


PAGE 19

## nVidia GeForce FX

Will nVidia's NV30  
outshine ATI's  
Radeon technically?  
PC Magazine finds out



THE LEADING INDEPENDENT LABS-BASED SOURCE

PAGE 16

## The future Office?


Our sneak preview of Office 11 reveals  
the enhancements Microsoft's next  
suite will offer business users

PAGE 12

REVIEW

## Dazzling colour

The Xerox 7300  
workgroup colour  
printer delivers  
crisp colour at  
high speed



PAGE 11

REVIEW

## A lower- cost PDA

We assess Dell's  
competitive new  
offering. But does the  
Axim X5 break any new  
technical ground?



PAGE 14

REVIEW

## Blackout Buster UPS

PAGE 5

NEWS

Mobile chips, wireless  
laptops, broadband  
via the mains, more  
XP & IE glitches...

PAGE 13

REVIEW

## Linksys EFG80



PAGE 23

USER TO USER

## Make your own free PDFs

How to create fully  
featured electronic  
documents for free

PAGE 28

CMS MASTERCLASS

## Going live

Our series concludes  
as we go online  
with our Web  
content system

PAGE 25

OFFICE TUTOR

## Preparing Excel data for Access


How to make the  
move From Excel to  
an Access database

PAGE 2, 3, 4 and 10

OPINION

Tony Westbrook  
Kelvyn Taylor  
Guy Kewney  
Alex Arias

# FROM THE FRONT



## Keep it coming

More news, views, reviews and previews in this latest issue of the *PC Magazine* e-book. And by the number of people registering with us, it seems that you like what we're doing

tony\_westbrook@vnu.co.uk

**A**s we saw in last month's issue where our VNU Labs experts did some crystal ball gazing, there's plenty to look forward to this year from a technical point of view. And with this in mind we've chosen to focus on a couple of those soon-to-arrive things in this issue.

First we've asked Geoff Eion, our Office Tutor author and all round productivity guru to take a close look at the beta version of Microsoft's Office 11 suite now being sent out to developers. So what is it going to offer when it ships? (This is important as there isn't much choice when it comes to office suites, so we want to make sure Microsoft gets it right). Office 11 is the first incarnation of Office to use XML in its native form to store and tag its data. This should add more weight to the general groundswell of support for XML as a common metadata standard. It's also interesting to see Microsoft turn to a true standard (defined by W3C) in place of the more traditional proprietary standards it has relied on until now. Turn to our feature on page 16 to find out what Geoff thinks we've got in store when Office 11 (or Office 2004 if Geoff has predicted the shipping name correctly) finally arrives.

Even as many areas of traditional PC hardware development seem to have peaked, one area where it's all still happening as fast as it ever is graphics, thanks to the insatiable demands of hardcore PC gamers. Today it's ATI, Matrox and nVidia who now seem to have become the main protagonists in supplying the next generation of 3D graphics chips. And, as yet, it's a pretty equal race with each company producing new chips to outflank the other—and all at record speed.

Today, ATI is on top with its Radeon 9700. But not for long if nVidia gets its way. We've sent our graphics guru Laurence Grayson on a mission to read through the product announcements, previews and product specification sheets to assess nVidia's chances of pushing ATI back in to second place when it starts shipping its soon to be released NV30 chip. See what he's has to say on page 19.

We also take a look at some of the latest batch of technology products. This month we've new reviews of everything from Xerox's latest printer, through Dell's cut price PDA, a network appliance and some drivers to the humble (but essential) UPS. As always, our VNU Labs team is giving them a full, hands-on assessment. Reviews start on page 11.

Our popular Solutions series continue in every issue and Alex Cruickshank has finally reached the end of his build your own CMS series. By the time you've finished this installment you should be running your own Web site and managing its content without having to recode every page in HTML. And keep subscribing to the PC magazine e-book to pick up on the great new Solutions series we've got planned over the coming months.

Many thanks to everyone for the great feedback on the e-book and the way we're distributing it. As a result of these comments we've made some changes to both the layout and distribution of this issue and, unlike a print publication, which is locked in to a fixed production schedule, we'll continue to respond to reader requests on an issue-by-issue basis. So keep your comments coming by emailing us at [feedback@pcmag.co.uk](mailto:feedback@pcmag.co.uk). If you're reading this and have not yet registered for your own copy, make sure you do at [www.pcmag.co.uk/pcm/next.jsp](http://www.pcmag.co.uk/pcm/next.jsp) ■

**One area where it's all still happening as fast as ever is graphics, thanks to the insatiable demands of PC gamers**



# TECH TALES

## Man versus machine

Kelvyn Taylor celebrates man's supremacy over machines, but struggles with Service Pack 1, and wonders if the latest integrated motherboards mark the end of his PC-tinkering days

kelvyn\_taylor@vnu.co.uk

I love fixing my PC. I even love fixing PCs belonging to my relatives or neighbours. And I absolutely adore fixing review PCs that won't run a benchmark test. Yet still I like to think of myself as a reasonably sane, well-adjusted individual.

