

RPC

COLLABORATORS

	<i>TITLE :</i> RPC	
<i>ACTION</i>	<i>NAME</i>	<i>DATE</i>
WRITTEN BY		October 27, 2024
<i>SIGNATURE</i>		

REVISION HISTORY

NUMBER	DATE	DESCRIPTION	NAME

Contents

1	RPC	1
1.1	RPC.guide	1
1.2	RPC.guide/Disclaimer & Copyright	1
1.3	RPC.guide/AmiTCP-IP Notes	2
1.4	RPC.guide/About the Port to AmiTCP-IP and Amiga	2
1.5	RPC.guide/Nonsupported Features	3
1.6	RPC.guide/About the Floating Point Format	3
1.7	RPC.guide/What Was Done to the Original SUNRPC 4.0 While Porting	4
1.8	RPC.guide/rpcgen Notes	4
1.9	RPC.guide/portmap and rpcinfo	5
1.10	RPC.guide/Using the Library	5
1.11	RPC.guide/Supported Compilers	5
1.12	RPC.guide/Compiling the library	6
1.13	RPC.guide/Choosing the Library to Use	6
1.14	RPC.guide/rpcgen Programming Guide	7
1.15	RPC.guide/The rpcgen Protocol Compiler	7
1.16	RPC.guide/Converting Local Procedures into Remote Procedures	8
1.17	RPC.guide/Generating XDR Routines	13
1.18	RPC.guide/The C-Preprocessor	17
1.19	RPC.guide/rpcgen Programming Notes	18
1.20	RPC.guide/Timeout Changes	18
1.21	RPC.guide/Handling Broadcast on the Server Side	19
1.22	RPC.guide/Other Information Passed to Server Procedures	20
1.23	RPC.guide/RPC Language	20
1.24	RPC.guide/Definitions	21
1.25	RPC.guide/Structures	21
1.26	RPC.guide/Unions	22
1.27	RPC.guide/Enumerations	23
1.28	RPC.guide/Typedef	23
1.29	RPC.guide/Constants	23

1.30	RPC.guide/Programs	24
1.31	RPC.guide/Declarations	25
1.32	RPC.guide/Simple Declarations	25
1.33	RPC.guide/Fixed-Length Array Declarations	25
1.34	RPC.guide/Variable-Length Array Declarations	26
1.35	RPC.guide/Pointer Declarations	26
1.36	RPC.guide/Special Cases	27
1.37	RPC.guide/XDR Booleans	27
1.38	RPC.guide/XDR Strings	27
1.39	RPC.guide/XDR Opaque Data	28
1.40	RPC.guide/XDR Voids	28
1.41	RPC.guide/Remote Procedure Call Programming Guide	28
1.42	RPC.guide/Introduction to RPC Programming Guide	29
1.43	RPC.guide/Layers of RPC	29
1.44	RPC.guide/The Highest Layer	29
1.45	RPC.guide/The Middle Layer	30
1.46	RPC.guide/The Lowest Layer	30
1.47	RPC.guide/The RPC Paradigm	31
1.48	RPC.guide/Higher Layers of RPC	31
1.49	RPC.guide/Highest Layer	32
1.50	RPC.guide/Intermediate Layer	33
1.51	RPC.guide/Assigning Program Numbers	36
1.52	RPC.guide/Passing Arbitrary Data Types	37
1.53	RPC.guide/Lowest Layer of RPC	40
1.54	RPC.guide/More on the Server Side	41
1.55	RPC.guide/Memory Allocation with XDR	43
1.56	RPC.guide/The Calling Side	44
1.57	RPC.guide/Other RPC Features	46
1.58	RPC.guide/Select on the Server Side	47
1.59	RPC.guide/Broadcast RPC	48
1.60	RPC.guide/Broadcast RPC Synopsis	48
1.61	RPC.guide/RPC Batching	49
1.62	RPC.guide/RPC Authentication	52
1.63	RPC.guide/UNIX Authentication	53
1.64	RPC.guide/DES Authentication	56
1.65	RPC.guide/Using Inetd	58
1.66	RPC.guide/More Examples	59
1.67	RPC.guide/Versions	59
1.68	RPC.guide/TCP	61

1.69	RPC.guide/Callback Procedures	64
1.70	RPC.guide/XDR Technical Notes	67
1.71	RPC.guide/Overview	67
1.72	RPC.guide/Justification	68
1.73	RPC.guide/A Canonical Standard	71
1.74	RPC.guide/The XDR Library	72
1.75	RPC.guide/XDR Library Primitives	74
1.76	RPC.guide/Number Filters	74
1.77	RPC.guide/Floating Point Filters	75
1.78	RPC.guide/Enumeration Filters	76
1.79	RPC.guide/No Data	76
1.80	RPC.guide/Constructed Data Type Filters	77
1.81	RPC.guide/Strings	77
1.82	RPC.guide/Byte Arrays	78
1.83	RPC.guide/Arrays	79
1.84	RPC.guide/XDR Examples	79
1.85	RPC.guide/Opaque Data	81
1.86	RPC.guide/Fixed Sized Arrays	82
1.87	RPC.guide/Discriminated Unions	83
1.88	RPC.guide/Pointers	84
1.89	RPC.guide/Non-filter Primitives	86
1.90	RPC.guide/XDR Operation Directions	86
1.91	RPC.guide/XDR Stream Access	87
1.92	RPC.guide/Standard I-O Streams	87
1.93	RPC.guide/Memory Streams	88
1.94	RPC.guide/Record (TCP-IP) Streams	88
1.95	RPC.guide/XDR Stream Implementation	89
1.96	RPC.guide/The XDR Object	90
1.97	RPC.guide/Advanced Topics	91
1.98	RPC.guide/Linked Lists	91
1.99	RPC.guide/XDR Protocol Specification	94
1.100	RPC.guide/Introduction to XDR	95
1.101	RPC.guide/Basic Block Size	95
1.102	RPC.guide/XDR Data Types	96
1.103	RPC.guide/Integer	97
1.104	RPC.guide/Unsigned Integer	97
1.105	RPC.guide/Enumeration	98
1.106	RPC.guide/Boolean	98
1.107	RPC.guide/Hyper Integer and Unsigned Hyper Integer	98

1.108RPC.guide/Floating-point	99
1.109RPC.guide/Double-precision Floating-point	100
1.110RPC.guide/Fixed-length Opaque Data	100
1.111RPC.guide/Variable-length Opaque Data	101
1.112RPC.guide/String	102
1.113RPC.guide/Fixed-length Array	103
1.114RPC.guide/Variable-length Array	103
1.115RPC.guide/Structure	104
1.116RPC.guide/Discriminated Union	104
1.117RPC.guide/Void	105
1.118RPC.guide/Constant	105
1.119RPC.guide/Typedef (type definition)	106
1.120RPC.guide/Optional-data	107
1.121RPC.guide/Areas for Future Enhancement	108
1.122RPC.guide/Discussion	108
1.123RPC.guide/Why a Language for Describing Data?	108
1.124RPC.guide/Why Only one Byte-Order for an XDR Unit?	109
1.125RPC.guide/Why does XDR use Big-Endian Byte-Order?	109
1.126RPC.guide/Why is the XDR Unit Four Bytes Wide?	109
1.127RPC.guide/Why must Variable-Length Data be Padded with Zeros?	110
1.128RPC.guide/Why is there No Explicit Data-Typing?	110
1.129RPC.guide/The XDR Language Specification	110
1.130RPC.guide/Notational Conventions	110
1.131RPC.guide/Lexical Notes	111
1.132RPC.guide/Syntax Information	112
1.133RPC.guide/XDR Syntax Notes	113
1.134RPC.guide/An Example of an XDR Data Description	114
1.135RPC.guide/XDR References	115
1.136RPC.guide/Remote Procedure Calls Protocol Specification	115
1.137RPC.guide/Introduction to RPC	116
1.138RPC.guide/Terminology	116
1.139RPC.guide/The RPC Model	116
1.140RPC.guide/Transports and Semantics	117
1.141RPC.guide/Binding and Rendezvous Independence	118
1.142RPC.guide/Authentication	118
1.143RPC.guide/RPC Protocol Requirements	119
1.144RPC.guide/Programs and Procedures	119
1.145RPC.guide/Authentication Requirements	120
1.146RPC.guide/Program Number Assignment	121

1.147RPC.guide/Other Uses of the RPC Protocol	121
1.148RPC.guide/Batching	122
1.149RPC.guide/RPC Broadcast	122
1.150RPC.guide/The RPC Message Protocol	122
1.151RPC.guide/Authentication Protocols	125
1.152RPC.guide/Null Authentication (protocol)	126
1.153RPC.guide/UNIX Authentication (protocol)	126
1.154RPC.guide/DES Authentication (protocol)	127
1.155RPC.guide/Naming	127
1.156RPC.guide/DES Authentication Verifiers	128
1.157RPC.guide/Nicknames and Clock Synchronization	129
1.158RPC.guide/DES Authentication Protocol (in XDR language)	129
1.159RPC.guide/Diffie-Hellman Encryption	131
1.160RPC.guide/Record Marking Standard	132
1.161RPC.guide/The RPC Language	133
1.162RPC.guide/An Example Service Described in the RPC Language	133
1.163RPC.guide/The RPC Language Specification	134
1.164RPC.guide/RPC Syntax Notes	134
1.165RPC.guide/Port Mapper Program Protocol	135
1.166RPC.guide/Port Mapper Protocol Specification (in RPC Language)	135
1.167RPC.guide/Port Mapper Operation	137
1.168RPC.guide/RPC References	138
1.169RPC.guide/Network File System Version 2 Protocol Specification	139
1.170RPC.guide/Introduction to NFS	139
1.171RPC.guide/Remote Procedure Call	139
1.172RPC.guide/External Data Representation	140
1.173RPC.guide/Stateless Servers	140
1.174RPC.guide/NFS Protocol Definition	141
1.175RPC.guide/File System Model	141
1.176RPC.guide/RPC Information	142
1.177RPC.guide/Sizes of XDR Structures	142
1.178RPC.guide/Basic Data Types	143
1.179RPC.guide/stat	143
1.180RPC.guide/ftype	145
1.181RPC.guide/fhandle	145
1.182RPC.guide/timeval	145
1.183RPC.guide/fattr	146
1.184RPC.guide/sattr	147
1.185RPC.guide/filename	148

1.186RPC.guide/path	148
1.187RPC.guide/attrstat	148
1.188RPC.guide/diopargs	148
1.189RPC.guide/diopres	149
1.190RPC.guide/Server Procedures	149
1.191RPC.guide/Do Nothing	150
1.192RPC.guide/Get File Attributes	151
1.193RPC.guide/Set File Attributes	151
1.194RPC.guide/Get Filesystem Root	151
1.195RPC.guide/Look Up File Name	152
1.196RPC.guide/Read From Symbolic Link	152
1.197RPC.guide/Read From File	152
1.198RPC.guide/Write to Cache	153
1.199RPC.guide/Write to File	153
1.200RPC.guide/Create File	154
1.201RPC.guide/Remove File	154
1.202RPC.guide/Rename File	154
1.203RPC.guide/Create Link to File	155
1.204RPC.guide/Create Symbolic Link	155
1.205RPC.guide/Create Directory	156
1.206RPC.guide/Remove Directory	156
1.207RPC.guide/Read From Directory	157
1.208RPC.guide/Get Filesystem Attributes	157
1.209RPC.guide/NFS Implementation Issues	158
1.210RPC.guide/Server-Client Relationship	159
1.211RPC.guide/Pathname Interpretation	159
1.212RPC.guide/Permission Issues	160
1.213RPC.guide/Setting RPC Parameters	160
1.214RPC.guide/Mount Protocol Definition	161
1.215RPC.guide/Mount Protocol Introduction	161
1.216RPC.guide/Mount Protocol RPC Information	162
1.217RPC.guide/Mount Protocol Sizes of XDR Structures	162
1.218RPC.guide/Mount Protocol Basic Data Types	162
1.219RPC.guide/fhandle (Mount)	163
1.220RPC.guide/fhstatus	163
1.221RPC.guide/dirpath	163
1.222RPC.guide/name	163
1.223RPC.guide/Mount Protocol Server Procedures	164
1.224RPC.guide/Do Nothing (Mount)	164

1.225RPC.guide/Add Mount Entry	165
1.226RPC.guide/Return Mount Entries	165
1.227RPC.guide/Remove Mount Entry	165
1.228RPC.guide/Remove All Mount Entries	165
1.229RPC.guide/Return Export List	166
1.230RPC.guide/Manual Pages	166
1.231RPC.guide/portmap Manual Page	167
1.232RPC.guide/rpcgen Manual Page	167
1.233RPC.guide/rpcinfo Manual Page	172
1.234RPC.guide/rpc Data Base File Format	173
1.235RPC.guide/Miscellaneous Utility Functions	174
1.236RPC.guide/bindresvport()	174
1.237RPC.guide/getrpcbyname()	175
1.238RPC.guide/getrpcbynumber()	176
1.239RPC.guide/getrpcport()	177
1.240RPC.guide/RPC Function Reference	177
1.241RPC.guide/auth_destroy()	179
1.242RPC.guide/authnone_create()	179
1.243RPC.guide/authunix_create()	179
1.244RPC.guide/authunix_create_default()	180
1.245RPC.guide/callrpc()	180
1.246RPC.guide/clnt_broadcast()	180
1.247RPC.guide/clnt_call()	181
1.248RPC.guide/clnt_destroy()	181
1.249RPC.guide/clnt_create()	182
1.250RPC.guide/clnt_control()	182
1.251RPC.guide/clnt_freeres()	183
1.252RPC.guide/clnt_geterr()	183
1.253RPC.guide/clnt_pcreateerror()	183
1.254RPC.guide/clnt_perrno()	184
1.255RPC.guide/clnt_perror()	184
1.256RPC.guide/clnt_screateerror()	184
1.257RPC.guide/clnt_sperno()	185
1.258RPC.guide/clnt_sperro()	185
1.259RPC.guide/clntraw_create()	185
1.260RPC.guide/clnttcp_create()	186
1.261RPC.guide/clntudp_create()	186
1.262RPC.guide/clntudp_bufcreate()	187
1.263RPC.guide/get_myaddress()	187

1.264RPC.guide/pmap_getmaps()	187
1.265RPC.guide/pmap_getport()	188
1.266RPC.guide/pmap_rmtcall()	188
1.267RPC.guide/pmap_set()	189
1.268RPC.guide/pmap_unset()	189
1.269RPC.guide/registerrpc()	189
1.270RPC.guide/rpc_createerr	190
1.271RPC.guide/svc_destroy()	190
1.272RPC.guide/svc_fdset	190
1.273RPC.guide/svc_fds	191
1.274RPC.guide/svc_freeargs()	191
1.275RPC.guide/svc_getargs()	191
1.276RPC.guide/svc_getcaller()	192
1.277RPC.guide/svc_getreqset()	192
1.278RPC.guide/svc_getreq()	192
1.279RPC.guide/svc_register()	192
1.280RPC.guide/svc_run()	193
1.281RPC.guide/svc_sendreply()	193
1.282RPC.guide/svc_unregister()	194
1.283RPC.guide/svcerr_auth()	194
1.284RPC.guide/svcerr_decode()	194
1.285RPC.guide/svcerr_noproc()	194
1.286RPC.guide/svcerr_noprogram()	195
1.287RPC.guide/svcerr_progvers()	195
1.288RPC.guide/svcerr_systemerr()	195
1.289RPC.guide/svcerr_weakauth()	196
1.290RPC.guide/svcraw_create()	196
1.291RPC.guide/svctcp_create()	196
1.292RPC.guide/svcfd_create()	197
1.293RPC.guide/svcudp_bufcreate()	197
1.294RPC.guide/svcudp_create()	197
1.295RPC.guide/xdr_accepted_reply()	198
1.296RPC.guide/xdr_authunix_parms()	198
1.297RPC.guide/xdr_callhdr()	198
1.298RPC.guide/xdr_callmsg()	198
1.299RPC.guide/xdr_opaque_auth()	199
1.300RPC.guide/xdr_pmap()	199
1.301RPC.guide/xdr_pmaplist()	199
1.302RPC.guide/xdr_rejected_reply()	200

1.303RPC.guide/xdr_replymsg()	200
1.304RPC.guide/xprt_register()	200
1.305RPC.guide/xprt_unregister()	201
1.306RPC.guide/XDR Function Reference	201
1.307RPC.guide/xdr_array()	202
1.308RPC.guide/xdr_bool()	202
1.309RPC.guide/xdr_bytes()	202
1.310RPC.guide/xdr_char()	203
1.311RPC.guide/xdr_destroy()	203
1.312RPC.guide/xdr_double()	203
1.313RPC.guide/xdr_enum()	204
1.314RPC.guide/xdr_float()	204
1.315RPC.guide/xdr_free()	204
1.316RPC.guide/xdr_getpos()	205
1.317RPC.guide/xdr_inline()	205
1.318RPC.guide/xdr_int()	205
1.319RPC.guide/xdr_long()	206
1.320RPC.guide/xdrmem_create()	206
1.321RPC.guide/xdr_opaque()	206
1.322RPC.guide/xdr_pointer()	207
1.323RPC.guide/xdrrec_create()	207
1.324RPC.guide/xdrrec_endofrecord()	208
1.325RPC.guide/xdrrec_eof()	208
1.326RPC.guide/xdrrec_skiprecord()	208
1.327RPC.guide/xdr_reference()	208
1.328RPC.guide/xdr_setpos()	209
1.329RPC.guide/xdr_short()	209
1.330RPC.guide/xdrstdio_create()	210
1.331RPC.guide/xdr_string()	210
1.332RPC.guide/xdr_u_char()	210
1.333RPC.guide/xdr_u_int()	211
1.334RPC.guide/xdr_u_long()	211
1.335RPC.guide/xdr_u_short()	211
1.336RPC.guide/xdr_union()	211
1.337RPC.guide/xdr_vector()	212
1.338RPC.guide/xdr_void()	212
1.339RPC.guide/xdr_wrapstring()	213
1.340RPC.guide/Function Index	213
1.341RPC.guide/Concept Index	216

Chapter 1

RPC

1.1 RPC.guide

AmiTCP/IP Remote Procedure Calls Manual

This document contains the original Sun RPC documents in a new format. All examples are as in the original ones, especially, they are meant to be run under UNIX. As I assume the potential readers of this manual to be familiar with various UNIXes, I didn't put my effort in converting the examples. So, do not expect to be able to follow the examples literally in AmigaOS. All RPC code should work as described, however.

Disclaimer & Copyright
AmiTCP-IP Notes
rpcgen Programming Guide
Remote Procedure Call Programming Guide
XDR Technical Notes
XDR Protocol Specification
Remote Procedure Calls Protocol Specification
Network File System Version 2 Protocol Specification
Manual Pages
Function Index
Concept Index

1.2 RPC.guide/Disclaimer & Copyright

Disclaimer & Copyright

Sun RPC is a product of Sun Microsystems, Inc. and is provided for unrestricted use provided that this legend is included on all tape media and as a part of the software program in whole or part. Users may copy or modify Sun RPC without charge, but are not authorized to license or distribute it to anyone else except as part of a product or

program developed by the user.

SUN RPC IS PROVIDED AS IS WITH NO WARRANTIES OF ANY KIND INCLUDING THE WARRANTIES OF DESIGN, MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, OR ARISING FROM A COURSE OF DEALING, USAGE OR TRADE PRACTICE.

Sun RPC is provided with no support and without any obligation on the part of Sun Microsystems, Inc. to assist in its use, correction, modification or enhancement.

SUN MICROSYSTEMS, INC. SHALL HAVE NO LIABILITY WITH RESPECT TO THE INFRINGEMENT OF COPYRIGHTS, TRADE SECRETS OR ANY PATENTS BY SUN RPC OR ANY PART THEREOF.

In no event will Sun Microsystems, Inc. be liable for any lost revenue or profits or other special, indirect and consequential damages, even if Sun has been advised of the possibility of such damages.

Sun Microsystems, Inc.
2550 Garcia Avenue
Mountain View, California 94043

This version of the manual is converted from Sun provided roff-files by Jarno Rajahalme 1993. All comments about this version should be sent to:

<AmiTCP-Group@hut.fi>

1.3 RPC.guide/AmiTCP-IP Notes

AmiTCP/IP RPC Notes

This chapter describes the AmiTCP/IP version of the Sun RPC 4.0.

Note: Read this chapter carefully and then do the "Hello World!" test with the provided demonstration protocols in 'demo' directory. Once you get these compiled and running you know that your setup is complete.

About the Port to AmiTCP-IP and Amiga
rpcgen Notes
portmap and rpcinfo
Using the Library

1.4 RPC.guide/About the Port to AmiTCP-IP and Amiga

About the Port to AmiTCP/IP and Amiga
=====

Nonsupported Features

About the Floating Point Format

What Was Done to the Original SUNRPC 4.0 While Porting

1.5 RPC.guide/Nonsupported Features

Nonsupported Features

Following features are not supported/modified in this version:

Secure RPC

This is not yet supported. I might decide to support it, if it is requested. This means that none of the DES authentication functions are included in the library.

Broadcast RPC

The current support for Broadcast RPC is limited in sense that the portmap blocks while serving the broadcasted RPC. This means that the 'portmap' services are not available during that time, which may prevent some RPC servers from functioning correctly.

'getrpcnt()'

Didn't include this, since AmiTCP/IP doesn't provide the '*ent()' functions in general. I might re-enable this if there is any demand for it.

'svcsfd_create()'

This is not included, since the library uses AmiTCP/IP socket descriptors. This might be added if there ever becomes a good, portable solution to the file vs. socket descriptor problem.

1.6 RPC.guide/About the Floating Point Format

About the Floating Point Format

The library supports only IEEE format floating point values ('xdr_float()' and 'xdr_double()'). The FFP format *must not* be used. Consult your compiler documentation about choosing the floating point format (SAS/C uses IEEE by default).

1.7 RPC.guide/What Was Done to the Original SUNRPC 4.0 While Porting

What Was Done to the Original SUNRPC 4.0 While Porting

1. Added ANSI C prototypes to all function and function pointer prototypes
2. ANSIfied the includes if possible ('sys/errno.h' -> 'errno.h', for example).
3. Removed dependency on select() retaining the value of the timeout 'timeval'.
4. Changed 'read()' and 'write()' calls to 'recv()' and 'send()', respectively. This prevents the usage of the real file descriptors (see the 'svcf_create()' in Nonsupported Features).
5. Added 'sys/param.h' as the first include file for every '.c' file for consistency.
6. Added 'XDRFUN' definition to 'rpc/types.h'. It is used in definition of every xdr-function.
7. Changed 'close()'s and 'ioctl()'s on sockets to 'CloseSocket()' and 'IoctlSocket()' if 'AMITCP' is defined.
8. Changed 'getpid()'s to 'FindTask(NULL)' if 'AMIGA' is defined.
9. Changed socket calls no to be retried after the 'EINTR' error return.

1.8 RPC.guide/rpcgen Notes

'rpcgen' Notes

=====

The 'rpcgen' included in this distribution is ported from the TIRPC(1) 'rpcgen' from Linux by Carsten Heyl <ch@irb.informatik.uni-dortmund.de>. This version of 'rpcgen' can produce K&R/ANSI-C/C++ compatible output. It also can generate 'inetd' compatible daemons, but this feature is not yet supported. Refer to the manual page of the 'rpcgen' for its usage.

The 'rpcgen' needs a C preprocessor to handle the macros and comments. One suitable C preprocessor ('cpp') is included in the distribution. It is based on an old DECUS cpp port by Fred Fish. Carsten added a '-P' option to suppress '#line' directives, which the 'cpp' left on the output.

See rpcgen Manual Page.

----- Footnotes -----

- (1) Transport Independent RPC

1.9 RPC.guide/portmap and rpcinfo

`'portmap'` and `'rpcinfo'`

=====

`'portmap'` is the RPC program to internet port mapper, see Port Mapper Program Protocol. The `'portmap'` must be started before any RPC services, and it is best to include line "`'run portmap'`" to your `'AmiTCP:bin/startnet'` script. See portmap Manual Page.

The `'rpcinfo'` can be used to view `'portmap'` bindings on any host. See the manual page for the usage. See rpcinfo Manual Page.

1.10 RPC.guide/Using the Library

Using the Library

=====

Supported Compilers
Compiling the library
Choosing the Library to Use

1.11 RPC.guide/Supported Compilers

Supported Compilers

The distribution contains precompiled library for the SAS/C 6.x, since it's the compiler I use. The library is highly portable though, and you should have no problems compiling it with DICE, GCC etc.

The only compiler dependent part I made there is the definition `'XDRFUN'` (in `'rpc/types.h'`), which is defined as `'__regargs'` for SAS/C and as nothing for others.

The makefiles are written for the SAS `'smake'`, but since it doesn't have any fancy features, they should be quite portable as well.

1.12 RPC.guide/Compiling the library

Compiling the library

The library code assumes that the 'bsdsocket.library' is opened by the caller, and the library base is in global variable named 'SocketBase'. Also the global 'errno' variable must be set with the AmiTCP/IP 'SocketBaseTagList()' function (1).

The 'Smakefile' of the library can compile eight different versions of the library, but only the 'rpcr.lib' version is compiled by default, see Choosing the Library to Use. This library is installed to the 'AmiTCP:netlib' directory by default.

----- Footnotes -----

(1) These are standard AmiTCP/IP practices and are provided by the 'net.lib' autoinitialization routine, so normally these are no concern.

1.13 RPC.guide/Choosing the Library to Use

Choosing the Library to Use

Eight different versions of the library are provided. The library is compiled either with or without the debugging information, with near or far data, and with register or stack arguments. Use the following table to select the right library:

Arguments		Data		Debugging		Library name
Stack	Reg.	Near	Far	No	Yes	
x		x		x		rpc.lib
	x	x		x		rpcr.lib
x			x	x		rpcnb.lib
	x		x	x		rpcnbr.lib
x		x			x	grpc.lib
	x	x			x	grpcr.lib
x			x		x	grpcnb.lib
	x		x		x	grpcnbr.lib

The use of registerized arguments and near data is highly recommended, since these options provide smaller and faster executables.

To compile the desired version of the library, issue 'smake <Library name>' or 'smake all' to compile all the versions. The makefile can only install the default version to the 'netlib:', you must move the other libraries yourself.

The reason why the library is not compiled with SAS/C 'ARGS=BOTH' is that there are some problems SAS/C not enforcing the function pointers prototyped as taking stack arguments actually being assigned with the

proper (stack) entry point of a function, if there are both stack and register versions of the function available (as is the case with the 'ARGS=BOTH' libraries).

Second reason is that if you want to use stack arguments, an 'ARGS=STACK' library is faster than the 'BOTH' equivalent.

The debugging info is compiled with SAS/C 'FullFlush' option.

1.14 RPC.guide/rpcgen Programming Guide

'rpcgen' Programming Guide

This chapter documents the Sun RPC Compiler, the 'rpcgen'.

The rpcgen Protocol Compiler
 Converting Local Procedures into Remote Procedures
 Generating XDR Routines
 The C-Preprocessor
 rpcgen Programming Notes
 RPC Language

1.15 RPC.guide/The rpcgen Protocol Compiler

The 'rpcgen' Protocol Compiler
=====

The details of programming applications to use Remote Procedure Calls can be overwhelming. Perhaps most daunting is the writing of the XDR routines necessary to convert procedure arguments and results into their network format and vice-versa.

Fortunately, 'rpcgen' exists to help programmers write RPC applications simply and directly. 'rpcgen' does most of the dirty work, allowing programmers to debug the main features of their application, instead of requiring them to spend most of their time debugging their network interface code.

'rpcgen' is a compiler. It accepts a remote program interface definition written in a language, called RPC Language, which is similar to C. It produces a C language output which includes stub versions of the client routines, a server skeleton, XDR filter routines for both parameters and results, and a header file that contains common definitions. The client stubs interface with the RPC library and effectively hide the network from their callers. The server stub similarly hides the network from the server procedures that are to be invoked by remote clients. 'rpcgen''s output files can be compiled and linked in the usual way. The developer writes server procedures -- in

any language that observes Sun calling conventions -- and links them with the server skeleton produced by 'rpcgen' to get an executable server program. To use a remote program, a programmer writes an ordinary main program that makes local procedure calls to the client stubs produced by 'rpcgen'. Linking this program with 'rpcgen''s stubs creates an executable program. (At present the main program must be written in C). 'rpcgen' options can be used to suppress stub generation and to specify the transport to be used by the server stub.

Like all compilers, 'rpcgen' reduces development time that would otherwise be spent coding and debugging low-level routines. All compilers, including 'rpcgen', do this at a small cost in efficiency and flexibility. However, many compilers allow escape hatches for programmers to mix low-level code with high-level code. 'rpcgen' is no exception. In speed-critical applications, hand-written routines can be linked with the 'rpcgen' output without any difficulty. Also, one may proceed by using 'rpcgen' output as a starting point, and then rewriting it as necessary. (If you need a discussion of RPC programming without 'rpcgen', see Remote Procedure Call Programming Guide).

1.16 RPC.guide/Converting Local Procedures into Remote Procedures

Converting Local Procedures into Remote Procedures

=====

Assume an application that runs on a single machine, one which we want to convert to run over the network. Here we will demonstrate such a conversion by way of a simple example -- a program that prints a message to the console:

```
/*
 * printmsg.c: print a message on the console
 */
#include <stdio.h>

main(argc, argv)
    int argc;
    char *argv[];
{
    char *message;

    if (argc < 2) {
        fprintf(stderr, "usage: %s <message>\n", argv[0]);
        exit(1);
    }
    message = argv[1];

    if (!printmessage(message)) {
        fprintf(stderr, "%s: couldn't print your message\n",
            argv[0]);
        exit(1);
    }
    printf("Message Delivered!\n");
}
```

```

        exit(0);
    }
    /*
    * Print a message to the console.
    * Return a boolean indicating whether the message was actually printed.
    */
    printmessage(msg)
        char *msg;
    {
        FILE *f;

        f = fopen("/dev/console", "w");
        if (f == NULL) {
            return (0);
        }
        fprintf(f, "%s\n", msg);
        fclose(f);
        return(1);
    }

```

And then, of course:

```

example% cc printmsg.c -o printmsg
example% printmsg "Hello, there."
Message delivered!
example%

```

If `'printmessage()'` was turned into a remote procedure, then it could be called from anywhere in the network. Ideally, one would just like to stick a keyword like `'remote'` in front of a procedure to turn it into a remote procedure. Unfortunately, we have to live within the constraints of the C language, since it existed long before RPC did. But even without language support, it's not very difficult to make a procedure remote.

In general, it's necessary to figure out what the types are for all procedure inputs and outputs. In this case, we have a procedure `'printmessage()'` which takes a string as input, and returns an integer as output. Knowing this, we can write a protocol specification in RPC language that describes the remote version of `'printmessage()'`. Here it is:

```

/*
 * msg.x: Remote message printing protocol
 */
program MESSAGEPROG {
    version MESSAGEVERS {
        int PRINTMESSAGE(string) = 1;
    } = 1;
} = 99;

```

Remote procedures are part of remote programs, so we actually declared an entire remote program here which contains the single procedure `'PRINTMESSAGE'`. This procedure was declared to be in version 1 of the remote program. No null procedure (procedure 0) is necessary because `'rpcgen'` generates it automatically.

Notice that everything is declared with all capital letters. This is not required, but is a good convention to follow.

Notice also that the argument type is `'string'` and not `'char *'`. This is because a `'char *'` in C is ambiguous. Programmers usually intend it to mean a null-terminated string of characters, but it could also represent a pointer to a single character or a pointer to an array of characters. In RPC language, a null-terminated string is unambiguously called a `'string'`.

There are just two more things to write. First, there is the remote procedure itself. Here's the definition of a remote procedure to implement the `'PRINTMESSAGE'` procedure we declared above:

```

/*
 * msg_proc.c: implementation of the remote procedure "printmessage"
 */
#include <stdio.h>
#include <rpc/rpc.h>    /* always needed */
#include "msg.h"        /* need this too: msg.h will be generated by rpcgen ←
 */

/*
 * Remote version of "printmessage"
 */
int *
printmessage_1(msg)
    char **msg;
{
    static int result; /* must be static! */
    FILE *f;

    f = fopen("/dev/console", "w");
    if (f == NULL) {
        result = 0;
        return (&result);
    }
    fprintf(f, "%s\n", *msg);
    fclose(f);
    result = 1;
    return (&result);
}

```

Notice here that the declaration of the remote procedure `'printmessage_1()'` differs from that of the local procedure `'printmessage()'` in three ways:

1. It takes a pointer to a string instead of a string itself. This is true of all remote procedures: they always take pointers to their arguments rather than the arguments themselves.
2. It returns a pointer to an integer instead of an integer itself. This is also generally true of remote procedures: they always return a pointer to their results.
3. It has an `'_1'` appended to its name. In general, all remote procedures called by `'rpcgen'` are named by the following rule: the

name in the program definition (here 'PRINTMESSAGE') is converted to all lower-case letters, an underbar ('_') is appended to it, and finally the version number (here '1') is appended.

The last thing to do is declare the main client program that will call the remote procedure. Here it is:

```

/*
 * rprintmsg.c: remote version of "printmsg.c"
 */
#include <stdio.h>
#include <rpc/rpc.h>      /* always needed */
#include "msg.h"         /* need this too: msg.h will be generated by rpcgen ←
                        */

main(argc, argv)
    int argc;
    char *argv[];
{
    CLIENT *cl;
    int *result;
    char *server;
    char *message;

    if (argc < 3) {
        fprintf(stderr, "usage: %s host message\n", argv[0]);
        exit(1);
    }

    /*
     * Save values of command line arguments
     */
    server = argv[1];
    message = argv[2];

    /*
     * Create client "handle" used for calling MESSAGEPROG on the
     * server designated on the command line. We tell the RPC package
     * to use the "tcp" protocol when contacting the server.
     */
    cl = clnt_create(server, MESSAGEPROG, MESSAGEEVERS, "tcp");
    if (cl == NULL) {
        /*
         * Couldn't establish connection with server.
         * Print error message and die.
         */
        clnt_pcreateerror(server);
        exit(1);
    }

    /*
     * Call the remote procedure "printmessage" on the server
     */
    result = printmessage_1(&message, cl);
    if (result == NULL) {
        /*
         * An error occurred while calling the server.

```

```

        * Print error message and die.
        */
        clnt_perror(cl, server);
        exit(1);
    }

    /*
     * Okay, we successfully called the remote procedure.
     */
    if (*result == 0) {
        /*
         * Server was unable to print our message.
         * Print error message and die.
         */
        fprintf(stderr, "%s: %s couldn't print your message\n",
                argv[0], server);
        exit(1);
    }

    /*
     * The message got printed on the server's console
     */
    printf("Message delivered to %s!\n", server);
}

```

There are two things to note here:

1. First a "client handle" is created using the RPC library routine 'clnt_create()'. This client handle will be passed to the stub routines which call the remote procedure.
2. The remote procedure 'printmessage_1()' is called exactly the same way as it is declared in 'msg_proc.c' except for the inserted client handle as the second argument.

Here's how to put all of the pieces together:

```

example% rpcgen msg.x
example% cc rprintmsg.c msg_clnt.c -o rprintmsg
example% cc msg_proc.c msg_svc.c -o msg_server

```

Two programs were compiled here: the client program 'rprintmsg' and the server program 'msg_server'. Before doing this though, 'rpcgen' was used to fill in the missing pieces.

Here is what 'rpcgen' did with the input file 'msg.x':

1. It created a header file called 'msg.h' that contained '#define's for 'MESSAGEPROG', 'MESSAGEVERS' and 'PRINTMESSAGE' for use in the other modules.
2. It created client 'stub' routines in the 'msg_clnt.c' file. In this case there is only one, the 'printmessage_1()' that was referred to from the 'rprintmsg' client program. The name of the output file for client stub routines is always formed in this way: if the name of the input file is 'FOO.x', the client stubs output file is called 'FOO_clnt.c'.

- It created the server program which calls `'printmessage_1()'` in `'msg_proc.c'`. This server program is named `'msg_svc.c'`. The rule for naming the server output file is similar to the previous one: for an input file called `'FOO.x'`, the output server file is named `'FOO_svc.c'`.

Now we're ready to have some fun. First, copy the server to a remote machine and run it. For this example, the machine is called `'moon'`. Server processes are run in the background, because they never exit.

```
moon% msg_server &
```

Then on our local machine (`'sun'`) we can print a message on `'moon'`'s console.

```
sun% printmsg moon "Hello, moon."
```

The message will get printed to `'moon'`'s console. You can print a message on anybody's console (including your own) with this program if you are able to copy the server to their machine and run it.

1.17 RPC.guide/Generating XDR Routines

Generating XDR Routines

=====

The previous example only demonstrated the automatic generation of client and server RPC code. `'rpcgen'` may also be used to generate XDR routines, that is, the routines necessary to convert local data structures into network format and vice-versa. This example presents a complete RPC service -- a remote directory listing service, which uses `'rpcgen'` not only to generate stub routines, but also to generate the XDR routines. Here is the protocol description file:

```
/*
 * dir.x: Remote directory listing protocol
 */
const MAXNAMELEN = 255;          /* maximum length of a directory entry */

typedef string nametype<MAXNAMELEN>; /* a directory entry */

typedef struct namenode *namelist; /* a link in the listing */

/*
 * A node in the directory listing
 */
struct namenode {
    nametype name;          /* name of directory entry */
    namelist next;         /* next entry */
};

/*
 * The result of a READDIR operation.
```

```

    */
union readdir_res switch (int errno) {
case 0:
    namelist list; /* no error: return directory listing */
default:
    void;          /* error occurred: nothing else to return */
};

/*
 * The directory program definition
 */
program DIRPROG {
    version DIRVERS {
        readdir_res
        READDIR(nametype) = 1;
    } = 1;
} = 76;

```

Note: *Types (like 'readdir_res' in the example above) can be defined using the 'struct', 'union' and 'enum' keywords, but those keywords should not be used in subsequent declarations of variables of those types. For example, if you define a union 'foo', you should declare using only 'foo' and not 'union foo'. In fact, 'rpcgen' compiles RPC unions into C structures and it is an error to declare them using the 'union' keyword.*

Running 'rpcgen' on 'dir.x' creates four output files. Three are the same as before: header file, client stub routines and server skeleton. The fourth are the XDR routines necessary for converting the data types we declared into XDR format and vice-versa. These are output in the file 'dir_xdr.c'.

Here is the implementation of the 'READDIR' procedure.

```

/*
 * dir_proc.c: remote readdir implementation
 */
#include <rpc/rpc.h>
#include <sys/dir.h>
#include "dir.h"

extern int errno;
extern char *malloc();
extern char *strdup();

readdir_res *
readdir_1(dirname)
    nametype *dirname;
{
    DIR *dirp;
    struct direct *d;
    namelist nl;
    namelist *nlp;
    static readdir_res res; /* must be static! */

    /*
     * Open directory

```

```

    */
    dirp = opendir(*dirname);
    if (dirp == NULL) {
        res.errno = errno;
        return (&res);
    }

    /*
     * Free previous result
     */
    xdr_free(xdr_readdir_res, &res);

    /*
     * Collect directory entries.
     * Memory allocated here will be freed by 'xdr_free()'
     * next time 'readdir_1()' is called
     */
    nlp = &res.readdir_res_u.list;
    while (d = readdir(dirp)) {
        nl = *nlp = (namenode *) malloc(sizeof(namenode));
        nl->name = strdup(d->d_name);
        nl->next = &nl->next;
    }
    *nlp = NULL;

    /*
     * Return the result
     */
    res.errno = 0;
    closedir(dirp);
    return (&res);
}

```

Finally, there is the client side program to call the server:

```

/*
 * rls.c: Remote directory listing client
 */
#include <stdio.h>
#include <rpc/rpc.h> /* always need this */
#include "dir.h" /* will be generated by rpcgen */

extern int errno;

main(argc, argv)
    int argc;
    char *argv[];
{
    CLIENT *cl;
    char *server;
    char *dir;
    readdir_res *result;
    namelist nl;

    if (argc != 3) {
        fprintf(stderr, "usage: %s host directory\n",

```

```
        argv[0]);
        exit(1);
    }

    /*
     * Remember what our command line arguments refer to
     */
    server = argv[1];
    dir = argv[2];

    /*
     * Create client "handle" used for calling MESSAGEPROG on the
     * server designated on the command line. We tell the RPC package
     * to use the "tcp" protocol when contacting the server.
     */
    cl = clnt_create(server, DIRPROG, DIRVERS, "tcp");
    if (cl == NULL) {
        /*
         * Couldn't establish connection with server.
         * Print error message and die.
         */
        clnt_pcreateerror(server);
        exit(1);
    }

    /*
     * Call the remote procedure readdir on the server
     */
    result = readdir_l(&dir, cl);
    if (result == NULL) {
        /*
         * An error occurred while calling the server.
         * Print error message and die.
         */
        clnt_perror(cl, server);
        exit(1);
    }

    /*
     * Okay, we successfully called the remote procedure.
     */
    if (result->errno != 0) {
        /*
         * A remote system error occurred.
         * Print error message and die.
         */
        errno = result->errno;
        perror(dir);
        exit(1);
    }

    /*
     * Successfully got a directory listing.
     * Print it out.
     */
    for (nl = result->readdir_res_u.list; nl != NULL;
         nl = nl->next) {
```

```

        printf("%s\n", nl->name);
    }
    exit(0);
}

```

Compile everything, and run:

```

sun% rpcgen dir.x
sun% cc rls.c dir_clnt.c dir_xdr.c -o rls
sun% cc dir_svc.c dir_proc.c dir_xdr.c -o dir_svc

sun% dir_svc &
moon% rls sun /usr/pub
.
..
ascii
eqnchar
greek
kbd
marg8
tabclr
tabs
tabs4
moon%

```

A final note about 'rpcgen': The client program and the server procedure can be tested together as a single program by simply linking them with each other rather than with the client and server stubs. The procedure calls will be executed as ordinary local procedure calls and the program can be debugged with a local debugger such as 'dbx'. When the program is working, the client program can be linked to the client stub produced by 'rpcgen' and the server procedures can be linked to the server stub produced by 'rpcgen'.

NOTE: *If you do this, you may want to comment out calls to RPC library routines, and have client-side routines call server routines directly.*

1.18 RPC.guide/The C-Preprocessor

The C-Preprocessor

=====

The C-preprocessor is run on all input files before they are compiled, so all the preprocessor directives are legal within a '.x' file. Four symbols may be defined, depending upon which output file is getting generated. The symbols are:

'Symbol'

Usage

'RPC_HDR'

for header-file output

'RPC_XDR'

```

    for XDR routine output

`RPC_SVC'
    for server-skeleton output

`RPC_CLNT'
    for client stub output

```

Also, ``rpcgen'` does a little preprocessing of its own. Any line that begins with a percent sign is passed directly into the output file, without any interpretation of the line. Here is a simple example that demonstrates the preprocessing features.

```

/*
 * time.x: Remote time protocol
 */
program TIMEPROG {
    version TIMEEVERS {
        unsigned int TIMEGET(void) = 1;
    } = 1;
} = 44;

#ifdef RPC_SVC
%int *
%timeget_1()
%{
%    static int thetime;
%
%    thetime = time(0);
%    return (&thetime);
%}
#endif

```

The ``%'` feature is not generally recommended, as there is no guarantee that the compiler will stick the output where you intended.

1.19 RPC.guide/rpcgen Programming Notes

``rpcgen'` Programming Notes

=====

Timeout Changes
 Handling Broadcast on the Server Side
 Other Information Passed to Server Procedures

1.20 RPC.guide/Timeout Changes

Timeout Changes

RPC sets a default timeout of 25 seconds for RPC calls when `'clnt_create()'` is used. This timeout may be changed using `'clnt_control()'`. Here is a small code fragment to demonstrate use of `'clnt_control()'`:

```

struct timeval tv;
CLIENT *cl;

cl = clnt_create("somehost", SOMEPROG, SOMEVERS, "tcp");
if (cl == NULL) {
    exit(1);
}
tv.tv_sec = 60; /* change timeout to 1 minute */
tv.tv_usec = 0;
clnt_control(cl, CLSET_TIMEOUT, &tv);

```

1.21 RPC.guide/Handling Broadcast on the Server Side

Handling Broadcast on the Server Side

When a procedure is known to be called via broadcast RPC, it is usually wise for the server to not reply unless it can provide some useful information to the client. This prevents the network from getting flooded by useless replies.

To prevent the server from replying, a remote procedure can return `'NULL'` as its result, and the server code generated by `'rpcgen'` will detect this and not send out a reply.

Here is an example of a procedure that replies only if it thinks it is an NFS server:

```

void *
reply_if_nfsserver()
{
    char notnull; /* just here so we can use its address */

    if (access("/etc/exports", F_OK) < 0) {
        return (NULL); /* prevent RPC from replying */
    }
    /*
     * return non-null pointer so RPC will send out a reply
     */
    return ((void *)&notnull);
}

```

Note that if procedure returns type `'void *'`, they must return a non-`'NULL'` pointer if they want RPC to reply for them.

1.22 RPC.guide/Other Information Passed to Server Procedures

Other Information Passed to Server Procedures

Server procedures will often want to know more about an RPC call than just its arguments. For example, getting authentication information is important to procedures that want to implement some level of security. This extra information is actually supplied to the server procedure as a second argument. Here is an example to demonstrate its use. What we've done here is rewrite the previous 'printmessage_1()' procedure to only allow root users to print a message to the console.

```
int *
printmessage_1(msg, rq)
    char **msg;
    struct svc_req *rq;
{
    static in result;          /* Must be static */
    FILE *f;
    struct authunix_parms *aup;

    aup = (struct authunix_parms *)rq->rq_clntcred;
    if (aup->aup_uid != 0) {
        result = 0;
        return (&result);
    }

    /*
     * Same code as before.
     */
}
```

1.23 RPC.guide/RPC Language

RPC Language

=====

RPC language is an extension of XDR language. The sole extension is the addition of the 'program' type. For a complete description of the XDR language syntax, See XDR Protocol Specification. For a description of the RPC extensions to the XDR language, See Remote Procedure Calls Protocol Specification.

However, XDR language is so close to C that if you know C, you know most of it already. We describe here the syntax of the RPC language, showing a few examples along the way. We also show how the various RPC and XDR type definitions get compiled into C type definitions in the output header file.

Definitions
Structures

Unions
Enumerations
Typedef
Constants
Programs
Declarations
Special Cases

1.24 RPC.guide/Definitions

Definitions

An RPC language file consists of a series of definitions.

```
definition-list:
    definition ";"
    definition ";" definition-list
```

It recognizes five types of definitions.

```
definition:
    enum-definition
    struct-definition
    union-definition
    typedef-definition
    const-definition
    program-definition
```

1.25 RPC.guide/Structures

Structures

An XDR struct is declared almost exactly like its C counterpart. It looks like the following:

```
struct-definition:
    "struct" struct-ident "{"
        declaration-list
    "}"

declaration-list:
    declaration ";"
    declaration ";" declaration-list
```

As an example, here is an XDR structure to a two-dimensional coordinate, and the C structure that it gets compiled into in the output header file.

```

struct coord {
    int x;
    int y;
};
-->
struct coord {
    int x;
    int y;
};
typedef struct coord coord;

```

The output is identical to the input, except for the added `'typedef'` at the end of the output. This allows one to use `'coord'` instead of `'struct coord'` when declaring items.

1.26 RPC.guide/Unions

Unions

XDR unions are discriminated unions, and look quite different from C unions. They are more analogous to Pascal variant records than they are to C unions.

```

union-definition:
    "union" union-ident "switch" "(" declaration ")" "{"
        case-list
    "}"

case-list:
    "case" value ":" declaration ";"
    "default" ":" declaration ";"
    "case" value ":" declaration ";" case-list

```

Here is an example of a type that might be returned as the result of a "read data" operation. If there is no error, return a block of data. Otherwise, don't return anything.

```

union read_result switch (int errno) {
case 0:
    opaque data[1024];
default:
    void;
};

```

It gets compiled into the following:

```

struct read_result {
    int errno;
    union {
        char data[1024];
    } read_result_u;
};
typedef struct read_result read_result;

```

Notice that the union component of the output struct has the name as the type name, except for the trailing `'_u'`.

1.27 RPC.guide/Enumerations

Enumerations

XDR enumerations have the same syntax as C enumerations.

```
enum-definition:
    "enum" enum-ident "{"
        enum-value-list
    "}"

enum-value-list:
    enum-value
    enum-value "," enum-value-list

enum-value:
    enum-value-ident
    enum-value-ident "=" value
```

Here is a short example of an XDR enum, and the C enum that it gets compiled into.

```
enum colortype {          enum colortype {
    RED = 0,              RED = 0,
    GREEN = 1,  -->     GREEN = 1,
    BLUE = 2             BLUE = 2,
};                       };
                        typedef enum colortype colortype;
```

1.28 RPC.guide/Typedef

Typedef

XDR typedefs have the same syntax as C typedefs.

```
typedef-definition:
    "typedef" declaration
```

Here is an example that defines a 'fname_type' used for declaring file name strings that have a maximum length of 255 characters.

```
typedef string fname_type<255>; --> typedef char *fname_type;
```

1.29 RPC.guide/Constants

Constants

XDR constants symbolic constants that may be used wherever a integer constant is used, for example, in array size specifications.

```
const-definition:
    "const" const-ident "=" integer
```

For example, the following defines a constant 'DOZEN' equal to 12.

```
const DOZEN = 12; --> #define DOZEN 12
```

1.30 RPC.guide/Programs

Programs

RPC programs are declared using the following syntax:

```
program-definition:
    "program" program-ident "{"
        version-list
    "}" "=" value

version-list:
    version ";"
    version ";" version-list

version:
    "version" version-ident "{"
        procedure-list
    "}" "=" value

procedure-list:
    procedure ";"
    procedure ";" procedure-list

procedure:
    type-ident procedure-ident "(" type-ident ")" "=" value
```

For example, here is the time protocol, revisited:

```
/*
 * time.x: Get or set the time. Time is represented as number of seconds
 * since 0:00, January 1, 1970.
 */
program TIMEPROG {
    version TIMEVERS {
        unsigned int TIMEGET(void) = 1;
        void TIMESET(unsigned) = 2;
    } = 1;
} = 44;
```

This file compiles into '#defines' in the output header file:

```
#define TIMEPROG 44
```

```
#define TIMEVERS 1
#define TIMEGET 1
#define TIMESET 2
```

1.31 RPC.guide/Declarations

Declarations

In XDR, there are only four kinds of declarations.

```
declaration:
    simple-declaration
    fixed-array-declaration
    variable-array-declaration
    pointer-declaration
```

Simple Declarations
Fixed-Length Array Declarations
Variable-Length Array Declarations
Pointer Declarations

1.32 RPC.guide/Simple Declarations

Simple Declarations
.....

Simple declarations are just like simple C declarations.

```
simple-declaration:
    type-ident variable-ident
```

Example:

```
colortype color;    --> colortype color;
```

1.33 RPC.guide/Fixed-Length Array Declarations

Fixed-Length Array Declarations
.....

Fixed-length array declarations are just like C array declarations:

```
fixed-array-declaration:
    type-ident variable-ident "[" value "]"
```

Example:

```
colortype palette[8];    --> colortype palette[8];
```

1.34 RPC.guide/Variable-Length Array Declarations

Variable-Length Array Declarations

.....

Variable-length array declarations have no explicit syntax in C, so XDR invents its own using angle-brackets.

```
variable-array-declaration:
    type-ident variable-ident "<" value ">"
    type-ident variable-ident "<" ">"
```

The maximum size is specified between the angle brackets. The size may be omitted, indicating that the array may be of any size.

```
int heights<12>;    /* at most 12 items */
int widths<>;      /* any number of items */
```

Since variable-length arrays have no explicit syntax in C, these declarations are actually compiled into 'struct's. For example, the 'heights' declaration gets compiled into the following struct:

```
struct {
    u_int heights_len; /* # of items in array */
    int *heights_val; /* pointer to array */
} heights;
```

Note that the number of items in the array is stored in the '_len' component and the pointer to the array is stored in the '_val' component. The first part of each of these component's names is the same as the name of the declared XDR variable.

1.35 RPC.guide/Pointer Declarations

Pointer Declarations

.....

Pointer declarations are made in XDR exactly as they are in C. You can't really send pointers over the network, but you can use XDR pointers for sending recursive data types such as lists and trees. The type is actually called 'optional-data', not 'pointer', in XDR language.

```
pointer-declaration:
    type-ident "*" variable-ident
```

Example:

```
listitem *next; --> listitem *next;
```

1.36 RPC.guide/Special Cases

Special Cases

There are a few exceptions to the rules described above.

- XDR Booleans
- XDR Strings
- XDR Opaque Data
- XDR Voids

1.37 RPC.guide/XDR Booleans

Booleans
.....

C has no built-in boolean type. However, the RPC library does a boolean type called 'bool_t' that is either 'TRUE' or 'FALSE'. Things declared as type 'bool' in XDR language are compiled into 'bool_t' in the output header file.

Example:

```
bool married; --> bool_t married;
```

1.38 RPC.guide/XDR Strings

Strings
.....

C has no built-in string type, but instead uses the null-terminated 'char *' convention. In XDR language, strings are declared using the 'string' keyword, and compiled into 'char *'s in the output header file. The maximum size contained in the angle brackets specifies the maximum number of characters allowed in the strings (not counting the 'NULL' character). The maximum size may be left off, indicating a string of arbitrary length.

Examples:

```
string name<32>; --> char *name;  
string longname<>; --> char *longname;
```

1.39 RPC.guide/XDR Opaque Data

Opaque Data

Opaque data is used in RPC and XDR to describe untyped data, that is, just sequences of arbitrary bytes. It may be declared either as a fixed or variable length array.

Examples:

```
opaque diskblock[512]; --> char diskblock[512];

opaque filedata<1024>; --> struct {
    u_int filedata_len;
    char *filedata_val;
} filedata;
```

1.40 RPC.guide/XDR Voids

Voids

In a void declaration, the variable is not named. The declaration is just 'void' and nothing else. Void declarations can only occur in two places: union definitions and program definitions (as the argument or result of a remote procedure).

1.41 RPC.guide/Remote Procedure Call Programming Guide

Remote Procedure Call Programming Guide

Introduction to RPC Programming Guide
 Layers of RPC
 Higher Layers of RPC
 Lowest Layer of RPC
 Other RPC Features
 More Examples

1.42 RPC.guide/Introduction to RPC Programming Guide

Introduction to RPC Programming Guide

=====

This document assumes a working knowledge of network theory. It is intended for programmers who wish to write network applications using remote procedure calls (explained below), and who want to understand the RPC mechanisms usually hidden by the 'rpcgen' protocol compiler. 'rpcgen' is described in detail in the previous chapter, the rpcgen Programming Guide. *Note:* *Before attempting to write a network application, or to convert an existing non-network application to run over the network, you may want to understand the material in this chapter. However, for most applications, you can circumvent the need to cope with the details presented here by using 'rpcgen'. Generating XDR Routines contains the complete source for a working RPC service -- a remote directory listing service which uses 'rpcgen' to generate XDR routines as well as client and server stubs.*

What are remote procedure calls? Simply put, they are the high-level communications paradigm used in the operating system. RPC presumes the existence of low-level networking mechanisms (such as TCP/IP and UDP/IP), and upon them it implements a logical client to server communications system designed specifically for the support of network applications. With RPC, the client makes a procedure call to send a data packet to the server. When the packet arrives, the server calls a dispatch routine, performs whatever service is requested, sends back the reply, and the procedure call returns to the client.

1.43 RPC.guide/Layers of RPC

Layers of RPC

=====

The RPC interface can be seen as being divided into three layers.

(1)

The Highest Layer
 The Middle Layer
 The Lowest Layer
 The RPC Paradigm

----- Footnotes -----

(1) For a complete specification of the routines in the remote procedure call Library, see the 'rpc(3N)' manual page.

1.44 RPC.guide/The Highest Layer

The Highest Layer

The highest layer is totally transparent to the operating system, machine and network upon which it is run. It's probably best to think of this level as a way of *using* RPC, rather than as a *part of* RPC proper. Programmers who write RPC routines should (almost) always make this layer available to others by way of a simple C front end that entirely hides the networking.

To illustrate, at this level a program can simply make a call to ``rnusers()'``, a C routine which returns the number of users on a remote machine. The user is not explicitly aware of using RPC -- they simply call a procedure, just as they would call ``malloc()'``.

1.45 RPC.guide/The Middle Layer

The Middle Layer

The middle layer is really "RPC proper." Here, the user doesn't need to consider details about sockets, the UNIX system, or other low-level implementation mechanisms. They simply make remote procedure calls to routines on other machines. The selling point here is simplicity. It's this layer that allows RPC to pass the "hello world" test -- simple things should be simple. The middle-layer routines are used for most applications.

RPC calls are made with the system routines ``registerrpc()'``, ``callrpc()'`` and ``svc_run()'``. The first two of these are the most fundamental: ``registerrpc()'`` obtains a unique system-wide procedure-identification number, and ``callrpc()'`` actually executes a remote procedure call. At the middle level, a call to ``rnusers()'`` is implemented by way of these two routines.

The middle layer is unfortunately rarely used in serious programming due to its inflexibility (simplicity). It does not allow timeout specifications or the choice of transport. It allows no UNIX process control or flexibility in case of errors. It doesn't support multiple kinds of call authentication. The programmer rarely needs all these kinds of control, but one or two of them is often necessary.

1.46 RPC.guide/The Lowest Layer

The Lowest Layer

The lowest layer does allow these details to be controlled by the programmer, and for that reason it is often necessary. Programs written

Highest Layer
 Intermediate Layer
 Assigning Program Numbers
 Passing Arbitrary Data Types

1.49 RPC.guide/Highest Layer

Highest Layer

Imagine you're writing a program that needs to know how many users are logged into a remote machine. You can do this by calling the RPC library routine `'rnusers()'` as illustrated below:

```
#include <stdio.h>

main(argc, argv)
    int argc;
    char **argv;
{
    int num;

    if (argc != 2) {
        fprintf(stderr, "usage: rnusers hostname\n");
        exit(1);
    }
    if ((num = rnusers(argv[1])) < 0) {
        fprintf(stderr, "error: rnusers\n");
        exit(-1);
    }
    printf("%d users on %s\n", num, argv[1]);
    exit(0);
}
```

RPC library routines such as `'rnusers()'` are in the RPC services library `'librpcsvc.a'` Thus, the program above should be compiled with:

```
% cc program.c -lrpcsvc
```

`'rnusers()'`, like the other RPC library routines, is documented in section 3R of the `'System Interface Manual for the Sun Workstation'`, the same section which documents the standard Sun RPC services. See the `'intro(3R)'` manual page for an explanation of the documentation strategy for these services and their RPC protocols.

Here are some of the RPC service library routines available to the C programmer:

<code>'Routine'</code>	Description
<code>'rnusers()'</code>	

Return number of users on remote machine

``rusers()'`

Return information about users on remote machine

``havedisk()'`

Determine if remote machine has disk

``rstats()'`

Get performance data from remote kernel

``rwall()'`

Write to specified remote machines

``yppasswd()'`

Update user password in Yellow Pages

★Table 1★ RPC Service Library Routines

Other RPC services -- for example ``ether()'` ``mount()'` ``rquota()'` and ``spray()'` -- are not available to the C programmer as library routines. They do, however, have RPC program numbers so they can be invoked with ``callrpc()'` which will be discussed in the next section. Most of them also have compilable ``rpcgen'` protocol description files. (The ``rpcgen'` protocol compiler radically simplifies the process of developing network applications. see `rpcgen Programming Guide` for detailed information about ``rpcgen'` and ``rpcgen'` protocol description files).

1.50 RPC.guide/Intermediate Layer

Intermediate Layer

The simplest interface, which explicitly makes RPC calls, uses the functions ``callrpc()'` and ``registerrpc()'` Using this method, the number of remote users can be gotten as follows:

```
#include <stdio.h>
#include <rpc/rpc.h>
#include <utmp.h>
#include <rpcsvc/rusers.h>

main(argc, argv)
    int argc;
    char **argv;
{
    unsigned long nusers;
    int stat;

    if (argc != 2) {
        fprintf(stderr, "usage: nusers hostname\n");
        exit(-1);
    }
}
```

```

    if (stat = callrpc(argv[1],
        RUSERSPROC, RUSERSVERS, RUSERSPROC_NUM,
        xdr_void, 0, xdr_u_long, &nusers) != 0) {
        clnt_perrno(stat);
        exit(1);
    }
    printf("%d users on %s\n", nusers, argv[1]);
    exit(0);
}

```

Each RPC procedure is uniquely defined by a program number, version number, and procedure number. The program number specifies a group of related remote procedures, each of which has a different procedure number. Each program also has a version number, so when a minor change is made to a remote service (adding a new procedure, for example), a new program number doesn't have to be assigned. When you want to call a procedure to find the number of remote users, you look up the appropriate program, version and procedure numbers in a manual, just as you look up the name of a memory allocator when you want to allocate memory.

The simplest way of making remote procedure calls is with the the RPC library routine 'callrpc()' It has eight parameters. The first is the name of the remote server machine. The next three parameters are the program, version, and procedure numbers -- together they identify the procedure to be called. The fifth and sixth parameters are an XDR filter and an argument to be encoded and passed to the remote procedure. The final two parameters are a filter for decoding the results returned by the remote procedure and a pointer to the place where the procedure's results are to be stored. Multiple arguments and results are handled by embedding them in structures. If 'callrpc()' completes successfully, it returns zero; else it returns a nonzero value. The return codes (of type cast into an integer) are found in '<rpc/clnt.h>'.

Since data types may be represented differently on different machines, 'callrpc()' needs both the type of the RPC argument, as well as a pointer to the argument itself (and similarly for the result). For 'RUSERSPROC_NUM', the return value is an 'unsigned long' so 'callrpc()' has 'xdr_u_long()' as its first return parameter, which says that the result is of type 'unsigned long' and '&nusers' as its second return parameter, which is a pointer to where the long result will be placed. Since 'RUSERSPROC_NUM' takes no argument, the argument parameter of 'callrpc()' is 'xdr_void()'.

After trying several times to deliver a message, if 'callrpc()' gets no answer, it returns with an error code. The delivery mechanism is UDP, which stands for User Datagram Protocol. Methods for adjusting the number of retries or for using a different protocol require you to use the lower layer of the RPC library, discussed later in this document. The remote server procedure corresponding to the above might look like this:

```

char *
nuser(indata)
    char *indata;
{
    unsigned long nusers;

```

```

/*
 * Code here to compute the number of users
 * and place result in variable nusers.
 */
return((char *)&nusers);
}

```

It takes one argument, which is a pointer to the input of the remote procedure call (ignored in our example), and it returns a pointer to the result. In the current version of C, character pointers are the generic pointers, so both the input argument and the return value are cast to `'char *'`.

Normally, a server registers all of the RPC calls it plans to handle, and then goes into an infinite loop waiting to service requests. In this example, there is only a single procedure to register, so the main body of the server would look like this:

```

#include <stdio.h>
#include <rpc/rpc.h>
#include <utmp.h>
#include <rpcsvc/rusers.h>

char *nuser();

main()
{
    registerrpc(RUSERSPROG, RUSERSVERS, RUSERSPROC_NUM,
               nuser, xdr_void, xdr_u_long);
    svc_run(); /* Never returns */
    fprintf(stderr, "Error: svc_run returned!\n");
    exit(1);
}

```

The `'registerrpc()'` routine registers a C procedure as corresponding to a given RPC procedure number. The first three parameters, `'RUSERPROG'`, `'RUSERSVERS'`, and `'RUSERSPROC_NUM'` are the program, version, and procedure numbers of the remote procedure to be registered; `'nuser()'` is the name of the local procedure that implements the remote procedure; and `'xdr_void()'` and `'xdr_u_long()'` are the XDR filters for the remote procedure's arguments and results, respectively. (Multiple arguments or multiple results are passed as structures).

Only the UDP transport mechanism can use `'registerrpc()'` thus, it is always safe in conjunction with calls generated by `'callrpc()'`.

***Warning:** the UDP transport mechanism can only deal with arguments and results less than 8K bytes in length.*

After registering the local procedure, the server program's main procedure calls `'svc_run()'`, the RPC library's remote procedure dispatcher. It is this function that calls the remote procedures in response to RPC call messages. Note that the dispatcher takes care of decoding remote procedure arguments and encoding results, using the XDR filters specified when the remote procedure was registered.

1.51 RPC.guide/Assigning Program Numbers

Assigning Program Numbers

Program numbers are assigned in groups of '0x20000000' according to the following chart:

```

0x0 - 0x1fffffff Defined by Sun
0x20000000 - 0x3fffffff Defined by user
0x40000000 - 0x5fffffff Transient
0x60000000 - 0x7fffffff Reserved
0x80000000 - 0x9fffffff Reserved
0xa0000000 - 0xbfffffff Reserved
0xc0000000 - 0xdfffffff Reserved
0xe0000000 - 0xffffffff Reserved

```

Sun Microsystems administers the first group of numbers, which should be identical for all Sun customers. If a customer develops an application that might be of general interest, that application should be given an assigned number in the first range. The second group of numbers is reserved for specific customer applications. This range is intended primarily for debugging new programs. The third group is reserved for applications that generate program numbers dynamically. The final groups are reserved for future use, and should not be used.

To register a protocol specification, send a request by network mail to 'rpc@sun' or write to:

```

RPC Administrator
Sun Microsystems
2550 Garcia Ave.
Mountain View, CA 94043

```

Please include a compilable 'rpcgen' '.x' file describing your protocol. You will be given a unique program number in return.

The RPC program numbers and protocol specifications of standard Sun RPC services can be found in the include files in '/usr/include/rpcsvc'. These services, however, constitute only a small subset of those which have been registered. The complete list of registered programs, as of the time when this manual was printed, is:

```

100000 PMAPPROG          portmapper
100001 RSTATPROG        remote stats
100002 RUSERSPROG       remote users
100003 NFSPROG          nfs
100004 YPPROG            Yellow Pages
100005 MOUNTPROG        mount demon
100006 DBXPROG          remote dbx
100007 YPBINDPROG       yp binder
100008 WALLPROG         shutdown msg
100009 YPPASSWDPROG     yppasswd server

```

100010	ETHERSTATPROG	ether stats
100011	RQUOTAPROG	disk quotas
100012	SPRAYPROG	spray packets
100013	IBM3270PROG	3270 mapper
100014	IBMRJEPROG	RJE mapper
100015	SELNSVCPROG	selection service
100016	RDATABASEPROG	remote database access
100017	REXECPROG	remote execution
100018	ALICEPROG	Alice Office Automation
100019	SCHEDPROG	scheduling service
100020	LOCKPROG	local lock manager
100021	NETLOCKPROG	network lock manager
100022	X25PROG	x.25 inr protocol
100023	STATMON1PROG	status monitor 1
100024	STATMON2PROG	status monitor 2
100025	SELNLIBPROG	selection library
100026	BOOTPARAMPROG	boot parameters service
100027	MAZEPROG	mazewars game
100028	YPUUPDATEPROG	yp update
100029	KEYSERVEPROG	key server
100030	SECURECMDPROG	secure login
100031	NETFWDIPROG	nfs net forwarder init
100032	NETFWDTPROG	nfs net forwarder trans
100033	SUNLINKMAP_PROG	sunlink MAP
100034	NETMONPROG	network monitor
100035	DBASEPROG	lightweight database
100036	PWDAUTHPROG	password authorization
100037	TFSPROG	translucent file svc
100038	NSEPROG	nse server
100039	NSE_ACTIVATE_PROG	nse activate daemon
150001	PCNFSDPROG	pc passwd authorization
200000	PYRAMIDLOCKINGPROG	Pyramid-locking
200001	PYRAMIDSYS5	Pyramid-sys5
200002	CADDS_IMAGE	CV cadds_image
300001	ADT_RFLOCKPROG	ADT file locking

Table 2: RPC Registered Programs

1.52 RPC.guide/Passing Arbitrary Data Types

Passing Arbitrary Data Types

In the previous example, the RPC call passes a single 'unsigned long' RPC can handle arbitrary data structures, regardless of different machines' byte orders or structure layout conventions, by always converting them to a network standard called "External Data Representation" (XDR) before sending them over the wire. The process of converting from a particular machine representation to XDR format is called *serializing*, and the reverse process is called *deserializing*. The type field parameters of 'callrpc()' and

'`registerrpc()`' can be a built-in procedure like '`xdr_u_long()`' in the previous example, or a user supplied one. XDR has these built-in type routines:

```
'xdr_int()'
'xdr_u_int()'
'xdr_enum()'
'xdr_long()'
'xdr_u_long()'
'xdr_bool()'
'xdr_short()'
'xdr_u_short()'
'xdr_wrapstring()'
'xdr_char()'
'xdr_u_char()'
```

Note that the routine '`xdr_string()`' exists, but cannot be used with '`callrpc()`' and '`registerrpc()`', which only pass two parameters to their XDR routines. '`xdr_wrapstring()`' has only two parameters, and is thus OK. It calls '`xdr_string()`'.

As an example of a user-defined type routine, if you wanted to send the structure

```
struct simple {
    int a;
    short b;
} simple;
```

then you would call '`callrpc()`' as

```
callrpc(hostname, PROGNUM, VERSNUM, PROCNUM,
        xdr_simple, &simple ...);
```

where '`xdr_simple()`' is written as:

```
#include <rpc/rpc.h>

xdr_simple(xdrsp, simplep)
    XDR *xdrsp;
    struct simple *simplep;
{
    if (!xdr_int(xdrsp, &simplep->a))
        return (0);
    if (!xdr_short(xdrsp, &simplep->b))
        return (0);
    return (1);
}
```

An XDR routine returns nonzero (true in the sense of C) if it completes successfully, and zero otherwise. A complete description of XDR is in the XDR Protocol Specification, only few implementation examples are given here.

In addition to the built-in primitives, there are also the prefabricated building blocks:

```
'xdr_array()'
```

```
'xdr_bytes()'
'xdr_reference()'
'xdr_vector()'
'xdr_union()'
'xdr_pointer()'
'xdr_string()'
'xdr_opaque()'
```

To send a variable array of integers, you might package them up as a structure like this

```
struct varintarr {
    int *data;
    int arrlen;
} arr;
```

and make an RPC call such as

```
callrpc(hostname, PROGNUM, VERSNUM, PROCNUM,
        xdr_varintarr, &arr...);
```

with 'xdr_varintarr()' defined as:

```
xdr_varintarr(xdrsp, arrp)
    XDR *xdrsp;
    struct varintarr *arrp;
{
    return (xdr_array(xdrsp, &arrp->data, &arrp->arrlen,
        MAXLEN, sizeof(int), xdr_int));
}
```

This routine takes as parameters the XDR handle, a pointer to the array, a pointer to the size of the array, the maximum allowable array size, the size of each array element, and an XDR routine for handling each array element.

If the size of the array is known in advance, one can use 'xdr_vector()', which serializes fixed-length arrays.

```
int intarr[SIZE];

xdr_intarr(xdrsp, intarr)
    XDR *xdrsp;
    int intarr[];
{
    int i;

    return (xdr_vector(xdrsp, intarr, SIZE, sizeof(int),
        xdr_int));
}
```

XDR always converts quantities to 4-byte multiples when serializing. Thus, if either of the examples above involved characters instead of integers, each character would occupy 32 bits. That is the reason for the XDR routine 'xdr_bytes()' which is like 'xdr_array()' except that it packs characters; 'xdr_bytes()' has four parameters, similar to the first four parameters of 'xdr_array()'. For null-terminated strings, there is also the 'xdr_string()' routine, which is the same as

'xdr_bytes()' without the length parameter. On serializing it gets the string length from 'strlen()', and on deserializing it creates a null-terminated string.

Here is a final example that calls the previously written 'xdr_simple()' as well as the built-in functions 'xdr_string()' and 'xdr_reference()', which chases pointers:

```

struct finalexample {
    char *string;
    struct simple *simplep;
} finalexample;

xdr_finalexample(xdrsp, finalp)
    XDR *xdrsp;
    struct finalexample *finalp;
{
    if (!xdr_string(xdrsp, &finalp->string, MAXSTRLEN))
        return (0);
    if (!xdr_reference(xdrsp, &finalp->simplep,
        sizeof(struct simple), xdr_simple);
        return (0);
    return (1);
}

```

Note that we could as easily call 'xdr_simple()' here instead of 'xdr_reference()'.

1.53 RPC.guide/Lowest Layer of RPC

Lowest Layer of RPC

=====

In the examples given so far, RPC takes care of many details automatically for you. In this section, we'll show you how you can change the defaults by using lower layers of the RPC library. It is assumed that you are familiar with sockets and the system calls for dealing with them.

More on the Server Side
 Memory Allocation with XDR
 The Calling Side

There are several occasions when you may need to use lower layers of RPC. First, you may need to use TCP, since the higher layer uses UDP, which restricts RPC calls to 8K bytes of data. Using TCP permits calls to send long streams of data. For an example, see the TCP. Second, you may want to allocate and free memory while serializing or deserializing with XDR routines. There is no call at the higher level to let you free memory explicitly. For more explanation, see the Memory Allocation with XDR.

Third, you may need to perform authentication on either the client or server side, by supplying credentials or verifying them. See the explanation in the Authentication.

1.54 RPC.guide/More on the Server Side

More on the Server Side

The server for the 'nusers()' program shown below does the same thing as the one using 'registerrpc()' above, but is written using a lower layer of the RPC package:

```
#include <stdio.h>
#include <rpc/rpc.h>
#include <utmp.h>
#include <rpcsvc/rusers.h>

main()
{
    SVCXPRT *transp;
    int nuser();

    transp = svcudp_create(RPC_ANYSOCK);
    if (transp == NULL) {
        fprintf(stderr, "can't create an RPC server\n");
        exit(1);
    }
    pmap_unset(RUSERSPROG, RUSERSVERS);
    if (!svc_register(transp, RUSERSPROG, RUSERSVERS,
                     nuser, IPPROTO_UDP)) {
        fprintf(stderr, "can't register RUSER service\n");
        exit(1);
    }
    svc_run(); /* Never returns */
    fprintf(stderr, "should never reach this point\n");
}

nuser(rqstp, transp)
    struct svc_req *rqstp;
    SVCXPRT *transp;
{
    unsigned long nusers;

    switch (rqstp->rq_proc) {
    case NULLPROC:
        if (!svc_sendreply(transp, xdr_void, 0))
            fprintf(stderr, "can't reply to RPC call\n");
        return;
    case RUSERSPROC_NUM:
        /*
         * Code here to compute the number of users
         * and assign it to the variable 'nusers'
         */
    }
```

```

        if (!svc_sendreply(transp, xdr_u_long, &nusers))
            fprintf(stderr, "can't reply to RPC call\n");
        return;
    default:
        svcerr_noproc(transp);
        return;
    }
}

```

First, the server gets a transport handle, which is used for receiving and replying to RPC messages. `'registerrpc()'` uses `'svcdp_create()'` to get a UDP handle. If you require a more reliable protocol, call `'svctcp_create()'` instead. If the argument to `'svcdp_create()'` is `'RPC_ANYSOCK'` the RPC library creates a socket on which to receive and reply to RPC calls. Otherwise, `'svcdp_create()'` expects its argument to be a valid socket number. If you specify your own socket, it can be bound or unbound. If it is bound to a port by the user, the port numbers of `'svcdp_create()'` and `'clnttcp_create()'` (the low-level client routine) must match.

If the user specifies the `'RPC_ANYSOCK'` argument, the RPC library routines will open sockets. Otherwise they will expect the user to do so. The routines `'svcdp_create()'` and `'clntudp_create()'` will cause the RPC library routines to `'bind()'` their socket if it is not bound already.

A service may choose to register its port number with the local portmapper service. This is done by specifying a non-zero protocol number in `'svc_register()'`. Incidentally, a client can discover the server's port number by consulting the portmapper on their server's machine. This can be done automatically by specifying a zero port number in `'clntudp_create()'` or `'clnttcp_create()'`.

After creating an `'SVCXPRT'`, the next step is to call `'pmap_unset()'` so that if the `'nusers()'` server crashed earlier, any previous trace of it is erased before restarting. More precisely, `'pmap_unset()'` erases the entry for `'RUSERSPROG'` from the port mapper's tables.

Finally, we associate the program number for `'nusers()'` with the procedure `'nuser()'`. The final argument to `'svc_register()'` is normally the protocol being used, which, in this case, is `'IPPROTO_UDP'`. Notice that unlike `'registerrpc()'`, there are no XDR routines involved in the registration process. Also, registration is done on the program, rather than procedure, level.

The user routine `'nuser()'` must call and dispatch the appropriate XDR routines based on the procedure number. Note that two things are handled by `'nuser()'` that `'registerrpc()'` handles automatically. The first is that procedure `'NULLPROC'` (currently zero) returns with no results. This can be used as a simple test for detecting if a remote program is running. Second, there is a check for invalid procedure numbers. If one is detected, `'svcerr_noproc()'` is called to handle the error.

The user service routine serializes the results and returns them to the RPC caller via `'svc_sendreply()'`. Its first parameter is the `'SVCXPRT'` handle, the second is the XDR routine, and the third is a

pointer to the data to be returned. Not illustrated above is how a server handles an RPC program that receives data. As an example, we can add a procedure 'RUSERSPROC_BOOL' which has an argument 'nusers()', and returns 'TRUE' or 'FALSE' depending on whether there are nusers logged on. It would look like this:

```

case RUSERSPROC_BOOL: {
    int bool;
    unsigned nuserquery;

    if (!svc_getargs(transp, xdr_u_int, &nuserquery) {
        svcerr_decode(transp);
        return;
    }
    /*
     * Code to set 'nusers' = number of users
     */
    if (nuserquery == nusers)
        bool = TRUE;
    else
        bool = FALSE;
    if (!svc_sendreply(transp, xdr_bool, &bool)) {
        fprintf(stderr, "can't reply to RPC call\n");
        return (1);
    }
    return;
}

```

The relevant routine is 'svc_getargs()' which takes an 'SVCXPRT' handle, the XDR routine, and a pointer to where the input is to be placed as arguments.

1.55 RPC.guide/Memory Allocation with XDR

Memory Allocation with XDR

XDR routines not only do input and output, they also do memory allocation. This is why the second parameter of 'xdr_array()' is a pointer to an array, rather than the array itself. If it is 'NULL', then 'xdr_array()' allocates space for the array and returns a pointer to it, putting the size of the array in the third argument. As an example, consider the following XDR routine 'xdr_chararr1()' which deals with a fixed array of bytes with length 'SIZE'.

```

xdr_chararr1(xdrsp, chararr)
    XDR *xdrsp;
    char chararr[];
{
    char *p;
    int len;

    p = chararr;
    len = SIZE;

```

```

    return (xdr_bytes(xdrsp, &p, &len, SIZE));
}

```

If space has already been allocated in 'chararr', it can be called from a server like this:

```

char chararr[SIZE];

svc_getargs(transp, xdr_chararr1, chararr);

```

If you want XDR to do the allocation, you would have to rewrite this routine in the following way:

```

xdr_chararr2(xdrsp, chararrp)
    XDR *xdrsp;
    char **chararrp;
{
    int len;

    len = SIZE;
    return (xdr_bytes(xdrsp, chararrp, &len, SIZE));
}

```

Then the RPC call might look like this:

```

char *arrptr;

arrptr = NULL;
svc_getargs(transp, xdr_chararr2, &arrptr);
/*
 * Use the result here
 */
svc_freeargs(transp, xdr_chararr2, &arrptr);

```

Note that, after being used, the character array can be freed with 'svc_freeargs()'. 'svc_freeargs()' will not attempt to free any memory if the variable indicating it is NULL. For example, in the routine 'xdr_finalexample()', given earlier, if 'finalp->string' was 'NULL', then it would not be freed. The same is true for 'finalp->simplep'.

To summarize, each XDR routine is responsible for serializing, deserializing, and freeing memory. When an XDR routine is called from 'callrpc()' the serializing part is used. When called from 'svc_getargs()' the deserializer is used. And when called from 'svc_freeargs()' the memory deallocator is used. When building simple examples like those in this section, a user doesn't have to worry about the three modes. see XDR Technical Notes for examples of more sophisticated XDR routines that determine which of the three modes they are in and adjust their behavior accordingly.

1.56 RPC.guide/The Calling Side

The Calling Side

When you use `callrpc()` you have no control over the RPC delivery mechanism or the socket used to transport the data. To illustrate the layer of RPC that lets you adjust these parameters, consider the following code to call the `nusers` service:

```
#include <stdio.h>
#include <rpc/rpc.h>
#include <utmp.h>
#include <rpcsvc/rusers.h>
#include <sys/socket.h>
#include <sys/time.h>
#include <netdb.h>

main(argc, argv)
    int argc;
    char **argv;
{
    struct hostent *hp;
    struct timeval pertry_timeout, total_timeout;
    struct sockaddr_in server_addr;
    int sock = RPC_ANYSOCK;
    register CLIENT *client;
    enum clnt_stat clnt_stat;
    unsigned long nusers;

    if (argc != 2) {
        fprintf(stderr, "usage: nusers hostname\n");
        exit(-1);
    }
    if ((hp = gethostbyname(argv[1])) == NULL) {
        fprintf(stderr, "can't get addr for %s\n", argv[1]);
        exit(-1);
    }
    pertry_timeout.tv_sec = 3;
    pertry_timeout.tv_usec = 0;
    bcopy(hp->h_addr, (caddr_t)&server_addr.sin_addr,
          hp->h_length);
    server_addr.sin_family = AF_INET;
    server_addr.sin_port = 0;
    if ((client = clntudp_create(&server_addr, RUSERSPROG,
                                RUSERSVERS, pertry_timeout, &sock)) == NULL) {
        clnt_pcreateerror("clntudp_create");
        exit(-1);
    }
    total_timeout.tv_sec = 20;
    total_timeout.tv_usec = 0;
    clnt_stat = clnt_call(client, RUSERSPROC_NUM, xdr_void,
                          0, xdr_u_long, &nusers, total_timeout);
    if (clnt_stat != RPC_SUCCESS) {
        clnt_perror(client, "rpc");
        exit(-1);
    }
    clnt_destroy(client);
    close(sock);
    exit(0);
}
```

The low-level version of `'callrpc()'` is `'clnt_call()'` which takes a `'CLIENT'` pointer rather than a host name. The parameters to `'clnt_call()'` are a `'CLIENT'` pointer, the procedure number, the XDR routine for serializing the argument, a pointer to the argument, the XDR routine for deserializing the return value, a pointer to where the return value will be placed, and the time in seconds to wait for a reply.

The `'CLIENT'` pointer is encoded with the transport mechanism. `'callrpc()'` uses UDP, thus it calls `'clntudp_create()'` to get a `'CLIENT'` pointer. To get TCP (Transmission Control Protocol), you would use `'clnttcp_create()'`.

The parameters to `'clntudp_create()'` are the server address, the program number, the version number, a timeout value (between tries), and a pointer to a socket. The final argument to `'clnt_call()'` is the total time to wait for a response. Thus, the number of tries is the `'clnt_call()'` timeout divided by the `'clntudp_create()'` timeout.

Note that the `'clnt_destroy()'` call always deallocates the space associated with the `'CLIENT'` handle. It closes the socket associated with the `'CLIENT'` handle, however, only if the RPC library opened it. If the socket was opened by the user, it stays open. This makes it possible, in cases where there are multiple client handles using the same socket, to destroy one handle without closing the socket that other handles are using.

To make a stream connection, the call to `'clntudp_create()'` is replaced with a call to `'clnttcp_create()'`.

```
clnttcp_create(&server_addr, prognum, versnum, &sock,
              inputsize, outputsize);
```

There is no timeout argument; instead, the receive and send buffer sizes must be specified. When the `'clnttcp_create()'` call is made, a TCP connection is established. All RPC calls using that `'CLIENT'` handle would use this connection. The server side of an RPC call using TCP has `'svcudp_create()'` replaced by `'svctcp_create()'`.

```
transp = svctcp_create(RPC_ANYSOCK, 0, 0);
```

The last two arguments to `'svctcp_create()'` are send and receive sizes respectively. If `'0'` is specified for either of these, the system chooses a reasonable default.

1.57 RPC.guide/Other RPC Features

Other RPC Features

=====

This section discusses some other aspects of RPC that are occasionally useful.

Select on the Server Side
 Broadcast RPC
 RPC Batching
 RPC Authentication
 Using Inetd

1.58 RPC.guide/Select on the Server Side

Select on the Server Side

Suppose a process is processing RPC requests while performing some other activity. If the other activity involves periodically updating a data structure, the process can set an alarm signal before calling 'svc_run()' But if the other activity involves waiting on a file descriptor, the 'svc_run()' call won't work. The code for 'svc_run()' is as follows:

```
void
svc_run()
{
    fd_set readfds;
    int dtbsz = getdtablesize();

    for (;;) {
        readfds = svc_fds;
        switch (select(dtbsz, &readfds, NULL, NULL, NULL)) {

            case -1:
                if (errno == EINTR)
                    continue;
                perror("select");
                return;
            case 0:
                break;
            default:
                svc_getreqset(&readfds);
        }
    }
}
```

You can bypass 'svc_run()' and call 'svc_getreqset()' yourself. All you need to know are the file descriptors of the socket(s) associated with the programs you are waiting on. Thus you can have your own 'select()' that waits on both the RPC socket, and your own descriptors. Note that 'svc_fds' is a bit mask of all the file descriptors that RPC is using for services. It can change everytime that *any* RPC library routine is called, because descriptors are constantly being opened and closed, for example for TCP connections.

1.59 RPC.guide/Broadcast RPC

Broadcast RPC

The 'portmapper' is a daemon that converts RPC program numbers into DARPA protocol port numbers; see the 'portmap' man page. You can't do broadcast RPC without the portmapper. Here are the main differences between broadcast RPC and normal RPC calls:

1. Normal RPC expects one answer, whereas broadcast RPC expects many answers (one or more answer from each responding machine).
2. Broadcast RPC can only be supported by packet-oriented (connectionless) transport protocols like UDP/IP.
3. The implementation of broadcast RPC treats all unsuccessful responses as garbage by filtering them out. Thus, if there is a version mismatch between the broadcaster and a remote service, the user of broadcast RPC never knows.
4. All broadcast messages are sent to the portmap port. Thus, only services that register themselves with their portmapper are accessible via the broadcast RPC mechanism.
5. Broadcast requests are limited in size to the MTU (Maximum Transfer Unit) of the local network. For Ethernet, the MTU is 1500 bytes.