Now before you write me off as a total lunatic, read that first paragraph again. I didn't say I enjoyed not being able to fix my PC, or being stumped by yet another obscure Windows incompatibility. You see, I'm sure you'll agree that solving a problem is immensely satisfying, yet there's nothing on this planet that compares to the frustration of being beaten by a man-made contraption. After all, if man made it, man should be able to fix it, surely? Unfortunately it's usually the case that the man who made it is never around when he's needed.

I had this experience the other weekend, when I decided to install Service Pack 1 on my daughter's Windows XP notebook PC. After a long download, I let it do its stuff, and then it rebooted and rebooted and rebooted. Somehow it managed to get stuck in an endless reboot cycle, with a momentary glimpse of a BSOD—not visible enough to be able to read the cause of the error. It would boot into Safe Mode, but there was precious little I could do there to rectify things—I hadn't let the Service Pack create an uninstall option as I was a bit short of disk space (as an aside, why on earth does a simple service pack need almost 1GB of disk space?).

After a little bit of thought (about 10 seconds), I decided to start afresh and repartition the hard disk drive (I'd been wanting to do that for years as the 2GB C: partition had precious little space for anything but the operating system). On went a shiny new installation of XP and then I had a moment of madness. I reinstalled Service Pack 1... with exactly the same results. So I went round the loop again, repartition, reformat, reinstall. This time I decided to leave well alone and not worry about any supposed 'fixes' or 'improvements' that Microsoft tried to offer me. The result? A perfectly functional system.

I never was able to get to the bottom of this problem, which is even more annoying. As I said earlier, I hate being defeated by inanimate objects, particularly PCs. I'm now even more determined than ever to install Service Pack 1 successfully on that machine, so keep an eye on this space.

Continuing the theme of upgrading PCs, I was recently accosted by Simon Crisp, one of our Labs Project Managers, raving about a new motherboard

he was testing. It was a Pentium 4 model, with a non-Intel chipset and integrated graphics. Nothing particularly innovative there, I thought. But it turned out that the graphics performance of this board was on a par with some of the budget add-on graphics cards you can buy these days. About the same, in fact, as ATI's Radeon 9000 chips. Now that's not particularly staggering but, traditionally, integrated graphic performance has been, well, to put it bluntly, rubbish. The Radeon 9000 is perfectly adequate for an occasional gamer such as me—especially if the games are a year or so old. You can pick these cards up for around £60. Which is why Simon was so impressed—the price of the motherboard is around £75. I'm not going to name names, as we'll no doubt be reviewing it soon. In addition to the graphics, it also has integrated 6-channel audio, a 10/100Base-T NIC, six USB 2.0 ports and a FireWire port.

This got me to thinking that maybe I should consider a new upgrade strategy. In the past, the motherboard is usually the last thing that goes into the bin—if you're anything like me, you try to eke out its life for as long as possible by adding memory, new CPUs, more expansion cards and so on. I reckon on average it's about two years before I'll replace the motherboard, whereas I've been through three graphics cards in the last 9 months—at around £75 to 100 a shot. If integrated graphics technology keeps improving at this rate, it's going to make more sense in economic terms to replace the motherboard every 6 months or so, and get the benefits of whatever new goodies the manufacturers are bolting on.

A 3GHz processor should be good for at least 12 to 18 months or so before it's struggling to feed the graphics processor. So perhaps now's the time to get started—my current motherboard is coming to the end of its life, as it's starting to lack some of things I'd find handy, such as FireWire, USB 2.0 and multichannel audio. When you tot up the price of add-in cards to give you these features, you're getting to the thick end of £100. For me, it sounds like a no-brainer, but I'd be interested to know your feelings. How do you upgrade your PCs? Do you just buy a complete new system, or, like me, enjoy mixing and matching components? It's certainly more of a challenge to keep adding bargain bits as and when you can afford it, but only when it saves you money. It sounds as though that might not be the case anymore. Let me know your thoughts. You can get in touch by emailing me at the address above. ■

**If man made it, then man should be able to fix it. But why is the man who made it never around?**



# ALEX ARIAS

## You've got snail mail

The UK's postal service is under severe threat, but don't count it out quite yet. Can it resurrect itself by offering a secure electronic mail service for the nation?

alex\_arias@vnu.co.uk

**E**mail is now something so fundamental to my working life that I would find it almost impossible to do my job without it. This dependence on electronic communication is beginning to have some serious repercussions for traditional mail.

The Royal Mail's recent financial struggles may not be so much mis-management as an indicator of a change in business culture. Advertisers and, more importantly, consumers are becoming more accustomed to sending and receiving information in electronic formats instead of physical ones.

A report by the Direct Mail Information Services (DMIS) indicates an overall increase in traffic moved by the Royal Mail. However, a closer look reveals a change in business patterns as the rate of growth of mail sent by companies begins to decline. Is snail mail's coffin being prepared?

There are significant and obvious advantages for businesses working electronically, with greatly reduced costs and shorter time frames needed to complete campaigns. But more significant is the probable shift to electronic mail for the average citizen. We could soon be receiving utility statements and paying bills online. This is significant because consumer postal mail still heavily outweighs business post.

So here lies the greatest threat to traditional mail carriers. I've noticed my own mail patterns changing: I now receive electronic statements, pay my itemised mobile bill online, send an e-card for a birthday and even the humble postcard may also have had its day. Why send a postcard when you can send your personalised photo straight from your mobile phone?

Will this mean that the red pillar box is to be consigned to history? I hope not, but progress is relentless and the postal service is one area where technology has continually revolutionised the service. But it's no longer being aided, but overtaken, by technology.

The false hope of keeping a postal service alive in its current guise is doomed. The imminent increase in the price of sending mail will decrease the volumes sent and the recent move to single daily deliveries will just accelerate that process. So should the Royal Mail and other postal services be shifting away from moving physical items to electronic bits and bytes?

This has already begun in some European countries. In Finland the Finish Post offers e-services for both companies and citizens. The Italians have access to national post and electronic postal services. So a hybrid model already exists.

A global electronic postal service already exists with services like Hotmail. The success of the service demonstrates that the need is present, although security and privacy problems limit the service. Microsoft's .NET Passport goes some way to address this, but so far it's not really aimed at helping you pay your gas bill securely. And although individual's confidential information is already in the hands of many private companies, there's still reluctance to hand over this type of information to a company such as Microsoft. A recent European ruling has forced Microsoft to relax its registration requirements, allowing European citizens more control over what information can be used by it or partner companies.

Globalisation is a bit of a dirty word these days, and there are definitely dangers in allowing private enterprise too much control of personal data. Perhaps governments would run such services better. There are still many problems, though, such as how to prevent fraud and identity snatching. And, of course, this does mean having to trust your government.

The widespread use of digital signatures should be the backbone of any system, although it has to be extremely simple to use. Security is all-important and maybe in the future we'll have to verify IDs for email with retinal scans, fingerprint checks, smart ID cards or conceivably DNA verification. Much of the technology is already available and some of it in use—Japan's government issues smart ID cards for access to local government information. It's not the technology but the political and moral issues that are preventing faster adoption of these measures.

Although the chances of re-nationalising the postal service and it then providing secure services, such as I've described, may seem remote, don't dismiss it entirely. The nuclear power industry and national railways have been effectively re-acquired by UK PLC. The postal service as a state monopoly did benefit the consumer by providing a national cheap postal service for all. In the future it should also include guaranteed secure electronic services for all.

Unfortunately in these crazy days of privatisation, freed of all constraints the postal service seems to have lost its way. Perhaps it should stick to its core business of enabling interpersonal communications, via whatever medium. And maybe there will come a point in the not too distant future where the concept of a device like the Amstrad emailer—released too early for the world—will be the first thing you check as you wake up and smell the coffee. ■

**Should the Royal Mail be shifting away from moving physical items to electronic bits and bytes?**



# TRENDS

NEWS | IN BRIEF | OPINION

## Wireless notebook PCs gain edge

Laptop PCs built around Intel's Centrino technology will start shipping early next month

**N**otebook PCs with Intel's Centrino mobile technology are expected to go on sale from its official launch date of 12 March, beginning a wave of mobile PC designs with wireless networking and other high-end features.

The developments, to be showcased at March's Cebit show, are likely to further shrink the closing gap between portable and desktop PC sales.

Centrino PCs are expected from vendors including Dell, Toshiba, HP, IBM and Acer. The units integrate Wi-Fi wireless LAN (WLAN) links and use new Pentium-M processors, which prolong battery life by anticipating user behaviour, deploying tactics like supplying power only to parts in use. Although Centrino is particularly applicable to ultra-portables, some vendors will deploy the technology across their mobile ranges.

Centrino's Pentium-M processor will initially be limited to 1.6GHz, but large caches should help boost performance. Integrating WLAN access will further expand the wireless market—BT last week said it would extend coverage to 120 hotspots throughout the country.

▶ Centrino is Intel's first brand to represent a combination of technologies

Other wireless advances are planned. IBM is already shipping the ThinkPad R40 with buttons linking to online help, for instance.

Sales growth for mobile PCs is outpacing that for desktops. In Europe, the Middle East and Africa, 2.9 million notebooks were sold last quarter, an annual increase of about 15 per cent, according to IDC. Desktop sales fell by about one per cent to 9.3 million.

Sales of Windows XP-based Tablet PCs are also currently strong. Though they've only been available since November last year, they made up about one per cent of European portable PC sales in the fourth quarter, according to analyst firm Context.

MARTIN VEITCH AND KARL FLINDERS



## Intel chip gives phones an IQ boost

GSM, GPRS and application processing are supported by a single chip

**I**ntel has launched a mobile chip that integrates all the major components for a smartphone. It said this will lead to smaller, more powerful devices.