Broadcast RPC Synopsis

1.60 RPC.guide/Broadcast RPC Synopsis

Broadcast RPC Synopsis
.....

```
#include <rpc/pmap_clnt.h>
. . .
enum clnt_stat  clnt_stat;
. . .
clnt_stat = clnt_broadcast(prognum, versnum, procnum,
    inproc, in, outproc, out, eachresult)
    u_long    prognum;        /* program number */
    u_long    versnum;       /* version number */
    u_long    procnum;       /* procedure number */
    xdrproc_t inproc;        /* xdr routine for args */
    caddr_t   in;            /* pointer to args */
    xdrproc_t outproc;       /* xdr routine for results */
    caddr_t   out;           /* pointer to results */
    bool_t    (*eachresult)(); /* call with each result gotten */
```

The procedure 'eachresult()' is called each time a valid result is obtained. It returns a boolean that indicates whether or not the user

wants more responses.

```
bool_t done;
. . .
done = eachresult(resultsp, raddr)
    caddr_t resultsp;
    struct sockaddr_in *raddr; /* Addr of responding machine */
```

If 'done' is 'TRUE', then broadcasting stops and 'clnt_broadcast()' returns successfully. Otherwise, the routine waits for another response. The request is rebroadcast after a few seconds of waiting. If no responses come back, the routine returns with 'RPC_TIMEDOUT'.

1.61 RPC.guide/RPC Batching

Batching

The RPC architecture is designed so that clients send a call message, and wait for servers to reply that the call succeeded. This implies that clients do not compute while servers are processing a call. This is inefficient if the client does not want or need an acknowledgement for every message sent. It is possible for clients to continue computing while waiting for a response, using RPC batch facilities.

RPC messages can be placed in a "pipeline" of calls to a desired server; this is called batching. Batching assumes that:

1. each RPC call in the pipeline requires no response from the server, and the server does not send a response message; and
2. the pipeline of calls is transported on a reliable byte stream transport such as TCP/IP.

Since the server does not respond to every call, the client can generate new calls in parallel with the server executing previous calls. Furthermore, the TCP/IP implementation can buffer up many call messages, and send them to the server in one 'send()' system call. This overlapped execution greatly decreases the interprocess communication overhead of the client and server processes, and the total elapsed time of a series of calls.

Since the batched calls are buffered, the client should eventually do a nonbatched call in order to flush the pipeline.

A contrived example of batching follows. Assume a string rendering service (like a window system) has two similar calls: one renders a string and returns void results, while the other renders a string and remains silent. The service (using the TCP/IP transport) may look like:

```
#include <stdio.h>
#include <rpc/rpc.h>
#include <suntool/windows.h>
```

```
void windowdispatch();

main()
{
    SVCXPRT *transp;

    transp = svctcp_create(RPC_ANYSOCK, 0, 0);
    if (transp == NULL) {
        fprintf(stderr, "can't create an RPC server\n");
        exit(1);
    }
    pmap_unset(WINDOWPROG, WINDOWVERS);
    if (!svc_register(transp, WINDOWPROG, WINDOWVERS,
        windowdispatch, IPPROTO_TCP)) {
        fprintf(stderr, "can't register WINDOW service\n");
        exit(1);
    }
    svc_run(); /* Never returns */
    fprintf(stderr, "should never reach this point\n");
}

void
windowdispatch(rqstp, transp)
    struct svc_req *rqstp;
    SVCXPRT *transp;
{
    char *s = NULL;

    switch (rqstp->rq_proc) {
    case NULLPROC:
        if (!svc_sendreply(transp, xdr_void, 0))
            fprintf(stderr, "can't reply to RPC call\n");
        return;
    case RENDERSTRING:
        if (!svc_getargs(transp, xdr_wrapstring, &s)) {
            fprintf(stderr, "can't decode arguments\n");
            /*
             * Tell caller he screwed up
             */
            svcerr_decode(transp);
            break;
        }
        /*
         * Code here to render the string 's'
         */
        if (!svc_sendreply(transp, xdr_void, NULL))
            fprintf(stderr, "can't reply to RPC call\n");
        break;
    case RENDERSTRING_BATCHED:
        if (!svc_getargs(transp, xdr_wrapstring, &s)) {
            fprintf(stderr, "can't decode arguments\n");
            /*
             * We are silent in the face of protocol errors
             */
            break;
        }
        /*

```

```

        * Code here to render string s, but send no reply!
        */
        break;
default:
    svcerr_noproc(transp);
    return;
}
/*
 * Now free string allocated while decoding arguments
 */
svc_freeargs(transp, xdr_wrapstring, &s);
}

```

Of course the service could have one procedure that takes the string and a boolean to indicate whether or not the procedure should respond.

In order for a client to take advantage of batching, the client must perform RPC calls on a TCP-based transport and the actual calls must have the following attributes:

1. the result's XDR routine must be zero ('NULL'), and
2. the RPC call's timeout must be zero.

Here is an example of a client that uses batching to render a bunch of strings; the batching is flushed when the client gets a null string (EOF):

```

#include <stdio.h>
#include <rpc/rpc.h>
#include <sys/socket.h>
#include <sys/time.h>
#include <netdb.h>
#include <suntool/windows.h>

main(argc, argv)
    int argc;
    char **argv;
{
    struct hostent *hp;
    struct timeval pertry_timeout, total_timeout;
    struct sockaddr_in server_addr;
    int sock = RPC_ANYSOCK;
    register CLIENT *client;
    enum clnt_stat clnt_stat;
    char buf[1000], *s = buf;

    if ((client = clnttcp_create(&server_addr,
        WINDOWPROG, WINDOWVERS, &sock, 0, 0)) == NULL) {
        perror("clnttcp_create");
        exit(-1);
    }
    total_timeout.tv_sec = 0;
    total_timeout.tv_usec = 0;
    while (scanf("%s", s) != EOF) {
        clnt_stat = clnt_call(client, RENDERSTRING_BATCHED,
            xdr_wrapstring, &s, NULL, NULL, total_timeout);
    }
}

```

```

        if (clnt_stat != RPC_SUCCESS) {
            clnt_perror(client, "batched rpc");
            exit(-1);
        }
    }

    /* Now flush the pipeline */

    total_timeout.tv_sec = 20;
    clnt_stat = clnt_call(client, NULLPROC, xdr_void, NULL,
        xdr_void, NULL, total_timeout);
    if (clnt_stat != RPC_SUCCESS) {
        clnt_perror(client, "rpc");
        exit(-1);
    }
    clnt_destroy(client);
    exit(0);
}

```

Since the server sends no message, the clients cannot be notified of any of the failures that may occur. Therefore, clients are on their own when it comes to handling errors.

The above example was completed to render all of the (2000) lines in the file `'/etc/termcap'`. The rendering service did nothing but throw the lines away. The example was run in the following four configurations:

1. machine to itself, regular RPC;
2. machine to itself, batched RPC;
3. machine to another, regular RPC; and
4. machine to another, batched RPC.

The results are as follows:

1. 50 seconds;
2. 16 seconds;
3. 52 seconds;
4. 10 seconds.

Running `'fscanf()'` on `'/etc/termcap'` only requires six seconds. These timings show the advantage of protocols that allow for overlapped execution, though these protocols are often hard to design.

1.62 RPC.guide/RPC Authentication

Authentication

In the examples presented so far, the caller never identified itself to the server, and the server never required an ID from the caller. Clearly, some network services, such as a network filesystem, require stronger security than what has been presented so far.

In reality, every RPC call is authenticated by the RPC package on the server, and similarly, the RPC client package generates and sends authentication parameters. Just as different transports (TCP/IP or UDP/IP) can be used when creating RPC clients and servers, different forms of authentication can be associated with RPC clients; the default authentication type used as a default is type 'none'.

The authentication subsystem of the RPC package is open ended. That is, numerous types of authentication are easy to support.

UNIX Authentication
DES Authentication

1.63 RPC.guide/UNIX Authentication

UNIX Authentication
.....

The Client Side
.....

When a caller creates a new RPC client handle as in:

```
clnt = clntudp_create(address, prognum, versnum,
                    wait, sockp)
```

the appropriate transport instance defaults the associate authentication handle to be

```
clnt->cl_auth = authnone_create();
```

The RPC client can choose to use 'UNIX' style authentication by setting 'clnt->cl_auth' after creating the RPC client handle:

```
clnt->cl_auth = authunix_create_default();
```

This causes each RPC call associated with 'clnt' to carry with it the following authentication credentials structure:

```
/*
 * UNIX style credentials.
 */
struct authunix_parms {
    u_long    aup_time;        /* credentials creation time */
    char      *aup_machname;   /* host name where client is */
    int       aup_uid;        /* client's UNIX effective uid */
    int       aup_gid;        /* client's current group id */
}
```

```

    u_int    aup_len;          /* element length of aup_gids */
    int      *aup_gids;       /* array of groups user is in */
};

```

These fields are set by `'authunix_create_default()'` by invoking the appropriate system calls. Since the RPC user created this new style of authentication, the user is responsible for destroying it with:

```
auth_destroy(clnt->cl_auth);
```

This should be done in all cases, to conserve memory.

The Server Side

.....

Service implementors have a harder time dealing with authentication issues since the RPC package passes the service dispatch routine a request that has an arbitrary authentication style associated with it. Consider the fields of a request handle passed to a service dispatch routine:

```

/*
 * An RPC Service request
 */
struct svc_req {
    u_long    rq_prog;        /* service program number */
    u_long    rq_vers;       /* service protocol vers num */
    u_long    rq_proc;       /* desired procedure number */
    struct opaque_auth rq_cred; /* raw credentials from wire */
    caddr_t   rq_clntcred;   /* credentials (read only) */
};

```

The `'rq_cred'` is mostly opaque, except for one field of interest: the style or flavor of authentication credentials:

```

/*
 * Authentication info. Mostly opaque to the programmer.
 */
struct opaque_auth {
    enum_t    oa_flavor;     /* style of credentials */
    caddr_t   oa_base;      /* address of more auth stuff */
    u_int     oa_length;    /* not to exceed MAX_AUTH_BYTES */
};

```

The RPC package guarantees the following to the service dispatch routine:

1. That the request's `'rq_cred'` is well formed. Thus the service implementor may inspect the request's `'rq_cred.oa_flavor'` to determine which style of authentication the caller used. The service implementor may also wish to inspect the other fields of `'rq_cred'` if the style is not one of the styles supported by the RPC package.
2. That the request's `'rq_clntcred'` field is either `'NULL'` or points to a well formed structure that corresponds to a supported style of authentication credentials. Remember that only `'UNIX'` style is

currently supported, so (currently) `'rq_clntcred'` could be cast to a pointer to an `'authunix_parms'` structure. If `'rq_clntcred'` is `'NULL'`, the service implementor may wish to inspect the other (opaque) fields of `'rq_cred'` in case the service knows about a new type of authentication that the RPC package does not know about.

Our remote users service example can be extended so that it computes results for all users except UID 16:

```
nuser(rqstp, transp)
    struct svc_req *rqstp;
    SVCXPRT *transp;
{
    struct authunix_parms *unix_cred;
    int uid;
    unsigned long nusers;

    /*
     * We don't care about authentication for null proc
     */
    if (rqstp->rq_proc == NULLPROC) {
        if (!svc_sendreply(transp, xdr_void, 0)) {
            fprintf(stderr, "can't reply to RPC call\n");
            return (1);
        }
        return;
    }
    /*
     * Now get the uid
     */
    switch (rqstp->rq_cred.oa_flavor) {
    case AUTH_UNIX:
        unix_cred =
            (struct authunix_parms *)rqstp->rq_clntcred;
        uid = unix_cred->aup_uid;
        break;
    case AUTH_NULL:
    default:
        svcerr_weakauth(transp);
        return;
    }
    switch (rqstp->rq_proc) {
    case RUSERSPROC_NUM:
        /*
         * Make sure caller is allowed to call this proc
         */
        if (uid == 16) {
            svcerr_systemerr(transp);
            return;
        }
        /*
         * Code here to compute the number of users
         * and assign it to the variable 'nusers'
         */
        if (!svc_sendreply(transp, xdr_u_long, &nusers)) {
            fprintf(stderr, "can't reply to RPC call\n");
            return (1);
        }
    }
}
```

```

        }
        return;
    default:
        svcerr_noproc(transp);
        return;
    }
}

```

A few things should be noted here. First, it is customary not to check the authentication parameters associated with the 'NULLPROC' (procedure number zero). Second, if the authentication parameter's type is not suitable for your service, you should call 'svcerr_weakauth()'. And finally, the service protocol itself should return status for access denied; in the case of our example, the protocol does not have such a status, so we call the service primitive 'svcerr_systemerr()' instead.

The last point underscores the relation between the RPC authentication package and the services; RPC deals only with *authentication* and not with individual services' *access control*. The services themselves must implement their own access control policies and reflect these policies as return statuses in their protocols.

1.64 RPC.guide/DES Authentication

DES Authentication

UNIX authentication is quite easy to defeat. Instead of using 'authunix_create_default()', one can call 'authunix_create()' and then modify the RPC authentication handle it returns by filling in whatever user ID and hostname they wish the server to think they have. DES authentication is thus recommended for people who want more security than UNIX authentication offers.

The details of the DES authentication protocol are complicated and are not explained here. see Remote Procedure Calls Protocol Specification for the details.

In order for DES authentication to work, the 'keyserv(8c)' daemon must be running on both the server and client machines. The users on these machines need public keys assigned by the network administrator in the 'publickey' database. And, they need to have decrypted their secret keys using their login password. This automatically happens when one logs in using 'login', or can be done manually using 'keylogin'. The 'Network Services' explains more how to setup secure networking.

Client Side

If a client wishes to use DES authentication, it must set its authentication handle appropriately. Here is an example:

```
cl->cl_auth =
    authdes_create(servername, 60, &server_addr, NULL);
```

The first argument is the network name or "netname" of the owner of the server process. Typically, server processes are root processes and their netname can be derived using the following call:

```
char servername[MAXNETNAMELEN];

host2netname(servername, rhostname, NULL);
```

Here, 'rhostname' is the hostname of the machine the server process is running on. 'host2netname()' fills in 'servername' to contain this root process's netname. If the server process was run by a regular user, one could use the call 'user2netname()' instead. Here is an example for a server process with the same user ID as the client:

```
char servername[MAXNETNAMELEN];

user2netname(servername, getuid(), NULL);
```

The last argument to both of these calls, 'user2netname()' and 'host2netname()', is the name of the naming domain where the server is located. The 'NULL' used here means "use the local domain name."

The second argument to 'authdes_create()' is a lifetime for the credential. Here it is set to sixty seconds. What that means is that the credential will expire 60 seconds from now. If some mischievous user tries to reuse the credential, the server RPC subsystem will recognize that it has expired and not grant any requests. If the same mischievous user tries to reuse the credential within the sixty second lifetime, he will still be rejected because the server RPC subsystem remembers which credentials it has already seen in the near past, and will not grant requests to duplicates.

The third argument to 'authdes_create()' is the address of the host to synchronize with. In order for DES authentication to work, the server and client must agree upon the time. Here we pass the address of the server itself, so the client and server will both be using the same time: the server's time. The argument can be 'NULL', which means "don't bother synchronizing." You should only do this if you are sure the client and server are already synchronized.

The final argument to 'authdes_create()' is the address of a DES encryption key to use for encrypting timestamps and data. If this argument is 'NULL', as it is in this example, a random key will be chosen. The client may find out the encryption key being used by consulting the 'ah_key' field of the authentication handle.

Server Side
.....

The server side is a lot simpler than the client side. Here is the previous example rewritten to use 'AUTH_DES' instead of 'AUTH_UNIX':

```
#include <sys/time.h>
#include <rpc/auth_des.h>
```

```

    . . .
    . . .
nuser(rqstp, transp)
    struct svc_req *rqstp;
    SVCXPRT *transp;
{
    struct authdes_cred *des_cred;
    int uid;
    int gid;
    int gidlen;
    int gidlist[10];
    /*
     * We don't care about authentication for null proc
     */

    if (rqstp->rq_proc == NULLPROC) {
        /* same as before */
    }

    /*
     * now get the uid
     */
    switch (rqstp->rq_cred.oa_flavor) {
    case AUTH_DES:
        des_cred =
            (struct authdes_cred *) rqstp->rq_clntcred;
        if (! netname2user(des_cred->adc_fullname.name,
            &uid, &gid, &gidlen, gidlist))
        {
            fprintf(stderr, "unknown user: %s\n",
                des_cred->adc_fullname.name);
            svcerr_systemerr(transp);
            return;
        }
        break;
    case AUTH_NULL:
    default:
        svcerr_weakauth(transp);
        return;
    }

    /*
     * The rest is the same as before
     */
}

```

Note the use of the routine `'netname2user()'`, the inverse of `'user2netname()'`: it takes a network ID and converts to a unix ID. `'netname2user()'` also supplies the group IDs which we don't use in this example, but which may be useful to other UNIX programs.

1.65 RPC.guide/Using Inetd

Using Inetd

An RPC server can be started from 'inetd'. The only difference from the usual code is that the service creation routine should be called in the following form:

```
transp = svcudp_create(0);      /* For UDP */
transp = svctcp_create(0,0,0); /* For listener TCP sockets */
transp = svcfd_create(0,0,0);  /* For connected TCP sockets */
```

since 'inet' passes a socket as file descriptor 0. Also, 'svc_register()' should be called as

```
svc_register(transp, PROGNUM, VERSNUM, service, 0);
```

with the final flag as 0, since the program would already be registered by 'inetd'. Remember that if you want to exit from the server process and return control to 'inetd' you need to explicitly exit, since 'svc_run()' never returns.

The format of entries in '/etc/inetd.conf' for RPC services is in one of the following two forms:

```
p_name/version dgram rpc/udp wait/nowait user server args
p_name/version stream rpc/tcp wait/nowait user server args
```

where 'p_name' is the symbolic name of the program as it appears in the 'rpc' data base, 'server' is the program implementing the server, and 'program' and 'version' are the program and version numbers of the service. For more information, see 'inetd.conf(5)'.

If the same program handles multiple versions, then the version number can be a range, as in this example:

```
rstatd/1-2 dgram rpc/udp wait root /usr/etc/rpc.rstatd
```

1.66 RPC.guide/More Examples

More Examples

=====

```
Versions
TCP
Callback Procedures
```

1.67 RPC.guide/Versions

Versions

By convention, the first version number of program 'PROG' is 'PROGVERS_ORIG' and the most recent version is 'PROGVERS'. Suppose there is a new version of the 'user' program that returns an 'unsigned short' rather than a 'long'. If we name this version 'RUSERSVERS_SHORT' then a server that wants to support both versions would do a double register.

```

if (!svc_register(transp, RUSERSPROG, RUSERSVERS_ORIG,
  nuser, IPPROTO_TCP)) {
    fprintf(stderr, "can't register RUSER service\n");
    exit(1);
}
if (!svc_register(transp, RUSERSPROG, RUSERSVERS_SHORT,
  nuser, IPPROTO_TCP)) {
    fprintf(stderr, "can't register RUSER service\n");
    exit(1);
}

```

Both versions can be handled by the same C procedure:

```

nuser(rqstp, transp)
    struct svc_req *rqstp;
    SVCXPRT *transp;
{
    unsigned long nusers;
    unsigned short nusers2;

    switch (rqstp->rq_proc) {
    case NULLPROC:
        if (!svc_sendreply(transp, xdr_void, 0)) {
            fprintf(stderr, "can't reply to RPC call\n");
            return (1);
        }
        return;
    case RUSERSPROC_NUM:
        /*
         * Code here to compute the number of users
         * and assign it to the variable 'nusers'
         */
        nusers2 = nusers;
        switch (rqstp->rq_vers) {
        case RUSERSVERS_ORIG:
            if (!svc_sendreply(transp, xdr_u_long,
                &nusers)) {
                fprintf(stderr, "can't reply to RPC call\n");
            }
            break;
        case RUSERSVERS_SHORT:
            if (!svc_sendreply(transp, xdr_u_short,
                &nusers2)) {
                fprintf(stderr, "can't reply to RPC call\n");
            }
            break;
        }
    default:
        svcerr_noproc(transp);
        return;
    }
}

```

```
}

```

1.68 RPC.guide/TCP

```
TCP
---
```

Here is an example that is essentially 'rcp'. The initiator of the RPC **snd** call takes its standard input and sends it to the server **rcv**, which prints it on standard output. The RPC call uses TCP. This also illustrates an XDR procedure that behaves differently on serialization than on deserialization.

```
/*
 * The xdr routine:
 *     on decode, read from wire, write onto fp
 *     on encode, read from fp, write onto wire
 */
#include <stdio.h>
#include <rpc/rpc.h>

xdr_rcp(xdrs, fp)
    XDR *xdrs;
    FILE *fp;
{
    unsigned long size;
    char buf[BUFSIZ], *p;

    if (xdrs->x_op == XDR_FREE)/* nothing to free */
        return 1;
    while (1) {
        if (xdrs->x_op == XDR_ENCODE) {
            if ((size = fread(buf, sizeof(char), BUFSIZ,
                fp)) == 0 && ferror(fp)) {
                fprintf(stderr, "can't fread\n");
                return (1);
            }
        }
        p = buf;
        if (!xdr_bytes(xdrs, &p, &size, BUFSIZ))
            return 0;
        if (size == 0)
            return 1;
        if (xdrs->x_op == XDR_DECODE) {
            if (fwrite(buf, sizeof(char), size,
                fp) != size) {
                fprintf(stderr, "can't fwrite\n");
                return (1);
            }
        }
    }
}
/*
```

```

    * The sender routines
    */
#include <stdio.h>
#include <netdb.h>
#include <rpc/rpc.h>
#include <sys/socket.h>
#include <sys/time.h>

main(argc, argv)
    int argc;
    char **argv;
{
    int xdr_rcp();
    int err;

    if (argc < 2) {
        fprintf(stderr, "usage: %s servername\n", argv[0]);
        exit(-1);
    }
    if ((err = callrpctcp(argv[1], RCPPROG, RCPPROC,
        RCPVERS, xdr_rcp, stdin, xdr_void, 0) != 0)) {
        clnt_perrno(err);
        fprintf(stderr, "can't make RPC call\n");
        exit(1);
    }
    exit(0);
}

callrpctcp(host, prognum, procnum, versnum,
    inproc, in, outproc, out)
    char *host, *in, *out;
    xdrproc_t inproc, outproc;
{
    struct sockaddr_in server_addr;
    int socket = RPC_ANYSOCK;
    enum clnt_stat clnt_stat;
    struct hostent *hp;
    register CLIENT *client;
    struct timeval total_timeout;

    if ((hp = gethostbyname(host)) == NULL) {
        fprintf(stderr, "can't get addr for '%s'\n", host);
        return (-1);
    }
    bcopy(hp->h_addr, (caddr_t)&server_addr.sin_addr,
        hp->h_length);
    server_addr.sin_family = AF_INET;
    server_addr.sin_port = 0;
    if ((client = clnttcp_create(&server_addr, prognum,
        versnum, &socket, BUFSIZ, BUFSIZ)) == NULL) {
        perror("rpctcp_create");
        return (-1);
    }
    total_timeout.tv_sec = 20;
    total_timeout.tv_usec = 0;
    clnt_stat = clnt_call(client, procnum,
        inproc, in, outproc, out, total_timeout);

```

```
    clnt_destroy(client);
    return (int)clnt_stat;
}

/*
 * The receiving routines
 */
#include <stdio.h>
#include <rpc/rpc.h>

main()
{
    register SVCXPRT *transp;
    int rcp_service(), xdr_rcp();

    if ((transp = svctcp_create(RPC_ANYSOCK,
        BUFSIZ, BUFSIZ)) == NULL) {
        fprintf("svctcp_create: error\n");
        exit(1);
    }
    pmap_unset(RCPPROG, RCPVERS);
    if (!svc_register(transp,
        RCPPROG, RCPVERS, rcp_service, IPPROTO_TCP)) {
        fprintf(stderr, "svc_register: error\n");
        exit(1);
    }
    svc_run(); /* never returns */
    fprintf(stderr, "svc_run should never return\n");
}

rcp_service(rqstp, transp)
register struct svc_req *rqstp;
register SVCXPRT *transp;
{
    switch (rqstp->rq_proc) {
    case NULLPROC:
        if (svc_sendreply(transp, xdr_void, 0) == 0) {
            fprintf(stderr, "err: rcp_service");
            return (1);
        }
        return;
    case RCPPROC_FP:
        if (!svc_getargs(transp, xdr_rcp, stdout)) {
            svcerr_decode(transp);
            return;
        }
        if (!svc_sendreply(transp, xdr_void, 0)) {
            fprintf(stderr, "can't reply\n");
            return;
        }
        return (0);
    default:
        svcerr_noproc(transp);
        return;
    }
}
```

1.69 RPC.guide/Callback Procedures

Callback Procedures

Occasionally, it is useful to have a server become a client, and make an RPC call back to the process which is its client. An example is remote debugging, where the client is a window system program, and the server is a debugger running on the remote machine. Most of the time, the user clicks a mouse button at the debugging window, which converts this to a debugger command, and then makes an RPC call to the server (where the debugger is actually running), telling it to execute that command. However, when the debugger hits a breakpoint, the roles are reversed, and the debugger wants to make an rpc call to the window program, so that it can inform the user that a breakpoint has been reached.

In order to do an RPC callback, you need a program number to make the RPC call on. Since this will be a dynamically generated program number, it should be in the transient range, '0x40000000 - 0x5fffffff'. The routine 'gettransient()' returns a valid program number in the transient range, and registers it with the portmapper. It only talks to the portmapper running on the same machine as the 'gettransient()' routine itself. The call to 'pmap_set()' is a test and set operation, in that it indivisibly tests whether a program number has already been registered, and if it has not, then reserves it. On return, the 'sockp' argument will contain a socket that can be used as the argument to an 'svcdp_create()' or 'svctcp_create()' call.

```
#include <stdio.h>
#include <rpc/rpc.h>
#include <sys/socket.h>

gettransient(proto, vers, sockp)
    int proto, vers, *sockp;
{
    static int prognum = 0x40000000;
    int s, len, socktype;
    struct sockaddr_in addr;

    switch(proto) {
        case IPPROTO_UDP:
            socktype = SOCK_DGRAM;
            break;
        case IPPROTO_TCP:
            socktype = SOCK_STREAM;
            break;
        default:
            fprintf(stderr, "unknown protocol type\n");
            return 0;
    }
    if (*sockp == RPC_ANYSOCK) {
        if ((s = socket(AF_INET, socktype, 0)) < 0) {
```

```

        perror("socket");
        return (0);
    }
    *sockp = s;
}
else
    s = *sockp;
addr.sin_addr.s_addr = 0;
addr.sin_family = AF_INET;
addr.sin_port = 0;
len = sizeof(addr);
/*
 * may be already bound, so don't check for error
 */
bind(s, &addr, len);
if (getsockname(s, &addr, &len) < 0) {
    perror("getsockname");
    return (0);
}
while (!pmap_set(prognum++, vers, proto,
    ntohs(addr.sin_port))) continue;
return (prognum-1);
}

```

Note: *The call to 'ntohs()' is necessary to ensure that the port number in 'addr.sin_port', which is in *network* byte order, is passed in *host* byte order (as 'pmap_set()' expects). See the 'byteorder(3N)' man page for more details on the conversion of network addresses from network to host byte order.*

The following pair of programs illustrate how to use the 'gettransient()' routine. The client makes an RPC call to the server, passing it a transient program number. Then the client waits around to receive a callback from the server at that program number. The server registers the program 'EXAMPLEPROG' so that it can receive the RPC call informing it of the callback program number. Then at some random time (on receiving an 'ALRM' signal in this example), it sends a callback RPC call, using the program number it received earlier.

```

/*
 * client
 */
#include <stdio.h>
#include <rpc/rpc.h>

int callback();
char hostname[256];

main()
{
    int x, ans, s;
    SVCXPRT *xpvt;

    gethostname(hostname, sizeof(hostname));
    s = RPC_ANYSOCK;
    x = gettransient(IPPROTO_UDP, 1, &s);
    fprintf(stderr, "client gets prognum %d\n", x);
}

```

```

    if ((xprt = svcudp_create(s)) == NULL) {
        fprintf(stderr, "rpc_server: svcudp_create\n");
        exit(1);
    }
    /* protocol is 0 - gettransient does registering
    */
    (void)svc_register(xprt, x, 1, callback, 0);
    ans = callrpc(hostname, EXAMPLEPROG, EXAMPLEVERS,
        EXAMPLEPROC_CALLBACK, xdr_int, &x, xdr_void, 0);
    if ((enum clnt_stat) ans != RPC_SUCCESS) {
        fprintf(stderr, "call: ");
        clnt_perrno(ans);
        fprintf(stderr, "\n");
    }
    svc_run();
    fprintf(stderr, "Error: svc_run shouldn't return\n");
}

callback(rqstp, transp)
register struct svc_req *rqstp;
register SVCXPRT *transp;
{
    switch (rqstp->rq_proc) {
        case 0:
            if (!svc_sendreply(transp, xdr_void, 0)) {
                fprintf(stderr, "err: exampleprog\n");
                return (1);
            }
            return (0);
        case 1:
            if (!svc_getargs(transp, xdr_void, 0)) {
                svcerr_decode(transp);
                return (1);
            }
            fprintf(stderr, "client got callback\n");
            if (!svc_sendreply(transp, xdr_void, 0)) {
                fprintf(stderr, "err: exampleprog");
                return (1);
            }
        }
    }
}

/*
 * server
 */
#include <stdio.h>
#include <rpc/rpc.h>
#include <sys/signal.h>

char *getnewprog();
char hostname[256];
int docallback();
int pnum;          /* program number for callback routine */

main()
{
    gethostname(hostname, sizeof(hostname));

```

```
    registerrpc(EXAMPLEPROG, EXAMPLEVERS,
        EXAMPLEPROC_CALLBACK, getnewprog, xdr_int, xdr_void);
    fprintf(stderr, "server going into svc_run\n");
    signal(SIGALRM, docallback);
    alarm(10);
    svc_run();
    fprintf(stderr, "Error: svc_run shouldn't return\n");
}

char *
getnewprog(pnum)
    char *pnum;
{
    pnum = *(int *)pnum;
    return NULL;
}

docallback()
{
    int ans;

    ans = callrpc(hostname, pnum, 1, 1, xdr_void, 0,
        xdr_void, 0);
    if (ans != 0) {
        fprintf(stderr, "server: ");
        clnt_perrno(ans);
        fprintf(stderr, "\n");
    }
}
```

1.70 RPC.guide/XDR Technical Notes

XDR Technical Notes

This chapter contains technical notes on Sun's implementation of the External Data Representation (XDR) standard, a set of library routines that allow a C programmer to describe arbitrary data structures in a machine-independent fashion. For a formal specification of the XDR standard, See XDR Protocol Specification.

- Overview
- Justification
- A Canonical Standard
- The XDR Library
- XDR Library Primitives
- Advanced Topics

1.71 RPC.guide/Overview

Overview

=====

This chapter contains technical notes on Sun's implementation of the External Data Representation ("XDR") standard, a set of library routines that allow a C programmer to describe arbitrary data structures in a machine-independent fashion. For a formal specification of the XDR standard, See XDR Protocol Specification. XDR is the backbone of Sun's Remote Procedure Call package, in the sense that data for remote procedure calls is transmitted using the standard. XDR library routines should be used to transmit data that is accessed (read or written) by more than one type of machine (1).

This chapter contains a short tutorial overview of the XDR library routines, a guide to accessing currently available XDR streams, and information on defining new streams and data types. XDR was designed to work across different languages, operating systems, and machine architectures. Most users (particularly RPC users) will only need the information in the Number Filters, Floating Point Filters, and Enumeration Filters. Programmers wishing to implement RPC and XDR on new machines will be interested in the rest of the chapter, as well as the XDR Protocol Specification, which will be their primary reference.

Note: * 'rpcgen' can be used to write XDR routines even in cases where no RPC calls are being made.*

On Sun systems, C programs that want to use XDR routines must include the file '<rpc/rpc.h>', which contains all the necessary interfaces to the XDR system. Since the C library 'libc.a' contains all the XDR routines, compile as normal:

```
example% cc program.c
```

```
----- Footnotes -----
```

(1) For a complete specification of the system External Data Representation routines, see the 'xdr(3N)' manual page.

1.72 RPC.guide/Justification

Justification

=====

Consider the following two programs, 'writer':

```
#include <stdio.h>

main()                /* writer.c */
{
    long i;

    for (i = 0; i < 8; i++) {
        if (fwrite((char *)&i, sizeof(i), 1, stdout) != 1) {
```

```

        fprintf(stderr, "failed!\n");
        exit(1);
    }
}
exit(0);
}

```

and 'reader':

```

#include <stdio.h>

main()                /* reader.c */
{
    long i, j;

    for (j = 0; j < 8; j++) {
        if (fread((char *)&i, sizeof (i), 1, stdin) != 1) {
            fprintf(stderr, "failed!\n");
            exit(1);
        }
        printf("%ld ", i);
    }
    printf("\n");
    exit(0);
}

```

The two programs appear to be portable, because

- a. they pass 'lint' checking, and
- b. they exhibit the same behavior when executed on two different hardware architectures, a Sun and a VAX.

Piping the output of the 'writer' program to the 'reader' program gives identical results on a Sun or a VAX.

```

sun% writer | reader
0 1 2 3 4 5 6 7
sun%

```

```

vax% writer | reader
0 1 2 3 4 5 6 7
vax%

```

With the advent of local area networks and 4.2BSD came the concept of "network pipes" -- a process produces data on one machine, and a second process consumes data on another machine. A network pipe can be constructed with 'writer' and 'reader'. Here are the results if the first produces data on a Sun, and the second consumes data on a VAX.

```

sun% writer | rsh vax reader
0 16777216 33554432 50331648 67108864 83886080 100663296 117440512
sun%

```

Identical results can be obtained by executing 'writer' on the VAX and 'reader' on the Sun. These results occur because the byte ordering of long integers differs between the VAX and the Sun, even though word

size is the same. Note that 16777216 is 2^{24} -- when four bytes are reversed, the 1 winds up in the 24th bit.

Whenever data is shared by two or more machine types, there is a need for portable data. Programs can be made data-portable by replacing the 'read()' and 'write()' calls with calls to an XDR library routine 'xdr_long()', a filter that knows the standard representation of a long integer in its external form. Here are the revised versions of 'writer':

```
#include <stdio.h>
#include <rpc/rpc.h>    /* xdr is a sub-library of rpc */

main()                /* writer.c */
{
    XDR xdrs;
    long i;

    xdrstdio_create(&xdrs, stdout, XDR_ENCODE);
    for (i = 0; i < 8; i++) {
        if (!xdr_long(&xdrs, &i)) {
            fprintf(stderr, "failed!\n");
            exit(1);
        }
    }
    exit(0);
}
```

and 'reader':

```
#include <stdio.h>
#include <rpc/rpc.h>    /* xdr is a sub-library of rpc */

main()                /* reader.c */
{
    XDR xdrs;
    long i, j;

    xdrstdio_create(&xdrs, stdin, XDR_DECODE);
    for (j = 0; j < 8; j++) {
        if (!xdr_long(&xdrs, &i)) {
            fprintf(stderr, "failed!\n");
            exit(1);
        }
        printf("%ld ", i);
    }
    printf("\n");
    exit(0);
}
```

The new programs were executed on a Sun, on a VAX, and from a Sun to a VAX; the results are shown below.

```
sun% writer | reader
0 1 2 3 4 5 6 7
sun%
```

```

vax% writer | reader
0 1 2 3 4 5 6 7
vax%

sun% writer | rsh vax reader
0 1 2 3 4 5 6 7
sun%

```

Note: *Integers are just the tip of the portable-data iceberg. Arbitrary data structures present portability problems, particularly with respect to alignment and pointers. Alignment on word boundaries may cause the size of a structure to vary from machine to machine. And pointers, which are very convenient to use, have no meaning outside the machine where they are defined.*

1.73 RPC.guide/A Canonical Standard

A Canonical Standard

=====

XDR's approach to standardizing data representations is *canonical*. That is, XDR defines a single byte order (Big Endian), a single floating-point representation (IEEE), and so on. Any program running on any machine can use XDR to create portable data by translating its local representation to the XDR standard representations; similarly, any program running on any machine can read portable data by translating the XDR standard representations to its local equivalents. The single standard completely decouples programs that create or send portable data from those that use or receive portable data. The advent of a new machine or a new language has no effect upon the community of existing portable data creators and users. A new machine joins this community by being "taught" how to convert the standard representations and its local representations; the local representations of other machines are irrelevant. Conversely, to existing programs running on other machines, the local representations of the new machine are also irrelevant; such programs can immediately read portable data produced by the new machine because such data conforms to the canonical standards that they already understand.

There are strong precedents for XDR's canonical approach. For example, TCP/IP, UDP/IP, XNS, Ethernet, and, indeed, all protocols below layer five of the ISO model, are canonical protocols. The advantage of any canonical approach is simplicity; in the case of XDR, a single set of conversion routines is written once and is never touched again. The canonical approach has a disadvantage, but it is unimportant in real-world data transfer applications. Suppose two Little-Endian machines are transferring integers according to the XDR standard. The sending machine converts the integers from Little-Endian byte order to XDR (Big-Endian) byte order; the receiving machine performs the reverse conversion. Because both machines observe the same byte order, their conversions are unnecessary. The point, however, is not necessity, but cost as compared to the alternative.

The time spent converting to and from a canonical representation is

insignificant, especially in networking applications. Most of the time required to prepare a data structure for transfer is not spent in conversion but in traversing the elements of the data structure. To transmit a tree, for example, each leaf must be visited and each element in a leaf record must be copied to a buffer and aligned there; storage for the leaf may have to be deallocated as well. Similarly, to receive a tree, storage must be allocated for each leaf, data must be moved from the buffer to the leaf and properly aligned, and pointers must be constructed to link the leaves together. Every machine pays the cost of traversing and copying data structures whether or not conversion is required. In networking applications, communications overhead -- the time required to move the data down through the sender's protocol layers, across the network and up through the receiver's protocol layers -- dwarfs conversion overhead.

1.74 RPC.guide/The XDR Library

The XDR Library

=====

The XDR library not only solves data portability problems, it also allows you to write and read arbitrary C constructs in a consistent, specified, well-documented manner. Thus, it can make sense to use the library even when the data is not shared among machines on a network.

The XDR library has filter routines for strings (null-terminated arrays of bytes), structures, unions, and arrays, to name a few. Using more primitive routines, you can write your own specific XDR routines to describe arbitrary data structures, including elements of arrays, arms of unions, or objects pointed at from other structures. The structures themselves may contain arrays of arbitrary elements, or pointers to other structures.

Let's examine the two programs more closely. There is a family of XDR stream creation routines in which each member treats the stream of bits differently. In our example, data is manipulated using standard I/O routines, so we use `'xdrstdio_create()'`.

The parameters to XDR stream creation routines vary according to their function. In our example, `'xdrstdio_create()'` takes a pointer to an XDR structure that it initializes, a pointer to a `'FILE'` that the input or output is performed on, and the operation. The operation may be `'XDR_ENCODE'` for serializing in the `'writer'` program, or `'XDR_DECODE'` for deserializing in the `'reader'` program.

Note: RPC users never need to create XDR streams; the RPC system itself creates these streams, which are then passed to the users.

The `'xdr_long()'` primitive is characteristic of most XDR library primitives and all client XDR routines. First, the routine returns `'FALSE'` (`'0'`) if it fails, and `'TRUE'` (`'1'`) if it succeeds. Second, for each data type, `'xxx'`, there is an associated XDR routine of the form:

```
xdr_xxx(xdrs, xp)
    XDR *xdrs;
    xxx *xp;
{
}
```

In our case, 'xxx' is long, and the corresponding XDR routine is a primitive, 'xdr_long()'. The client could also define an arbitrary structure 'xxx' in which case the client would also supply the routine 'xdr_xxx()', describing each field by calling XDR routines of the appropriate type. In all cases the first parameter, 'xdrs' can be treated as an opaque handle, and passed to the primitive routines.

XDR routines are direction independent; that is, the same routines are called to serialize or deserialize data. This feature is critical to software engineering of portable data. The idea is to call the same routine for either operation -- this almost guarantees that serialized data can also be deserialized. One routine is used by both producer and consumer of networked data. This is implemented by always passing the address of an object rather than the object itself -- only in the case of deserialization is the object modified. This feature is not shown in our trivial example, but its value becomes obvious when nontrivial data structures are passed among machines. If needed, the user can obtain the direction of the XDR operation. See the XDR Operation Directions for details.

Let's look at a slightly more complicated example. Assume that a person's gross assets and liabilities are to be exchanged among processes. Also assume that these values are important enough to warrant their own data type:

```
struct gnumbers {
    long g_assets;
    long g_liabilities;
};
```

The corresponding XDR routine describing this structure would be:

```
bool_t          /* TRUE is success, FALSE is failure */
xdr_gnumbers(xdrs, gp)
    XDR *xdrs;
    struct gnumbers *gp;
{
    if (xdr_long(xdrs, &gp->g_assets) &&
        xdr_long(xdrs, &gp->g_liabilities))
        return(TRUE);
    return(FALSE);
}
```

Note that the parameter 'xdrs' is never inspected or modified; it is only passed on to the subcomponent routines. It is imperative to inspect the return value of each XDR routine call, and to give up immediately and return 'FALSE' if the subroutine fails.

This example also shows that the type 'bool_t' is declared as an integer whose only values are 'TRUE' (1) and 'FALSE' (0). This document uses the following definitions:

```
#define bool_t  int
#define TRUE    1
#define FALSE   0
```

Keeping these conventions in mind, 'xdr_gnumbers()' can be rewritten as follows:

```
xdr_gnumbers(xdrs, gp)
    XDR *xdrs;
    struct gnumbers *gp;
{
    return(xdr_long(xdrs, &gp->g_assets) &&
           xdr_long(xdrs, &gp->g_liabilities));
}
```

This document uses both coding styles.

1.75 RPC.guide/XDR Library Primitives

XDR Library Primitives

=====

This section gives a synopsis of each XDR primitive. It starts with basic data types and moves on to constructed data types. Finally, XDR utilities are discussed. The interface to these primitives and utilities is defined in the include file '<rpc/xdr.h>', automatically included by '<rpc/rpc.h>'.

- Number Filters
- Floating Point Filters
- Enumeration Filters
- No Data
- Constructed Data Type Filters
- Non-filter Primitives
- XDR Operation Directions
- XDR Stream Access
- XDR Stream Implementation

1.76 RPC.guide/Number Filters

Number Filters

The XDR library provides primitives to translate between numbers and their corresponding external representations. Primitives cover the set of numbers in:

```
[signed, unsigned] * [short, int, long]
```

Specifically, the eight primitives are:

```
bool_t xdr_char(xdrs, cp)
    XDR *xdrs;
    char *cp;

bool_t xdr_u_char(xdrs, ucp)
    XDR *xdrs;
    unsigned char *ucp;

bool_t xdr_int(xdrs, ip)
    XDR *xdrs;
    int *ip;

bool_t xdr_u_int(xdrs, up)
    XDR *xdrs;
    unsigned *up;

bool_t xdr_long(xdrs, lip)
    XDR *xdrs;
    long *lip;

bool_t xdr_u_long(xdrs, lup)
    XDR *xdrs;
    u_long *lup;

bool_t xdr_short(xdrs, sip)
    XDR *xdrs;
    short *sip;

bool_t xdr_u_short(xdrs, sup)
    XDR *xdrs;
    u_short *sup;
```

The first parameter, 'xdrs', is an XDR stream handle. The second parameter is the address of the number that provides data to the stream or receives data from it. All routines return 'TRUE' if they complete successfully, and 'FALSE' otherwise.

1.77 RPC.guide/Floating Point Filters

Floating Point Filters

The XDR library also provides primitive routines for C's floating point types:

```
bool_t xdr_float(xdrs, fp)
    XDR *xdrs;
    float *fp;

bool_t xdr_double(xdrs, dp)
    XDR *xdrs;
```

```
double *dp;
```

The first parameter, 'xdrs' is an XDR stream handle. The second parameter is the address of the floating point number that provides data to the stream or receives data from it. Both routines return 'TRUE' if they complete successfully, and 'FALSE' otherwise.