The Intel XScale PXA800F chip combines an XScale applications chip with a communications processor, GSM/GPRS functions and memory.

It is aimed at mainstream phones much like Orange's SPV rather than high-end devices, according to Tony Sica, Intel vice president for wireless communications and computing.

'The PXA800F offers the capability to run enhanced applications on mainstream phones, so they are able to handle things like Multimedia Messaging Service (MMS),' he said. Unlike previous mobile phone chips, the new XScale includes

performance headroom for running multiple applications, he added.

Handsets based on the PXA800F are not expected until the end of the year but some will be announced at this week's 3GSM conference in Cannes, including at least one Linux-based device. Low-cost PDAs will also probably feature the chip.

The PXA800F includes an XScale core running at 312MHz with 4MB of Flash memory for firmware and 512kB of RAM. A 104MHz digital signal processor (DSP) based on Intel's Micro Signal Architecture (MSA) handles comms functions, and has its own 512KB of Flash and 64KB of RAM. The PXA800F will cost \$35 (£22) in volume. DANIEL ROBINSON, IT WEEK

## WEB KITEMARKS STILL IN BUSINESS

Logos flag safe shopping sites for consumers

The Which? Web Trader scheme may have closed down in January but the concept of the safe shopping kitemark lives on.

The Interactive Media in Retail Group (IMRG) said it will continue to run its own larger but less well-known scheme featuring the Internet Shopping Is Safe



▲ The IMRG will continue to run its own scheme

(ISIS) logo, which runs on similar principles to Which? Web Trader.

James Roper, chief executive of the IMRG, said: 'The consumer has a raft of regulations to protect them but without a kitemark there is no way of knowing if a site is good or bad.'

Members of the IMRG sign up voluntarily for the right to use the logo and agree to be bound by certain criteria. These include promising to run their e-commerce services in a legal, decent, fair, honest and truthful manner, with any disputes handled promptly and efficiently.

The IMRG runs regular and random spot checks of members using the kitemark. Breaches of compliance, or any activity that threatens the reputation of the industry, may result in withdrawal of the right to display the kitemark. DINAH GREEK

# IBM builds dual-WLAN notebook PCs

Latest ThinkPad R series offers built-in support for old and new 802.11 WLAN standards

IBM has unveiled a new range of ThinkPad notebook PCs, which integrate both old and new 802.11 wireless LAN (WLAN) standards, and come with easier access to service and support. They're designed to integrate seamlessly into changing corporate infrastructures, the firm said.


The new ThinkPad R40 range, shipping now, offers dual 802.11a and 802.11b WLAN capability integrated into selected models, providing better network performance and backwards compatibility with current WLANs.

IBM's ThinkPad R series is a budget corporate range, designed to offer businesses a good blend of portability and essential features. The R40 models will be available with Mobile Pentium 4-M processors up to 2.2GHz or Mobile Celeron chips up to 1.6GHz and up to 1GB of memory.

Versions of the R40 with dual WLAN capability will ship with IBM's Access Connections location-manager utility. This tool can automatically switch the notebook between networks.

ThinkPad R40s will come with Access IBM, described by the company as a portal to IBM support. This also provides onboard diagnostic tools, accessed via a dedicated button.

IBM's traditional TrackPoint has also been given an update in the new ThinkPad R40. All



## ▲ WLAN support comes as standard

models will be fitted with IBM's UltraNav, which combines a TrackPoint with the TouchPad device found on most other notebook PCs. IBM has also introduced two new ThinkPad docking ports.

Prices start at £1,069 (ex. VAT) for a ThinkPad R40 with 1.6GHz Mobile Celeron processor, and £1,191 (ex. VAT) for a 2GHz Mobile Pentium 4-M model. *DANIEL ROBINSON, IT WEEK*

## Power lines span digital divide

Scottish Hydro-Electric is trialling broadband electricity lines



Electricity power lines are being used to deliver broadband access for the first time, in a trial that could ultimately help beat the digital divide in rural areas.

Power company Scottish Hydro-Electric (SHE) has been running technical trials since last July delivering speeds of up to 2MB/s over its existing power network. It is now planning two

larger trials, said telecoms infrastructure manager Antony Lowe.

Ultimately the service could be made available to all the company's electricity customers across northern Scotland and the south of England. 'What we have done so far is very much a technical trial. We have established we like what the technology does, we are happy to see it on the power network and there are no technical reasons that we can't deliver broadband,' said Lowe.

The cost of rolling out telecoms infrastructure to sparse rural populations is hampering the availability of high-speed access in remote areas. Electricity suppliers face high costs, but the costs look more attractive than for traditional telecoms networks, says Lowe.

'The economics we have been looking at can deliver broadband to smaller communities than it would appear BT can,' he said.

*SARAH ARNOTT*

## NEWS IN BRIEF

- Sony's latest TFT flat panels offer higher-quality images and include a 19in. model, the SDM-S91, for under £600. The Sony S series also includes 18in., 17in. and 15in. models. All are on sale now at £595, £549, £425 and £255 respectively.

- Toshiba has quietly moved to selling direct to customers. IT buyers in need of quick delivery of peripherals and even notebook PCs can now buy them directly from Toshiba's Web site, [www.computers.toshiba.co.uk](http://www.computers.toshiba.co.uk)

- TrustUK, the government-backed online hallmark, is close to signing a deal with a third party that will issue quality seals to approved Web traders. Robert Dirskovski, TrustUK secretary, hopes to attract around half of the 2,700 companies accredited under Which? Web Trader.

- Processor giant Intel has announced that it will formally launch its mobile chip, code-named Banias, on 12 March in New York. The Pentium-M processor and associated technologies (chipsets and 802.11 Wi-Fi) will be marketed under the brand name Centrino. This is the first time Intel has branded a combination of technologies under one name.

- Microsoft has filed a counter-suit against UK mobile phone maker Sendo in its dispute over the failure of the Smartphone Z100. The suit claims that Sendo 'consistently failed to meet its contractual obligations to design and develop' the phone. In December, Sendo filed a claim in a Texas court against Microsoft.

- Apple has lowered prices and simplified its flat-panel iMac product line in the US. Two base models—one with a 15in. screen and a CD-burning DVD drive at \$1,299 and another with a 17in. screen and a DVD-burning drive at \$1,799—are now \$200 cheaper than previous models. The prices of eMac desktop computers have also been reduced.

## USB DRIVE READS FLASH CARDS

**LaCie's Universal Media Drive reads wide variety of solid-state storage media**

LaCie has launched a new USB media reader that allows users to read and write to most formats of Flash storage from any PC or Macintosh system with a USB port. The drive could help users to move files more easily and quickly between handheld PCs and desktop or notebook PCs.

LaCie's Universal Media Drive, available now, supports CompactFlash, Smartmedia, Secure Digital, MultiMedia Card and Sony's Memory Stick Flash



▲ LaCie's compact reader recognises six formats

memory formats. It can also be used to access the Microdrive miniature hard disks designed by IBM—now sold by Hitachi Global Storage Technologies—which fit in to a CompactFlash slot.

'The USB Universal Media Drive is perfect for those who are working in environments where different kinds of digital media are used,' said Olivier Moulin, LaCie's product manager.

Flash memory is commonly used for non-volatile storage of data on PDAs. The most recent Palm models feature a Secure Digital slot, while devices based on Microsoft's Pocket PC platform sometimes have both CompactFlash and Secure Digital.

The Universal Media Drive is about the size of a credit card and draws power from the host computer's USB port. The drive is compatible with Windows 98/98 SE, 2000, ME and XP, Mac OS 8.6, 9.x and OS X.

The LaCie Universal Media Drive costs £19 (ex. VAT).

**DANIEL ROBINSON, IT WEEK**

# Consortium to promote Linux on PCs

Linux vendors to form association to promote open-source software on desktops

The Desktop Linux Consortium recently announced its founding membership, which includes several Linux distributors as well as organisations producing supporting applications and tools. The new association's declared goal is to shape the future of Linux as an office desktop environment, the consortium said.

Linux developer Bruce Perens said: 'The consortium will assure that there is fairness in all desktop-related issues and events. All vendors will be fully represented and the open-source ethos will be respected.'

Linux is widely used as a server platform, where its low cost and reliability have attracted many large firms, particularly for operating Web sites. Because it's open-source, the software has also been touted as an alternative to Windows on the desktop, especially as many firms object to Microsoft's recent licensing and pricing changes.



**DANIEL ROBINSON, IT WEEK**

## Desktop Linux Consortium

▲ Hoping to see Linux on the desktop

However, Linux is still seen as lacking the ease of use and choice of productivity applications that Windows boasts. These are issues that the Desktop Linux Consortium hopes to address.

Consortium members include Linux vendors Suse, MandrakeSoft, Debian, Xandros, Ximian and ArkLinux, plus OpenOffice.org, which develops an open-source productivity suite, and the KDE Project, developer of a desktop environment for Linux.

The consortium said its goals include raising awareness of the role that Linux can play on the desktop as well as facilitating communications between users and developers.

## Microsoft finds more glitches in XP and IE

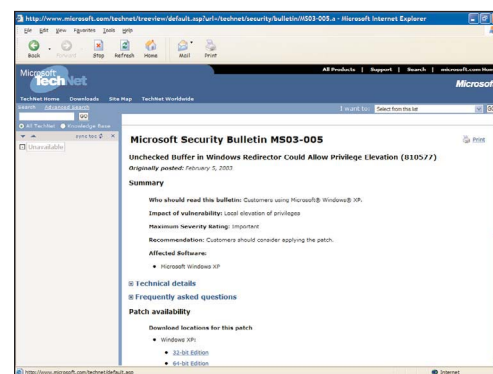
Fifth security advisory this year warns of IE and XP bugs

Two more Microsoft security advisories have appeared concerning Windows XP and Internet Explorer (IE). The latest Windows XP bug brings the total number of Microsoft security advisories issued this year to five.