Note: Since the numbers are represented in IEEE floating point, routines may fail when decoding a valid IEEE representation into a machine-specific representation, or vice-versa.

1.78 RPC.guide/Enumeration Filters

Enumeration Filters

The XDR library provides a primitive for generic enumerations. The primitive assumes that a C 'enum' has the same representation inside the machine as a C integer. The boolean type is an important instance of the 'enum'. The external representation of a boolean is always 'TRUE' (1) or 'FALSE' (0).

```
#define bool_t    int
#define FALSE    0
#define TRUE     1

#define enum_t   int

bool_t xdr_enum(xdrs, ep)
    XDR *xdrs;
    enum_t *ep;

bool_t xdr_bool(xdrs, bp)
    XDR *xdrs;
    bool_t *bp;
```

The second parameters 'ep' and 'bp' are addresses of the associated type that provides data to, or receives data from, the stream 'xdrs'.

1.79 RPC.guide/No Data

No Data

Occasionally, an XDR routine must be supplied to the RPC system, even when no data is passed or required. The library provides such a routine:

```
bool_t xdr_void(); /* always returns TRUE */
```

1.80 RPC.guide/Constructed Data Type Filters

Constructed Data Type Filters

Constructed or compound data type primitives require more parameters and perform more complicated functions than the primitives discussed above. This section includes primitives for strings, arrays, unions, and pointers to structures.

Constructed data type primitives may use memory management. In many cases, memory is allocated when deserializing data with 'XDR_DECODE'. Therefore, the XDR package must provide means to deallocate memory. This is done by an XDR operation, 'XDR_FREE'. To review, the three XDR directional operations are 'XDR_ENCODE', 'XDR_DECODE' and 'XDR_FREE'.

- Strings
- Byte Arrays
- Arrays
- XDR Examples
- Opaque Data
- Fixed Sized Arrays
- Discriminated Unions
- Pointers

1.81 RPC.guide/Strings

Strings
.....

In C, a string is defined as a sequence of bytes terminated by a null byte, which is not considered when calculating string length. However, when a string is passed or manipulated, a pointer to it is employed. Therefore, the XDR library defines a string to be a 'char *' and not a sequence of characters. The external representation of a string is drastically different from its internal representation. Externally, strings are represented as sequences of ASCII characters, while internally, they are represented with character pointers. Conversion between the two representations is accomplished with the routine 'xdr_string()':

```
bool_t xdr_string(xdrs, sp, maxlength)
    XDR *xdrs;
    char **sp;
    u_int maxlength;
```

The first parameter 'xdrs' is the XDR stream handle. The second parameter 'sp' is a pointer to a string (type 'char **'). The third parameter 'maxlength' specifies the maximum number of bytes allowed during encoding or decoding. Its value is usually specified by a protocol. For example, a protocol specification may say that a file name may be no longer than 255 characters.

The routine returns 'FALSE' if the number of characters exceeds 'maxlength', and 'TRUE' if it doesn't. *Keep* 'maxlength' small. If it is too big you can blow the heap, since 'xdr_string()' will call 'malloc()' for space.

The behavior of 'xdr_string()' is similar to the behavior of other routines discussed in this section. The direction 'XDR_ENCODE' is easiest to understand. The parameter 'sp' points to a string of a certain length; if the string does not exceed 'maxlength', the bytes are serialized.

The effect of deserializing a string is subtle. First the length of the incoming string is determined; it must not exceed 'maxlength'. Next 'sp' is dereferenced; if the the value is 'NULL', then a string of the appropriate length is allocated and '*sp' is set to this string. If the original value of '*sp' is non-null, then the XDR package assumes that a target area has been allocated, which can hold strings no longer than 'maxlength'. In either case, the string is decoded into the target area. The routine then appends a null character to the string.

In the 'XDR_FREE' operation, the string is obtained by dereferencing 'sp'. If the string is not 'NULL', it is freed and '*sp' is set to 'NULL'. In this operation, 'xdr_string()' ignores the 'maxlength' parameter.

1.82 RPC.guide/Byte Arrays

Byte Arrays

.....

Often variable-length arrays of bytes are preferable to strings. Byte arrays differ from strings in the following three ways:

1. the length of the array (the byte count) is explicitly located in an unsigned integer,
2. the byte sequence is not terminated by a null character, and
3. the external representation of the bytes is the same as their internal representation.

The primitive 'xdr_bytes()' converts between the internal and external representations of byte arrays:

```
bool_t xdr_bytes(xdrs, bpp, lp, maxlength)
    XDR *xdrs;
    char **bpp;
    u_int *lp;
    u_int maxlength;
```

The usage of the first, second and fourth parameters are identical to the first, second and third parameters of 'xdr_string()', respectively.

The length of the byte area is obtained by dereferencing 'lp' when serializing; '*lp' is set to the byte length when deserializing.

1.83 RPC.guide/Arrays

Arrays

.....

The XDR library package provides a primitive for handling arrays of arbitrary elements. The 'xdr_bytes()' routine treats a subset of generic arrays, in which the size of array elements is known to be 1, and the external description of each element is built-in. The generic array primitive, 'xdr_array()', requires parameters identical to those of 'xdr_bytes()' plus two more: the size of array elements, and an XDR routine to handle each of the elements. This routine is called to encode or decode each element of the array.

```
bool_t
xdr_array(xdrs, ap, lp, maxlength, elementsiz, xdr_element)
    XDR *xdrs;
    char **ap;
    u_int *lp;
    u_int maxlength;
    u_int elementsiz;
    bool_t (*xdr_element)();
```

The parameter 'ap' is the address of the pointer to the array. If '*ap' is 'NULL' when the array is being deserialized, XDR allocates an array of the appropriate size and sets '*ap' to that array. The element count of the array is obtained from '*lp' when the array is serialized; '*lp' is set to the array length when the array is deserialized. The parameter 'maxlength' is the maximum number of elements that the array is allowed to have; 'elementsiz' is the byte size of each element of the array (the C function 'sizeof()' can be used to obtain this value). The 'xdr_element()' routine is called to serialize, deserialize, or free each element of the array.

1.84 RPC.guide/XDR Examples

XDR Examples

.....

Before defining more constructed data types, it is appropriate to present three examples.

Example A: Network User

.....

A user on a networked machine can be identified by

- a. the machine name, such as 'krypton': see the 'gethostname()' man page;
- b. the user's UID: see the 'geteuid()' man page; and
- c. the group numbers to which the user belongs: see the 'getgroups()' man page.

A structure with this information and its associated XDR routine could be coded like this:

```

struct netuser {
    char    *nu_machinename;
    int     nu_uid;
    u_int   nu_glen;
    int     *nu_gids;
};
#define NLEN 255      /* machine names < 256 chars */
#define NGRPS 20     /* user can't be in > 20 groups */

bool_t
xdr_netuser(xdrs, nup)
    XDR *xdrs;
    struct netuser *nup;
{
    return(xdr_string(xdrs, &nup->nu_machinename, NLEN) &&
           xdr_int(xdrs, &nup->nu_uid) &&
           xdr_array(xdrs, &nup->nu_gids, &nup->nu_glen,
                     NGRPS, sizeof (int), xdr_int));
}

```

Example B: Party of Network Users

.....

A party of network users could be implemented as an array of 'netuser' structure. The declaration and its associated XDR routines are as follows:

```

struct party {
    u_int p_len;
    struct netuser *p_nusers;
};
#define PLEN 500     /* max number of users in a party */

bool_t
xdr_party(xdrs, pp)
    XDR *xdrs;
    struct party *pp;
{
    return(xdr_array(xdrs, &pp->p_nusers, &pp->p_len, PLEN,
                    sizeof (struct netuser), xdr_netuser));
}

```

Example C: History of Commands

.....

The well-known parameters to 'main()', 'argc' and 'argv' can be

combined into a structure. An array of these structures can make up a history of commands. The declarations and XDR routines might look like:

```

struct cmd {
    u_int c_argc;
    char **c_argv;
};
#define ALEN 1000 /* args cannot be > 1000 chars */
#define NARGC 100 /* commands cannot have > 100 args */

struct history {
    u_int h_len;
    struct cmd *h_cmds;
};
#define NCMDS 75 /* history is no more than 75 commands */

bool_t
xdr_wrap_string(xdrs, sp)
    XDR *xdrs;
    char **sp;
{
    return(xdr_string(xdrs, sp, ALEN));
}

bool_t
xdr_cmd(xdrs, cp)
    XDR *xdrs;
    struct cmd *cp;
{
    return(xdr_array(xdrs, &cp->c_argv, &cp->c_argc, NARGC,
        sizeof(char *), xdr_wrap_string));
}

bool_t
xdr_history(xdrs, hp)
    XDR *xdrs;
    struct history *hp;
{
    return(xdr_array(xdrs, &hp->h_cmds, &hp->h_len, NCMDS,
        sizeof(struct cmd), xdr_cmd));
}

```

The most confusing part of this example is that the routine `'xdr_wrap_string()'` is needed to package the `'xdr_string()'` routine, because the implementation of `'xdr_array()'` only passes two parameters to the array element description routine; `'xdr_wrap_string()'` supplies the third parameter to `'xdr_string()'`.

By now the recursive nature of the XDR library should be obvious. Let's continue with more constructed data types.

1.85 RPC.guide/Opaque Data

Opaque Data

In some protocols, handles are passed from a server to client. The client passes the handle back to the server at some later time. Handles are never inspected by clients; they are obtained and submitted. That is to say, handles are opaque. The 'xdr_opaque()' primitive is used for describing fixed sized, opaque bytes.

```
bool_t xdr_opaque(xdrs, p, len)
    XDR *xdrs;
    char *p;
    u_int len;
```

The parameter 'p' is the location of the bytes; 'len' is the number of bytes in the opaque object. By definition, the actual data contained in the opaque object are not machine portable.

1.86 RPC.guide/Fixed Sized Arrays

Fixed Sized Arrays

The XDR library provides a primitive, 'xdr_vector()', for fixed-length arrays.

```
#define NLEN 255      /* machine names must be < 256 chars */
#define NGRPS 20     /* user belongs to exactly 20 groups */

struct netuser {
    char *nu_machinename;
    int nu_uid;
    int nu_gids[NGRPS];
};

bool_t
xdr_netuser(xdrs, nup)
    XDR *xdrs;
    struct netuser *nup;
{
    int i;

    if (!xdr_string(xdrs, &nup->nu_machinename, NLEN))
        return(FALSE);
    if (!xdr_int(xdrs, &nup->nu_uid))
        return(FALSE);
    if (!xdr_vector(xdrs, nup->nu_gids, NGRPS, sizeof(int),
        xdr_int)) {
        return(FALSE);
    }
    return(TRUE);
}
```

1.87 RPC.guide/Discriminated Unions

Discriminated Unions

.....

The XDR library supports discriminated unions. A discriminated union is a C union and an 'enum_t' value that selects an "arm" of the union.

```

struct xdr_discrim {
    enum_t value;
    bool_t (*proc)();
};

bool_t xdr_union(xdrs, dscmp, unp, arms, defaultarm)
    XDR *xdrs;
    enum_t *dscmp;
    char *unp;
    struct xdr_discrim *arms;
    bool_t (*defaultarm)(); /* may equal NULL */

```

First the routine translates the discriminant of the union located at '*dscmp'. The discriminant is always an 'enum_t'. Next the union located at '*unp' is translated. The parameter 'arms' is a pointer to an array of 'xdr_discrim' structures. Each structure contains an ordered pair of *[value,proc]*. If the union's discriminant is equal to the associated *value*, then the *proc* is called to translate the union. The end of the 'xdr_discrim' structure array is denoted by a routine of value 'NULL' ('0'). If the discriminant is not found in the 'arms' array, then the 'defaultarm' procedure is called if it is non-null; otherwise the routine returns 'FALSE'.

Example D: Union

.....

Suppose the type of a union may be integer, character pointer (a string), or a 'gnumbers' structure. Also, assume the union and its current type are declared in a structure. The declaration is:

```

enum utype { INTEGER=1, STRING=2, GNUMBERS=3 };

struct u_tag {
    enum utype utype; /* the union's discriminant */
    union {
        int ival;
        char *pval;
        struct gnumbers gn;
    } uval;
};

```

The following constructs and XDR procedure (de)serialize the discriminated union:

```

struct xdr_discrim u_tag_arms[4] = {
    { INTEGER, xdr_int },
    { GNUMBERS, xdr_gnumbers }

```

```

    { STRING, xdr_wrap_string },
    { __dontcare__, NULL }
    /* always terminate arms with a NULL xdr_proc */
}

bool_t
xdr_u_tag(xdrs, utp)
    XDR *xdrs;
    struct u_tag *utp;
{
    return(xdr_union(xdrs, &utp->utype, &utp->uval,
                    u_tag_arms, NULL));
}

```

The routine `'xdr_gnumbers()'` was presented above in The XDR Library. `'xdr_wrap_string()'` was presented in example C. The default `'arm'` parameter to `'xdr_union()'` (the last parameter) is `'NULL'` in this example. Therefore the value of the union's discriminant may legally take on only values listed in the `'u_tag_arms'` array. This example also demonstrates that the elements of the arm's array do not need to be sorted.

It is worth pointing out that the values of the discriminant may be sparse, though in this example they are not. It is always good practice to assign explicitly integer values to each element of the discriminant's type. This practice both documents the external representation of the discriminant and guarantees that different C compilers emit identical discriminant values.

Exercise: Implement `'xdr_union()'` using the other primitives in this section.

1.88 RPC.guide/Pointers

Pointers

In C it is often convenient to put pointers to another structure within a structure. The `'xdr_reference()'` primitive makes it easy to serialize, deserialize, and free these referenced structures.

```

bool_t xdr_reference(xdrs, pp, size, proc)
    XDR *xdrs;
    char **pp;
    u_int ssize;
    bool_t (*proc)();

```

Parameter `'pp'` is the address of the pointer to the structure; parameter `'ssize'` is the size in bytes of the structure (use the C function `'sizeof()'` to obtain this value); and `'proc'` is the XDR routine that describes the structure. When decoding data, storage is allocated if `'*pp'` is `'NULL'`.

There is no need for a primitive `'xdr_struct()'` to describe

structures within structures, because pointers are always sufficient.

Exercise: Implement `'xdr_reference()'` using `'xdr_array()'`. Warning: `'xdr_reference()'` and `'xdr_array()'` are NOT interchangeable external representations of data.

Example E: Pointers

.....

Suppose there is a structure containing a person's name and a pointer to a 'gnumbers' structure containing the person's gross assets and liabilities. The construct is:

```
struct pgn {
    char *name;
    struct gnumbers *gnp;
};
```

The corresponding XDR routine for this structure is:

```
bool_t
xdr_pgn(xdrs, pp)
    XDR *xdrs;
    struct pgn *pp;
{
    if (xdr_string(xdrs, &pp->name, NLEN) &&
        xdr_reference(xdrs, &pp->gnp,
            sizeof(struct gnumbers), xdr_gnumbers))
        return(TRUE);
    return(FALSE);
}
```

Pointer Semantics and XDR

.....

In many applications, C programmers attach double meaning to the values of a pointer. Typically the value 'NULL' (or zero) means data is not needed, yet some application-specific interpretation applies. In essence, the C programmer is encoding a discriminated union efficiently by overloading the interpretation of the value of a pointer. For instance, in example E a 'NULL' pointer value for 'gnp' could indicate that the person's assets and liabilities are unknown. That is, the pointer value encodes two things: whether or not the data is known; and if it is known, where it is located in memory. Linked lists are an extreme example of the use of application-specific pointer interpretation.

The primitive `'xdr_reference()'` cannot and does not attach any special meaning to a null-value pointer during serialization. That is, passing an address of a pointer whose value is 'NULL' to `'xdr_reference()'` when serializing data will most likely cause a memory fault and, on the UNIX system, a core dump.

`'xdr_pointer()'` correctly handles 'NULL' pointers. For more information about its use, see the Linked Lists.

Exercise:

.....

After reading the Linked Lists, return here and extend example E so that it can correctly deal with 'NULL' pointer values.

Exercise:

.....

Using the 'xdr_union()', 'xdr_reference()' and 'xdr_void()' primitives, implement a generic pointer handling primitive that implicitly deals with 'NULL' pointers. That is, implement 'xdr_pointer()'.

1.89 RPC.guide/Non-filter Primitives

Non-filter Primitives

XDR streams can be manipulated with the primitives discussed in this section.

```
u_int xdr_getpos(xdrs)
    XDR *xdrs;

bool_t xdr_setpos(xdrs, pos)
    XDR *xdrs;
    u_int pos;

xdr_destroy(xdrs)
    XDR *xdrs;
```

The routine 'xdr_getpos()' returns an unsigned integer that describes the current position in the data stream. Warning: In some XDR streams, the returned value of 'xdr_getpos()' is meaningless; the routine returns a -1 in this case (though -1 should be a legitimate value).

The routine 'xdr_setpos()' sets a stream position to 'pos'. Warning: In some XDR streams, setting a position is impossible; in such cases, 'xdr_setpos()' will return 'FALSE'. This routine will also fail if the requested position is out-of-bounds. The definition of bounds varies from stream to stream.

The 'xdr_destroy()' primitive destroys the XDR stream. Usage of the stream after calling this routine is undefined.

1.90 RPC.guide/XDR Operation Directions

XDR Operation Directions

At times you may wish to optimize XDR routines by taking advantage of the direction of the operation -- 'XDR_ENCODE', 'XDR_DECODE' or 'XDR_FREE'. The value 'xdrs->x_op' always contains the direction of the XDR operation. Programmers are not encouraged to take advantage of this information. Therefore, no example is presented here. However, an example in the Linked Lists, demonstrates the usefulness of the 'xdrs->x_op' field.

1.91 RPC.guide/XDR Stream Access

XDR Stream Access

An XDR stream is obtained by calling the appropriate creation routine. These creation routines take arguments that are tailored to the specific properties of the stream.

Streams currently exist for (de)serialization of data to or from standard I/O 'FILE' streams, TCP/IP connections and UNIX files, and memory.

Standard I-O Streams
 Memory Streams
 Record (TCP-IP) Streams

1.92 RPC.guide/Standard I-O Streams

Standard I/O Streams

.....

XDR streams can be interfaced to standard I/O using the 'xdrstdio_create()' routine as follows:

```
#include <stdio.h>
#include <rpc/rpc.h>    /* xdr streams part of rpc */

void
xdrstdio_create(xdrs, fp, x_op)
    XDR *xdrs;
    FILE *fp;
    enum xdr_op x_op;
```

The routine 'xdrstdio_create()' initializes an XDR stream pointed to by 'xdrs'. The XDR stream interfaces to the standard I/O library. Parameter 'fp' is an open file, and 'x_op' is an XDR direction.

1.93 RPC.guide/Memory Streams

Memory Streams

.....

Memory streams allow the streaming of data into or out of a specified area of memory:

```
#include <rpc/rpc.h>

void
xdrmem_create(xdrs, addr, len, x_op)
    XDR *xdrs;
    char *addr;
    u_int len;
    enum xdr_op x_op;
```

The routine `'xdrmem_create()'` initializes an XDR stream in local memory. The memory is pointed to by parameter `'addr'`; parameter `'len'` is the length in bytes of the memory. The parameters `'xdrs'` and `'x_op'` are identical to the corresponding parameters of `'xdrstdio_create()'`. Currently, the UDP/IP implementation of RPC uses `'xdrmem_create()'`. Complete call or result messages are built in memory before calling the `'sendto()'` system routine.

1.94 RPC.guide/Record (TCP-IP) Streams

Record (TCP/IP) Streams

.....

A record stream is an XDR stream built on top of a record marking standard that is built on top of the UNIX file or 4.2 BSD connection interface.

```
#include <rpc/rpc.h> /* xdr streams part of rpc */

xdrrec_create(xdrs,
    sendsize, recvsize, iohandle, readproc, writeproc)
    XDR *xdrs;
    u_int sendsize, recvsize;
    char *iohandle;
    int (*readproc)(), (*writeproc)();
```

The routine `'xdrrec_create()'` provides an XDR stream interface that allows for a bidirectional, arbitrarily long sequence of records. The contents of the records are meant to be data in XDR form. The stream's primary use is for interfacing RPC to TCP connections. However, it can be used to stream data into or out of normal UNIX files.

The parameter `'xdrs'` is similar to the corresponding parameter described above. The stream does its own data buffering similar to that of standard I/O. The parameters `'sendsize'` and `'recvsize'` determine the size in bytes of the output and input buffers, respectively; if

their values are zero ('0'), then predetermined defaults are used. When a buffer needs to be filled or flushed, the routine 'readproc()' or 'writeproc()' is called, respectively. The usage and behavior of these routines are similar to the UNIX system calls 'read()' and 'write()'. However, the first parameter to each of these routines is the opaque parameter 'iohandle'. The other two parameters ('buf' and 'nbytes') and the results (byte count) are identical to the system routines. If 'xxx' is 'readproc()' or 'writeproc()', then it has the following form:

```
/*
 * returns the actual number of bytes transferred.
 * -1 is an error
 */
int
xxx(iohandle, buf, len)
    char *iohandle;
    char *buf;
    int nbytes;
```

The XDR stream provides means for delimiting records in the byte stream. The implementation details of delimiting records in a stream are discussed in the Advanced Topics. The primitives that are specific to record streams are as follows:

```
bool_t
xdrrec_endofrecord(xdrs, flushnow)
    XDR *xdrs;
    bool_t flushnow;

bool_t
xdrrec_skiprecord(xdrs)
    XDR *xdrs;

bool_t
xdrrec_eof(xdrs)
    XDR *xdrs;
```

The routine 'xdrrec_endofrecord()' causes the current outgoing data to be marked as a record. If the parameter 'flushnow' is 'TRUE', then the stream's 'writeproc' will be called; otherwise, 'writeproc' will be called when the output buffer has been filled.

The routine 'xdrrec_skiprecord()' causes an input stream's position to be moved past the current record boundary and onto the beginning of the next record in the stream.

If there is no more data in the stream's input buffer, then the routine 'xdrrec_eof()' returns 'TRUE'. That is not to say that there is no more data in the underlying file descriptor.

1.95 RPC.guide/XDR Stream Implementation

XDR Stream Implementation

This section provides the abstract data types needed to implement new instances of XDR streams.

The XDR Object

1.96 RPC.guide/The XDR Object

The XDR Object

.....

The following structure defines the interface to an XDR stream:

```
enum xdr_op { XDR_ENCODE=0, XDR_DECODE=1, XDR_FREE=2 };

typedef struct {
    enum xdr_op x_op;          /* operation; fast added param */
    struct xdr_ops {
        bool_t (*x_getlong)(); /* get long from stream */
        bool_t (*x_putlong)(); /* put long to stream */
        bool_t (*x_getbytes)(); /* get bytes from stream */
        bool_t (*x_putbytes)(); /* put bytes to stream */
        u_int (*x_getpostn)(); /* return stream offset */
        bool_t (*x_setpostn)(); /* reposition offset */
        caddr_t (*x_inline)(); /* ptr to buffered data */
        VOID (*x_destroy)(); /* free private area */
    } *x_ops;
    caddr_t x_public;         /* users' data */
    caddr_t x_private;        /* pointer to private data */
    caddr_t x_base;           /* private for position info */
    int x_handy;             /* extra private word */
} XDR;
```

The 'x_op' field is the current operation being performed on the stream. This field is important to the XDR primitives, but should not affect a stream's implementation. That is, a stream's implementation should not depend on this value. The fields 'x_private', 'x_base', and 'x_handy' are private to the particular stream's implementation. The field 'x_public' is for the XDR client and should never be used by the XDR stream implementations or the XDR primitives. 'x_getpostn()', 'x_setpostn()' and 'x_destroy()' are macros for accessing operations. The operation 'x_inline()' takes two parameters: an 'XDR *', and an unsigned integer, which is a byte count. The routine returns a pointer to a piece of the stream's internal buffer. The caller can then use the buffer segment for any purpose. From the stream's point of view, the bytes in the buffer segment have been consumed or put. The routine may return 'NULL' if it cannot return a buffer segment of the requested size. (The 'x_inline()' routine is for cycle squeezers. Use of the resulting buffer is not data-portable. Users are encouraged not to use this feature.)

The operations `'x_getbytes()'` and `'x_putbytes()'` blindly get and put sequences of bytes from or to the underlying stream; they return `'TRUE'` if they are successful, and `'FALSE'` otherwise. The routines have identical parameters (replace `'xxx'`):

```
bool_t
xxxbytes(xdrs, buf, bytecount)
    XDR *xdrs;
    char *buf;
    u_int bytecount;
```

The operations `'x_getlong()'` and `'x_putlong()'` receive and put long numbers from and to the data stream. It is the responsibility of these routines to translate the numbers between the machine representation and the (standard) external representation. The UNIX primitives `'htonl()'` and `'ntohl()'` can be helpful in accomplishing this. The higher-level XDR implementation assumes that signed and unsigned long integers contain the same number of bits, and that nonnegative integers have the same bit representations as unsigned integers. The routines return `'TRUE'` if they succeed, and `'FALSE'` otherwise. They have identical parameters:

```
bool_t
xxxlong(xdrs, lp)
    XDR *xdrs;
    long *lp;
```

Implementors of new XDR streams must make an XDR structure (with new operation routines) available to clients, using some kind of create routine.

1.97 RPC.guide/Advanced Topics

Advanced Topics

=====

This section describes techniques for passing data structures that are not covered in the preceding sections. Such structures include linked lists (of arbitrary lengths). Unlike the simpler examples covered in the earlier sections, the following examples are written using both the XDR C library routines and the XDR data description language. The XDR Protocol Specification describes this language in complete detail.

Linked Lists

1.98 RPC.guide/Linked Lists

Linked Lists

The last example in the Pointers topic earlier in this chapter presented a C data structure and its associated XDR routines for a individual's gross assets and liabilities. The example is duplicated below:

```

struct gnumbers {
    long g_assets;
    long g_liabilities;
};

bool_t
xdr_gnumbers(xdrs, gp)
    XDR *xdrs;
    struct gnumbers *gp;
{
    if (xdr_long(xdrs, &(gp->g_assets)))
        return(xdr_long(xdrs, &(gp->g_liabilities)));
    return(FALSE);
}

```

Now assume that we wish to implement a linked list of such information. A data structure could be constructed as follows:

```

struct gnumbers_node {
    struct gnumbers gn_numbers;
    struct gnumbers_node *gn_next;
};

typedef struct gnumbers_node *gnumbers_list;

```

The head of the linked list can be thought of as the data object; that is, the head is not merely a convenient shorthand for a structure. Similarly the 'gn_next' field is used to indicate whether or not the object has terminated. Unfortunately, if the object continues, the 'gn_next' field is also the address of where it continues. The link addresses carry no useful information when the object is serialized.

The XDR data description of this linked list is described by the recursive declaration of 'gnumbers_list':

```

struct gnumbers {
    int g_assets;
    int g_liabilities;
};

struct gnumbers_node {
    gnumbers gn_numbers;
    gnumbers_node *gn_next;
};

```

In this description, the boolean indicates whether there is more data following it. If the boolean is 'FALSE', then it is the last data field of the structure. If it is 'TRUE', then it is followed by a gnumbers

structure and (recursively) by a 'gnumbers_list'. Note that the C declaration has no boolean explicitly declared in it (though the 'gn_next' field implicitly carries the information), while the XDR data description has no pointer explicitly declared in it.

Hints for writing the XDR routines for a 'gnumbers_list' follow easily from the XDR description above. Note how the primitive 'xdr_pointer()' is used to implement the XDR union above.

```

bool_t
xdr_gnumbers_node(xdrs, gn)
    XDR *xdrs;
    gnumbers_node *gn;
{
    return(xdr_gnumbers(xdrs, &gn->gn_numbers) &&
           xdr_gnumbers_list(xdrs, &gn->gn_next));
}

bool_t
xdr_gnumbers_list(xdrs, gnp)
    XDR *xdrs;
    gnumbers_list *gnp;
{
    return(xdr_pointer(xdrs, gnp,
                      sizeof(struct gnumbers_node),
                      xdr_gnumbers_node));
}

```

The unfortunate side effect of XDR'ing a list with these routines is that the C stack grows linearly with respect to the number of node in the list. This is due to the recursion. The following routine collapses the above two mutually recursive into a single, non-recursive one.

```

bool_t
xdr_gnumbers_list(xdrs, gnp)
    XDR *xdrs;
    gnumbers_list *gnp;
{
    bool_t more_data;
    gnumbers_list *nextp;

    for (;;) {
        more_data = (*gnp != NULL);
        if (!xdr_bool(xdrs, &more_data)) {
            return(FALSE);
        }
        if (! more_data) {
            break;
        }
        if (xdrs->x_op == XDR_FREE) {
            nextp = &(*gnp)->gn_next;
        }
        if (!xdr_reference(xdrs, gnp,
                          sizeof(struct gnumbers_node), xdr_gnumbers)) {
            return(FALSE);
        }
    }
}

```

```

        gnp = (xdrs->x_op == XDR_FREE) ?
            nextp : &(*gnp)->gn_next;
    }
    *gnp = NULL;
    return(TRUE);
}

```

The first task is to find out whether there is more data or not, so that this boolean information can be serialized. Notice that this statement is unnecessary in the 'XDR_DECODE' case, since the value of 'more_data' is not known until we deserialize it in the next statement.

The next statement XDR's the more_data field of the XDR union. Then if there is truly no more data, we set this last pointer to 'NULL' to indicate the end of the list, and return 'TRUE' because we are done. Note that setting the pointer to 'NULL' is only important in the 'XDR_DECODE' case, since it is already 'NULL' in the 'XDR_ENCODE' and 'XDR_FREE' cases.

Next, if the direction is 'XDR_FREE', the value of 'nextp' is set to indicate the location of the next pointer in the list. We do this now because we need to dereference gnp to find the location of the next item in the list, and after the next statement the storage pointed to by 'gnp' will be freed up and no be longer valid. We can't do this for all directions though, because in the 'XDR_DECODE' direction the value of 'gnp' won't be set until the next statement.

Next, we XDR the data in the node using the primitive 'xdr_reference()'. 'xdr_reference()' is like 'xdr_pointer()' which we used before, but it does not send over the boolean indicating whether there is more data. We use it instead of 'xdr_pointer()' because we have already XDR'd this information ourselves. Notice that the xdr routine passed is not the same type as an element in the list. The routine passed is 'xdr_gnumbers()', for XDR'ing gnumbers, but each element in the list is actually of type 'gnumbers_node'. We don't pass 'xdr_gnumbers_node()' because it is recursive, and instead use 'xdr_gnumbers()' which XDR's all of the non-recursive part. Note that this trick will work only if the 'gn_numbers' field is the first item in each element, so that their addresses are identical when passed to 'xdr_reference()'.

Finally, we update 'gnp' to point to the next item in the list. If the direction is 'XDR_FREE', we set it to the previously saved value, otherwise we can dereference 'gnp' to get the proper value. Though harder to understand than the recursive version, this non-recursive routine is far less likely to blow the C stack. It will also run more efficiently since a lot of procedure call overhead has been removed. Most lists are small though (in the hundreds of items or less) and the recursive version should be sufficient for them.

1.99 RPC.guide/XDR Protocol Specification

External Data Representation Standard: Protocol Specification

This chapter specifies a protocol that Sun Microsystems, Inc., and others are using. It has been designated RFC1014 by the ARPA Network Information Center.

Introduction to XDR
XDR Data Types
Discussion
The XDR Language Specification
An Example of an XDR Data Description
XDR References

1.100 RPC.guide/Introduction to XDR

Introduction to XDR
=====

XDR is a standard for the description and encoding of data. It is useful for transferring data between different computer architectures, and has been used to communicate data between such diverse machines as the Sun Workstation, VAX, IBM-PC, and Cray. XDR fits into the ISO presentation layer, and is roughly analogous in purpose to X.409, ISO Abstract Syntax Notation. The major difference between these two is that XDR uses implicit typing, while X.409 uses explicit typing.

XDR uses a language to describe data formats. The language can only be used only to describe data; it is not a programming language. This language allows one to describe intricate data formats in a concise manner. The alternative of using graphical representations (itself an informal language) quickly becomes incomprehensible when faced with complexity. The XDR language itself is similar to the C language [1], just as Courier [4] is similar to Mesa. Protocols such as Sun RPC (Remote Procedure Call) and the NFS (Network File System) use XDR to describe the format of their data.

The XDR standard makes the following assumption: that bytes (or octets) are portable, where a byte is defined to be 8 bits of data. A given hardware device should encode the bytes onto the various media in such a way that other hardware devices may decode the bytes without loss of meaning. For example, the Ethernet standard suggests that bytes be encoded in "little-endian" style [2], or least significant bit first.

Basic Block Size

1.101 RPC.guide/Basic Block Size

Basic Block Size

The representation of all items requires a multiple of four bytes (or 32 bits) of data. The bytes are numbered 0 through $n-1$. The bytes are read or written to some byte stream such that byte m always precedes byte $m+1$. If the n bytes needed to contain the data are not a multiple of four, then the n bytes are followed by enough (0 to 3) residual zero bytes, r , to make the total byte count a multiple of 4.

We include the familiar graphic box notation for illustration and comparison. In most illustrations, each box (delimited by a plus sign at the 4 corners and vertical bars and dashes) depicts a byte. Ellipses (...) between boxes show zero or more additional bytes where required.

A Block

```
+-----+-----+...+-----+-----+...+-----+
| byte 0 | byte 1 |...|byte n-1|    0  |...|    0  |
+-----+-----+...+-----+-----+...+-----+
|<-----n bytes----->|<-----r bytes----->|
|<-----n+r (where (n+r) mod 4 = 0)>----->|
```

1.102 RPC.guide/XDR Data Types

XDR Data Types

=====

Each of the sections that follow describes a data type defined in the XDR standard, shows how it is declared in the language, and includes a graphic illustration of its encoding.

- Integer
- Unsigned Integer
- Enumeration
- Boolean
- Hyper Integer and Unsigned Hyper Integer
- Floating-point
- Double-precision Floating-point
- Fixed-length Opaque Data
- Variable-length Opaque Data
- String
- Fixed-length Array
- Variable-length Array
- Structure
- Discriminated Union
- Void
- Constant
- Typedef (type definition)
- Optional-data
- Areas for Future Enhancement

For each data type in the language we show a general paradigm declaration. Note that angle brackets ('<' and '>') denote variable length sequences of data and square brackets ('[' and ']') denote fixed-length sequences of data. 'n', 'm' and 'r' denote integers. For the full language specification and more formal definitions of terms such as "identifier" and "declaration", refer to The XDR Language Specification.

For some data types, more specific examples are included. A more extensive example of a data description is in An Example of an XDR Data Description.

1.103 RPC.guide/Integer

Integer

An XDR signed integer is a 32-bit datum that encodes an integer in the range [-2147483648,2147483647]. The integer is represented in two's complement notation. The most and least significant bytes are 0 and 3, respectively. Integers are declared as follows:

Integer

```

(MSB)                                (LSB)
+-----+-----+-----+-----+
|byte 0 |byte 1 |byte 2 |byte 3 |
+-----+-----+-----+-----+
<-----32 bits----->
```

1.104 RPC.guide/Unsigned Integer

Unsigned Integer

An XDR unsigned integer is a 32-bit datum that encodes a nonnegative integer in the range [0,4294967295]. It is represented by an unsigned binary number whose most and least significant bytes are 0 and 3, respectively. An unsigned integer is declared as follows:

Unsigned Integer

```

(MSB)                                (LSB)
+-----+-----+-----+-----+
|byte 0 |byte 1 |byte 2 |byte 3 |
+-----+-----+-----+-----+
<-----32 bits----->
```

1.105 RPC.guide/Enumeration

Enumeration

Enumerations have the same representation as signed integers. Enumerations are handy for describing subsets of the integers. Enumerated data is declared as follows:

```
enum { name-identifier = constant, ... } identifier;
```

For example, the three colors red, yellow, and blue could be described by an enumerated type:

```
enum { RED = 2, YELLOW = 3, BLUE = 5 } colors;
```

It is an error to encode as an enum any other integer than those that have been given assignments in the enum declaration.

1.106 RPC.guide/Boolean

Boolean

Booleans are important enough and occur frequently enough to warrant their own explicit type in the standard. Booleans are declared as follows:

```
bool identifier;
```

This is equivalent to:

```
enum { FALSE = 0, TRUE = 1 } identifier;
```

1.107 RPC.guide/Hyper Integer and Unsigned Hyper Integer

Hyper Integer and Unsigned Hyper Integer

The standard also defines 64-bit (8-byte) numbers called hyper integer and unsigned hyper integer. Their representations are the obvious extensions of integer and unsigned integer defined above. They are represented in two's complement notation. The most and least significant bytes are 0 and 7, respectively. Their declarations:

Hyper Integer

Unsigned Hyper Integer

(MSB)

(LSB)

+-----+-----+-----+-----+-----+-----+-----+

```

|byte 0 |byte 1 |byte 2 |byte 3 |byte 4 |byte 5 |byte 6 |byte 7 |
+-----+-----+-----+-----+-----+-----+-----+-----+
<-----64 bits----->

```

1.108 RPC.guide/Floating-point

Floating-point

The standard defines the floating-point data type 'float' (32 bits or 4 bytes). The encoding used is the IEEE standard for normalized single-precision floating-point numbers [3]. The following three fields describe the single-precision floating-point number:

S:

The sign of the number. Values 0 and 1 represent positive and negative, respectively. One bit.

E:

The exponent of the number, base 2. 8 bits are devoted to this field. The exponent is biased by 127.

F:

The fractional part of the number's mantissa, base 2. 23 bits are devoted to this field.

Therefore, the floating-point number is described by:

$$(-1)^S * 2^{(E-Bias)} * 1.F$$

It is declared as follows:

Single-Precision Floating-Point

```

+-----+-----+-----+-----+
|byte 0 |byte 1 |byte 2 |byte 3 |
S|  E  |      F      |
+-----+-----+-----+
1|<- 8 ->|<-----23 bits----->|
<-----32 bits----->

```

Just as the most and least significant bytes of a number are 0 and 3, the most and least significant bits of a single-precision floating-point number are 0 and 31. The beginning bit (and most significant bit) offsets of S, E, and F are 0, 1, and 9, respectively. Note that these numbers refer to the mathematical positions of the bits, and NOT to their actual physical locations (which vary from medium to medium).

The IEEE specifications should be consulted concerning the encoding for signed zero, signed infinity (overflow), and denormalized numbers (underflow) [3]. According to IEEE specifications, the 'NaN' (not a number) is system dependent and should not be used externally.

1.109 RPC.guide/Double-precision Floating-point

Double-precision Floating-point

The standard defines the encoding for the double-precision floating-point data type 'double' (64 bits or 8 bytes). The encoding used is the IEEE standard for normalized double-precision floating-point numbers [3]. The standard encodes the following three fields, which describe the double-precision floating-point number:

S:

The sign of the number. Values 0 and 1 represent positive and negative, respectively. One bit.

E:

The exponent of the number, base 2. 11 bits are devoted to this field. The exponent is biased by 1023.

F:

The fractional part of the number's mantissa, base 2. 52 bits are devoted to this field.

Therefore, the floating-point number is described by:

$$(-1)^S * 2^{(E-Bias)} * 1.F$$

It is declared as follows:

Double-Precision Floating-Point

```
+-----+-----+-----+-----+-----+-----+-----+-----+
|byte 0|byte 1|byte 2|byte 3|byte 4|byte 5|byte 6|byte 7|
S|     E  |           F           |
+-----+-----+-----+-----+-----+-----+-----+-----+
1|<--11-->|<-----52 bits----->|
<-----64 bits----->
```

Just as the most and least significant bytes of a number are 0 and 3, the most and least significant bits of a double-precision floating-point number are 0 and 63. The beginning bit (and most significant bit) offsets of S, E, and F are 0, 1, and 12, respectively. Note that these numbers refer to the mathematical positions of the bits, and NOT to their actual physical locations (which vary from medium to medium).

The IEEE specifications should be consulted concerning the encoding for signed zero, signed infinity (overflow), and denormalized numbers (underflow) [3]. According to IEEE specifications, the 'NaN' (not a number) is system dependent and should not be used externally.

1.110 RPC.guide/Fixed-length Opaque Data

Fixed-length Opaque Data

At times, fixed-length uninterpreted data needs to be passed among machines. This data is called 'opaque' and is declared as follows:

```
opaque identifier[n];
```

where the constant n is the (static) number of bytes necessary to contain the opaque data. If n is not a multiple of four, then the n bytes are followed by enough (0 to 3) residual zero bytes, r , to make the total byte count of the opaque object a multiple of four.

Fixed-Length Opaque

```

0          1          ...
+-----+-----+...+-----+-----+...+-----+
| byte 0 | byte 1 |...|byte n-1|    0   |...|    0   |
+-----+-----+...+-----+-----+...+-----+
|<-----n bytes----->|<-----r bytes----->|
|<-----n+r (where (n+r) mod 4 = 0)----->|

```

1.111 RPC.guide/Variable-length Opaque Data

Variable-length Opaque Data

The standard also provides for variable-length (counted) opaque data, defined as a sequence of n (numbered 0 through $n-1$) arbitrary bytes to be the number n encoded as an unsigned integer (as described below), and followed by the n bytes of the sequence.

Byte m of the sequence always precedes byte $m+1$ of the sequence, and byte 0 of the sequence always follows the sequence's length (count). enough (0 to 3) residual zero bytes, r , to make the total byte count a multiple of four. Variable-length opaque data is declared in the following way:

```
opaque identifier<m>;
```

or

```
opaque identifier<>;
```

The constant m denotes an upper bound of the number of bytes that the sequence may contain. If m is not specified, as in the second declaration, it is assumed to be $(2^{32}) - 1$, the maximum length. The constant m would normally be found in a protocol specification. For example, a filing protocol may state that the maximum data transfer size is 8192 bytes, as follows:

```
opaque filedata<8192>;
```

This can be illustrated as follows:

Variable-Length Opaque

```

0      1      2      3      4      5      ...
+-----+-----+-----+-----+-----+-----+...+-----+-----+...+-----+
|           length n           |byte0|byte1|...| n-1 | 0 |...| 0 |
+-----+-----+-----+-----+-----+-----+...+-----+-----+...+-----+
|<-----4 bytes----->|<-----n bytes----->|<---r bytes--->|
|<-----n+r (where (n+r) mod 4 = 0)---->|

```

It is an error to encode a length greater than the maximum described in the specification.

1.112 RPC.guide/String

String

The standard defines a string of n (numbered 0 through $n-1$) ASCII bytes to be the number n encoded as an unsigned integer (as described above), and followed by the n bytes of the string. Byte m of the string always precedes byte $m+1$ of the string, and byte 0 of the string always follows the string's length. If n is not a multiple of four, then the n bytes are followed by enough (0 to 3) residual zero bytes, r , to make the total byte count a multiple of four. Counted byte strings are declared as follows:

```
string object<m>;
```

or

```
string object<>;
```

The constant m denotes an upper bound of the number of bytes that a string may contain. If m is not specified, as in the second declaration, it is assumed to be $(2^{32}) - 1$, the maximum length. The constant m would normally be found in a protocol specification. For example, a filing protocol may state that a file name can be no longer than 255 bytes, as follows:

```
string filename<255>;
```

Which can be illustrated as:

A String

```

0      1      2      3      4      5      ...
+-----+-----+-----+-----+-----+-----+...+-----+-----+...+-----+
|           length n           |byte0|byte1|...| n-1 | 0 |...| 0 |
+-----+-----+-----+-----+-----+-----+...+-----+-----+...+-----+
|<-----4 bytes----->|<-----n bytes----->|<---r bytes--->|
|<-----n+r (where (n+r) mod 4 = 0)---->|

```

It is an error to encode a length greater than the maximum described in the specification.

1.113 RPC.guide/Fixed-length Array

Fixed-length Array

Declarations for fixed-length arrays of homogeneous elements are in the following form:

```
type-name identifier[n];
```

Fixed-length arrays of elements numbered 0 through $n-1$ are encoded by individually encoding the elements of the array in their natural order, 0 through $n-1$. Each element's size is a multiple of four bytes. Though all elements are of the same type, the elements may have different sizes. For example, in a fixed-length array of strings, all elements are of type 'string', yet each element will vary in its length.

Fixed-Length Array

```
+---+---+---+---+---+---+---+---+...+---+---+---+---+
| element 0 | element 1 |...| element n-1 |
+---+---+---+---+---+---+---+---+...+---+---+---+---+
|<-----n elements----->|
```

1.114 RPC.guide/Variable-length Array

Variable-length Array

Counted arrays provide the ability to encode variable-length arrays of homogeneous elements. The array is encoded as the element count n (an unsigned integer) followed by the encoding of each of the array's elements, starting with element 0 and progressing through element $n-1$. The declaration for variable-length arrays follows this form:

```
type-name identifier<m>;
```

or

```
type-name identifier<>;
```

The constant m specifies the maximum acceptable element count of an array; if m is not specified, as in the second declaration, it is assumed to be $(2^{32}) - 1$.

Counted Array

```
0 1 2 3
+---+---+---+---+---+---+---+---+...+---+---+---+---+
```

```

|      n      | element 0 | element 1 |...|element n-1|
+---+---+---+---+---+---+---+---+---+---+...+---+---+---+
|<-4 bytes->|<-----n elements----->|

```

It is an error to encode a value of *n* that is greater than the maximum described in the specification.

1.115 RPC.guide/Structure

Structure

Structures are declared as follows:

```

struct {
    component-declaration-A;
    component-declaration-B;
    ...
} identifier;

```

The components of the structure are encoded in the order of their declaration in the structure. Each component's size is a multiple of four bytes, though the components may be different sizes.

Structure

```

+-----+-----+...
| component A | component B |...
+-----+-----+...

```

1.116 RPC.guide/Discriminated Union

Discriminated Union

A discriminated union is a type composed of a discriminant followed by a type selected from a set of prearranged types according to the value of the discriminant. The type of discriminant is either 'int', 'unsigned int', or an enumerated type, such as 'bool'. The component types are called "arms" of the union, and are preceded by the value of the discriminant which implies their encoding. Discriminated unions are declared as follows:

```

union switch (discriminant-declaration) {
    case discriminant-value-A:
        arm-declaration-A;
    case discriminant-value-B:
        arm-declaration-B;
    ...
    default: default-declaration;
}

```

```
    } identifier;
```

Each 'case' keyword is followed by a legal value of the discriminant. The default arm is optional. If it is not specified, then a valid encoding of the union cannot take on unspecified discriminant values. The size of the implied arm is always a multiple of four bytes.

The discriminated union is encoded as its discriminant followed by the encoding of the implied arm.

Discriminated Union

```
0   1   2   3
+---+---+---+---+---+---+---+---+
| discriminant | implied arm |
+---+---+---+---+---+---+---+---+
|<---4 bytes--->|
```

1.117 RPC.guide/Void

Void

An XDR void is a 0-byte quantity. Voids are useful for describing operations that take no data as input or no data as output. They are also useful in unions, where some arms may contain data and others do not. The declaration is simply as follows:

```
void;
```

Voids are illustrated as follows:

```
Void
++
||
++
--><-- 0 bytes
```

1.118 RPC.guide/Constant

Constant

The data declaration for a constant follows this form:

```
const name-identifier = n;
```

'const' is used to define a symbolic name for a constant; it does

not declare any data. The symbolic constant may be used anywhere a regular constant may be used. For example, the following defines a symbolic constant DOZEN, equal to 12.

```
const DOZEN = 12;
```

1.119 RPC.guide/Typedef (type definition)

Typedef

'typedef' does not declare any data either, but serves to define new identifiers for declaring data. The syntax is:

```
typedef declaration;
```

The new type name is actually the variable name in the declaration part of the typedef. For example, the following defines a new type called 'eggbox' using an existing type called 'egg':

```
typedef egg eggbox[DOZEN];
```

Variables declared using the new type name have the same type as the new type name would have in the typedef, if it was considered a variable. For example, the following two declarations are equivalent in declaring the variable 'fresheggs':

```
eggbox fresheggs;  
egg    fresheggs[DOZEN];
```

When a typedef involves a struct, enum, or union definition, there is another (preferred) syntax that may be used to define the same type. In general, a typedef of the following form:

```
typedef <<struct, union, or enum definition>> identifier;
```

may be converted to the alternative form by removing the 'typedef' part and placing the identifier after the 'struct', 'union', or 'enum' keyword, instead of at the end. For example, here are the two ways to define the type 'bool':

```
typedef enum {      /* using typedef */  
    FALSE = 0,  
    TRUE  = 1  
} bool;  
  
enum bool {        /* preferred alternative */  
    FALSE = 0,  
    TRUE  = 1  
};
```

The reason this syntax is preferred is one does not have to wait until the end of a declaration to figure out the name of the new type.

1.120 RPC.guide/Optional-data

Optional-data

Optional-data is one kind of union that occurs so frequently that we give it a special syntax of its own for declaring it. It is declared as follows:

```
type-name *identifier;
```

This is equivalent to the following union:

```
union switch (bool opted) {
    case TRUE:
        type-name element;
    case FALSE:
        void;
} identifier;
```

It is also equivalent to the following variable-length array declaration, since the boolean 'opted' can be interpreted as the length of the array:

```
type-name identifier<1>;
```

Optional-data is not so interesting in itself, but it is very useful for describing recursive data-structures such as linked-lists and trees. For example, the following defines a type 'stringlist' that encodes lists of arbitrary length strings:

```
struct *stringlist {
    string item<>;
    stringlist next;
};
```

It could have been equivalently declared as the following union:

```
union stringlist switch (bool opted) {
    case TRUE:
        struct {
            string item<>;
            stringlist next;
        } element;
    case FALSE:
        void;
};
```

or as a variable-length array:

```
struct stringlist<1> {
    string item<>;
    stringlist next;
```

```
};
```

Both of these declarations obscure the intention of the stringlist type, so the optional-data declaration is preferred over both of them. The optional-data type also has a close correlation to how recursive data structures are represented in high-level languages such as Pascal or C by use of pointers. In fact, the syntax is the same as that of the C language for pointers.

1.121 RPC.guide/Areas for Future Enhancement

Areas for Future Enhancement

The XDR standard lacks representations for bit fields and bitmaps, since the standard is based on bytes. Also missing are packed (or binary-coded) decimals.

The intent of the XDR standard was not to describe every kind of data that people have ever sent or will ever want to send from machine to machine. Rather, it only describes the most commonly used data-types of high-level languages such as Pascal or C so that applications written in these languages will be able to communicate easily over some medium.

One could imagine extensions to XDR that would let it describe almost any existing protocol, such as TCP. The minimum necessary for this are support for different block sizes and byte-orders. The XDR discussed here could then be considered the 4-byte big-endian member of a larger XDR family.

1.122 RPC.guide/Discussion

Discussion
=====

Why a Language for Describing Data?
Why Only one Byte-Order for an XDR Unit?
Why does XDR use Big-Endian Byte-Order?
Why is the XDR Unit Four Bytes Wide?
Why must Variable-Length Data be Padded with Zeros?
Why is there No Explicit Data-Typing?

1.123 RPC.guide/Why a Language for Describing Data?

Why a Language for Describing Data?

There are many advantages in using a data-description language such as XDR versus using diagrams. Languages are more formal than diagrams and lead to less ambiguous descriptions of data. Languages are also easier to understand and allow one to think of other issues instead of the low-level details of bit-encoding. Also, there is a close analogy between the types of XDR and a high-level language such as C or Pascal. This makes the implementation of XDR encoding and decoding modules an easier task. Finally, the language specification itself is an ASCII string that can be passed from machine to machine to perform on-the-fly data interpretation.

1.124 RPC.guide/Why Only one Byte-Order for an XDR Unit?

Why Only one Byte-Order for an XDR Unit?

Supporting two byte-orderings requires a higher level protocol for determining in which byte-order the data is encoded. Since XDR is not a protocol, this can't be done. The advantage of this, though, is that data in XDR format can be written to a magnetic tape, for example, and any machine will be able to interpret it, since no higher level protocol is necessary for determining the byte-order.

1.125 RPC.guide/Why does XDR use Big-Endian Byte-Order?

Why does XDR use Big-Endian Byte-Order?

Yes, it is unfair, but having only one byte-order means you have to be unfair to somebody. Many architectures, such as the Motorola 68000 and IBM 370, support the big-endian byte-order.

1.126 RPC.guide/Why is the XDR Unit Four Bytes Wide?

Why is the XDR Unit Four Bytes Wide?

There is a tradeoff in choosing the XDR unit size. Choosing a small size such as two makes the encoded data small, but causes alignment problems for machines that aren't aligned on these boundaries. A large size such as eight means the data will be aligned on virtually every machine, but causes the encoded data to grow too big. We chose four as a compromise. Four is big enough to support most architectures

efficiently, except for rare machines such as the eight-byte aligned Cray. Four is also small enough to keep the encoded data restricted to a reasonable size.

1.127 RPC.guide/Why must Variable-Length Data be Padded with Zeros?

Why must Variable-Length Data be Padded with Zeros?

It is desirable that the same data encode into the same thing on all machines, so that encoded data can be meaningfully compared or checksummed. Forcing the padded bytes to be zero ensures this.

1.128 RPC.guide/Why is there No Explicit Data-Typing?

Why is there No Explicit Data-Typing?

Data-typing has a relatively high cost for what small advantages it may have. One cost is the expansion of data due to the inserted type fields. Another is the added cost of interpreting these type fields and acting accordingly. And most protocols already know what type they expect, so data-typing supplies only redundant information. However, one can still get the benefits of data-typing using XDR. One way is to encode two things: first a string which is the XDR data description of the encoded data, and then the encoded data itself. Another way is to assign a value to all the types in XDR, and then define a universal type which takes this value as its discriminant and for each value, describes the corresponding data type.

1.129 RPC.guide/The XDR Language Specification

The XDR Language Specification
=====

Notational Conventions
Lexical Notes
Syntax Information

1.130 RPC.guide/Notational Conventions

Notational Conventions

This specification uses an extended Backus–Naur Form notation for describing the XDR language. Here is a brief description of the notation:

- * The characters `'`, `(`, `)`, `[`, `]`, `"`, and `*` are special.
- * Terminal symbols are strings of any characters surrounded by double quotes.
- * Non-terminal symbols are strings of non-special characters.
- * Alternative items are separated by a vertical bar (`|`).
- * Optional items are enclosed in brackets.
- * Items are grouped together by enclosing them in parentheses.
- * A `*` following an item means 0 or more occurrences of that item.

For example, consider the following pattern:

```
"a " "very" (" , " " very")* [" cold " "and"] " rainy " ("day" | "night")
```

An infinite number of strings match this pattern. A few of them are:

```
"a very rainy day"
"a very, very rainy day"
"a very cold and rainy day"
"a very, very, very cold and rainy night"
```

1.131 RPC.guide/Lexical Notes

Lexical Notes

1. Comments begin with `/*` and terminate with `*/`.
 2. White space serves to separate items and is otherwise ignored.
 3. An identifier is a letter followed by an optional sequence of letters, digits or underbar (`'_'`). The case of identifiers is not ignored.
 4. A constant is a sequence of one or more decimal digits, optionally preceded by a minus-sign (`'-'`).
-

1.132 RPC.guide/Syntax Information

Syntax Information

XDR Syntax Notes

```

declaration:
    type-specifier identifier
    | type-specifier identifier "[" value "]"
    | type-specifier identifier "<" [ value ] ">"
    | "opaque" identifier "[" value "]"
    | "opaque" identifier "<" [ value ] ">"
    | "string" identifier "<" [ value ] ">"
    | type-specifier "*" identifier
    | "void"

value:
    constant
    | identifier

type-specifier:
    [ "unsigned" ] "int"
    | [ "unsigned" ] "hyper"
    | "float"
    | "double"
    | "bool"
    | enum-type-spec
    | struct-type-spec
    | union-type-spec
    | identifier

enum-type-spec:
    "enum" enum-body

enum-body:
    "{"
    ( identifier "=" value )
    ( "," identifier "=" value ) *
    "}"

struct-type-spec:
    "struct" struct-body

struct-body:
    "{"
    ( declaration ";" )
    ( declaration ";" ) *
    "}"

union-type-spec:
    "union" union-body

union-body:
    "switch" "(" declaration ")" "{"

```

```
( "case" value ":" declaration ";" )
( "case" value ":" declaration ";" ) *
[ "default" ":" declaration ";" ]
"}"
```

```
constant-def:
    "const" identifier "=" constant ";"
```

```
type-def:
    "typedef" declaration ";"
    | "enum" identifier enum-body ";"
    | "struct" identifier struct-body ";"
    | "union" identifier union-body ";"
```

```
definition:
    type-def
    | constant-def
```

```
specification:
    definition *
```

1.133 RPC.guide/XDR Syntax Notes

Syntax Notes

.....

1. The following are keywords and cannot be used as identifiers:

```
'bool', 'case', 'const', 'default', 'double',
'enum', 'float', 'hyper', 'opaque', 'string',
'struct', 'switch', 'typedef', 'union',
'unsigned' and 'void'.
```

2. Only unsigned constants may be used as size specifications for arrays. If an identifier is used, it must have been declared previously as an unsigned constant in a 'const' definition.
3. Constant and type identifiers within the scope of a specification are in the same name space and must be declared uniquely within this scope.
4. Similarly, variable names must be unique within the scope of struct and union declarations. Nested struct and union declarations create new scopes.
5. The discriminant of a union must be of a type that evaluates to an integer. That is, 'int', 'unsigned int', 'bool', an enumerated type or any typedefed type that evaluates to one of these is legal. Also, the case values must be one of the legal values of the discriminant. Finally, a case value may not be specified more than once within the scope of a union declaration.

1.134 RPC.guide/An Example of an XDR Data Description

An Example of an XDR Data Description

Here is a short XDR data description of a thing called a "file", which might be used to transfer files from one machine to another.

```

const MAXUSERNAME = 32;      /* max length of a user name */
const MAXFILELEN = 65535;   /* max length of a file      */
const MAXNAMELEN = 255;    /* max length of a file name */

/*
 * Types of files:
 */
enum filekind {
    TEXT = 0,      /* ascii data */
    DATA = 1,    /* raw data   */
    EXEC = 2      /* executable */
};

/*
 * File information, per kind of file:
 */
union filetype switch (filekind kind) {
    case TEXT:
        void;      /* no extra information */
    case DATA:
        string creator<MAXNAMELEN>; /* data creator      */
    case EXEC:
        string interpreter<MAXNAMELEN>; /* program interpreter */
};

/*
 * A complete file:
 */
struct file {
    string filename<MAXNAMELEN>; /* name of file */
    filetype type;              /* info about file */
    string owner<MAXUSERNAME>; /* owner of file  */
    opaque data<MAXFILELEN>;   /* file data      */
};

```

Suppose now that there is a user named "john" who wants to store his lisp program 'sillyprog' that contains just the data "(quit)". His file would be encoded as follows:

Offset	Hex Bytes	ASCII	Description
0	00 00 00 09	Length of filename = 9
4	73 69 6c 6c	sill	Filename characters
8	79 70 72 6f	ypro	... and more characters ...
12	67 00 00 00	g...	... and 3 zero-bytes of fill
16	00 00 00 02	Filekind is EXEC = 2
20	00 00 00 04	Length of interpreter = 4

```

24      6c 69 73 70 lisp   Interpretor characters
28      00 00 00 04 ....   Length of owner = 4
32      6a 6f 68 6e john   Owner characters
36      00 00 00 06 ....   Length of file data = 6
40      28 71 75 69 (qui   File data bytes ...
44      74 29 00 00 t)..   ... and 2 zero-bytes of fill

```

1.135 RPC.guide/XDR References

References

=====

- [1] Brian W. Kernighan & Dennis M. Ritchie, 'The C Programming Language', Bell Laboratories, Murray Hill, New Jersey, 1978.
- [2] Danny Cohen, 'On Holy Wars and a Plea for Peace', IEEE Computer, October 1981.
- [3] 'IEEE Standard for Binary Floating-Point Arithmetic', ANSI/IEEE Standard 754-1985, Institute of Electrical and Electronics Engineers, August 1985.
- [4] 'Courier: The Remote Procedure Call Protocol', XEROX Corporation, X SIS 038112, December 1981.

1.136 RPC.guide/Remote Procedure Calls Protocol Specification

Remote Procedure Calls: Protocol Specification

This chapter specifies a protocol that Sun Microsystems, Inc., and others are using. It has been designated RFC1050 by the ARPA Network Information Center.

- Introduction to RPC
- RPC Protocol Requirements
- The RPC Message Protocol
- Authentication Protocols
- Record Marking Standard
- The RPC Language
- Port Mapper Program Protocol
- RPC References

1.137 RPC.guide/Introduction to RPC

Introduction to RPC

=====

This chapter specifies a message protocol used in implementing Sun's Remote Procedure Call (RPC) package. (The message protocol is specified with the External Data Representation (XDR) language. see XDR Protocol Specification for the details. Here, we assume that the reader is familiar with XDR and do not attempt to justify it or its uses). The paper by Birrell and Nelson [1] is recommended as an excellent background to and justification of RPC.

Terminology

The RPC Model

Transports and Semantics

Binding and Rendezvous Independence

Authentication

1.138 RPC.guide/Terminology

Terminology

This chapter discusses servers, services, programs, procedures, clients, and versions. A "server" is a piece of software where network services are implemented. A "network service" is a collection of one or more remote programs. A "remote program" implements one or more remote procedures; the procedures, their parameters, and results are documented in the specific program's "protocol specification" (see the Port Mapper Program Protocol, below, for an example). "Network clients" are pieces of software that initiate remote procedure calls to services. A server may support more than one version of a remote program in order to be forward compatible with changing protocols.

For example, a network file service may be composed of two programs. One program may deal with high-level applications such as file system access control and locking. The other may deal with low-level file IO and have procedures like 'read' and 'write'. A client machine of the network file service would call the procedures associated with the two programs of the service on behalf of some user on the client machine.

1.139 RPC.guide/The RPC Model

The RPC Model

The remote procedure call model is similar to the local procedure call model. In the local case, the caller places arguments to a

procedure in some well-specified location (such as a result register). It then transfers control to the procedure, and eventually gains back control. At that point, the results of the procedure are extracted from the well-specified location, and the caller continues execution.

The remote procedure call is similar, in that one thread of control logically winds through two processes -- one is the caller's process, the other is a server's process. That is, the caller process sends a call message to the server process and waits (blocks) for a reply message. The call message contains the procedure's parameters, among other things. The reply message contains the procedure's results, among other things. Once the reply message is received, the results of the procedure are extracted, and caller's execution is resumed.

On the server side, a process is dormant awaiting the arrival of a call message. When one arrives, the server process extracts the procedure's parameters, computes the results, sends a reply message, and then awaits the next call message.

Note that in this model, only one of the two processes is active at any given time. However, this model is only given as an example. The RPC protocol makes no restrictions on the concurrency model implemented, and others are possible. For example, an implementation may choose to have RPC calls be asynchronous, so that the client may do useful work while waiting for the reply from the server. Another possibility is to have the server create a task to process an incoming request, so that the server can be free to receive other requests.

1.140 RPC.guide/Transports and Semantics

Transports and Semantics

The RPC protocol is independent of transport protocols. That is, RPC does not care how a message is passed from one process to another. The protocol deals only with specification and interpretation of messages.

It is important to point out that RPC does not try to implement any kind of reliability and that the application must be aware of the type of transport protocol underneath RPC. If it knows it is running on top of a reliable transport such as TCP/IP[6], then most of the work is already done for it. On the other hand, if it is running on top of an unreliable transport such as UDP/IP[7], it must implement its own retransmission and time-out policy as the RPC layer does not provide this service.

Because of transport independence, the RPC protocol does not attach specific semantics to the remote procedures or their execution. Semantics can be inferred from (but should be explicitly specified by) the underlying transport protocol. For example, consider RPC running on top of an unreliable transport such as UDP/IP. If an application retransmits RPC messages after short time-outs, the only thing it can infer if it receives no reply is that the procedure was executed zero or more times. If it does receive a reply, then it can infer that the

procedure was executed at least once.

A server may wish to remember previously granted requests from a client and not regrant them in order to insure some degree of execute-at-most-once semantics. A server can do this by taking advantage of the transaction ID that is packaged with every RPC request. The main use of this transaction is by the client RPC layer in matching replies to requests. However, a client application may choose to reuse its previous transaction ID when retransmitting a request. The server application, knowing this fact, may choose to remember this ID after granting a request and not regrant requests with the same ID in order to achieve some degree of execute-at-most-once semantics. The server is not allowed to examine this ID in any other way except as a test for equality.

On the other hand, if using a reliable transport such as TCP/IP, the application can infer from a reply message that the procedure was executed exactly once, but if it receives no reply message, it cannot assume the remote procedure was not executed. Note that even if a connection-oriented protocol like TCP is used, an application still needs time-outs and reconnection to handle server crashes.

There are other possibilities for transports besides datagram- or connection-oriented protocols. For example, a request-reply protocol such as VMTP[2] is perhaps the most natural transport for RPC.

NOTE: At Sun, RPC is currently implemented on top of both TCP/IP and UDP/IP transports.

1.141 RPC.guide/Binding and Rendezvous Independence

Binding and Rendezvous Independence

The act of binding a client to a service is NOT part of the remote procedure call specification. This important and necessary function is left up to some higher-level software. (The software may use RPC itself -- see the Port Mapper Program Protocol, below).

Implementors should think of the RPC protocol as the jump-subroutine instruction ('JSR') of a network; the loader (binder) makes 'JSR' useful, and the loader itself uses 'JSR' to accomplish its task. Likewise, the network makes RPC useful, using RPC to accomplish this task.

1.142 RPC.guide/Authentication

Authentication

The RPC protocol provides the fields necessary for a client to identify itself to a service and vice-versa. Security and access control mechanisms can be built on top of the message authentication. Several different authentication protocols can be supported. A field in the RPC header indicates which protocol is being used. More information on specific authentication protocols can be found in the Authentication Protocols, below.

1.143 RPC.guide/RPC Protocol Requirements

RPC Protocol Requirements

=====

Programs and Procedures
Authentication Requirements
Program Number Assignment
Other Uses of the RPC Protocol

The RPC protocol must provide for the following:

1. Unique specification of a procedure to be called.
2. Provisions for matching response messages to request messages.
3. Provisions for authenticating the caller to service and vice-versa.

Besides these requirements, features that detect the following are worth supporting because of protocol roll-over errors, implementation bugs, user error, and network administration:

1. RPC protocol mismatches.
2. Remote program protocol version mismatches.
3. Protocol errors (such as misspecification of a procedure's parameters).
4. Reasons why remote authentication failed.
5. Any other reasons why the desired procedure was not called.

1.144 RPC.guide/Programs and Procedures

Programs and Procedures

The RPC call message has three unsigned fields: remote program number, remote program version number, and remote procedure number. The three fields uniquely identify the procedure to be called. Program

numbers are administered by some central authority (like Sun). Once an implementor has a program number, he can implement his remote program; the first implementation would most likely have the version number of 1. Because most new protocols evolve into better, stable, and mature protocols, a version field of the call message identifies which version of the protocol the caller is using. Version numbers make speaking old and new protocols through the same server process possible.

The procedure number identifies the procedure to be called. These numbers are documented in the specific program's protocol specification. For example, a file service's protocol specification may state that its procedure number 5 is 'read' and procedure number 12 is 'write'.

Just as remote program protocols may change over several versions, the actual RPC message protocol could also change. Therefore, the call message also has in it the RPC version number, which is always equal to two for the version of RPC described here.

The reply message to a request message has enough information to distinguish the following error conditions:

1. The remote implementation of RPC does speak protocol version 2. The lowest and highest supported RPC version numbers are returned.
2. The remote program is not available on the remote system.
3. The remote program does not support the requested version number. The lowest and highest supported remote program version numbers are returned.
4. The requested procedure number does not exist. (This is usually a caller side protocol or programming error.)
5. The parameters to the remote procedure appear to be garbage from the server's point of view. (Again, this is usually caused by a disagreement about the protocol between client and service.)

1.145 RPC.guide/Authentication Requirements

Authentication

Provisions for authentication of caller to service and vice-versa are provided as a part of the RPC protocol. The call message has two authentication fields, the credentials and verifier. The reply message has one authentication field, the response verifier. The RPC protocol specification defines all three fields to be the following opaque type:

```
enum auth_flavor {
    AUTH_NULL      = 0,
    AUTH_UNIX      = 1,
    AUTH_SHORT      = 2,
    AUTH_DES        = 3
```

```

        /* and more to be defined */
};

struct opaque_auth {
    auth_flavor flavor;
    opaque body<400>;
};

```

In simple English, any 'opaque_auth' structure is an 'auth_flavor' enumeration followed by bytes which are opaque to the RPC protocol implementation.

The interpretation and semantics of the data contained within the authentication fields is specified by individual, independent authentication protocol specifications. (See Authentication Protocols, below, for definitions of the various authentication protocols.)

If authentication parameters were rejected, the response message contains information stating why they were rejected.

1.146 RPC.guide/Program Number Assignment

Program Number Assignment

Program numbers are given out in groups of '0x20000000' (decimal 536870912) according to the following chart:

Program Numbers	Description
0 - 1fffffff	Defined by Sun
20000000 - 3fffffff	Defined by user
40000000 - 5fffffff	Transient
60000000 - 7fffffff	Reserved
80000000 - 9fffffff	Reserved
a0000000 - bfffffff	Reserved
c0000000 - dfffffff	Reserved
e0000000 - ffffffff	Reserved

The first group is a range of numbers administered by Sun Microsystems and should be identical for all sites. The second range is for applications peculiar to a particular site. This range is intended primarily for debugging new programs. When a site develops an application that might be of general interest, that application should be given an assigned number in the first range. The third group is for applications that generate program numbers dynamically. The final groups are reserved for future use, and should not be used.

1.147 RPC.guide/Other Uses of the RPC Protocol

Other Uses of the RPC Protocol

The intended use of this protocol is for calling remote procedures. That is, each call message is matched with a response message. However, the protocol itself is a message-passing protocol with which other (non-RPC) protocols can be implemented. Sun currently uses, or perhaps abuses, the RPC message protocol for the following two (non-RPC) protocols: batching (or pipelining) and broadcast RPC. These two protocols are discussed but not defined below.

Batching
RPC Broadcast

1.148 RPC.guide/Batching

Batching
.....

Batching allows a client to send an arbitrarily large sequence of call messages to a server; batching typically uses reliable byte stream protocols (like TCP/IP) for its transport. In the case of batching, the client never waits for a reply from the server, and the server does not send replies to batch requests. A sequence of batch calls is usually terminated by a legitimate RPC in order to flush the pipeline (with positive acknowledgement).

1.149 RPC.guide/RPC Broadcast

Broadcast RPC
.....

In broadcast RPC-based protocols, the client sends a broadcast packet to the network and waits for numerous replies. Broadcast RPC uses unreliable, packet-based protocols (like UDP/IP) as its transports. Servers that support broadcast protocols only respond when the request is successfully processed, and are silent in the face of errors. Broadcast RPC uses the Port Mapper RPC service to achieve its semantics. See the Port Mapper Program Protocol, below, for more information.

1.150 RPC.guide/The RPC Message Protocol

The RPC Message Protocol
=====

This section defines the RPC message protocol in the XDR data description language. The message is defined in a top-down style.

```

enum msg_type {
    CALL = 0,
    REPLY = 1
};

/*
 * A reply to a call message can take on two forms:
 * The message was either accepted or rejected.
 */
enum reply_stat {
    MSG_ACCEPTED = 0,
    MSG_DENIED = 1
};

/*
 * Given that a call message was accepted, the following is the
 * status of an attempt to call a remote procedure.
 */
enum accept_stat {
    SUCCESS = 0, /* RPC executed successfully */
    PROG_UNAVAIL = 1, /* remote hasn't exported program */
    PROG_MISMATCH = 2, /* remote can't support version # */
    PROC_UNAVAIL = 3, /* program can't support procedure */
    GARBAGE_ARGS = 4 /* procedure can't decode params */
};

/*
 * Reasons why a call message was rejected:
 */
enum reject_stat {
    RPC_MISMATCH = 0, /* RPC version number != 2 */
    AUTH_ERROR = 1 /* remote can't authenticate caller */
};

/*
 * Why authentication failed:
 */
enum auth_stat {
    AUTH_BADCRED = 1, /* bad credentials */
    AUTH_REJECTEDCRED = 2, /* client must begin new session */
    AUTH_BADVERF = 3, /* bad verifier */
    AUTH_REJECTEDVERF = 4, /* verifier expired or replayed */
    AUTH_TOOWEAK = 5 /* rejected for security reasons */
};

/*
 * The RPC message:
 * All messages start with a transaction identifier, 'xid',
 * followed by a two-armed discriminated union. The union's discriminant
 * is a 'msg_type' which switches to one of the two types of the
 * message. The 'xid' of a 'REPLY' message always matches that of the
 * initiating 'CALL' message. NB: The 'xid' field is only used
 * for clients matching reply messages with call messages or for servers
 * detecting retransmissions; the service side cannot treat this id as

```

```

    * any type of sequence number.
    */
struct rpc_msg {
    unsigned int xid;
    union switch (msg_type mtype) {
        case CALL:
            call_body cbody;
        case REPLY:
            reply_body rbody;
    } body;
};

/*
 * Body of an RPC request call: In version 2 of the RPC protocol
 * specification, 'rpcvers' must be equal to 2. The fields 'prog', 'vers', ←
 * and
 * 'proc' specify the remote program, its version number, and the procedure
 * within the remote program to be called. After these fields are two
 * authentication parameters: 'cred' (authentication credentials) and 'verf'
 * (authentication verifier). The two authentication parameters are
 * followed by the parameters to the remote procedure, which are
 * specified by the specific program protocol.
 */
struct call_body {
    unsigned int rpcvers; /* must be equal to two (2) */
    unsigned int prog;
    unsigned int vers;
    unsigned int proc;
    opaque_auth cred;
    opaque_auth verf;
    /* procedure specific parameters start here */
};

/*
 * Body of a reply to an RPC request:
 * The call message was either accepted or rejected.
 */
union reply_body switch (reply_stat stat) {
    case MSG_ACCEPTED:
        accepted_reply areply;
    case MSG_DENIED:
        rejected_reply rreply;
} reply;

/*
 * Reply to an RPC request that was accepted by the server:
 * there could be an error even though the request was accepted.
 * The first field is an authentication verifier that the server
 * generates in order to validate itself to the caller. It is
 * followed by a union whose discriminant is an enum
 * 'accept_stat'. The 'SUCCESS' arm of the union is protocol
 * specific. The 'PROG_UNAVAIL', 'PROC_UNAVAIL', and 'GARBAGE_ARGP'
 * arms of the union are void. The 'PROG_MISMATCH' arm specifies
 * the lowest and highest version numbers of the remote program
 * supported by the server.
 */
struct accepted_reply {

```

```

opaque_auth verf;
union switch (accept_stat stat) {
    case SUCCESS:
        opaque results[0];
        /* procedure-specific results start here */
    case PROG_MISMATCH:
        struct {
            unsigned int low;
            unsigned int high;
        } mismatch_info;
    default:
        /*
         * Void. Cases include 'PROG_UNAVAIL, PROC_UNAVAIL',
         * and 'GARBAGE_ARGS'.
         */
        void;
    } reply_data;
};

/*
 * Reply to an RPC request that was rejected by the server: The
 * request can be rejected for two reasons: either the server is not
 * running a compatible version of the RPC protocol ('RPC_MISMATCH'),
 * or the server refuses to authenticate the caller ('AUTH_ERROR'). In
 * case of an RPC version mismatch, the server returns the lowest and
 * highest supported RPC version numbers. In case of refused
 * authentication, failure status is returned.
 */
union rejected_reply switch (reject_stat stat) {
    case RPC_MISMATCH:
        struct {
            unsigned int low;
            unsigned int high;
        } mismatch_info;
    case AUTH_ERROR:
        auth_stat stat;
};

```

1.151 RPC.guide/Authentication Protocols

Authentication Protocols

=====

As previously stated, authentication parameters are opaque, but open-ended to the rest of the RPC protocol. This section defines some "flavors" of authentication implemented at (and supported by) Sun. Other sites are free to invent new authentication types, with the same rules of flavor number assignment as there is for program number assignment.

Null Authentication (protocol)
 UNIX Authentication (protocol)
 DES Authentication (protocol)

1.152 RPC.guide/Null Authentication (protocol)

Null Authentication

Often calls must be made where the caller does not know who he is or the server does not care who the caller is. In this case, the flavor value (the discriminant of the 'opaque_auth''s union) of the RPC message's credentials, verifier, and response verifier is 'AUTH_NULL'. The bytes of the 'opaque_auth''s body are undefined. It is recommended that the opaque length be zero.

1.153 RPC.guide/UNIX Authentication (protocol)

UNIX Authentication

The caller of a remote procedure may wish to identify himself as he is identified on a UNIX system. The value of the credential's discriminant of an RPC call message is 'AUTH_UNIX'. The bytes of the credential's opaque body encode the following structure:

```
struct auth_unix {
    unsigned int stamp;
    string machinename<255>;
    unsigned int uid;
    unsigned int gid;
    unsigned int gids<10>;
};
```

The 'stamp' is an arbitrary ID which the caller machine may generate. The 'machinename' is the name of the caller's machine (like "krypton"). The 'uid' is the caller's effective user ID. The 'gid' is the caller's effective group ID. The 'gids' is a counted array of groups which contain the caller as a member. The verifier accompanying the credentials should be of 'AUTH_NULL' (defined above).

The value of the discriminant of the response verifier received in the reply message from the server may be 'AUTH_NULL' or 'AUTH_SHORT'. In the case of 'AUTH_SHORT', the bytes of the response verifier's string encode an opaque structure. This new opaque structure may now be passed to the server instead of the original 'AUTH_UNIX' flavor credentials. The server keeps a cache which maps shorthand opaque structures (passed back by way of an 'AUTH_SHORT' style response verifier) to the original credentials of the caller. The caller can save network bandwidth and server cpu cycles by using the new credentials.

The server may flush the shorthand opaque structure at any time. If

this happens, the remote procedure call message will be rejected due to an authentication error. The reason for the failure will be 'AUTH_REJECTEDCRED'. At this point, the caller may wish to try the original 'AUTH_UNIX' style of credentials.

1.154 RPC.guide/DES Authentication (protocol)

DES Authentication

UNIX authentication suffers from two major problems:

1. The naming is too UNIX-system oriented.
2. There is no verifier, so credentials can easily be faked.

DES authentication attempts to fix these two problems.

Naming

DES Authentication Verifiers

Nicknames and Clock Synchronization

DES Authentication Protocol (in XDR language)

Diffie-Hellman Encryption

1.155 RPC.guide/Naming

Naming

.....

The first problem is handled by addressing the caller by a simple string of characters instead of by an operating system specific integer. This string of characters is known as the "netname" or network name of the caller. The server is not allowed to interpret the contents of the caller's name in any other way except to identify the caller. Thus, netnames should be unique for every caller in the internet.

It is up to each operating system's implementation of DES authentication to generate netnames for its users that insure this uniqueness when they call upon remote servers. Operating systems already know how to distinguish users local to their systems. It is usually a simple matter to extend this mechanism to the network. For example, a UNIX user at Sun with a user ID of 515 might be assigned the following netname: "unix.515@sun.com". This netname contains three items that serve to insure it is unique. Going backwards, there is only one naming domain called "sun.com" in the internet. Within this domain, there is only one UNIX user with user ID 515. However, there may be another user on another operating system, for example VMS, within the same naming domain that, by coincidence, happens to have the

same user ID. To insure that these two users can be distinguished we add the operating system name. So one user is "unix.515@sun.com" and the other is "vms.515@sun.com".

The first field is actually a naming method rather than an operating system name. It just happens that today there is almost a one-to-one correspondence between naming methods and operating systems. If the world could agree on a naming standard, the first field could be the name of that standard, instead of an operating system name.

1.156 RPC.guide/DES Authentication Verifiers

DES Authentication Verifiers

.....

Unlike UNIX authentication, DES authentication does have a verifier so the server can validate the client's credential (and vice-versa). The contents of this verifier is primarily an encrypted timestamp. The server can decrypt this timestamp, and if it is close to what the real time is, then the client must have encrypted it correctly. The only way the client could encrypt it correctly is to know the "conversation key" of the RPC session. And if the client knows the conversation key, then it must be the real client.

The conversation key is a DES [5] key which the client generates and notifies the server of in its first RPC call. The conversation key is encrypted using a public key scheme in this first transaction. The particular public key scheme used in DES authentication is Diffie-Hellman [3] with 192-bit keys. The details of this encryption method are described later.

The client and the server need the same notion of the current time in order for all of this to work. If network time synchronization cannot be guaranteed, then client can synchronize with the server before beginning the conversation, perhaps by consulting the Internet Time Server (TIME[4]).

The way a server determines if a client timestamp is valid is somewhat complicated. For any other transaction but the first, the server just checks for two things:

1. the timestamp is greater than the one previously seen from the same client.
2. the timestamp has not expired.

A timestamp is expired if the server's time is later than the sum of the client's timestamp plus what is known as the client's "window". The "window" is a number the client passes (encrypted) to the server in its first transaction. You can think of it as a lifetime for the credential.

This explains everything but the first transaction. In the first transaction, the server checks only that the timestamp has not expired.

If this was all that was done though, then it would be quite easy for the client to send random data in place of the timestamp with a fairly good chance of succeeding. As an added check, the client sends an encrypted item in the first transaction known as the "window verifier" which must be equal to the window minus 1, or the server will reject the credential.

The client too must check the verifier returned from the server to be sure it is legitimate. The server sends back to the client the encrypted timestamp it received from the client, minus one second. If the client gets anything different than this, it will reject it.

1.157 RPC.guide/Nicknames and Clock Synchronization

Nicknames and Clock Synchronization

.....

After the first transaction, the server's DES authentication subsystem returns in its verifier to the client an integer "nickname" which the client may use in its further transactions instead of passing its netname, encrypted DES key and window every time. The nickname is most likely an index into a table on the server which stores for each client its netname, decrypted DES key and window.

Though they originally were synchronized, the client's and server's clocks can get out of sync again. When this happens the client RPC subsystem most likely will get back 'RPC_AUTHERROR' at which point it should resynchronize.

A client may still get the 'RPC_AUTHERROR' error even though it is synchronized with the server. The reason is that the server's nickname table is a limited size, and it may flush entries whenever it wants. A client should resend its original credential in this case and the server will give it a new nickname. If a server crashes, the entire nickname table gets flushed, and all clients will have to resend their original credentials.

1.158 RPC.guide/DES Authentication Protocol (in XDR language)

DES Authentication Protocol (in XDR language)

.....

```
/*
 * There are two kinds of credentials: one in which the client uses
 * its full network name, and one in which it uses its "nickname"
 * (just an unsigned integer) given to it by the server. The
 * client must use its fullname in its first transaction with the
 * server, in which the server will return to the client its
 * nickname. The client may use its nickname in all further
 * transactions with the server. There is no requirement to use the
```

```

    * nickname, but it is wise to use it for performance reasons.
    */
enum authdes_namekind {
    ADN_FULLNAME = 0,
    ADN_NICKNAME = 1
};

/*
 * A 64-bit block of encrypted DES data
 */
typedef opaque des_block[8];

/*
 * Maximum length of a network user's name
 */
const MAXNETNAMELEN = 255;

/*
 * A fullname contains the network name of the client, an encrypted
 * conversation key and the window. The window is actually a
 * lifetime for the credential. If the time indicated in the
 * verifier timestamp plus the window has past, then the server
 * should expire the request and not grant it. To insure that
 * requests are not replayed, the server should insist that
 * timestamps are greater than the previous one seen, unless it is
 * the first transaction. In the first transaction, the server
 * checks instead that the window verifier is one less than the
 * window.
 */
struct authdes_fullname {
    string name<MAXNETNAMELEN>; /* name of client */
    des_block key; /* encrypted conversation key */
    unsigned int window; /* encrypted window */
};

/*
 * A credential is either a fullname or a nickname
 */
union authdes_cred switch (authdes_namekind adc_namekind) {
    case ADN_FULLNAME:
        authdes_fullname adc_fullname;
    case ADN_NICKNAME:
        unsigned int adc_nickname;
};

/*
 * A timestamp encodes the time since midnight, January 1, 1970.
 */
struct timestamp {
    unsigned int seconds; /* seconds */
    unsigned int useconds; /* and microseconds */
};

/*
 * Verifier: client variety
 * The window verifier is only used in the first transaction. In
 * conjunction with a fullname credential, these items are packed

```

```

* into the following structure before being encrypted:
*
* struct {
*     adv_timestamp;           -- one DES block
*     adc_fullname.window;    -- one half DES block
*     adv_winverf;           -- one half DES block
* }
* This structure is encrypted using CBC mode encryption with an
* input vector of zero. All other encryptions of timestamps use
* ECB mode encryption.
*/
struct authdes_verf_clnt {
    timestamp adv_timestamp;    /* encrypted timestamp */
    unsigned int adv_winverf;  /* encrypted window verifier */
};

/*
* Verifier: server variety
* The server returns (encrypted) the same timestamp the client
* gave it minus one second. It also tells the client its nickname
* to be used in future transactions (unencrypted).
*/
struct authdes_verf_svr {
    timestamp adv_timeverf;    /* encrypted verifier */
    unsigned int adv_nickname; /* new nickname for client */
};

```

1.159 RPC.guide/Diffie-Hellman Encryption

Diffie-Hellman Encryption

.....

In this scheme, there are two constants, 'BASE' and 'MODULUS'. The particular values Sun has chosen for these for the DES authentication protocol are:

```

const BASE = 3;
const MODULUS =
    "d4a0ba0250b6fd2ec626e7efd637df76c716e22d0944b88b"; /* hex */

```

The way this scheme works is best explained by an example. Suppose there are two people 'A' and 'B' who want to send encrypted messages to each other. So, A and B both generate "secret" keys at random which they do not reveal to anyone. Let these keys be represented as 'SK(A)' and 'SK(B)'. They also publish in a public directory their "public" keys. These keys are computed as follows:

```

PK(A) = ( BASE ^ SK(A) ) mod MODULUS
PK(B) = ( BASE ^ SK(B) ) mod MODULUS

```

The '^' notation is used here to represent exponentiation. Now, both A and B can arrive at the "common" key between them, represented here as 'CK(A, B)', without revealing their secret keys.

A computes:

$$CK(A, B) = (PK(B) ^ SK(A)) \text{ mod } MODULUS$$

while B computes:

$$CK(A, B) = (PK(A) ^ SK(B)) \text{ mod } MODULUS$$

These two can be shown to be equivalent:

$$(PK(B) ^ SK(A)) \text{ mod } MODULUS = (PK(A) ^ SK(B)) \text{ mod } MODULUS$$

We drop the "mod MODULUS" parts and assume modulo arithmetic to simplify things:

$$PK(B) ^ SK(A) = PK(A) ^ SK(B)$$

Then, replace 'PK(B)' by what B computed earlier and likewise for 'PK(A)'.

$$((BASE ^ SK(B)) ^ SK(A) = (BASE ^ SK(A)) ^ SK(B)$$

which leads to:

$$BASE ^ (SK(A) * SK(B)) = BASE ^ (SK(A) * SK(B))$$

This common key 'CK(A, B)' is not used to encrypt the timestamps used in the protocol. Rather, it is used only to encrypt a conversation key which is then used to encrypt the timestamps. The reason for doing this is to use the common key as little as possible, for fear that it could be broken. Breaking the conversation key is a far less serious offense, since conversations are relatively short-lived.