The XP vulnerability has been caused by an unchecked buffer in the Windows Redirector function on the operating system, Microsoft said. The company explained that an attacker exploiting the vulnerability could crash the system or run their own code with system privileges. This could allow them to take any action on the machine, such as adding, deleting, or modifying data on the system, and creating or deleting user accounts.

According to Microsoft the vulnerability can't be exploited remotely as calls to the Windows Redirector can only be made locally. Attackers would therefore need to log on to the system using an interactive logon in order to attempt to exploit this vulnerability.

The IE glitch has been found in versions 5.01, 5.5 and 6.0. And Microsoft warned that since it no longer supports IE 5.0 and earlier versions, these could also be vulnerable. The security issue could enable an attacker to read files or run programs on a PC used to view the attacker's Web site.



▲ The advisories warn that an attacker could crash the system or run their own code

In a recent interview, Mark Greatorex, director of .Net Developer Group in the UK, said the company hadn't experienced any new vulnerabilities in the past 10 months. Two thirds of the company's developers, he added, are actually involved in software testing, not development. Security, he said, is an industry wide problem and one where 'stones cannot be continually thrown at Microsoft without dispassionately looking at what is happening elsewhere'. **NICK FARRELL**

## NEWS IN BRIEF

● Dell has announced plans to stop installing floppy disk drives on its high-end Dimension computers, offering them only as an option from next month. According to the Associated Press, the trend is toward larger capacity technology, such as portable hard drives and rewritable CDs.

Dell will first offer its own 16MB USB flash memory drives as standard in its Dimension high-end models. The company said it would then consider making flash memory available on all of its desktops, depending on demand.

● PC users have been warned to ignore an email that invites users to download a 'vital' security patch. The email invites recipients to call a premium rate number so they can find out more about downloading a patch for Windows. However, technicalsupport@microsoft.com, is non-existent and Microsoft doesn't use premium rate numbers for customer support. A Microsoft representative advised users to simply delete the email.

● NTL recently told customers of its broadband service that they will be limited to 1GB of downloads per a day. Download capping, already introduced in Europe, the US and Australia, will let operators plan future investment in networks and infrastructure. NTL has no plans to introduce draconian penalties if heavy users exceed the new limits, although it would not rule out further restrictions in the future.

● A weakness in the widely used Concurrent Versions System (CVS) development aid has left Linux and open-source code vulnerable to attack. A Computer Emergency Response Team advisory has warned the flaw could allow hackers to alter the operation of the CVS program, read sensitive information or launch denial of service attacks. Users should download a patch from their Linux supplier, or find out more from E-Matters at <http://security.e-matters.de/advisories/>

## Integrated chips to boost mobiles

Mobile chips will integrate much of the circuitry needed to build wireless mobile devices

New mobile chips from Texas Instruments (TI) and Samsung Electronics could lead to smaller wireless clients with better performance than current models. Both firms have announced ARM-compatible chips that integrate most or all of the functions required for a handheld PDA or smartphone into a single component.

TI has announced forthcoming additions to its Open Multimedia Applications Protocol (Omap) family, while Samsung launched a device, code-named Oyster, that combines a processor with Flash storage. Both companies are using system-in-package designs where two or more silicon chips are 'stacked' on top of each other.

TI's Omap 1610, 1611 and 1612 are designed to work with existing baseband radio chipsets. How-

ever, the Omap 730 and 732 integrate a radio for GSM and GPRS mobile networks. Both families include hardware acceleration for processing security algorithms and are due to ship in the fourth quarter of this year.

The Omaps offer up to 256MB of stacked low-power SDRAM. This will reduce the size and power consumption of smartphones and PDAs.

The new Omap chips have hardware support for Secure Digital (SD) slots, can interface directly to LCD displays, and have a high-speed serial bus for connecting to TI wireless LAN adapter chips.

Samsung's Oyster, available in June, combines the firm's S3C2410 ARM-based processor core stacked with 256MB of Flash memory for firmware and 256MB of SDRAM in a single package. 'Our solution offers the world's first combination of an application processor with Nand Flash, which will be the leading non-volatile storage solution for next-generation handhelds,' said Samsung's Hyung Lae Roh. Nand Flash is a memory design optimised for file storage.

Intel is expected to introduce its XScale PXA261 and PXA262 processors this quarter, combining ARM-compatible processors with Flash memory. *DANIEL ROBINSON, IT WEEK*



▲ The Oyster combines processor with Flash memory

## Dell takes PDAs in hand

Two new models spearhead push into the handheld market

Dell has sounded a wake-up call to other manufacturers in the competitive handheld sector with the launch of its first own-brand device, the Axim X5 (see our full review).