The conversation key is encrypted using 56-bit DES keys, yet the common key is 192 bits. To reduce the number of bits, 56 bits are selected from the common key as follows. The middle-most 8-bytes are selected from the common key, and then parity is added to the lower order bit of each byte, producing a 56-bit key with 8 bits of parity.

1.160 RPC.guide/Record Marking Standard

Record Marking Standard

=====

When RPC messages are passed on top of a byte stream protocol (like TCP/IP), it is necessary, or at least desirable, to delimit one message from another in order to detect and possibly recover from user protocol errors. This is called record marking (RM). Sun uses this RM/TCP/IP transport for passing RPC messages on TCP streams. One RPC message fits into one RM record.

A record is composed of one or more record fragments. A record fragment is a four-byte header followed by 0 to $(2^{31}) - 1$ bytes of fragment data. The bytes encode an unsigned binary number; as with XDR

integers, the byte order is from highest to lowest. The number encodes two values -- a boolean which indicates whether the fragment is the last fragment of the record (bit value 1 implies the fragment is the last fragment) and a 31-bit unsigned binary value which is the length in bytes of the fragment's data. The boolean value is the highest-order bit of the header; the length is the 31 low-order bits. (Note that this record specification is NOT in XDR standard form!)

1.161 RPC.guide/The RPC Language

The RPC Language

=====

Just as there was a need to describe the XDR data-types in a formal language, there is also need to describe the procedures that operate on these XDR data-types in a formal language as well. We use the RPC Language for this purpose. It is an extension to the XDR language. The following example is used to describe the essence of the language.

An Example Service Described in the RPC Language
 The RPC Language Specification
 RPC Syntax Notes

1.162 RPC.guide/An Example Service Described in the RPC Language

An Example Service Described in the RPC Language

Here is an example of the specification of a simple ping program.

```

/*
 * Simple ping program
 */
program PING_PROG {
    /* Latest and greatest version */
    version PING_VERS_PINGBACK {
        void
        PINGPROC_NULL(void) = 0;

        /*
         * Ping the caller, return the round-trip time
         * (in microseconds). Returns -1 if the operation
         * timed out.
         */
        int
        PINGPROC_PINGBACK(void) = 1;
    } = 2;
}
/*

```

```

    * Original version
    */
    version PING_VERS_ORIG {
        void
        PINGPROC_NULL(void) = 0;
    } = 1;
} = 1;

const PING_VERS = 2;      /* latest version */

```

The first version described is 'PING_VERS_PINGBACK' with two procedures, 'PINGPROC_NULL' and 'PINGPROC_PINGBACK'. 'PINGPROC_NULL' takes no arguments and returns no results, but it is useful for computing round-trip times from the client to the server and back again. By convention, procedure 0 of any RPC protocol should have the same semantics, and never require any kind of authentication. The second procedure is used for the client to have the server do a reverse ping operation back to the client, and it returns the amount of time (in microseconds) that the operation used. The next version, 'PING_VERS_ORIG', is the original version of the protocol and it does not contain 'PINGPROC_PINGBACK' procedure. It is useful for compatibility with old client programs, and as this program matures it may be dropped from the protocol entirely.

1.163 RPC.guide/The RPC Language Specification

The RPC Language Specification

The RPC language is identical to the XDR language, except for the added definition of a 'program-def' described below.

```

program-def:
    "program" identifier "{"
        version-def
        version-def *
    "}" "=" constant ";"

version-def:
    "version" identifier "{"
        procedure-def
        procedure-def *
    "}" "=" constant ";"

procedure-def:
    type-specifier identifier "(" type-specifier ")"
    "=" constant ";"

```

1.164 RPC.guide/RPC Syntax Notes

Syntax Notes

1. The following keywords are added and cannot be used as identifiers: `'program'` and `'version'`;
2. A version name cannot occur more than once within the scope of a program definition. Nor can a version number occur more than once within the scope of a program definition.
3. A procedure name cannot occur more than once within the scope of a version definition. Nor can a procedure number occur more than once within the scope of version definition.
4. Program identifiers are in the same name space as constant and type identifiers.
5. Only unsigned constants can be assigned to programs, versions and procedures.

1.165 RPC.guide/Port Mapper Program Protocol

Port Mapper Program Protocol

=====

The port mapper program maps RPC program and version numbers to transport-specific port numbers. This program makes dynamic binding of remote programs possible.

This is desirable because the range of reserved port numbers is very small and the number of potential remote programs is very large. By running only the port mapper on a reserved port, the port numbers of other remote programs can be ascertained by querying the port mapper.

The port mapper also aids in broadcast RPC. A given RPC program will usually have different port number bindings on different machines, so there is no way to directly broadcast to all of these programs. The port mapper, however, does have a fixed port number. So, to broadcast to a given program, the client actually sends its message to the port mapper located at the broadcast address. Each port mapper that picks up the broadcast then calls the local service specified by the client. When the port mapper gets the reply from the local service, it sends the reply on back to the client.

Port Mapper Protocol Specification (in RPC Language)
Port Mapper Operation

1.166 RPC.guide/Port Mapper Protocol Specification (in RPC Language)

Port Mapper Protocol Specification (in RPC Language)

```
-----

const PMAP_PORT = 111;          /* portmapper port number */

/*
 * A mapping of (program, version, protocol) to port number
 */
struct mapping {
    unsigned int prog;
    unsigned int vers;
    unsigned int prot;
    unsigned int port;
};

/*
 * Supported values for the "prot" field
 */
const IPPROTO_TCP = 6;          /* protocol number for TCP/IP */
const IPPROTO_UDP = 17;        /* protocol number for UDP/IP */

/*
 * A list of mappings
 */
struct *pmaplist {
    mapping map;
    pmaplist next;
};

/*
 * Arguments to callit
 */
struct call_args {
    unsigned int prog;
    unsigned int vers;
    unsigned int proc;
    opaque args<>;
};

/*
 * Results of callit
 */
struct call_result {
    unsigned int port;
    opaque res<>;
};

/*
 * Port mapper procedures
 */
program PMAP_PROG {
    version PMAP_VERS {
        void
        PMAPPROC_NULL(void)          = 0;

        bool
```

```

    PMAPPROC_SET(mapping)          = 1;

    bool
    PMAPPROC_UNSET(mapping)        = 2;

    unsigned int
    PMAPPROC_GETPORT(mapping)      = 3;

    pmaplist
    PMAPPROC_DUMP(void)            = 4;

    call_result
    PMAPPROC_CALLIT(call_args)     = 5;
} = 2;
} = 100000;
```

1.167 RPC.guide/Port Mapper Operation

Port Mapper Operation

The portmapper program currently supports two protocols (UDP/IP and TCP/IP). The portmapper is contacted by talking to it on assigned port number 111 (SUNRPC [8]) on either of these protocols. The following is a description of each of the portmapper procedures:

PMAPPROC_NULL:

This procedure does no work. By convention, procedure zero of any protocol takes no parameters and returns no results.

PMAPPROC_SET:

When a program first becomes available on a machine, it registers itself with the port mapper program on the same machine. The program passes its program number 'prog', version number 'vers', transport protocol number 'prot', and the port 'port' on which it awaits service request. The procedure returns a boolean response whose value is 'TRUE' if the procedure successfully established the mapping and 'FALSE' otherwise. The procedure refuses to establish a mapping if one already exists for the tuple '(prog, vers, prot)'.

PMAPPROC_UNSET:

When a program becomes unavailable, it should unregister itself with the port mapper program on the same machine. The parameters and results have meanings identical to those of 'PMAPPROC_SET'. The protocol and port number fields of the argument are ignored.

PMAPPROC_GETPORT:

Given a program number 'prog', version number 'vers', and transport protocol number 'prot', this procedure returns the port number on which the program is awaiting call requests. A port value of zeros means the program has not been registered. The 'port' field of the argument is ignored.

PMAPPROC_DUMP:

This procedure enumerates all entries in the port mapper's database. The procedure takes no parameters and returns a list of program, version, protocol, and port values.

PMAPPROC_CALLIT:

This procedure allows a caller to call another remote procedure on the same machine without knowing the remote procedure's port number. It is intended for supporting broadcasts to arbitrary remote programs via the well-known port mapper's port. The parameters 'prog', 'vers', 'proc', and the bytes of 'args' are the program number, version number, procedure number, and parameters of the remote procedure.

Note:

1. This procedure only sends a response if the procedure was successfully executed and is silent (no response) otherwise.
2. The port mapper communicates with the remote program using UDP/IP only.

The procedure returns the remote program's port number, and the bytes of results are the results of the remote procedure.

1.168 RPC.guide/RPC References

References

=====

- [1] Birrell, Andrew D. & Nelson, Bruce Jay; 'Implementing Remote Procedure Calls'; XEROX CSL-83-7, October 1983.
 - [2] Cheriton, D.; 'VMTP: Versatile Message Transaction Protocol', Preliminary Version 0.3; Stanford University, January 1987.
 - [3] Diffie & Hellman; 'New Directions in Cryptography'; IEEE Transactions on Information Theory IT-22, November 1976.
 - [4] Harrenstien, K.; 'Time Server', RFC 738; Information Sciences Institute, October 1977.
 - [5] National Bureau of Standards; 'Data Encryption Standard'; Federal Information Processing Standards Publication 46, January 1977.
 - [6] Postel, J.; 'Transmission Control Protocol - DARPA Internet Program Protocol Specification', RFC 793; Information Sciences Institute, September 1981.
-

- [7] Postel, J.; 'User Datagram Protocol', RFC 768; Information Sciences Institute, August 1980.
- [8] Reynolds, J. & Postel, J.; 'Assigned Numbers', RFC 923; Information Sciences Institute, October 1984.

1.169 RPC.guide/Network File System Version 2 Protocol Specification

Network File System: Version 2 Protocol Specification

Note: This document specifies a protocol that Sun Microsystems, Inc., and others are using. It specifies it in standard ARPA RFC form.

Introduction to NFS
NFS Protocol Definition
NFS Implementation Issues
Mount Protocol Definition

1.170 RPC.guide/Introduction to NFS

Introduction to NFS
=====

The Sun Network Filesystem (NFS) protocol provides transparent remote access to shared filesystems over local area networks. The NFS protocol is designed to be machine, operating system, network architecture, and transport protocol independent. This independence is achieved through the use of Remote Procedure Call (RPC) primitives built on top of an External Data Representation (XDR). Implementations exist for a variety of machines, from personal computers to supercomputers.

The supporting mount protocol allows the server to hand out remote access privileges to a restricted set of clients. It performs the operating system-specific functions that allow, for example, to attach remote directory trees to some local file system.

Remote Procedure Call
External Data Representation
Stateless Servers

1.171 RPC.guide/Remote Procedure Call

Remote Procedure Call

Sun's remote procedure call specification provides a procedure-oriented interface to remote services. Each server supplies a program that is a set of procedures. NFS is one such "program". The combination of host address, program number, and procedure number specifies one remote service procedure. RPC does not depend on services provided by specific protocols, so it can be used with any underlying transport protocol. See Remote Procedure Calls Protocol Specification.

1.172 RPC.guide/External Data Representation

External Data Representation

The External Data Representation (XDR) standard provides a common way of representing a set of data types over a network. The NFS Protocol Specification is written using the RPC data description language. For more information, See XDR Protocol Specification. Sun provides implementations of XDR and RPC, but NFS does not require their use. Any software that provides equivalent functionality can be used, and if the encoding is exactly the same it can interoperate with other implementations of NFS.

1.173 RPC.guide/Stateless Servers

Stateless Servers

The NFS protocol is stateless. That is, a server does not need to maintain any extra state information about any of its clients in order to function correctly. Stateless servers have a distinct advantage over stateful servers in the event of a failure. With stateless servers, a client need only retry a request until the server responds; it does not even need to know that the server has crashed, or the network temporarily went down. The client of a stateful server, on the other hand, needs to either detect a server crash and rebuild the server's state when it comes back up, or cause client operations to fail.

This may not sound like an important issue, but it affects the protocol in some unexpected ways. We feel that it is worth a bit of extra complexity in the protocol to be able to write very simple servers that do not require fancy crash recovery.

On the other hand, NFS deals with objects such as files and directories that inherently have state - what good would a file be if it did not keep its contents intact? The goal is to not introduce any

extra state in the protocol itself. Another way to simplify recovery is by making operations "idempotent" whenever possible (so that they can potentially be repeated).

1.174 RPC.guide/NFS Protocol Definition

NFS Protocol Definition

=====

Servers have been known to change over time, and so can the protocol that they use. So RPC provides a version number with each RPC request. This RFC describes version two of the NFS protocol. Even in the second version, there are various obsolete procedures and parameters, which will be removed in later versions. An RFC for version three of the NFS protocol is currently under preparation(1).

File System Model
 RPC Information
 Sizes of XDR Structures
 Basic Data Types
 Server Procedures

----- Footnotes -----

(1) Editor's note: The version three specification is ready by now.

1.175 RPC.guide/File System Model

File System Model

NFS assumes a file system that is hierarchical, with directories as all but the bottom-level files. Each entry in a directory (file, directory, device, etc.) has a string name. Different operating systems may have restrictions on the depth of the tree or the names used, as well as using different syntax to represent the "pathname", which is the concatenation of all the "components" (directory and file names) in the name. A "file system" is a tree on a single server (usually a single disk or physical partition) with a specified "root". Some operating systems provide a "mount" operation to make all file systems appear as a single tree, while others maintain a "forest" of file systems. Files are unstructured streams of uninterpreted bytes. Version 3 of NFS uses a slightly more general file system model.

NFS looks up one component of a pathname at a time. It may not be obvious why it does not just take the whole pathname, traipse down the directories, and return a file handle when it is done. There are several good reasons not to do this. First, pathnames need separators between the directory components, and different operating systems use

different separators. We could define a Network Standard Pathname Representation, but then every pathname would have to be parsed and converted at each end. Other issues are discussed in NFS Implementation Issues below.

Although files and directories are similar objects in many ways, different procedures are used to read directories and files. This provides a network standard format for representing directories. The same argument as above could have been used to justify a procedure that returns only one directory entry per call. The problem is efficiency. Directories can contain many entries, and a remote call to return each would be just too slow.

1.176 RPC.guide/RPC Information

RPC Information

Authentication

The NFS service uses 'AUTH_UNIX', 'AUTH_DES', or 'AUTH_SHORT' style authentication, except in the NULL procedure where 'AUTH_NONE' is also allowed.

Transport Protocols

NFS currently is supported on UDP/IP only.

Port Number

The NFS protocol currently uses the UDP port number 2049. This is not an officially assigned port, so later versions of the protocol use the "Portmapping" facility of RPC.

1.177 RPC.guide/Sizes of XDR Structures

Sizes of XDR Structures

These are the sizes, given in decimal bytes, of various XDR structures used in the protocol:

```
/* The maximum number of bytes of data in a READ or WRITE request */
const MAXDATA = 8192;
```

```
/* The maximum number of bytes in a pathname argument */
const MAXPATHLEN = 1024;
```

```
/* The maximum number of bytes in a file name argument */
const MAXNAMLEN = 255;
```

```
/* The size in bytes of the opaque "cookie" passed by REaddir */
const COOKIESIZE = 4;
```

```
/* The size in bytes of the opaque file handle */  
const FHSIZE = 32;
```

1.178 RPC.guide/Basic Data Types

Basic Data Types

The following XDR definitions are basic structures and types used in other structures described further on.

```
stat  
ftype  
fhandle  
timeval  
fattr  
sattr  
filename  
path  
attrstat  
diropargs  
diropres
```

1.179 RPC.guide/stat

```
stat  
....  
  
enum stat {  
    NFS_OK = 0,  
    NFSERR_PERM=1,  
    NFSERR_NOENT=2,  
    NFSERR_IO=5,  
    NFSERR_NXIO=6,  
    NFSERR_ACCES=13,  
    NFSERR_EXIST=17,  
    NFSERR_NODEV=19,  
    NFSERR_NOTDIR=20,  
    NFSERR_ISDIR=21,  
    NFSERR_FBIG=27,  
    NFSERR_NOSPC=28,  
    NFSERR_ROFS=30,  
    NFSERR_NAMETOOLONG=63,  
    NFSERR_NOTEMPTY=66,  
    NFSERR_DQUOT=69,  
    NFSERR_STALE=70,  
    NFSERR_WFLUSH=99  
};
```

The 'stat' type is returned with every procedure's results. A value of 'NFS_OK' indicates that the call completed successfully and the results are valid. The other values indicate some kind of error occurred on the server side during the servicing of the procedure. The error values are derived from UNIX error numbers.

'NFSERR_PERM'

Not owner. The caller does not have correct ownership to perform the requested operation.

'NFSERR_NOENT'

No such file or directory. The file or directory specified does not exist.

'NFSERR_IO'

Some sort of hard error occurred when the operation was in progress. This could be a disk error, for example.

'NFSERR_NXIO'

No such device or address.

'NFSERR_ACCES'

Permission denied. The caller does not have the correct permission to perform the requested operation.

'NFSERR_EXIST'

File exists. The file specified already exists.

'NFSERR_NODEV'

No such device.

'NFSERR_NOTDIR'

Not a directory. The caller specified a non-directory in a directory operation.

'NFSERR_ISDIR'

Is a directory. The caller specified a directory in a non-directory operation.

'NFSERR_FBIG'

File too large. The operation caused a file to grow beyond the server's limit.

'NFSERR_NOSPC'

No space left on device. The operation caused the server's filesystem to reach its limit.

'NFSERR_ROFS'

Read-only filesystem. Write attempted on a read-only filesystem.

'NFSERR_NAMETOOLONG'

File name too long. The file name in an operation was too long.

'NFSERR_NOTEMPTY'

Directory not empty. Attempted to remove a directory that was not empty.

``NFSEERR_DQUOT'`

Disk quota exceeded. The client's disk quota on the server has been exceeded.

``NFSEERR_STALE'`

The `'fhandle'` given in the arguments was invalid. That is, the file referred to by that file handle no longer exists, or access to it has been revoked.

``NFSEERR_WFLUSH'`

The server's write cache used in the `'WRITECACHE'` call got flushed to disk.

1.180 RPC.guide/ftype

`ftype`

.....

```
enum ftype {
    NFNON = 0,
    NFREG = 1,
    NFDIR = 2,
    NFBLK = 3,
    NFCHR = 4,
    NFLNK = 5
};
```

The enumeration `'ftype'` gives the type of a file. The type `'NFNON'` indicates a non-file, `'NFREG'` is a regular file, `'NFDIR'` is a directory, `'NFBLK'` is a block-special device, `'NFCHR'` is a character-special device, and `'NFLNK'` is a symbolic link.

1.181 RPC.guide/fhandle

`fhandle`

.....

```
typedef opaque fhandle[FHSIZE];
```

The `'fhandle'` is the file handle passed between the server and the client. All file operations are done using file handles to refer to a file or directory. The file handle can contain whatever information the server needs to distinguish an individual file.

1.182 RPC.guide/timeval

timeval
.....

```
struct timeval {
    unsigned int seconds;
    unsigned int useconds;
};
```

The 'timeval' structure is the number of seconds and microseconds since midnight January 1, 1970, Greenwich Mean Time. It is used to pass time and date information.

1.183 RPC.guide/fattr

fattr
.....

```
struct fattr {
    ftype      type;
    unsigned int mode;
    unsigned int nlink;
    unsigned int uid;
    unsigned int gid;
    unsigned int size;
    unsigned int blocksize;
    unsigned int rdev;
    unsigned int blocks;
    unsigned int fsid;
    unsigned int fileid;
    timeval     atime;
    timeval     mtime;
    timeval     ctime;
};
```

The 'fattr' structure contains the attributes of a file; 'type' is the type of the file; 'nlink' is the number of hard links to the file (the number of different names for the same file); 'uid' is the user identification number of the owner of the file; 'gid' is the group identification number of the group of the file; 'size' is the size in bytes of the file; 'blocksize' is the size in bytes of a block of the file; 'rdev' is the device number of the file if it is type 'NFCHR' or 'NFBLK'; 'blocks' is the number of blocks the file takes up on disk; 'fsid' is the file system identifier for the filesystem containing the file; 'fileid' is a number that uniquely identifies the file within its filesystem; 'atime' is the time when the file was last accessed for either read or write; 'mtime' is the time when the file data was last modified (written); and 'ctime' is the time when the status of the file was last changed. Writing to the file also changes 'ctime' if the size of the file changes.

'mode' is the access mode encoded as a set of bits. Notice that the file type is specified both in the mode bits and in the file type. This is really a bug in the protocol and will be fixed in future

versions. The descriptions given below specify the bit positions using octal numbers.

Bit	Description
0040000	This is a directory; 'type' field should be NFDIR.
0020000	This is a character special file; 'type' field should be NFCHR.
0060000	This is a block special file; 'type' field should be NFBLK.
0100000	This is a regular file; 'type' field should be NFREG.
0120000	This is a symbolic link file; 'type' field should be NFLNK.
0140000	This is a named socket; 'type' field should be NFNON.
0004000	Set user id on execution.
0002000	Set group id on execution.
0001000	Save swapped text even after use.
0000400	Read permission for owner.
0000200	Write permission for owner.
0000100	Execute and search permission for owner.
0000040	Read permission for group.
0000020	Write permission for group.
0000010	Execute and search permission for group.
0000004	Read permission for others.
0000002	Write permission for others.
0000001	Execute and search permission for others.

Notes:

1. The bits are the same as the mode bits returned by the 'stat' system call in the UNIX system. The file type is specified both in the mode bits and in the file type. This is fixed in future versions.
2. The "rdev" field in the attributes structure is an operating system specific device specifier. It will be removed and generalized in the next revision of the protocol.

1.184 RPC.guide/sattr

```
sattr
.....
```

```
struct sattr {
    unsigned int mode;
    unsigned int uid;
    unsigned int gid;
    unsigned int size;
    timeval      atime;
    timeval      mtime;
};
```

The 'sattr' structure contains the file attributes which can be set from the client. The fields are the same as for 'fattr' above. A 'size' of zero means the file should be truncated. A value of '-1' indicates a field that should be ignored.

1.185 RPC.guide/filename

filename
.....

```
typedef string filename<MAXNAMLEN>;
```

The type 'filename' is used for passing file names or pathname components.

1.186 RPC.guide/path

path
....

```
typedef string path<MAXPATHLEN>;
```

The type 'path' is a pathname. The server considers it as a string with no internal structure, but to the client it is the name of a node in a filesystem tree.

1.187 RPC.guide/attrstat

attrstat
.....

```
union attrstat switch (stat status) {  
    case NFS_OK:  
        fattr attributes;  
    default:  
        void;  
};
```

The 'attrstat' structure is a common procedure result. It contains a 'status' and, if the call succeeded, it also contains the attributes of the file on which the operation was done.

1.188 RPC.guide/diopargs

diopargs
.....

```
struct diopargs {  
    fhandle dir;  
    filename name;  
};
```

The 'diropargs' structure is used in directory operations. The 'fhandle dir' is the directory in which to find the file 'name'. A directory operation is one in which the directory is affected.

1.189 RPC.guide/diopres

diopres
.....

```
union diopres switch (stat status) {
    case NFS_OK:
        struct {
            fhandle file;
            fattr  attributes;
        } diopok;
    default:
        void;
};
```

The results of a directory operation are returned in a 'diopres' structure. If the call succeeded, a new file handle 'file' and the 'attributes' associated with that file are returned along with the 'status'.

1.190 RPC.guide/Server Procedures

Server Procedures

The protocol definition is given as a set of procedures with arguments and results defined using the RPC language. A brief description of the function of each procedure should provide enough information to allow implementation.

- Do Nothing
- Get File Attributes
- Set File Attributes
- Get Filesystem Root
- Look Up File Name
- Read From Symbolic Link
- Read From File
- Write to Cache
- Write to File
- Create File
- Remove File
- Rename File
- Create Link to File
- Create Symbolic Link

```

Create Directory
Remove Directory
Read From Directory
Get Filesystem Attributes

```

All of the procedures in the NFS protocol are assumed to be synchronous. When a procedure returns to the client, the client can assume that the operation has completed and any data associated with the request is now on stable storage. For example, a client 'WRITE' request may cause the server to update data blocks, filesystem information blocks (such as indirect blocks), and file attribute information (size and modify times). When the 'WRITE' returns to the client, it can assume that the write is safe, even in case of a server crash, and it can discard the data written. This is a very important part of the statelessness of the server. If the server waited to flush data from remote requests, the client would have to save those requests so that it could resend them in case of a server crash.

```

/*
 * Remote file service routines
 */
program NFS_PROGRAM {
    version NFS_VERSION {
        void          NFSPROC_NULL(void)           = 0;
        attrstat      NFSPROC_GETATTR(fhandle)     = 1;
        attrstat      NFSPROC_SETATTR(sattrargs)   = 2;
        void          NFSPROC_ROOT(void)           = 3;
        diropres      NFSPROC_LOOKUP(diropargs)    = 4;
        readlinkres   NFSPROC_READLINK(fhandle)   = 5;
        readres       NFSPROC_READ(readargs)       = 6;
        void          NFSPROC_WRITECACHE(void)     = 7;
        attrstat      NFSPROC_WRITE(writeargs)    = 8;
        diropres      NFSPROC_CREATE(createargs)   = 9;
        stat          NFSPROC_REMOVE(diropargs)    = 10;
        stat          NFSPROC_RENAME(renameargs)   = 11;
        stat          NFSPROC_LINK(linkargs)       = 12;
        stat          NFSPROC_SYMLINK(symargs)     = 13;
        diropres      NFSPROC_MKDIR(createargs)    = 14;
        stat          NFSPROC_RMDIR(diropargs)     = 15;
        readdirres    NFSPROC_READDIR(readdirargs) = 16;
        statfsres     NFSPROC_STATFS(fhandle)     = 17;
    } = 2;
} = 100003;

```

1.191 RPC.guide/Do Nothing

Do Nothing

.....

```

void
NFSPROC_NULL(void) = 0;

```

This procedure does no work. It is made available in all RPC services to allow server response testing and timing.

1.192 RPC.guide/Get File Attributes

Get File Attributes


```
attrstat
NFSPROC_GETATTR (fhandle) = 1;
```

If the reply status is 'NFS_OK', then the reply attributes contains the attributes for the file given by the input fhandle.

1.193 RPC.guide/Set File Attributes

Set File Attributes


```
struct sattrargs {
    fhandle file;
    sattr attributes;
};

attrstat
NFSPROC_SETATTR (sattrargs) = 2;
```

The 'attributes' argument contains fields which are either '-1' or are the new value for the attributes of 'file'. If the reply status is 'NFS_OK', then the reply attributes have the attributes of the file after the 'SETATTR' operation has completed.

Note: The use of '-1' to indicate an unused field in 'attributes' is changed in the next version of the protocol.

1.194 RPC.guide/Get Filesystem Root

Get Filesystem Root


```
void
NFSPROC_ROOT(void) = 3;
```

Obsolete. This procedure is no longer used because finding the root file handle of a filesystem requires moving pathnames between client and server. To do this right we would have to define a network standard representation of pathnames. Instead, the function of looking up the root file handle is done by the 'MNTPROC_MNT()' procedure. (For details, see Mount Protocol Definition).

1.195 RPC.guide/Look Up File Name

Look Up File Name


```
diopres
NFSPROC_LOOKUP(diopargs) = 4;
```

If the reply 'status' is 'NFS_OK', then the reply 'file' and reply 'attributes' are the file handle and attributes for the file 'name' in the directory given by 'dir' in the argument.

1.196 RPC.guide/Read From Symbolic Link

Read From Symbolic Link


```
union readlinkres switch (stat status) {
    case NFS_OK:
        path data;
    default:
        void;
};

readlinkres
NFSPROC_READLINK(fhandle) = 5;
```

If 'status' has the value 'NFS_OK', then the reply 'data' is the data in the symbolic link given by the file referred to by the 'fhandle' argument.

Note: since NFS always parses pathnames on the client, the pathname in a symbolic link may mean something different (or be meaningless) on a different client or on the server if a different pathname syntax is used.

1.197 RPC.guide/Read From File

Read From File


```
struct readargs {
    fhandle file;
    unsigned offset;
    unsigned count;
    unsigned totalcount;
```

```

};

union readres switch (stat status) {
    case NFS_OK:
        fattr attributes;
        opaque data<NFS_MAXDATA>;
    default:
        void;
};

readres
NFSPROC_READ(readargs) = 6;

```

Returns up to 'count' bytes of 'data' from the file given by 'file', starting at 'offset' bytes from the beginning of the file. The first byte of the file is at offset zero. The file attributes after the read takes place are returned in 'attributes'.

Note: The argument 'totalcount' is unused, and is removed in the next protocol revision.

1.198 RPC.guide/Write to Cache

Write to Cache

.....

```

void
NFSPROC_WRITECACHE(void) = 7;

```

To be used in the next protocol revision.

1.199 RPC.guide/Write to File

Write to File

.....

```

struct writeargs {
    fhandle file;
    unsigned beginoffset;
    unsigned offset;
    unsigned totalcount;
    opaque data<NFS_MAXDATA>;
};

attrstat
NFSPROC_WRITE(writeargs) = 8;

```

Writes 'data' beginning 'offset' bytes from the beginning of 'file'. The first byte of the file is at offset zero. If the reply 'status' is 'NFS_OK', then the reply 'attributes' contains the attributes of the

file after the write has completed. The write operation is atomic. Data from this call to 'WRITE' will not be mixed with data from another client's calls.

Note: The arguments 'beginoffset' and 'totalcount' are ignored and are removed in the next protocol revision.

1.200 RPC.guide/Create File

Create File

.....

```
struct createargs {
    diropargs where;
    sattr attributes;
};

diropres
NFSPROC_CREATE(createargs) = 9;
```

The file 'name' is created in the directory given by 'dir'. The initial attributes of the new file are given by 'attributes'. A reply 'status' of 'NFS_OK' indicates that the file was created, and reply 'file' and reply 'attributes' are its file handle and attributes. Any other reply 'status' means that the operation failed and no file was created.

Note: This routine should pass an exclusive create flag, meaning "create the file only if it is not already there".

1.201 RPC.guide/Remove File

Remove File

.....

```
stat
NFSPROC_REMOVE(diropargs) = 10;
```

The file 'name' is removed from the directory given by 'dir'. A reply of 'NFS_OK' means the directory entry was removed.

Note: possibly non-idempotent operation.

1.202 RPC.guide/Rename File

Rename File

.....

```

struct renameargs {
    diropargs from;
    diropargs to;
};

stat
NFSPROC_RENAME(renameargs) = 11;

```

The existing file 'from.name' in the directory given by 'from.dir' is renamed to 'to.name' in the directory given by 'to.dir'. If the reply is 'NFS_OK', the file was renamed. The RENAME operation is atomic on the server; it cannot be interrupted in the middle.

Note: possibly non-idempotent operation.

1.203 RPC.guide/Create Link to File

Create Link to File

.....

```

struct linkargs {
    fhandle from;
    diropargs to;
};

stat
NFSPROC_LINK(linkargs) = 12;

```

Creates the file 'to.name' in the directory given by 'to.dir', which is a hard link to the existing file given by 'from'. If the return value is 'NFS_OK', a link was created. Any other return value indicates an error, and the link was not created.

A hard link should have the property that changes to either of the linked files are reflected in both files. When a hard link is made to a file, the attributes for the file should have a value for 'nlink' that is one greater than the value before the link.

Note: possibly non-idempotent operation.

1.204 RPC.guide/Create Symbolic Link

Create Symbolic Link

.....

```

struct symlinkargs {

```

```

    diropargs from;
    path to;
    sattr attributes;
};

stat
NFSPROC_SYMLINK(symlinkargs) = 13;

```

Creates the file 'from.name' with ftype 'NFLNK' in the directory given by 'from.dir'. The new file contains the pathname 'to' and has initial attributes given by 'attributes'. If the return value is 'NFS_OK', a link was created. Any other return value indicates an error, and the link was not created.

A symbolic link is a pointer to another file. The name given in 'to' is not interpreted by the server, only stored in the newly created file. When the client references a file that is a symbolic link, the contents of the symbolic link are normally transparently reinterpreted as a pathname to substitute. A 'READLINK' operation returns the data to the client for interpretation.

Note: On UNIX servers the attributes are never used, since symbolic links always have mode '0777'.

1.205 RPC.guide/Create Directory

Create Directory


```

diropres
NFSPROC_MKDIR (createargs) = 14;

```

The new directory 'where.name' is created in the directory given by 'where.dir'. The initial attributes of the new directory are given by 'attributes'. A reply 'status' of 'NFS_OK' indicates that the new directory was created, and reply 'file' and reply 'attributes' are its file handle and attributes. Any other reply 'status' means that the operation failed and no directory was created.

Note: possibly non-idempotent operation.

1.206 RPC.guide/Remove Directory

Remove Directory


```

stat
NFSPROC_RMDIR(diropargs) = 15;

```

The existing empty directory 'name' in the directory given by 'dir'

is removed. If the reply is 'NFS_OK', the directory was removed.

Note: possibly non-idempotent operation.

1.207 RPC.guide/Read From Directory

Read From Directory

.....

```

struct readdirargs {
    fhandle dir;
    nfscookie cookie;
    unsigned count;
};

struct entry {
    unsigned fileid;
    filename name;
    nfscookie cookie;
    entry *nextentry;
};

union readdirres switch (stat status) {
    case NFS_OK:
        struct {
            entry *entries;
            bool eof;
        } readdirok;
    default:
        void;
};

readdirres
NFSPROC_READDIR (readdirargs) = 16;

```

Returns a variable number of directory entries, with a total size of up to 'count' bytes, from the directory given by 'dir'. If the returned value of 'status' is 'NFS_OK', then it is followed by a variable number of 'entry's. Each 'entry' contains a 'fileid' which consists of a unique number to identify the file within a filesystem, the 'name' of the file, and a 'cookie' which is an opaque pointer to the next entry in the directory. The cookie is used in the next 'READDIR' call to get more entries starting at a given point in the directory. The special cookie zero (all bits zero) can be used to get the entries starting at the beginning of the directory. The 'fileid' field should be the same number as the 'fileid' in the the attributes of the file. (see Basic Data Types.) The 'eof' flag has a value of 'TRUE' if there are no more entries in the directory.

1.208 RPC.guide/Get Filesystem Attributes

Get Filesystem Attributes


```

union statfsres (stat status) {
    case NFS_OK:
        struct {
            unsigned tsize;
            unsigned bsize;
            unsigned blocks;
            unsigned bfree;
            unsigned bavail;
        } info;
    default:
        void;
};

statfsres
NFSPROC_STATFS(fhandle) = 17;

```

If the reply 'status' is 'NFS_OK', then the reply 'info' gives the attributes for the filesystem that contains file referred to by the input fhandle. The attribute fields contain the following values:

'tsize'

The optimum transfer size of the server in bytes. This is the number of bytes the server would like to have in the data part of READ and WRITE requests.

'bsize'

The block size in bytes of the filesystem.

'blocks'

The total number of 'bsize' blocks on the filesystem.

'bfree'

The number of free 'bsize' blocks on the filesystem.

'bavail'

The number of 'bsize' blocks available to non-privileged users.

Note: This call does not work well if a filesystem has variable size blocks.

1.209 RPC.guide/NFS Implementation Issues

NFS Implementation Issues
 =====

The NFS protocol is designed to be operating system independent, but since this version was designed in a UNIX environment, many operations have semantics similar to the operations of the UNIX file system. This section discusses some of the implementation-specific semantic issues.

Server-Client Relationship
Pathname Interpretation
Permission Issues
Setting RPC Parameters

1.210 RPC.guide/Server-Client Relationship

Server/Client Relationship

The NFS protocol is designed to allow servers to be as simple and general as possible. Sometimes the simplicity of the server can be a problem, if the client wants to implement complicated filesystem semantics.

For example, some operating systems allow removal of open files. A process can open a file and, while it is open, remove it from the directory. The file can be read and written as long as the process keeps it open, even though the file has no name in the filesystem. It is impossible for a stateless server to implement these semantics. The client can do some tricks such as renaming the file on remove, and only removing it on close. We believe that the server provides enough functionality to implement most file system semantics on the client.

Every NFS client can also potentially be a server, and remote and local mounted filesystems can be freely intermixed. This leads to some interesting problems when a client travels down the directory tree of a remote filesystem and reaches the mount point on the server for another remote filesystem. Allowing the server to follow the second remote mount would require loop detection, server lookup, and user revalidation. Instead, we decided not to let clients cross a server's mount point. When a client does a 'LOOKUP' on a directory on which the server has mounted a filesystem, the client sees the underlying directory instead of the mounted directory. A client can do remote mounts that match the server's mount points to maintain the server's view.

1.211 RPC.guide/Pathname Interpretation

Pathname Interpretation

There are a few complications to the rule that pathnames are always parsed on the client. For example, symbolic links could have different interpretations on different clients. Another common problem for non-UNIX implementations is the special interpretation of the pathname '..' to mean the parent of a given directory. The next revision of the protocol uses an explicit flag to indicate the parent instead.

1.212 RPC.guide/Permission Issues

Permission Issues

The NFS protocol, strictly speaking, does not define the permission checking used by servers. However, it is expected that a server will do normal operating system permission checking using 'AUTH_UNIX' style authentication as the basis of its protection mechanism. The server gets the client's effective "uid", effective "gid", and groups on each call and uses them to check permission. There are various problems with this method that can be resolved in interesting ways.

Using "uid" and "gid" implies that the client and server share the same "uid" list. Every server and client pair must have the same mapping from user to "uid" and from group to "gid". Since every client can also be a server, this tends to imply that the whole network shares the same "uid/gid" space. 'AUTH_DES' (and the next revision of the NFS protocol) uses string names instead of numbers, but there are still complex problems to be solved.

Another problem arises due to the usually stateful open operation. Most operating systems check permission at open time, and then check that the file is open on each read and write request. With stateless servers, the server has no idea that the file is open and must do permission checking on each read and write call. On a local filesystem, a user can open a file and then change the permissions so that no one is allowed to touch it, but will still be able to write to the file because it is open. On a remote filesystem, by contrast, the write would fail. To get around this problem, the server's permission checking algorithm should allow the owner of a file to access it regardless of the permission setting.

A similar problem has to do with paging in from a file over the network. The operating system usually checks for execute permission before opening a file for demand paging, and then reads blocks from the open file. The file may not have read permission, but after it is opened it doesn't matter. An NFS server can not tell the difference between a normal file read and a demand page-in read. To make this work, the server allows reading of files if the "uid" given in the call has execute or read permission on the file.

In most operating systems, a particular user (on the user ID zero) has access to all files no matter what permission and ownership they have. This "super-user" permission may not be allowed on the server, since anyone who can become super-user on their workstation could gain access to all remote files. The UNIX server by default maps user id 0 to -2 before doing its access checking. This works except for NFS root filesystems, where super-user access cannot be avoided.

1.213 RPC.guide/Setting RPC Parameters

Setting RPC Parameters

Various file system parameters and options should be set at mount time. The mount protocol is described in the appendix below. For example, "Soft" mounts as well as "Hard" mounts are usually both provided. Soft mounted file systems return errors when RPC operations fail (after a given number of optional retransmissions), while hard mounted file systems continue to retransmit forever. Clients and servers may need to keep caches of recent operations to help avoid problems with non-idempotent operations.

1.214 RPC.guide/Mount Protocol Definition

Mount Protocol Definition

=====

Mount Protocol Introduction
Mount Protocol RPC Information
Mount Protocol Sizes of XDR Structures
Mount Protocol Basic Data Types
Mount Protocol Server Procedures

1.215 RPC.guide/Mount Protocol Introduction

Introduction

The mount protocol is separate from, but related to, the NFS protocol. It provides operating system specific services to get the NFS off the ground - looking up server path names, validating user identity, and checking access permissions. Clients use the mount protocol to get the first file handle, which allows them entry into a remote filesystem.

The mount protocol is kept separate from the NFS protocol to make it easy to plug in new access checking and validation methods without changing the NFS server protocol.

Notice that the protocol definition implies stateful servers because the server maintains a list of client's mount requests. The mount list information is not critical for the correct functioning of either the client or the server. It is intended for advisory use only, for example, to warn possible clients when a server is going down.

Version one of the mount protocol is used with version two of the NFS protocol. The only connecting point is the 'fhandle' structure, which is the same for both protocols.

1.216 RPC.guide/Mount Protocol RPC Information

RPC Information

Authentication

The mount service uses 'AUTH_UNIX' and 'AUTH_DES' style authentication only.

Transport Protocols

The mount service is currently supported on UDP/IP only.

Port Number

Consult the server's portmapper, described in Remote Procedure Calls Protocol Specification, to find the port number on which the mount service is registered.

1.217 RPC.guide/Mount Protocol Sizes of XDR Structures

Sizes of XDR Structures

These are the sizes, given in decimal bytes, of various XDR structures used in the protocol:

```
/* The maximum number of bytes in a pathname argument */  
const MNTPATHLEN = 1024;
```

```
/* The maximum number of bytes in a name argument */  
const MNTNAMLEN = 255;
```

```
/* The size in bytes of the opaque file handle */  
const FHSIZE = 32;
```

1.218 RPC.guide/Mount Protocol Basic Data Types

Basic Data Types

This section presents the data types used by the mount protocol. In many cases they are similar to the types used in NFS.

```
fhandle (Mount)  
fhstatus  
dirpath  
name
```

1.219 RPC.guide/fhandle (Mount)

fhandle
.....

```
typedef opaque fhandle[FHSIZE];
```

The type 'fhandle' is the file handle that the server passes to the client. All file operations are done using file handles to refer to a file or directory. The file handle can contain whatever information the server needs to distinguish an individual file.

This is the same as the 'fhandle' XDR definition in version 2 of the NFS protocol; see Basic Data Types in the definition of the NFS protocol, above.

1.220 RPC.guide/fhstatus

fhstatus
.....

```
union fhstatus switch (unsigned status) {  
    case 0:  
        fhandle directory;  
    default:  
        void;  
};
```

The type 'fhstatus' is a union. If a 'status' of zero is returned, the call completed successfully, and a file handle for the 'directory' follows. A non-zero status indicates some sort of error. In this case the status is a UNIX error number.

1.221 RPC.guide/dirpath

dirpath
.....

```
typedef string dirpath<MNTPATHLEN>;
```

The type 'dirpath' is a server pathname of a directory.

1.222 RPC.guide/name

name
....

```
typedef string name<MNTNAMLEN>;
```

The type 'name' is an arbitrary string used for various names.

1.223 RPC.guide/Mount Protocol Server Procedures

Server Procedures

The following sections define the RPC procedures supplied by a mount server.

```
Do Nothing (Mount)
Add Mount Entry
Return Mount Entries
Remove Mount Entry
Remove All Mount Entries
Return Export List
```

```
/*
 * Protocol description for the mount program
 */
program MOUNTPROG {
/*
 * Version 1 of the mount protocol used with
 * version 2 of the NFS protocol.
 */
    version MOUNTVERS {
        void          MOUNTPROC_NULL(void)      = 0;
        fhstatus      MOUNTPROC_MNT(dirpath)    = 1;
        mountlist     MOUNTPROC_DUMP(void)      = 2;
        void          MOUNTPROC_UMNT(dirpath)   = 3;
        void          MOUNTPROC_UMNTALL(void)   = 4;
        exportlist    MOUNTPROC_EXPORT(void)    = 5;
    } = 1;
} = 100005;
```

1.224 RPC.guide/Do Nothing (Mount)

Do Nothing
.....

```
void
MNTPROC_NULL(void) = 0;
```

This procedure does no work. It is made available in all RPC services to allow server response testing and timing.

1.225 RPC.guide/Add Mount Entry

Add Mount Entry

.....

```
fhstatus
MNTPROC_MNT(dirpath) = 1;
```

If the reply 'status' is '0', then the reply 'directory' contains the file handle for the directory 'dirname'. This file handle may be used in the NFS protocol. This procedure also adds a new entry to the mount list for this client mounting 'dirname'.

1.226 RPC.guide/Return Mount Entries

Return Mount Entries

.....

```
struct *mountlist {
    name      hostname;
    dirpath   directory;
    mountlist nextentry;
};
```

```
mountlist
MNTPROC_DUMP(void) = 2;
```

Returns the list of remote mounted filesystems. The 'mountlist' contains one entry for each 'hostname' and 'directory' pair.

1.227 RPC.guide/Remove Mount Entry

Remove Mount Entry

.....

```
void
MNTPROC_UMNT(dirpath) = 3;
```

Removes the mount list entry for the input 'dirpath'.

1.228 RPC.guide/Remove All Mount Entries

Remove All Mount Entries

.....

```
void
```

```
MNTPROC_UMNTALL(void) = 4;
```

Removes all of the mount list entries for this client.

1.229 RPC.guide/Return Export List

Return Export List

.....

```
struct *groups {
    name grname;
    groups grnext;
};

struct *exportlist {
    dirpath filesys;
    groups groups;
    exportlist next;
};

exportlist
MNTPROC_EXPORT(void) = 5;
```

Returns a variable number of export list entries. Each entry contains a filesystem name and a list of groups that are allowed to import it. The filesystem name is in 'filesys', and the group name is in the list 'groups'.

Note: The exportlist should contain more information about the status of the filesystem, such as a read-only flag.

1.230 RPC.guide/Manual Pages

Manual Pages

All functions in these manual pages need the include file 'rpc/rpc.h' to be included.

Utilities:

portmap	The RPC program-to-port mapper
rpcgen	The RPC protocol compiler
rpcinfo	View RPC program information

'rpc' Data Base File Format:

rpc Data Base File Format	rpc program number database
---------------------------	-----------------------------

Library Functions:

Miscellaneous Utility Functions

RPC Function Reference
XDR Function Reference

1.231 RPC.guide/portmap Manual Page

`portmap' Manual Page
=====

NAME

`portmap' - DARPA port to RPC program number mapper

SYNOPSIS

`AmiTCP:bin/portmap'

DESCRIPTION

`portmap' is a server that converts RPC program numbers into DARPA protocol port numbers. It must be running in order to make RPC calls. When an RPC server is started, it will tell `portmap' what port number it is listening to, and what RPC program numbers it is prepared to serve. When a client wishes to make an RPC call to a given program number, it will first contact `portmap' on the server machine to determine the port number where RPC packets should be sent. Normally, standard RPC servers are started by `inetd', so `portmap' must be started before `inetd' is invoked.

SEE ALSO

rpcinfo.

BUGS

If `portmap' crashes, all servers must be restarted.

1.232 RPC.guide/rpcgen Manual Page

`rpcgen' Manual Page
=====

NAME

`rpcgen' - an RPC protocol compiler

SYNOPSIS

`rpcgen infile'

`rpcgen [-Dname[=value]] [-T] [-K secs] infile'

`rpcgen -c | -h | -l | -m | -t [-o outfile] infile'

`rpcgen [-I] -s nettype [-o outfile] infile'

`rpcgen -n netid [-o outfile] infile'

DESCRIPTION

'rpcgen' is a tool that generates C code to implement an RPC protocol. The input to 'rpcgen' is a language similar to C known as RPC Language (Remote Procedure Call Language).

'rpcgen' is normally used as in the first synopsis where it takes an input file and generates up to four output files. If the infile is named 'proto.x', then 'rpcgen' will generate a header file in 'proto.h', XDR routines in 'proto_xdr.c', server-side stubs in 'proto_svc.c', and client-side stubs in 'proto_clnt.c'. With the '-T' option, it will also generate the RPC dispatch table in 'proto_tbl.i'. With the '-Sc' option, it will also generate sample code which would illustrate how to use the remote procedures on the client side. This code would be created in 'proto_client.c'. With the '-Ss' option, it will also generate a sample server code which would illustrate how to write the remote procedures. This code would be created in 'proto_server.c'.

The server created can be started both by the port monitors (for example, inetd or listen) or by itself. When it is started by a port monitor, it creates servers only for the transport for which the file descriptor 0 was passed. The name of the transport must be specified by setting up the environmental variable PM_TRANSPORT. When the server generated by 'rpcgen' is executed, it creates server handles for all the transports specified in NETPATH environment variable, or if it is unset, it creates server handles for all the visible transports from '/etc/netconfig' file. Note: the transports are chosen at run time and not at compile time. When the server is self-started, it backgrounds itself by default. A special define symbol RPC_SVC_FG can be used to run the server process in foreground.

The second synopsis provides special features which allow for the creation of more sophisticated RPC servers. These features include support for user provided #defines and RPC dispatch tables. The entries in the RPC dispatch table contain:

- + pointers to the service routine corresponding to that procedure,
- + a pointer to the input and output arguments
- + the size of these routines

A server can use the dispatch table to check authorization and then to execute the service routine; a client library may use it to deal with the details of storage management and XDR data conversion.

The other three synopses shown above are used when one does not want to generate all the output files, but only a particular one. Some examples of their usage is described in the EXAMPLE section below. When 'rpcgen' is executed with the '-s' option, it creates servers for that particular class of transports. When executed with the '-n' option, it creates a server for the transport specified by netid. If infile is not specified, 'rpcgen' accepts the standard input.

The C preprocessor, `ccp`, is run on the input file before it is actually interpreted by `'rpcgen'`. For each type of output file, `'rpcgen'` defines a special preprocessor symbol for use by the `'rpcgen'` programmer:

```
RPC_HDR    defined when compiling into header files
RPC_XDR    defined when compiling into XDR routines
RPC_SVC    defined when compiling into server-side stubs
RPC_CLNT   defined when compiling into client-side stubs
RPC_TBL    defined when compiling into RPC dispatch tables
```

Any line beginning with `'%'` is passed directly into the output file, uninterpreted by `'rpcgen'`.

For every data type referred to in `infile`, `'rpcgen'` assumes that there exists a routine with the string `xdr_` prepended to the name of the data type. If this routine does not exist in the RPC/XDR library, it must be provided. Providing an undefined data type allows customization of XDR routines.

The following options are available:

- `'-a'`
Generate all the files including sample code for client and server side.
 - `'-b'`
This generates code for the SunOS4.1 style of rpc. It is for backward compatibility. This is the default.
 - `'-5'`
This generates code for the SysVr4 style of rpc. It is used by the Transport Independent RPC that is in Svr4 systems. By default `'rpcgen'` generates code for SunOS4.1 style of rpc.
 - `'-c'`
Compile into XDR routines.
 - `'-C'`
Generate code in ANSI C. This option also generates code that could be compiled with the C++ compiler. This is the default.
 - `'-k'`
Generate code in K&R C. The default is ANSI C.
 - `'-Dname[=value]'`
Define a symbol name. Equivalent to the `#define` directive in the source. If no value is given, value is defined as 1. This option may be specified more than once.
 - `'-h'`
Compile into C data-definitions (a header file). `'-T'` option can be used in conjunction to produce a header file which supports RPC dispatch tables.
 - `'-I'`
Generate a service that can be started from `inetd`. The
-

default is to generate a static service that handles transports selected with `'-s'`. Using `'-I'` allows starting a service by either method.

`'-K secs'`

By default, services created using `'rpcgen'` wait 120 seconds after servicing a request before exiting. That interval can be changed using the `'-K'` flag. To create a server that exits immediately upon servicing a request, `'-K 0'` can be used. To create a server that never exits, the appropriate argument is `'-K -1'`.

When monitoring for a server, some portmonitors, like `listen(1M)`, always spawn a new process in response to a service request. If it is known that a server will be used with such a monitor, the server should exit immediately on completion. For such servers, `'rpcgen'` should be used with `'-K -1'`.

`'-l'`

Compile into client-side stubs.

`'-m'`

Compile into server-side stubs, but do not generate a main routine. This option is useful for doing callback-routines and for users who need to write their own main routine to do initialization.

`'-n netid'`

Compile into server-side stubs for the transport specified by `netid`. There should be an entry for `netid` in the `netconfig` database. This option may be specified more than once, so as to compile a server that serves multiple transports.

`'-N'`

Use the newstyle of `'rpcgen'`. This allows procedures to have multiple arguments. It also uses the style of parameter passing that closely resembles C. So, when passing an argument to a remote procedure you do not have to pass a pointer to the argument but the argument itself. This behaviour is different from the oldstyle of `'rpcgen'` generated code. The newstyle is not the default case because of backward compatibility.

`'-o outfile'`

Specify the name of the output file. If none is specified, standard output is used (`'-c'`, `'-h'`, `'-l'`, `'-m'`, `'-n'`, `'-s'`, `'-sSc'`, `'-sSs'` and `'-t'` modes only).

`'-s nettype'`

Compile into server-side stubs for all the transports belonging to the class `nettype`. The supported classes are `netpath`, `visible`, `circuit_n`, `circuit_v`, `datagram_n`, `datagram_v`, `tcp`, and `udp` [see `rpc(3N)` for the meanings associated with these classes]. This option may be specified more than once. Note: the transports are chosen at run time and not at compile time.

- `'-Sc'`
Generate sample code to show the use of remote procedure and how to bind to the server before calling the client side stubs generated by `'rpcgen'`.
- `'-Ss'`
Generate skeleton code for the remote procedures on the server side. You would need to fill in the actual code for the remote procedures.
- `'-t'`
Compile into RPC dispatch table.
- `'-T'`
Generate the code to support RPC dispatch tables.

The options `'-c'`, `'-h'`, `'-l'`, `'-m'`, `'-s'` and `'-t'` are used exclusively to generate a particular type of file, while the options `'-D'` and `'-T'` are global and can be used with the other options.

NOTES

The RPC Language does not support nesting of structures. As a work-around, structures can be declared at the top-level, and their name used inside other structures in order to achieve the same effect.

Name clashes can occur when using program definitions, since the apparent scoping does not really apply. Most of these can be avoided by giving unique names for programs, versions, procedures and types.

The server code generated with `'-n'` option refers to the transport indicated by `netid` and hence is very site specific.

EXAMPLE

The following example:

```
$ rpcgen -T prot.x
```

generates the five files: `'prot.h'`, `'prot_clnt.c'`, `'prot_svc.c'`, `'prot_xdr.c'` and `'prot_tbl.i'`.

The following example sends the C data-definitions (header file) to the standard output.

```
$ rpcgen -h prot.x
```

To send the test version of the `'-DTEST'`, server side stubs for all the transport belonging to the class `datagram_n` to standard output, use:

```
$ rpcgen -s datagram_n -DTEST prot.x
```

To create the server side stubs for the transport indicated by `netid tcp`, use:

```
$ rpcgen -n tcp -o prot_svc.c prot.x
```

SEE ALSO

'cpp' manual

1.233 RPC.guide/rpcinfo Manual Page

'rpcinfo' Manual Page

=====

NAME

rpcinfo - report RPC information

SYNOPSIS

'rpcinfo -p [host]'

'rpcinfo [-n portnum] -u host program [version]'

'rpcinfo [-n portnum] -t host program [version]'

'rpcinfo -b program version'

'rpcinfo -d program version'

DESCRIPTION

'rpcinfo' makes an RPC call to an RPC server and reports what it finds.

OPTIONS

'-p'

Probe the portmapper on host, and print a list of all registered RPC programs. If host is not specified, it defaults to the value returned by 'gethostname()'.

'-u'

Make an RPC call to procedure 0 of program on the specified host using UDP, and report whether a response was received.

'-t'

Make an RPC call to procedure 0 of program on the specified host using TCP, and report whether a response was received.

'-n'

Use 'portnum' as the port number for the '-t' and '-u' options instead of the port number given by the portmapper.

'-b'

Make an RPC broadcast to procedure 0 of the specified program and version using UDP and report all hosts that respond.

'-d'

Delete registration for the RPC service of the specified program and version. This option can be exercised only by

the superuser. The program argument can be either a name or a number. If a version is specified, 'rpcinfo' attempts to call that version of the specified program. Otherwise, rpcinfo attempts to find all the registered version numbers for the specified program by calling version 0 (which is presumed not to exist; if it does exist, rpcinfo attempts to obtain this information by calling an extremely high version number instead) and attempts to call each registered version.

Note: the version number is required for -b and -d options.

EXAMPLES

To show all of the RPC services registered on the local machine use:

```
example% rpcinfo -p
```

To show all of the RPC services registered on the machine named klaxon use:

```
example% rpcinfo -p klaxon
```

To show all machines on the local net that are running the Yellow Pages service use:

```
example% rpcinfo -b ypserv 'version' | uniq
```

where 'version' is the current Yellow Pages version obtained from the results of the '-p' switch above.

To delete the registration for version 1 of the 'walld' service use:

```
example% rpcinfo -d walld 1
```

SEE ALSO

rpc, portmap.

1.234 RPC.guide/rpc Data Base File Format

rpc Data Base File Format

=====

The 'AmiTCP:db/rpc' file contains user readable names that can be used in place of rpc program numbers. Each line has the following information:

1. name of server for the rpc program
2. rpc program number
3. aliases

Items are separated by any number of blanks and/or tab characters. A

"#" indicates the beginning of a comment; characters up to the end of the line are not interpreted by routines which search the file. Here is an example of the 'rpc' file from the Sun RPC Source distribution.

```
#
# rpc 88/08/01 4.0 RPCSRC; from 1.12 88/02/07 SMI
#
portmapper          100000      portmap sunrpc
rstatd              100001      rstat rstat_svc rup perfmeter
rusersd             100002      rusers
nfs                 100003      nfsprog
ypserv              100004      ypprog
mountd              100005      mount showmount
ypbind              100007
walld               100008      rwall shutdown
yppasswdd           100009      yppasswd
etherstatd          100010      etherstat
rquotad             100011      rquotaprogram quota rquota
sprayd              100012      spray
3270_mapper         100013
rje_mapper          100014
selection_svc       100015      selnsvc
database_svc        100016
rex                 100017      rex
alis                100018
sched               100019
llockmgr            100020
nlockmgr            100021
x25.inr             100022
statmon             100023
status              100024
bootparam           100026
ypupdated           100028      ypupdate
keyserv             100029      keyserver
tfsd                100037
nsd                 100038
nsemntd             100039
```

1.235 RPC.guide/Miscellaneous Utility Functions

Miscellaneous Utility Functions

=====

```
bindresvport ()
getrpcbyname ()
getrpcbynumber ()
getrpcport ()
```

1.236 RPC.guide/bindresvport()

bindresvport ()

Bind a socket to a privileged IP port.

```
#include <sys/types.h>
#include <netinet/in.h>

int
bindresvport(sd, sin)
int sd;
struct sockaddr_in *sin;
```

bindresvport() is used to bind a socket descriptor to a privileged IP port, that is, a port number in the range 0-1023. The routine returns 0 if it is successful, otherwise -1 is returned and errno set to reflect the cause of the error. Only root can bind to a privileged port; this call will fail for any other users.

1.237 RPC.guide/getrpcbyname()

getrpcbyname ()

Get an RPC entry.

```
#include <netdb.h>

struct rpcent *
getrpcbyname(name)
char *name;
```

Return a pointer to an object with the following structure containing the broken-out fields of a line in the rpc program number data base, 'AmiTCP:db/rpc'.

```
struct    rpcent {
    char *r_name;          /* name of server for this rpc program */
    char **r_aliases;     /* alias list */
    long r_number;        /* rpc program number */
};
```

The members of this structure are:

'r_name'

The name of the server for this rpc program.

'r_aliases'

A zero terminated list of alternate names for the rpc program.

'number'

The rpc program number for this service.

'getrpcbyname()' sequentially searches from the beginning of the

file until a matching rpc program 'name' is found, or until end-of-file is encountered.

A 'NULL' pointer is returned on 'EOF' or error.

NOTE: All information is contained in a static area so it must be copied if it is to be saved.

See also:

=====

getrpcbynumber(), rpc, rpcinfo.

1.238 RPC.guide/getrpcbynumber()

getrpcbynumber()

Get an RPC entry.

```
#include <netdb.h>

struct rpcent *
getrpcbynumber(number)
int number;
```

Return a pointer to an object with the following structure containing the broken-out fields of a line in the rpc program number data base, 'AmiTCP:db/rpc'.

```
struct    rpcent {
    char *r_name;          /* name of server for this rpc program */
    char **r_aliases;     /* alias list */
    long r_number;        /* rpc program number */
};
```

The members of this structure are:

'r_name'

The name of the server for this rpc program.

'r_aliases'

A zero terminated list of alternate names for the rpc program.

'number'

The rpc program number for this service.

'getrpcbynumber()' sequentially searches from the beginning of the file until a matching rpc program 'number' is found, or until end-of-file is encountered.

A 'NULL' pointer is returned on 'EOF' or error.

NOTE: All information is contained in a static area so it must be copied if it is to be saved.

See also:

=====

getrpcbyname(), rpc, rpcinfo.

1.239 RPC.guide/getrpcport()

getrpcport()

Get RPC port number.

```
int
getrpcport(host, prognum, versnum, proto)
char *host;
int prognum, versnum, proto;
```

Returns the port number for version 'versnum' of the RPC program 'prognum' running on 'host' and using protocol 'proto'. It returns 0 if it cannot contact the portmapper, or if 'prognum' is not registered. If 'prognum' is registered but not with version 'versnum', it will still return a port number (for some version of the program) indicating that the program is indeed registered. The version mismatch will be detected upon the first call to the service.

1.240 RPC.guide/RPC Function Reference

RPC Function Reference

=====

These routines allow C programs to make procedure calls on other machines across the network. First, the client calls a procedure to send a data packet to the server. Upon receipt of the packet, the server calls a dispatch routine to perform the requested service, and then sends back a reply. Finally, the procedure call returns to the client.

Routines that are used for Secure RPC (DES authentication) are not available in this release.

```
auth_destroy()
authnone_create()
authunix_create()
authunix_create_default()
callrpc()
clnt_broadcast()
clnt_call()
clnt_destroy()
```

```
clnt_create()
clnt_control()
clnt_freeres()
clnt_geterr()
clnt_pcreateerror()
clnt_perrno()
clnt_perror()
clnt_spcreateerror()
clnt_sperrno()
clnt_sperror()
clntraw_create()
clnttcp_create()
clntudp_create()
clntudp_bufcreate()
get_myaddress()
pmap_getmaps()
pmap_getport()
pmap_rmtcall()
pmap_set()
pmap_unset()
registerrpc()
rpc_createerr
svc_destroy()
svc_fdset
svc_fds
svc_freeargs()
svc_getargs()
svc_getcaller()
svc_getreqset()
svc_getreq()
svc_register()
svc_run()
svc_sendreply()
svc_unregister()
svcerr_auth()
svcerr_decode()
svcerr_noproc()
svcerr_noprog()
svcerr_progvers()
svcerr_systemerr()
svcerr_weakauth()
svcrw_create()
svctcp_create()
svcfid_create()
svcudp_bufcreate()
svcudp_create()
xdr_accepted_reply()
xdr_authunix_parms()
xdr_callhdr()
xdr_callmsg()
xdr_opaque_auth()
xdr_pmap()
xdr_pmaplist()
xdr_rejected_reply()
xdr_replymsg()
xpirt_register()
xpirt_unregister()
```

1.241 RPC.guide/auth_destroy()

auth_destroy()

```
void
auth_destroy(auth)
AUTH *auth;
```

A macro that destroys the authentication information associated with 'auth'. Destruction usually involves deallocation of private data structures. The use of 'auth' is undefined after calling 'auth_destroy()'.

1.242 RPC.guide/authnone_create()

authnone_create()

```
AUTH *
authnone_create()
```

Create and returns an 'RPC' authentication handle that passes nonusable authentication information with each remote procedure call. This is the default authentication used by 'RPC.'

1.243 RPC.guide/authunix_create()

authunix_create()

```
AUTH *
authunix_create(host, uid, gid, len, aup_gids)
char *host;
int uid, gid, len, *aup_gids;
```

Create and return an 'RPC' authentication handle that contains Unix authentication information. The parameter 'host' is the name of the machine on which the information was created; 'uid' is the user's user 'ID'; 'gid' is the user's current group 'ID'; 'len' and 'aup_gids' refer to a counted array of groups to which the user belongs. It is easy to impersonate a user.

1.244 RPC.guide/authunix_create_default()

```
authunix_create_default()
-----
```

```
AUTH *
authunix_create_default()
```

Calls `authunix_create()` with the appropriate parameters.

1.245 RPC.guide/callrpc()

```
callrpc()
-----
```

```
int
callrpc(host, prognum, versnum, procnum, inproc, in, outproc, out)
char *host;
u_long prognum, versnum, procnum;
char *in, *out;
xdrproc_t inproc, outproc;
```

Call the remote procedure associated with `'prognum'`, `'versnum'`, and `'procnum'` on the machine, `'host'`. The parameter `'in'` is the address of the procedure's argument(s), and `'out'` is the address of where to place the result(s); `'inproc'` is used to encode the procedure's parameters, and `'outproc'` is used to decode the procedure's results. This routine returns zero if it succeeds, or the value of `'enum clnt_stat'` cast to an integer if it fails. The routine `clnt_perrno()` is handy for translating failure statuses into messages.

**Warning:* calling remote procedures with this routine uses `'UDP/IP'` as a transport; see `clntudp_create()` for restrictions. You do not have control of timeouts or authentication using this routine.

1.246 RPC.guide/clnt_broadcast()

```
clnt_broadcast()
-----
```

```
enum clnt_stat
clnt_broadcast(prognum, versnum, procnum, inproc, in, outproc, out, ←
eachresult)
u_long prognum, versnum, procnum;
char *in, *out;
xdrproc_t inproc, outproc;
resultproc_t eachresult;
```

Like `callrpc()`, except the call message is broadcast to all locally connected broadcast nets. Each time it receives a response, this

routine calls `'eachresult()'`, whose form is:

```
int
eachresult(out, addr)
char *out;
struct sockaddr_in *addr;
```

where `'out'` is the same as `'out'` passed to `'clnt_broadcast()'`, except that the remote procedure's output is decoded there; `'addr'` points to the address of the machine that sent the results. If `'eachresult()'` returns zero, `'clnt_broadcast()'` waits for more replies; otherwise it returns with appropriate status.

**Warning:* broadcast sockets are limited in size to the maximum transfer unit of the data link. For ethernet, this value is 1500 bytes.

1.247 RPC.guide/clnt_call()

`clnt_call()`

```
enum clnt_stat
clnt_call(clnt, procnum, inproc, in, outproc, out, tout)
CLIENT *clnt;
u_long procnum;
xdrproc_t inproc, outproc;
char *in, *out;
struct timeval tout;
```

A macro that calls the remote procedure `'procnum'` associated with the client handle, `'clnt'`, which is obtained with an `'RPC'` client creation routine such as `clnt_create()`. The parameter `'in'` is the address of the procedure's argument(s), and `'out'` is the address of where to place the result(s); `'inproc'` is used to encode the procedure's parameters, and `'outproc'` is used to decode the procedure's results; `'tout'` is the time allowed for results to come back.

1.248 RPC.guide/clnt_destroy()

`clnt_destroy()`

```
void
clnt_destroy(clnt)
CLIENT *clnt;
```

A macro that destroys the client's `'RPC'` handle. Destruction usually involves deallocation of private data structures, including `'clnt'` itself. Use of `'clnt'` is undefined after calling `'clnt_destroy()'`. If the `'RPC'` library opened the associated socket, it will close it also.

Otherwise, the socket remains open.

1.249 RPC.guide/clnt_create()

clnt_create()

```
CLIENT *
clnt_create(host, prog, vers, proto)
char *host;
u_long prog, vers;
char *proto;
```

Generic client creation routine. 'host' identifies the name of the remote host where the server is located. 'proto' indicates which kind of transport protocol to use. The currently supported values for this field are "udp" and "tcp". Default timeouts are set, but can be modified using clnt_control().

Warning: Using 'UDP' has its shortcomings. Since 'UDP'-based 'RPC' messages can only hold up to 8 Kbytes of encoded data, this transport cannot be used for procedures that take large arguments or return huge results.

1.250 RPC.guide/clnt_control()

clnt_control()

```
bool_t
clnt_control(cl, req, info)
CLIENT *cl;
char *info;
```

A macro used to change or retrieve various information about a client object. 'req' indicates the type of operation, and 'info' is a pointer to the information. For both 'UDP' and 'TCP', the supported values of 'req' and their argument types and what they do are:

CLSET_TIMEOUT	struct timeval	set total timeout
CLGET_TIMEOUT	struct timeval	get total timeout

Note: if you set the timeout using 'clnt_control()', the timeout parameter passed to clnt_call() will be ignored in all future calls.

CLGET_SERVER_ADDR	struct sockaddr_in	get server's address
-------------------	--------------------	----------------------

The following operations are valid for 'UDP' only:

CLSET_RETRY_TIMEOUT	struct timeval	set the retry timeout
---------------------	----------------	-----------------------

```
CLGET_RETRY_TIMEOUT struct timeval    get the retry timeout
```

The retry timeout is the time that 'UDP RPC' waits for the server to reply before retransmitting the request.

1.251 RPC.guide/clnt_freeres()

```
clnt_freeres()
```

```
-----
```

```
bool_t
clnt_freeres(clnt, outproc, out)
CLIENT *clnt;
xdrproc_t outproc;
char *out;
```

A macro that frees any data allocated by the 'RPC/XDR' system when it decoded the results of an 'RPC' call. The parameter 'out' is the address of the results, and 'outproc' is the 'XDR' routine describing the results. This routine returns one if the results were successfully freed, and zero otherwise.

1.252 RPC.guide/clnt_geterr()

```
clnt_geterr()
```

```
-----
```

```
void
clnt_geterr(clnt, errp)
CLIENT *clnt;
struct rpc_err *errp;
```

A macro that copies the error structure out of the client handle to the structure at address 'errp'.

1.253 RPC.guide/clnt_pcreateerror()

```
clnt_pcreateerror()
```

```
-----
```

```
void
clnt_pcreateerror(s)
char *s;
```

Print a message to standard error indicating why a client 'RPC' handle could not be created. The message is prepended with string 's' and a colon. Used when a `clnt_create()`, `clntraw_create()`, `clnttcp_create()`,

or `clntudp_create()` call fails.

1.254 RPC.guide/clnt_perrno()

`clnt_perrno()`

```
void
clnt_perrno(stat)
enum clnt_stat stat;
```

Print a message to standard error corresponding to the condition indicated by 'stat'. Used after `callrpc()`.

1.255 RPC.guide/clnt_perror()

`clnt_perror()`

```
void
clnt_perror(clnt, s)
CLIENT *clnt;
char *s;
```

Print a message to standard error indicating why an 'RPC' call failed; 'clnt' is the handle used to do the call. The message is prepended with string 's' and a colon. Used after `clnt_call()`.

1.256 RPC.guide/clnt_screateerror()

`clnt_screateerror()`

```
char *
clnt_screateerror(s)
char *s;
```

Like `clnt_pcreateerror()`, except that it returns a string instead of printing to the standard error.

Bugs: returns pointer to static data that is overwritten on each call.

1.257 RPC.guide/clnt_sperrno()

```
clnt_sperrno()
-----
```

```
char *
clnt_sperrno(stat)
enum clnt_stat stat;
```

Take the same arguments as `clnt_perrno()`, but instead of sending a message to the standard error indicating why an 'RPC' call failed, return a pointer to a string which contains the message. The string ends with a '\n'. `'clnt_sperrno()'` is used instead of `clnt_perrno()` if the program does not have a standard error (as a program running as a server quite likely does not), or if the programmer does not want the message to be output with `'printf'`, or if a message format different than that supported by `clnt_perrno()` is to be used.

**Note:* unlike `clnt_sperror()` and `clnt_spcreateerror()`, `'clnt_sperrno()'` returns pointer to static data, but the result will not get overwritten on each call.

1.258 RPC.guide/clnt_sperror()

```
clnt_sperror()
-----
```

```
char *
clnt_sperror(rpch, s)
CLIENT *rpch;
char *s;
```

Like `clnt_perror()`, except that (like `clnt_sperrno()`) it returns a string instead of printing to standard error.

**Bugs:* returns pointer to static data that is overwritten on each call.

1.259 RPC.guide/clntraw_create()

```
clntraw_create()
-----
```

```
CLIENT *
clntraw_create(prognum, versnum)
u_long prognum, versnum;
```

This routine creates a toy 'RPC' client for the remote program `'prognum'`, version `'versnum'`. The transport used to pass messages to the service is actually a buffer within the process's address space, so

the corresponding 'RPC' server should live in the same address space; see `svcrw_create()`. This allows simulation of 'RPC' and acquisition of 'RPC' overheads, such as round trip times, without any kernel interference. This routine returns 'NULL' if it fails.

1.260 RPC.guide/clnttcp_create()

`clnttcp_create()`

```
CLIENT *
clnttcp_create(addr, prognum, versnum, sockp, sendsz, recvsz)
struct sockaddr_in *addr;
u_long prognum, versnum;
int *sockp;
u_int sendsz, recvsz;
```

This routine creates an 'RPC' client for the remote program 'prognum', version 'versnum'; the client uses 'TCP/IP' as a transport. The remote program is located at Internet address '*addr'. If 'addr->sin_port' is zero, then it is set to the actual port that the remote program is listening on (the remote 'portmap' service is consulted for this information). The parameter 'sockp' is a socket; if it is 'RPC_ANYSOCK', then this routine opens a new one and sets 'sockp'. Since 'TCP'-based 'RPC' uses buffered 'I/O', the user may specify the size of the send and receive buffers with the parameters 'sendsz' and 'recvsz'; values of zero choose suitable defaults. This routine returns 'NULL' if it fails.

1.261 RPC.guide/clntudp_create()

`clntudp_create()`

```
CLIENT *
clntudp_create(addr, prognum, versnum, wait, sockp)
struct sockaddr_in *addr;
u_long prognum, versnum;
struct timeval wait;
int *sockp;
```

This routine creates an 'RPC' client for the remote program 'prognum', version 'versnum'; the client uses 'UDP/IP' as a transport. The remote program is located at Internet address 'addr'. If 'addr->sin_port' is zero, then it is set to actual port that the remote program is listening on (the remote 'portmap' service is consulted for this information). The parameter 'sockp' is a socket; if it is 'RPC_ANYSOCK', then this routine opens a new one and sets 'sockp'. The 'UDP' transport resends the call message in intervals of 'wait' time until a response is received or until the call times out. The total

time for the call to time out is specified by `clnt_call()`.

**Warning:* since 'UDP'-based 'RPC' messages can only hold up to 8 Kbytes of encoded data, this transport cannot be used for procedures that take large arguments or return huge results.

1.262 RPC.guide/clntudp_bufcreate()

`clntudp_bufcreate()`

```
CLIENT *
clntudp_bufcreate(addr, prognum, versnum, wait, sockp, sendsize, recosize)
struct sockaddr_in *addr;
u_long prognum, versnum;
struct timeval wait;
int *sockp;
unsigned int sendsize;
unsigned int recosize;
```

This routine creates an 'RPC' client for the remote program 'prognum', on 'versnum'; the client uses 'UDP/IP' as a transport. The remote program is located at Internet address 'addr'. If 'addr->sin_port' is zero, then it is set to actual port that the remote program is listening on (the remote 'portmap' service is consulted for this information). The parameter 'sockp' is a socket; if it is 'RPC_ANYSOCK', then this routine opens a new one and sets 'sockp'. The 'UDP' transport resends the call message in intervals of 'wait' time until a response is received or until the call times out. The total time for the call to time out is specified by `clnt_call()`. This allows the user to specify the maximum packet size for sending and receiving 'UDP-based' 'RPC' messages.

1.263 RPC.guide/get_myaddress()

`get_myaddress()`

```
void
get_myaddress(addr)
struct sockaddr_in *addr;
```

Stuff the machine's 'IP' address into '*addr', without consulting the library routines that deal with host database or name service. The port number is always set to 'htons(PMAPPORT)'.

1.264 RPC.guide/pmap_getmaps()

pmap_getmaps ()

```
struct pmaplist *
pmap_getmaps(addr)
struct sockaddr_in *addr;
```

A user interface to the 'portmap' service, which returns a list of the current 'RPC' program-to-port mappings on the host located at 'IP' address '*addr'. This routine can return 'NULL'. The command portmap uses this routine.

1.265 RPC.guide/pmap_getport()

pmap_getport ()

```
u_short
pmap_getport(addr, prognum, versnum, protocol)
struct sockaddr_in *addr;
u_long prognum, versnum, protocol;
```

A user interface to the 'portmap' service, which returns the port number on which waits a service that supports program number 'prognum', version 'versnum', and speaks the transport protocol associated with 'protocol'. The value of 'protocol' is most likely 'IPPROTO_UDP' or 'IPPROTO_TCP'. A return value of zero means that the mapping does not exist or that the 'RPC' system failed to contact the remote 'portmap' service. In the latter case, the global variable rpc_createerr contains the 'RPC' status.

1.266 RPC.guide/pmap_rmtcall()

pmap_rmtcall ()

```
enum clnt_stat
pmap_rmtcall(addr, prognum, versnum, procnum,
             inproc, in, outproc, out, tout, portp)
struct sockaddr_in *addr;
u_long prognum, versnum, procnum;
char *in, *out;
xdrproc_t inproc, outproc;
struct timeval tout;
u_long *portp;
```

A user interface to the 'portmap' service, which instructs 'portmap' on the host at 'IP' address '*addr' to make an 'RPC' call on your behalf to a procedure on that host. The parameter '*portp' will be modified

to the program's port number if the procedure succeeds. The definitions of other parameters are discussed in `callrpc()` and `clnt_call()`. This procedure should be used for a "ping" and nothing else. See also `clnt_broadcast()`.

1.267 RPC.guide/pmap_set()

`pmap_set()`

```
bool_t
pmap_set(prognum, versnum, protocol, port)
u_long prognum, versnum, protocol;
u_short port;
```

A user interface to the 'portmap' service, which establishes a mapping between the triple [prognum,versnum,protocol] and 'port' on the machine's 'portmap' service. The value of 'protocol' is most likely 'IPPROTO_UDP' or 'IPPROTO_TCP'. This routine returns one if it succeeds, zero otherwise. Automatically done by `svc_register()`.

1.268 RPC.guide/pmap_unset()

`pmap_unset()`

```
bool_t
pmap_unset(prognum, versnum)
u_long prognum, versnum;
```

A user interface to the 'portmap' service, which destroys all mapping between the triple [prognum,versnum,*] and 'ports' on the machine's 'portmap' service. This routine returns one if it succeeds, zero otherwise.

1.269 RPC.guide/registerrpc()

`registerrpc()`

```
int
registerrpc(prognum, versnum, procnum, procname, inproc, outproc)
u_long prognum, versnum, procnum;
char *(*procname) ();
xdrproc_t inproc, outproc;
```

Register procedure 'procname' with the 'RPC' service package. If a

request arrives for program 'prognum', version 'versnum', and procedure 'procnum', 'procname' is called with a pointer to its parameter(s); 'progname' should return a pointer to its static result(s); 'inproc' is used to decode the parameters while 'outproc' is used to encode the results. This routine returns zero if the registration succeeded, -1 otherwise.

Warning: remote procedures registered in this form are accessed using the 'UDP/IP' transport; see svcudp_create() for restrictions.

1.270 RPC.guide/rpc_createerr

rpc_createerr

```
struct rpc_createerr    rpc_createerr;
```

A global variable whose value is set by any 'RPC' client creation routine that does not succeed. Use the routine clnt_pcreateerror() to print the reason why.

1.271 RPC.guide/svc_destroy()

svc_destroy()

```
void
svc_destroy(xprt)
SVCXPRT *xprt;
```

A macro that destroys the 'RPC' service transport handle, 'xprt'. Destruction usually involves deallocation of private data structures, including 'xprt' itself. Use of 'xprt' is undefined after calling this routine.

1.272 RPC.guide/svc_fdset

svc_fdset

```
fd_set svc_fdset;
```

A global variable reflecting the 'RPC' service side's read file descriptor bit mask; it is suitable as a parameter to the 'select' system call. This is only of interest if a service implementor does not call svc_run(), but rather does his own asynchronous event processing. This variable is read-only (do not pass its address to 'select'!), yet

it may change after calls to `svc_getreqset()` or any creation routines.

1.273 RPC.guide/svc_fds

`svc_fds`

```
int svc_fds;
```

Similar to `svc_fdset`, but limited to 32 descriptors. This interface is obsoleted by `svc_fdset`.

1.274 RPC.guide/svc_freeargs()

`svc_freeargs()`

```
bool_t
svc_freeargs(xprt, inproc, in)
SVCXPRT *xprt;
xdrproc_t inproc;
char *in;
```

A macro that frees any data allocated by the 'RPC/XDR' system when it decoded the arguments to a service procedure using `svc_getargs()`. This routine returns 1 if the results were successfully freed, and zero otherwise.

1.275 RPC.guide/svc_getargs()

`svc_getargs()`

```
bool_t
svc_getargs(xprt, inproc, in)
SVCXPRT *xprt;
xdrproc_t inproc;
char *in;
```

A macro that decodes the arguments of an 'RPC' request associated with the 'RPC' service transport handle, 'xprt'. The parameter 'in' is the address where the arguments will be placed; 'inproc' is the 'XDR' routine used to decode the arguments. This routine returns one if decoding succeeds, and zero otherwise.

1.276 RPC.guide/svc_getcaller()

svc_getcaller()

```
struct sockaddr_in *
svc_getcaller(xprt)
SVCXPRT *xprt;
```

The approved way of getting the network address of the caller of a procedure associated with the 'RPC' service transport handle, 'xprt'.

1.277 RPC.guide/svc_getreqset()

svc_getreqset()

```
void
svc_getreqset(rdfds)
fd_set *rdfds;
```

This routine is only of interest if a service implementor does not call `svc_run()`, but instead implements custom asynchronous event processing. It is called when the 'select()' system call has determined that an 'RPC' request has arrived on some 'RPC' 'socket(s)'; 'rdfds' is the resultant read file descriptor bit mask. The routine returns when all sockets associated with the value of 'rdfds' have been serviced.

1.278 RPC.guide/svc_getreq()

svc_getreq()

```
void
svc_getreq(rdfds)
int rdfds;
```

Similar to `svc_getreqset()`, but limited to 32 descriptors. This interface is obsoleted by `svc_getreqset()`.

1.279 RPC.guide/svc_register()

svc_register()

```
bool_t
```

```

svc_register(xprt, prognum, versnum, dispatch, protocol)
SVCXPRT *xprt;
u_long prognum, versnum;
void (*dispatch) ();
u_long protocol;

```

Associates 'prognum' and 'versnum' with the service dispatch procedure, 'dispatch'. If 'protocol' is zero, the service is not registered with the 'portmap' service. If 'protocol' is non-zero, then a mapping of the triple [prognum,versnum,protocol] to 'xprt->xp_port' is established with the local 'portmap' service (generally 'protocol' is zero, 'IPPROTO_UDP' or 'IPPROTO_TCP'). The procedure 'dispatch' has the following form:

```

void
dispatch(request, xprt)
struct svc_req *request;
SVCXPRT *xprt;

```

The 'svc_register()' routine returns one if it succeeds, and zero otherwise.

1.280 RPC.guide/svc_run()

```

svc_run()
-----

```

```

void
svc_run()

```

This routine never returns. It waits for 'RPC' requests to arrive, and calls the appropriate service procedure using svc_getreqset() when one arrives. This procedure is usually waiting for a 'select()' system call to return.

1.281 RPC.guide/svc_sendreply()

```

svc_sendreply()
-----

```

```

bool_t
svc_sendreply(xprt, outproc, out)
SVCXPRT *xprt;
xdrproc_t outproc;
char *out;

```

Called by an 'RPC' service's dispatch routine to send the results of a remote procedure call. The parameter 'xprt' is the request's associated transport handle; 'outproc' is the 'XDR' routine which is used to encode the results; and 'out' is the address of the results.

This routine returns one if it succeeds, zero otherwise.

1.282 RPC.guide/svc_unregister()

svc_unregister()

```
void
svc_unregister(prognum, versnum)
u_long prognum, versnum;
```

Remove all mapping of the double [prognum,versnum] to dispatch routines, and of the triple [prognum,versnum,*] to port number.

1.283 RPC.guide/svcerr_auth()

svcerr_auth()

```
void
svcerr_auth(xprt, why)
SVCXPRT *xprt;
enum auth_stat why;
```

Called by a service dispatch routine that refuses to perform a remote procedure call due to an authentication error.

1.284 RPC.guide/svcerr_decode()

svcerr_decode()

```
void
svcerr_decode(xprt)
SVCXPRT *xprt;
```

Called by a service dispatch routine that cannot successfully decode its parameters. Also See svc_getargs().

1.285 RPC.guide/svcerr_noproc()

```
svcerr_noproc ()
```

```
-----
```

```
void
svcerr_noproc(xprt)
SVCXPRT *xprt;
```

Called by a service dispatch routine that does not implement the procedure number that the caller requests.

1.286 RPC.guide/svcerr_noprog()

```
svcerr_noprog ()
```

```
-----
```

```
void
svcerr_noprog(xprt)
SVCXPRT *xprt;
```

Called when the desired program is not registered with the 'RPC' package. Service implementors usually do not need this routine.

1.287 RPC.guide/svcerr_progvers()

```
svcerr_progvers ()
```

```
-----
```

```
void
svcerr_progvers(xprt)
SVCXPRT *xprt;
```

Called when the desired version of a program is not registered with the 'RPC' package. Service implementors usually do not need this routine.

1.288 RPC.guide/svcerr_systemerr()

```
svcerr_systemerr ()
```

```
-----
```

```
void
svcerr_systemerr(xprt)
SVCXPRT *xprt;
```

Called by a service dispatch routine when it detects a system error not covered by any particular protocol. For example, if a service can no longer allocate storage, it may call this routine.

1.289 RPC.guide/svcerr_weakauth()

svcerr_weakauth()

```
void
svcerr_weakauth(xprt)
SVCXPRT *xprt;
```

Called by a service dispatch routine that refuses to perform a remote procedure call due to insufficient authentication parameters. The routine calls 'svcerr_auth(xprt, AUTH_TOOWEAK)'.

1.290 RPC.guide/svcraw_create()

svcraw_create()

```
SVCXPRT *
svcraw_create()
```

This routine creates a toy 'RPC' service transport, to which it returns a pointer. The transport is really a buffer within the process's address space, so the corresponding 'RPC' client should live in the same address space; see clntraw_create(). This routine allows simulation of 'RPC' and acquisition of 'RPC' overheads (such as round trip times), without any kernel interference. This routine returns 'NULL' if it fails.

1.291 RPC.guide/svctcp_create()

svctcp_create()

```
SVCXPRT *
svctcp_create(sock, send_buf_size, recv_buf_size)
int sock;
u_int send_buf_size, recv_buf_size;
```

This routine creates a 'TCP/IP'-based 'RPC' service transport, to which it returns a pointer. The transport is associated with the socket 'sock', which may be 'RPC_ANYSOCK', in which case a new socket is created. If the socket is not bound to a local 'TCP' port, then this routine binds it to an arbitrary port. Upon completion, 'xprt->xp_sock' is the transport's socket descriptor, and 'xprt->xp_port' is the transport's port number. This routine returns 'NULL' if it fails. Since 'TCP'-based 'RPC' uses buffered 'I/O', users may specify the size of buffers; values of zero choose suitable defaults.