Despite sluggish sales over the last year, analyst company Gartner said it expected Dell to pick up 25 per cent of global PDA sales with the X5.

While other devices feature selling points like looks and being light weight, it is felt that Dell's aggressive pricing poses a real threat to other manufacturers such as HP. Gartner predicted that even Palm, which is dominant in the lower end of the market, will feel the pressure to cut prices or come up with its own unique selling points.


Based on Microsoft Pocket PC 2002 software, Dell has launched two versions of the X5.

The higher end device will cost £229 and features a 400MHz processor, 64MB of SDRAM and 48MB of Intel StrataFlash ROM. The cradle sup-

plied with the device has an extra battery charge, so that battery life can be extended by up to as much as 25 hours.

For users with lower memory and speed requirements Dell is offering a cheaper X5, priced at £169, which uses a 300MHz Intel processor, 32MB of SDRAM, 32MB of Flash ROM and a sync cable instead of a cradle.

You can choose from two wireless options, Bluetooth or Wi-Fi 802.11b, with either unit. Both models will also come with a wide variety of accessories, such as folding and snap-on keyboards. *DINAH GREEK*





## NEWS IN BRIEF

● Europe's mobile phones will jump from General Packet Radio Service (GPRS) straight to third-generation (3G), leapfrogging the intermediate Edge technology, according to analysts. Edge has regularly been promoted as the last step towards 3G, but analysts from The Shostek Group, which looks at the cellular and wireless Internet markets, now forecast that it will only be adopted in the US and the Americas, where progress towards 3G is well behind Europe.

● Rambus had been suing memory chip makers Infineon, Micron Technology and Hynix for non-payment of royalties when Infineon launched a fraud counter claim. But a three-judge US Court of Appeals panel has now ruled by two to one that a previous jury was mistaken in its verdict of fraud and that the evidence didn't support it. The company had been accused of amending pending secret patent applications to take into account future memory technology standards agreed at standards body meetings it attended.

● Microsoft and nVidia have settled their dispute over the price of chips for the Xbox games console. Last spring Microsoft took nVidia to court for alleged violations of a supply agreement. But nVidia has now said that the companies have agreed to collaborate on future cost reductions for the Xbox. The company announced that it would give more details after its quarterly earnings report.

● Cheaper alternatives to 2Mbit/s E1 leased lines may soon be more widely available, with BT set to increase the number of exchanges supporting Symmetric DSL (SDSL) services to around 180. BT's 2Mbit/s SDSL will be leased wholesale to other service providers at a fixed price of £200 per month, plus a connection fee of £450 per customer. SDSL bandwidth is shared rather than dedicated, with contention rates up to 10:1, and SLAs are rarely offered.

## Good for another decade, says creator of Moore's law

The sun will rise and chip speeds will double every two years

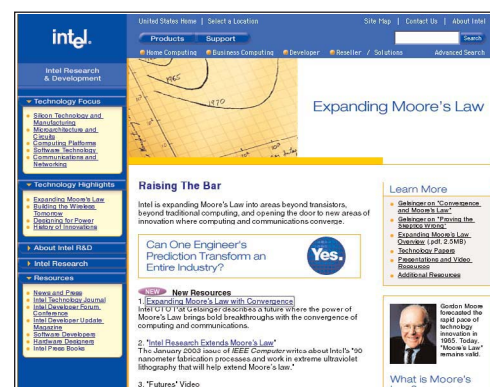
Computer pioneer Gordon Moore recently predicted that the law to which he gave his name will last for at least another decade.

The 74-year-old co-founder of Intel, the world's biggest chipmaker, predicted in 1965 that the number of transistors on a semiconductor, and thus overall chip performance, would double every two years.

According to Reuters, the creator of Moore's Law told a meeting of many of the world's pre-eminent chip designers that engineers must now concentrate on overcoming power leakage and reducing heat levels as more and more circuits are crammed closer together.

Asked if Moore's Law would run its course in the foreseeable future, he replied: 'Another decade is probably straightforward. None of these things hits an abrupt wall.'

Current limitations on the law include electrical power leakage and heat dissipation that



▲ Moore remains chairman emeritus at Intel

increases each time more transistors are packed into a smaller area.

This makes chips 'not far from the power density of a nuclear reactor', said Moore. **NICK FARRELL**

## E-traders flout legislation

Almost half of the UK's top e-traders breach consumer rights legislation

Leading retail Web sites are breaching laws that protect the rights of consumers when shopping online and risk fines as a result, according to a recent survey.

Almost half of Internet retailers are currently breaking the law on consumer rights, according to a study of 50 leading e-commerce sites carried out by law firm Browne Jacobson. It found that while the UK's top 10 retailers comply with legislation, many firms in the top 50 of the Retail Week top 500 fail to do so.

Sites were assessed for compliance with key e-trade legislation like the Data Protection Act.

Firms breaching the law included one that failed to provide an appropriate data protection consent form. Another told customers that personal data would be passed to third parties unless they sent an email opting out, but consumers must opt in to receive marketing information.