1.292 RPC.guide/svcfd_create()

svcfd_create()

```
SVCXPRT *
svcfd_create(fd, sendsize, recvsize)
int fd;
u_int sendsize;
u_int recvsize;
```

Create a service on top of any open descriptor. Typically, this descriptor is a connected socket for a stream protocol such as 'TCP.' 'sendsize' and 'recvsize' indicate sizes for the send and receive buffers. If they are zero, a reasonable default is chosen.

1.293 RPC.guide/svcudp_bufcreate()

svcudp_bufcreate()

```
SVCXPRT *
svcudp_bufcreate(sock, sendsize, recvsize)
int sock;
u_int sendsize, recvsize;
```

This routine creates a 'UDP/IP'-based 'RPC' service transport, to which it returns a pointer. The transport is associated with the socket 'sock', which may be 'RPC_ANYSOCK', in which case a new socket is created. If the socket is not bound to a local 'UDP' port, then this routine binds it to an arbitrary port. Upon completion, 'xpirt->xp_sock' is the transport's socket descriptor, and 'xpirt->xp_port' is the transport's port number. This routine returns 'NULL' if it fails. This allows the user to specify the maximum packet size for sending and receiving 'UDP'-based 'RPC' messages.

1.294 RPC.guide/svcudp_create()

svcudp_create()

```
SVCXPRT *
svcudp_create(sock)
int sock;
```

This routine creates a 'UDP/IP'-based 'RPC' service transport, to which it returns a pointer. This is done by calling the see `svcudp_bufcreate()` with the default buffer size 'UDPMSSIZE'.

1.295 RPC.guide/xdr_accepted_reply()

xdr_accepted_reply()

```
bool_t XDRFUN
xdr_accepted_reply(xdrs, ar)
XDR *xdrs;
struct accepted_reply *ar;
```

Used for encoding 'RPC' reply messages. This routine is useful for users who wish to generate RPC-style messages without using the 'RPC' package.

1.296 RPC.guide/xdr_authunix_parms()

xdr_authunix_parms()

```
bool_t XDRFUN
xdr_authunix_parms(xdrs, aupp)
XDR *xdrs;
struct authunix_parms *aupp;
```

Used for describing 'UNIX' credentials. This routine is useful for users who wish to generate these credentials without using the 'RPC' authentication package.

1.297 RPC.guide/xdr_callhdr()

xdr_callhdr()

```
void
xdr_callhdr(xdrs, chdr)
XDR *xdrs;
struct rpc_msg *chdr;
```

Used for describing 'RPC' call header messages. This routine is useful for users who wish to generate 'RPC-style' messages without using the 'RPC' package.

1.298 RPC.guide/xdr_callmsg()

xdr_callmsg()

```
bool_t XDRFUN
xdr_callmsg(xdrs, cmsg)
XDR *xdrs;
struct rpc_msg *cmsg;
```

Used for describing 'RPC' call messages. This routine is useful for users who wish to generate 'RPC-style' messages without using the 'RPC' package.

1.299 RPC.guide/xdr_opaque_auth()

```
xdr_opaque_auth()
```

```
-----
```

```
bool_t XDRFUN
xdr_opaque_auth(xdrs, ap)
XDR *xdrs;
struct opaque_auth *ap;
```

Used for describing 'RPC' authentication information messages. This routine is useful for users who wish to generate 'RPC-style' messages without using the 'RPC' package.

1.300 RPC.guide/xdr_pmap()

```
xdr_pmap()
```

```
-----
```

```
bool_t XDRFUN
xdr_pmap(xdrs, regs)
XDR *xdrs;
struct pmap *regs;
```

Used for describing parameters to various 'portmap' procedures, externally. This routine is useful for users who wish to generate these parameters without using the 'pmap' interface.

1.301 RPC.guide/xdr_pmaplist()

```
xdr_pmaplist()
```

```
-----
```

```
bool_t XDRFUN
xdr_pmaplist(xdrs, rp)
XDR *xdrs;
struct pmaplist **rp;
```

Used for describing a list of port mappings, externally. This routine is useful for users who wish to generate these parameters without using the 'pmap' interface.

1.302 RPC.guide/xdr_rejected_reply()

xdr_rejected_reply()

```
bool_t XDRFUN
xdr_rejected_reply(xdrs, rr)
XDR *xdrs;
struct rejected_reply *rr;
```

Used for describing 'RPC' reply messages. This routine is useful for users who wish to generate 'RPC-style' messages without using the 'RPC' package.

1.303 RPC.guide/xdr_replymsg()

xdr_replymsg()

```
bool_t XDRFUN
xdr_replymsg(xdrs, rmsg)
XDR *xdrs;
struct rpc_msg *rmsg;
```

Used for describing 'RPC' reply messages. This routine is useful for users who wish to generate 'RPC' style messages without using the 'RPC' package.

1.304 RPC.guide/xprt_register()

xprt_register()

```
void
xprt_register(xprt)
SVCXPRT *xprt;
```

After 'RPC' service transport handles are created, they should register themselves with the 'RPC' service package. This routine modifies the global variable svc_fds. Service implementors usually do not need this routine.

1.305 RPC.guide/xprt_unregister()

```
xprt_unregister()  
-----
```

```
void  
xprt_unregister(xprt)  
SVCXPRT *xprt;
```

Before an 'RPC' service transport handle is destroyed, it should unregister itself with the 'RPC' service package. This routine modifies the global variable `svc_fds`. Service implementors usually do not need this routine.

1.306 RPC.guide/XDR Function Reference

```
XDR Function Reference  
=====
```

These routines allow C programmers to describe arbitrary data structures in a machine-independent fashion. Data for remote procedure calls are transmitted using these routines.

```
xdr_array()  
xdr_bool()  
xdr_bytes()  
xdr_char()  
xdr_destroy()  
xdr_double()  
xdr_enum()  
xdr_float()  
xdr_free()  
xdr_getpos()  
xdr_inline()  
xdr_int()  
xdr_long()  
xdrmem_create()  
xdr_opaque()  
xdr_pointer()  
xdrrec_create()  
xdrrec_endofrecord()  
xdrrec_eof()  
xdrrec_skiprecord()  
xdr_reference()  
xdr_setpos()  
xdr_short()  
xdrstdio_create()  
xdr_string()  
xdr_u_char()  
xdr_u_int()  
xdr_u_long()  
xdr_u_short()
```

```
xdr_union()  
xdr_vector()  
xdr_void()  
xdr_wrapstring()
```

1.307 RPC.guide/xdr_array()

```
xdr_array()  
-----
```

```
bool_t XDRFUN  
xdr_array(xdrs, arrp, sizep, maxsize, elsize, elproc)  
XDR *xdrs;  
char **arrp;  
u_int *sizep, maxsize, elsize;  
xdrproc_t elproc;
```

A filter primitive that translates between variable-length arrays and their corresponding external representations. The parameter 'arrp' is the address of the pointer to the array, while 'sizep' is the address of the element count of the array; this element count cannot exceed 'maxsize'. The parameter 'elsize' is the 'sizeof' each of the array's elements, and 'elproc' is an 'XDR' filter that translates between the array elements' C form, and their external representation. This routine returns one if it succeeds, zero otherwise.

1.308 RPC.guide/xdr_bool()

```
xdr_bool()  
-----
```

```
bool_t XDRFUN  
xdr_bool(xdrs, bp)  
XDR *xdrs;  
bool_t *bp;
```

A filter primitive that translates between booleans (C integers) and their external representations. When encoding data, this filter produces values of either one or zero. This routine returns one if it succeeds, zero otherwise.

1.309 RPC.guide/xdr_bytes()

```
xdr_bytes()  
-----
```

```
bool_t XDRFUN
```

```
xdr_bytes(xdrs, sp, sizep, maxsize)
XDR *xdrs;
char **sp;
u_int *sizep, maxsize;
```

A filter primitive that translates between counted byte strings and their external representations. The parameter 'sp' is the address of the string pointer. The length of the string is located at address 'sizep'; strings cannot be longer than 'maxsize'. This routine returns one if it succeeds, zero otherwise.

1.310 RPC.guide/xdr_char()

```
xdr_char()
-----
```

```
bool_t XDRFUN
xdr_char(xdrs, cp)
XDR *xdrs;
char *cp;
```

A filter primitive that translates between C characters and their external representations. This routine returns one if it succeeds, zero otherwise. Note: encoded characters are not packed, and occupy 4 bytes each. For arrays of characters, it is worthwhile to consider `xdr_bytes()`, `xdr_opaque()` or `xdr_string()`.

1.311 RPC.guide/xdr_destroy()

```
xdr_destroy()
-----
```

```
void
xdr_destroy(xdrs)
XDR *xdrs;
```

A macro that invokes the destroy routine associated with the 'XDR' stream, 'xdrs'. Destruction usually involves freeing private data structures associated with the stream. Using 'xdrs' after invoking 'xdr_destroy()' is undefined.

1.312 RPC.guide/xdr_double()

```
xdr_double()
-----
```

```
bool_t XDRFUN
```

```
xdr_double(xdrs, dp)
XDR *xdrs;
double *dp;
```

A filter primitive that translates between C 'double' precision numbers and their external representations. This routine returns one if it succeeds, zero otherwise.

1.313 RPC.guide/xdr_enum()

```
xdr_enum()
-----
```

```
bool_t XDRFUN
xdr_enum(xdrs, ep)
XDR *xdrs;
enum_t *ep;
```

A filter primitive that translates between C 'enum's (actually integers) and their external representations. This routine returns one if it succeeds, zero otherwise.

1.314 RPC.guide/xdr_float()

```
xdr_float()
-----
```

```
bool_t XDRFUN
xdr_float(xdrs, fp)
XDR *xdrs;
float *fp;
```

A filter primitive that translates between C 'float's and their external representations. This routine returns one if it succeeds, zero otherwise.

1.315 RPC.guide/xdr_free()

```
xdr_free()
-----
```

```
void
xdr_free(proc, objp)
xdrproc_t proc;
char *objp;
```

Generic freeing routine. The first argument is the 'XDR' routine for

the object being freed. The second argument is a pointer to the object itself. Note: the pointer passed to this routine is 'not' freed, but what it points to 'is' freed (recursively).

1.316 RPC.guide/xdr_getpos()

xdr_getpos()

```
u_int
xdr_getpos(xdrs)
XDR *xdrs;
```

A macro that invokes the get-position routine associated with the 'XDR' stream, 'xdrs'. The routine returns an unsigned integer, which indicates the position of the 'XDR' byte stream. A desirable feature of 'XDR' streams is that simple arithmetic works with this number, although the 'XDR' stream instances need not guarantee this.

1.317 RPC.guide/xdr_inline()

xdr_inline()

```
long *
xdr_inline(xdrs, len)
XDR *xdrs;
int len;
```

A macro that invokes the in-line routine associated with the 'XDR' stream, 'xdrs'. The routine returns a pointer to a contiguous piece of the stream's buffer; 'len' is the byte length of the desired buffer. Note: pointer is cast to 'long *'. *Warning:* 'xdr_inline()' may return 'NULL' (0) if it cannot allocate a contiguous piece of a buffer. Therefore the behavior may vary among stream instances; it exists for the sake of efficiency.

1.318 RPC.guide/xdr_int()

xdr_int()

```
bool_t XDRFUN
xdr_int(xdrs, ip)
XDR *xdrs;
int *ip;
```

A filter primitive that translates between C integers and their external representations. This routine returns one if it succeeds, zero otherwise.

1.319 RPC.guide/xdr_long()

xdr_long()

```
bool_t XDRFUN
xdr_long(xdrs, lp)
XDR *xdrs;
long *lp;
```

A filter primitive that translates between C 'long' integers and their external representations. This routine returns one if it succeeds, zero otherwise.

1.320 RPC.guide/xdrmem_create()

xdrmem_create()

```
void
xdrmem_create(xdrs, addr, size, op)
XDR *xdrs;
char *addr;
u_int size;
enum xdr_op op;
```

This routine initializes the 'XDR' stream object pointed to by 'xdrs'. The stream's data is written to, or read from, a chunk of memory at location 'addr' whose length is no more than 'size' bytes long. The 'op' determines the direction of the 'XDR' stream (either 'XDR_ENCODE', 'XDR_DECODE', or 'XDR_FREE').

1.321 RPC.guide/xdr_opaque()

xdr_opaque()

```
bool_t XDRFUN
xdr_opaque(xdrs, cp, cnt)
XDR *xdrs;
char *cp;
u_int cnt;
```

A filter primitive that translates between fixed size opaque data and its external representation. The parameter 'cp' is the address of the opaque object, and 'cnt' is its size in bytes. This routine returns one if it succeeds, zero otherwise.

1.322 RPC.guide/xdr_pointer()

xdr_pointer()

```
bool_t XDRFUN
xdr_pointer(xdrs, objpp, objsize, xdrobj)
XDR *xdrs;
char **objpp;
u_int objsize;
xdrproc_t xdrobj;
```

Like xdr_reference() except that it serializes 'NULL' pointers, whereas xdr_reference() does not. Thus, 'xdr_pointer()' can represent recursive data structures, such as binary trees or linked lists.

1.323 RPC.guide/xdrrec_create()

xdrrec_create()

```
void
xdrrec_create(xdrs, sendsize, recvsize, handle, readit, writeit)
XDR *xdrs;
u_int sendsize, recvsize;
char *handle;
int (*readit) (), (*writeit) ();
```

This routine initializes the 'XDR' stream object pointed to by 'xdrs'. The stream's data is written to a buffer of size 'sendsize'; a value of zero indicates the system should use a suitable default. The stream's data is read from a buffer of size 'recvsize'; it too can be set to a suitable default by passing a zero value. When a stream's output buffer is full, 'writeit' is called. Similarly, when a stream's input buffer is empty, 'readit' is called. The behavior of these two routines is similar to the system calls 'read' and 'write', except that 'handle' is passed to the former routines as the first parameter.

Note: the 'XDR' stream's 'op' field must be set by the caller.

Warning: this 'XDR' stream implements an intermediate record stream. Therefore there are additional bytes in the stream to provide record boundary information.

1.324 RPC.guide/xdrrec_endofrecord()

```
xdrrec_endofrecord()  
-----
```

```
bool_t  
xdrrec_endofrecord(xdrs, sendnow)  
XDR *xdrs;  
bool_t sendnow;
```

This routine can be invoked only on streams created by `xdrrec_create()`. The data in the output buffer is marked as a completed record, and the output buffer is optionally written out if 'sendnow' is non-zero. This routine returns one if it succeeds, zero otherwise.

1.325 RPC.guide/xdrrec_eof()

```
xdrrec_eof()  
-----
```

```
bool_t  
xdrrec_eof(xdrs)  
XDR *xdrs;
```

This routine can be invoked only on streams created by `xdrrec_create()`. After consuming the rest of the current record in the stream, this routine returns one if the stream has no more input, zero otherwise.

1.326 RPC.guide/xdrrec_skiprecord()

```
xdrrec_skiprecord()  
-----
```

```
bool_t  
xdrrec_skiprecord(xdrs)  
XDR *xdrs;
```

This routine can be invoked only on streams created by `xdrrec_create()`. It tells the 'XDR' implementation that the rest of the current record in the stream's input buffer should be discarded. This routine returns one if it succeeds, zero otherwise.

1.327 RPC.guide/xdr_reference()

```
xdr_reference()  
-----
```

```
bool_t XDRFUN
xdr_reference(xdrs, pp, size, proc)
XDR *xdrs;
char **pp;
u_int size;
xdrproc_t proc;
```

A primitive that provides pointer chasing within structures. The parameter 'pp' is the address of the pointer; 'size' is the 'sizeof' the structure that '*pp' points to; and 'proc' is an 'XDR' procedure that filters the structure between its C form and its external representation. This routine returns one if it succeeds, zero otherwise. Warning: this routine does not understand 'NULL' pointers. Use xdr_pointer() instead.

1.328 RPC.guide/xdr_setpos()

```
xdr_setpos()
```

```
-----
```

```
bool_t XDRFUN
xdr_setpos(xdrs, pos)
XDR *xdrs;
u_int pos;
```

A macro that invokes the set position routine associated with the 'XDR' stream 'xdrs'. The parameter 'pos' is a position value obtained from xdr_getpos(). This routine returns one if the 'XDR' stream could be repositioned, and zero otherwise. Warning: it is difficult to reposition some types of 'XDR' streams, so this routine may fail with one type of stream and succeed with another.

1.329 RPC.guide/xdr_short()

```
xdr_short()
```

```
-----
```

```
bool_t XDRFUN
xdr_short(xdrs, sp)
XDR *xdrs;
short *sp;
```

A filter primitive that translates between C 'short' integers and their external representations. This routine returns one if it succeeds, zero otherwise.

1.330 RPC.guide/xdrstdio_create()

xdrstdio_create()

```
void
xdrstdio_create(xdrs, file, op)
XDR *xdrs;
FILE *file;
enum xdr_op op;
```

This routine initializes the 'XDR' stream object pointed to by 'xdrs'. The 'XDR' stream data is written to, or read from, the Standard 'I/O' stream 'file'. The parameter 'op' determines the direction of the 'XDR' stream (either 'XDR_ENCODE', 'XDR_DECODE', or 'XDR_FREE').
Warning: the destroy routine associated with such 'XDR' streams calls 'fflush()' on the 'file' stream, but never 'fclose()'.

1.331 RPC.guide/xdr_string()

xdr_string()

```
bool_t XDRFUN
xdr_string(xdrs, sp, maxsize)
XDR *xdrs;
char **sp;
u_int maxsize;
```

A filter primitive that translates between C strings and their corresponding external representations. Strings cannot be longer than 'maxsize'. Note: 'sp' is the address of the string's pointer. This routine returns one if it succeeds, zero otherwise.

1.332 RPC.guide/xdr_u_char()

xdr_u_char()

```
bool_t XDRFUN
xdr_u_char(xdrs, ucp)
XDR *xdrs;
unsigned char *ucp;
```

A filter primitive that translates between 'unsigned' C characters and their external representations. This routine returns one if it succeeds, zero otherwise.

1.333 RPC.guide/xdr_u_int()

xdr_u_int()

```
bool_t XDRFUN
xdr_u_int(xdrs, up)
XDR *xdrs;
unsigned *up;
```

A filter primitive that translates between C 'unsigned' integers and their external representations. This routine returns one if it succeeds, zero otherwise.

1.334 RPC.guide/xdr_u_long()

xdr_u_long()

```
bool_t XDRFUN
xdr_u_long(xdrs, ulp)
XDR *xdrs;
unsigned long *ulp;
```

A filter primitive that translates between C 'unsigned long' integers and their external representations. This routine returns one if it succeeds, zero otherwise.

1.335 RPC.guide/xdr_u_short()

xdr_u_short()

```
bool_t XDRFUN
xdr_u_short(xdrs, usp)
XDR *xdrs;
unsigned short *usp;
```

A filter primitive that translates between C 'unsigned short' integers and their external representations. This routine returns one if it succeeds, zero otherwise.

1.336 RPC.guide/xdr_union()

xdr_union()

```

bool_t XDRFUN
xdr_union(xdrs, dscmp, unp, choices, default)
XDR *xdrs;
int *dscmp;
char *unp;
struct xdr_discrim *choices;
bool_t (*defaultarm) (); /* may equal NULL */

```

A filter primitive that translates between a discriminated C 'union' and its corresponding external representation. It first translates the discriminant of the union located at 'dscmp'. This discriminant is always an 'enum_t'. Next the union located at 'unp' is translated. The parameter 'choices' is a pointer to an array of 'xdr_discrim' structures. Each structure contains an ordered pair of ['value', 'proc']. If the union's discriminant is equal to the associated 'value', then the 'proc' is called to translate the union. The end of the 'xdr_discrim' structure array is denoted by a routine of value 'NULL'. If the discriminant is not found in the 'choices' array, then the 'defaultarm' procedure is called (if it is not 'NULL'). Returns one if it succeeds, zero otherwise.

1.337 RPC.guide/xdr_vector()

```

xdr_vector()
-----

```

```

bool_t XDRFUN
xdr_vector(xdrs, arrp, size, elsize, elproc)
XDR *xdrs;
char *arrp;
u_int size, elsize;
xdrproc_t elproc;

```

A filter primitive that translates between fixed-length arrays and their corresponding external representations. The parameter 'arrp' is the address of the pointer to the array, while 'size' is the element count of the array. The parameter 'elsize' is the 'sizeof' each of the array's elements, and 'elproc' is an 'XDR' filter that translates between the array elements' C form, and their external representation. This routine returns one if it succeeds, zero otherwise.

1.338 RPC.guide/xdr_void()

```

xdr_void()
-----

```

```

bool_t XDRFUN
xdr_void(xdrs, dummy)
XDR *xdrs;

```

```
void *dummy;
```

This routine always returns one. It may be passed to 'RPC' routines that require a function parameter, where nothing is to be done. 'NULL' can be passed as argument pointer.

1.339 RPC.guide/xdr_wrapstring()

```
xdr_wrapstring()
```

```
-----
```

```
bool_t XDRFUN
xdr_wrapstring(xdrs, sp)
XDR *xdrs;
char **sp;
```

A primitive that calls 'xdr_string(xdrs, sp, UINT_MAX);' where 'UINT_MAX' is the maximum value of an unsigned integer. 'xdr_wrapstring()' is handy because the 'RPC' package passes a maximum of two parameters to 'XDR' routines, and xdr_string(), one of the most frequently used primitives, requires three. Returns one if it succeeds, zero otherwise.

1.340 RPC.guide/Function Index

Function Index

```
*****
```

authdes_create()	DES Authentication
authnone_create()	authnone_create()
authnone_create()	UNIX Authentication
authunix_create()	authunix_create()
authunix_create()	DES Authentication
authunix_create_default()	authunix_create_default()
authunix_create_default()	UNIX Authentication
auth_destroy()	auth_destroy()
bindresvport()	bindresvport()
callrpc()	Intermediate Layer
callrpc()	callrpc()
clntraw_create()	clntraw_create()
clnttcp_create()	The Calling Side
clnttcp_create()	clnttcp_create()
clntudp_bufcreate()	clntudp_bufcreate()
clntudp_create()	The Calling Side
clntudp_create()	clntudp_create()
clnt_broadcast()	Broadcast RPC Synopsis
clnt_broadcast()	clnt_broadcast()
clnt_call()	The Calling Side

clnt_call()	clnt_call()
clnt_control()	clnt_control()
clnt_create()	clnt_create()
clnt_destroy()	clnt_destroy()
clnt_destroy()	The Calling Side
clnt_freeres()	clnt_freeres()
clnt_geterr()	clnt_geterr()
clnt_pcreateerror()	clnt_pcreateerror()
clnt_perrno()	clnt_perrno()
clnt_perror()	clnt_perror()
clnt_spcreateerror()	clnt_spcreateerror()
clnt_sperrno()	clnt_sperrno()
clnt_sperror()	clnt_sperror()
getrpcbyname()	getrpcbyname()
getrpcbynumber()	getrpcbynumber()
getrpcport()	getrpcport()
gettransient()	Callback Procedures
get_myaddress()	get_myaddress()
host2netname()	DES Authentication
netname2user()	DES Authentication
ntohs()	Callback Procedures
pmap_getmaps()	pmap_getmaps()
pmap_getport()	pmap_getport()
pmap_rmtcall()	pmap_rmtcall()
pmap_set()	Callback Procedures
pmap_set()	pmap_set()
pmap_unset()	More on the Server Side
pmap_unset()	pmap_unset()
registerrpc()	Intermediate Layer
registerrpc()	registerrpc()
rpc_createerr	rpc_createerr
select()	Select on the Server Side
svcerr_auth()	svcerr_auth()
svcerr_decode()	svcerr_decode()
svcerr_noproc()	svcerr_noproc()
svcerr_noproc()	More on the Server Side
svcerr_noprog()	svcerr_noprog()
svcerr_progvers()	svcerr_progvers()
svcerr_systemerr()	svcerr_systemerr()
svcerr_systemerr()	UNIX Authentication
svcerr_weakauth()	UNIX Authentication
svcerr_weakauth()	svcerr_weakauth()
svcfld_create()	svcfld_create()
svccraw_create()	svccraw_create()
svctcp_create()	More on the Server Side
svctcp_create()	svctcp_create()
svcudp_bufcreate()	svcudp_bufcreate()
svcudp_create()	More on the Server Side
svcudp_create()	svcudp_create()
svc_destroy()	svc_destroy()
svc_fds	Select on the Server Side
svc_fds	svc_fds
svc_fdset	svc_fdset
svc_freeargs()	svc_freeargs()
svc_freeargs()	Memory Allocation with XDR
svc_getargs()	Memory Allocation with XDR
svc_getargs()	svc_getargs()

svc_getcaller()	svc_getcaller()
svc_getreq()	svc_getreq()
svc_getreqset()	Select on the Server Side
svc_getreqset()	svc_getreqset()
svc_register()	More on the Server Side
svc_register()	svc_register()
svc_run()	Intermediate Layer
svc_run()	svc_run()
svc_run()	Select on the Server Side
svc_sendreply()	svc_sendreply()
svc_sendreply()	More on the Server Side
svc_unregister()	svc_unregister()
user2netname()	DES Authentication
xdrmem_create()	Memory Streams
xdrmem_create()	xdrmem_create()
xdrrec_create()	xdrrec_create()
xdrrec_endofrecord()	Record (TCP-IP) Streams
xdrrec_endofrecord()	xdrrec_endofrecord()
xdrrec_eof()	Record (TCP-IP) Streams
xdrrec_eof()	xdrrec_eof()
xdrrec_skiprecord()	Record (TCP-IP) Streams
xdrrec_skiprecord()	xdrrec_skiprecord()
xdrstdio_create()	xdrstdio_create()
xdrstdio_create()	Standard I-O Streams
xdr_accepted_reply()	xdr_accepted_reply()
xdr_array()	Arrays
xdr_array()	xdr_array()
xdr_authunix_parms()	xdr_authunix_parms()
xdr_bool()	xdr_bool()
xdr_bytes()	xdr_bytes()
xdr_bytes()	Byte Arrays
xdr_callhdr()	xdr_callhdr()
xdr_callmsg()	xdr_callmsg()
xdr_char()	xdr_char()
xdr_destroy()	Non-filter Primitives
xdr_destroy()	xdr_destroy()
xdr_double()	xdr_double()
xdr_element()	Arrays
xdr_enum()	xdr_enum()
xdr_float()	xdr_float()
xdr_free()	xdr_free()
xdr_getpos()	xdr_getpos()
xdr_getpos()	Non-filter Primitives
xdr_inline()	xdr_inline()
xdr_int()	xdr_int()
xdr_long()	xdr_long()
xdr_opaque()	Opaque Data
xdr_opaque()	xdr_opaque()
xdr_opaque_auth()	xdr_opaque_auth()
xdr_pmap()	xdr_pmap()
xdr_pmaplist()	xdr_pmaplist()
xdr_pointer()	xdr_pointer()
xdr_reference()	xdr_reference()
xdr_reference()	Pointers
xdr_rejected_reply()	xdr_rejected_reply()
xdr_replymsg()	xdr_replymsg()
xdr_setpos()	xdr_setpos()

xdr_setpos()	Non-filter Primitives
xdr_short()	xdr_short()
xdr_string()	Strings
xdr_string()	xdr_string()
xdr_union()	xdr_union()
xdr_union()	Discriminated Unions
xdr_u_char()	xdr_u_char()
xdr_u_int()	xdr_u_int()
xdr_u_long()	xdr_u_long()
xdr_u_short()	xdr_u_short()
xdr_vector()	Fixed Sized Arrays
xdr_vector()	xdr_vector()
xdr_void()	xdr_void()
xdr_wrapstring()	xdr_wrapstring()
xprt_register()	xprt_register()
xprt_unregister()	xprt_unregister()

1.341 RPC.guide/Concept Index

Concept Index

About the floating point format	About the Floating Point Format
About the Port to AmiTCP/IP and Amiga	About the Port to AmiTCP-IP and Amiga
Administration of RPC	Assigning Program Numbers
AmiTCP/IP RPC Notes	AmiTCP-IP Notes
Arbitrary data types	Passing Arbitrary Data Types
Assigning program numbers	Assigning Program Numbers
Authentication	RPC Authentication
Authentication protocols	Authentication Protocols
Authentication, DES	DES Authentication (protocol)
Authentication, general	Authentication
Authentication, Null	Null Authentication (protocol)
Authentication, requirements	Authentication Requirements
Authentication, UNIX	UNIX Authentication (protocol)
Batching	RPC Batching
Batching (RFC)	Batching
Binding	Binding and Rendezvous Independence
bool	XDR Booleans
Booleans	XDR Booleans
bool_t	XDR Booleans
Broadcast RPC	Broadcast RPC
Broadcast RPC, (RFC)	RPC Broadcast
Broadcast RPC, server response	Handling Broadcast on the Server Side
Broadcast RPC, synopsis	Broadcast RPC Synopsis
C-preprocessor	The C-Preprocessor
Changes from SUNRPC 4.0	What Was Done to the Original SUNRPC 4.0 ↔
While Porting	
Choosing the Library to Use	Choosing the Library to Use
Client handle, used by rpcgen	Converting Local Procedures into Remote ↔
Procedures	
Compiling the library	Compiling the library
Constants	Constants
Debugging with rpcgen	Generating XDR Routines

Declarations	Declarations
Definitions	Definitions
DES authentication protocol language)	DES Authentication Protocol (in XDR ←
DES authentication verifiers	DES Authentication Verifiers
DES authentication, protocol	DES Authentication (protocol)
DES clock synchronization	Nicknames and Clock Synchronization
DES naming	Naming
DES nicknames	Nicknames and Clock Synchronization
Diffie-Hellman encryption	Diffie-Hellman Encryption
direction of XDR operations	XDR Operation Directions
Disclaimer & Copyright	Disclaimer & Copyright
Encryption, Diffie-Hellman	Diffie-Hellman Encryption
enum clnt_stat (in RPC programming)	Intermediate Layer
Enumerations	Enumerations
External Data Representation: protocol specification	XDR Protocol Specification
file format, rpc	rpc Data Base File Format
File system model	File System Model
Fixed-length array declarations	Fixed-Length Array Declarations
Function reference, RPC	RPC Function Reference
Function reference, XDR	XDR Function Reference
Generating XDR routines	Generating XDR Routines
Highest layer of RPC	Highest Layer
IEEE vs. FFP floating point format	About the Floating Point Format
inetd	Using Inetd
Intermediate layer of RPC	Intermediate Layer
Introction to RPC protocol	Introduction to RPC
Introduction to RPC programming guide	Introduction to RPC Programming Guide
Introduction to XDR protocol	Introduction to XDR
Introduction, mount protocol	Mount Protocol Introduction
Introduction, NFS	Introduction to NFS
layers of RPC	Layers of RPC
library primitives for XDR	XDR Library Primitives
Local procedures	Converting Local Procedures into Remote ←
Procedures	
Lowest layer of RPC	Lowest Layer of RPC
Manual Pages	Manual Pages
Memory allocation with XDR	Memory Allocation with XDR
Message protocol	The RPC Message Protocol
Miscellaneous RPC features	Other RPC Features
Miscellaneous utility functions	Miscellaneous Utility Functions
Mount data types	Mount Protocol Basic Data Types
Mount data types, dirpath	dirpath
Mount data types, fhandle	fhandle (Mount)
Mount data types, fhstatus	fhstatus
Mount data types, name	name
Mount protocol	Mount Protocol Definition
Mount protocol server procedures	Mount Protocol Server Procedures
Mount protocol, basic data types	Mount Protocol Basic Data Types
Mount protocol, introduction	Mount Protocol Introduction
Mount protocol, RPC information	Mount Protocol RPC Information
Mount protocol, XDR structure sizes	Mount Protocol Sizes of XDR Structures
Mount server procedures, MNTPROC_DUMP()	Return Mount Entries
Mount server procedures, MNTPROC_EXPORT()	Return Export List
Mount server procedures, MNTPROC_MNT()	Add Mount Entry
Mount server procedures, MNTPROC_NULL()	Do Nothing (Mount)
Mount server procedures, MNTPROC_UMNT()	Remove Mount Entry

Mount server procedures, MNTPROC_UMNTALL() Remove All Mount Entries
Network File System: version 2 protocol specification Network File System ↔
Version 2 Protocol Specification

NFS data types	Basic Data Types
NFS data types, attrstat	attrstat
NFS data types, diropargs	diropargs
NFS data types, diropres	diropres
NFS data types, fattr	fattr
NFS data types, fhandle	fhandle
NFS data types, filename	filename
NFS data types, ftype	ftype
NFS data types, path	path
NFS data types, sattr	sattr
NFS data types, stat	stat
NFS data types, timeval	timeval
NFS implementation	NFS Implementation Issues
NFS introduction	Introduction to NFS
NFS Mount protocol	Mount Protocol Definition
NFS protocol	NFS Protocol Definition
NFS server procedures	Server Procedures
NFS server procedures, NFSPROC_CREATE()	Create File
NFS server procedures, NFSPROC_GETATTR()	Get File Attributes
NFS server procedures, NFSPROC_LINK()	Create Link to File
NFS server procedures, NFSPROC_LOOKUP()	Look Up File Name
NFS server procedures, NFSPROC_MKDIR()	Create Directory
NFS server procedures, NFSPROC_NULL()	Do Nothing
NFS server procedures, NFSPROC_READ()	Read From File
NFS server procedures, NFSPROC_READDIR()	Read From Directory
NFS server procedures, NFSPROC_READLINK()	Read From Symbolic Link
NFS server procedures, NFSPROC_REMOVE()	Remove File
NFS server procedures, NFSPROC_RENAME()	Rename File
NFS server procedures, NFSPROC_RMDIR()	Remove Directory
NFS server procedures, NFSPROC_ROOT()	Get Filesystem Root
NFS server procedures, NFSPROC_SETATTR()	Set File Attributes
NFS server procedures, NFSPROC_STATFS()	Get Filesystem Attributes
NFS server procedures, NFSPROC_SYMLINK()	Create Symbolic Link
NFS server procedures, NFSPROC_WRITE()	Write to File
NFS server procedures, NFSPROC_WRITECACHE()	Write to Cache
NFS, basic data types	Basic Data Types
NFS, external data representation	External Data Representation
NFS, pathname interpretation	Pathname Interpretation
NFS, permission issues	Permission Issues
NFS, protocol definition	NFS Protocol Definition
NFS, remote procedure call	Remote Procedure Call
NFS, RPC information	RPC Information
NFS, server/client relationship	Server-Client Relationship
NFS, setting RPC parameters	Setting RPC Parameters
NFS, version 2 protocol specification	Network File System Version 2 Protocol ↔ Specification
NFS, XDR structure sizes	Sizes of XDR Structures
Nonsupported features	Nonsupported Features
Null authentication, protocol	Null Authentication (protocol)
opaque	XDR Opaque Data
Opaque data	XDR Opaque Data
Pointer declarations	Pointer Declarations
Pointer semantics and XDR	Pointers
Port Mapper operation	Port Mapper Operation

Port Mapper program protocol	Port Mapper Program Protocol
Port Mapper protocol specification Language)	Port Mapper Protocol Specification (in RPC ↔
portmap (the program)	portmap and rpcinfo
portmap manual page	portmap Manual Page
Preprocessor	The C-Preprocessor
Program number assignment	Assigning Program Numbers
Program number assignment (RFC)	Program Number Assignment
Programs	Programs
Protocol Requirements (RPC)	RPC Protocol Requirements
Protocol specification, NFS version 2 Specification	Network File System Version 2 Protocol ↔
Protocol specification, RPC Specification	Remote Procedure Calls Protocol ↔
Protocol specification, XDR	XDR Protocol Specification
Record Marking standard	Record Marking Standard
References (RPC)	RPC References
References (XDR)	XDR References
Remote Procedure Call programming guide	Remote Procedure Call Programming Guide
Remote Procedure Calls: protocol specification Specification	Remote Procedure Calls Protocol ↔
Remote procedures Procedures	Converting Local Procedures into Remote ↔
Rendezvous independence	Binding and Rendezvous Independence
RFC, XDR	XDR Protocol Specification
RPC administration	Assigning Program Numbers
RPC authentication, library functions	DES Authentication
RPC binding	Binding and Rendezvous Independence
rpc data base file format	rpc Data Base File Format
RPC DES	DES Authentication
RPC function reference	RPC Function Reference
RPC guarantees	UNIX Authentication
RPC language	RPC Language
RPC language specification	The RPC Language Specification
RPC language, example service Language	An Example Service Described in the RPC ↔
RPC language, ext. to the XDR Language	The RPC Language
RPC language, syntax notes	RPC Syntax Notes
RPC message protocol	The RPC Message Protocol
RPC model	The RPC Model
RPC paradigm	The RPC Paradigm
RPC procedures	Programs and Procedures
RPC programming guide, introduction	Introduction to RPC Programming Guide
RPC programs	Programs and Procedures
RPC Protocol Requirements	RPC Protocol Requirements
RPC protocol, introduction	Introduction to RPC
RPC protocol, other uses	Other Uses of the RPC Protocol
RPC select ()	Select on the Server Side
RPC services	Highest Layer
RPC terminology	Terminology
RPC transport semantics	Transports and Semantics
RPC, authentication	RPC Authentication
RPC, batching	RPC Batching
RPC, broadcast	Broadcast RPC
RPC, broadcast synopsis	Broadcast RPC Synopsis
RPC, built-in routines	Passing Arbitrary Data Types
RPC, callback procedures	Callback Procedures

RPC, calling side	The Calling Side
RPC, highest layer	Highest Layer
RPC, intermediate layer	Intermediate Layer
RPC, layers	Layers of RPC
RPC, lowest layer	Lowest Layer of RPC
RPC, miscellaneous features	Other RPC Features
RPC, protocol specification	Remote Procedure Calls Protocol ←
Specification	
RPC, server side	More on the Server Side
RPC, the highest layer, introduction	The Highest Layer
RPC, the lowest layer, introduction	The Lowest Layer
RPC, the middle layer, introduction	The Middle Layer
RPC, versions	Versions
rpcgen	Introduction to RPC Programming Guide
rpcgen C-preprocessor	The C-Preprocessor
rpcgen constants	Constants
rpcgen declarations	Declarations
rpcgen definitions	Definitions
rpcgen enumerations	Enumerations
rpcgen manual page	rpcgen Manual Page
rpcgen Notes	rpcgen Notes
rpcgen other operations	rpcgen Programming Notes
rpcgen programming guide	rpcgen Programming Guide
rpcgen programs	Programs
rpcgen special cases	Special Cases
rpcgen structures	Structures
rpcgen typedef	Typedef
rpcgen unions	Unions
rpcgen, broadcast RPC	Handling Broadcast on the Server Side
rpcgen, introduction	The rpcgen Protocol Compiler
rpcgen, RPC language	RPC Language
rpcinfo (the program)	portmap and rpcinfo
rpcinfo manual page	rpcinfo Manual Page
RPCL	RPC Language
SAS/C 6.x	Supported Compilers
select() on the server side	Select on the Server Side
Servers, stateless	Stateless Servers
Simple declarations	Simple Declarations
Stateless servers	Stateless Servers
stream implementation in XDR	XDR Stream Implementation
string	XDR Strings
Strings	XDR Strings
Structures	Structures
Supported Compilers	Supported Compilers
TCP	TCP
Terminology	Terminology
The rpcgen protocol compiler	The rpcgen Protocol Compiler
Timeout changes	Timeout Changes
Transport semantics	Transports and Semantics
Typedef	Typedef
UDP 8K warning	Intermediate Layer
Unions	Unions
UNIX authentication, library functions	UNIX Authentication
UNIX authentication, protocol	UNIX Authentication (protocol)
Using inetd	Using Inetd
Using the Library	Using the Library
Variable-length array declarations	Variable-Length Array Declarations

Versions	Versions
void	XDR Voids
Voids	XDR Voids
XDR array, fixed length	Fixed-length Array
XDR array, variable length	Variable-length Array
XDR basic block size	Basic Block Size
XDR block size	Basic Block Size
XDR boolean	Boolean
XDR byte order	Why Only one Byte-Order for an XDR Unit?
XDR constant	Constant
XDR data types	XDR Data Types
XDR data, optional	Optional-data
XDR discriminated union	Discriminated Union
XDR double-precision floating-point integer	Double-precision Floating-point
XDR enumeration	Enumeration
XDR fixed-length array	Fixed-length Array
XDR fixed-length opaque data	Fixed-length Opaque Data
XDR floating-point integer	Floating-point
XDR function reference	XDR Function Reference
XDR futures	Areas for Future Enhancement
XDR hyper integer	Hyper Integer and Unsigned Hyper Integer
XDR integer	Integer
XDR integer, double-precision floating point	Double-precision Floating-point
XDR integer, floating point	Floating-point
XDR integer, hyper	Hyper Integer and Unsigned Hyper Integer
XDR integer, unsigned	Unsigned Integer
XDR justification	Justification
XDR language	Why a Language for Describing Data?
XDR language notation	Notational Conventions
XDR language specification	The XDR Language Specification
XDR language syntax	Syntax Information
XDR language syntax notes	XDR Syntax Notes
XDR library, arrays	Arrays
XDR library, byte arrays	Byte Arrays
XDR library, constructed data type filters	Constructed Data Type Filters
XDR library, constructed data type filters: examples	XDR Examples
XDR library, discriminated unions	Discriminated Unions
XDR library, enumeration filters	Enumeration Filters
XDR library, fixed sized arrays	Fixed Sized Arrays
XDR library, floating point filters	Floating Point Filters
XDR library, no data	No Data
XDR library, number filters	Number Filters
XDR library, opaque data	Opaque Data
XDR library, pointers	Pointers
XDR library, strings	Strings
XDR object	The XDR Object
XDR opaque data, fixed length	Fixed-length Opaque Data
XDR opaque data, variable length	Variable-length Opaque Data
XDR optional data	Optional-data
XDR protocol specification	XDR Protocol Specification
XDR protocol, introduction	Introduction to XDR
XDR RFC	XDR Protocol Specification
XDR routines, generating	Generating XDR Routines
XDR string	String
XDR Strings	XDR Strings
XDR structure	Structure
XDR Technical Notes	XDR Technical Notes

XDR typedef	Typedef (type definition)
XDR union discriminated	Discriminated Union
XDR unsigned integer	Unsigned Integer
XDR variable-length array	Variable-length Array
XDR variable-length data with Zeros?	Why must Variable-Length Data be Padded ↔
XDR variable-length opaque data	Variable-length Opaque Data
XDR void	Void
XDR, advanced topics	Advanced Topics
XDR, canonical standard	A Canonical Standard
XDR, library	The XDR Library
XDR, library primitives	XDR Library Primitives
XDR, linked lists	Linked Lists
XDR, memory allocation	Memory Allocation with XDR
XDR, memory streams	Memory Streams
XDR, non-filter primitives	Non-filter Primitives
XDR, operation directions	XDR Operation Directions
XDR, overview	Overview
XDR, portable data	Justification
XDR, record (TCP/IP) streams	Record (TCP-IP) Streams
XDR, standard I/O streams	Standard I-O Streams
XDR, stream access	XDR Stream Access
XDR, stream implementation	XDR Stream Implementation
XDR, Sun Technical Notes	XDR Technical Notes
XDR, system routines	Overview
xdrstdio_create(), example	The XDR Library
xdr_double(), note	About the Floating Point Format
xdr_float(), note	About the Floating Point Format
xdr_reference() with NULL pointers	Pointers