Firms must ensure compliance with the law or could face fines, said Simon White, a solicitor at

Browne Jacobson. For example, firms that process consumer data without obtaining consent and registering under the DPA are committing a criminal act and could be fined £5,000. 'One of the biggest failings of sites was not complying in spirit or in fact with the act,' said White.

Firms could also face fines if they breach the Distance Selling Regulations. If consumers aren't given a clear indication of the returns policy, they can claim a full refund for up to three months. 'If companies are selling expensive goods, this is quite risky,' said White.

Several steps must be taken to ensure full compliance. Firms must check whether they collect and process personal data. If so, they must ensure they're registered under the DPA and have a system set up to request and record users' consent. Companies also need to make sure terms and conditions are easy to access and understand, and that they provide consumers with contact details and returns policies. **MADELINE BENNETT, IT WEEK**

**'One of the biggest failings of sites was not complying in spirit with the act.'**



OPINION  
GUY KEWNEY

## GAMES AND PCS

JAVA ISN'T A GAMES PLATFORM, or rather it wasn't intended to be one. It has, however, become one on phones and portable PDAs. But it may be about to lose its crown—unwanted and unplanned though it was.

You can get into trouble if you have a good idea and tell lots of people about it.

Thirty years ago, I went through that with 'micro-computers', which I thought were a pretty good idea. And at some point over the next three decades, people discovered I had another little obsession—computer games.

Unfortunately I took them seriously, not just as things to do when you should be doing important work, but as a powerful force in the computer industry, an instrument for change. I believed that they would force the development of smarter hardware and innovative software, because they always pushed the capabilities of machines. And so I suppose it wasn't terribly surprising that some people thought Tao Systems was a bit of a joke of mine.

Tao started out as the way-out idea of a truly eccentric games programmer. He

came up with the idea of a universal emulator, which would run on any processor. And even better than that, it could multi-task, so that you could have two processors sharing the same application. That meant if you took one away, the application wouldn't care—it would just pull in the remote task to a new virtual machine on the first processor.

At the time, this idea was the preserve of the very rich. It was networked multi-processor arrays and, generally speaking, it couldn't be done. So the claim to have done it wasn't taken seriously by many.

Today, a decade after I first heard about it, this platform is a threat to Java. In fact, you may already have a Tao product in your pocket, as the biggest users of the platform are games providers for phones. Hardly a threat to Java, you might think. But you'd be wrong, because the force of convergence, so glibly used by seminar pundits and columnists, is genuinely changing the world.

You will find Tao in things like JVC digicams, an iPronto set-top box, a Kyocera PDA

and a growing number of 'multimedia platforms' for the simple reason that it truly does run everywhere. The Tao engine does what the Java virtual machine promises: you write the application and then just port the Tao engine to your platform. And because it was created as a full games platform, it isn't just an audio streaming plug-in, a pixel-shifter extra, or just a pocket browser. It's a complete multi-tasking, multiprocessor networking operating system.

It is also blindingly fast, which is the important part.

Tao has been selling a trojan horse. It's been going to hardware makers and saying: 'Use our Java interpreter and it will run faster.' It does. What they haven't said is: 'If you use our engine, which we use to speed up Java, but use it directly, it will go even quicker. Ten times quicker, at least.' And it does.

From being a derided hobby of an eccentric games writer, Tao is now a company which has investment from a wide array of mobile and phone companies.

It's not just a threat to Java, though no doubt

people will smile when I say that it's a threat to any mainstream operating system built on client-server ancestry, such as Unix or Windows. On the desktop, it has always been the case that if you need new features, you can plug in more hardware.

It's got beyond a joke, though. I can remember running Desqview on a 10MHz 286 processor with full multi-tasking and getting virtually instant response from half a dozen applications, out of a machine with only one megabyte of internal RAM, and less than 50MB of disk space. Today, bloat has taken things to the point where a quarter of a gigabyte of RAM is regarded as small. You can get away with that sort of bloat on a desktop, but not in a battery-powered, mobile machine. And what you can do comfortably on a pocket PC, you can do too fast to see on a 3GHz P4.

It won't happen this year and it'll be ignored for a lot longer than that, but Tao is about to eat Microsoft alive—from the inside. Windows is simply no way to run a mobile device and Tao is. And from there it will grow...



**PC Magazine forum** If you have any comments regarding this e-book, why not visit the *PC Magazine* forum.

It can be found at [www.pcmag.co.uk](http://www.pcmag.co.uk). All you need to do is click the reader forum link. To add your comments you'll need to register, which is quick and free.

**Enquiries** *PC Magazine*, VNU House, 32-34 Broadwick Street, London W1A 2HG. **Tel** 020-7316 9000 **Fax** 020-7316 9089

**Web address** [www.pcmag.co.uk](http://www.pcmag.co.uk) **Editorial email** [feedback@pcmag.co.uk](mailto:feedback@pcmag.co.uk) **Press release email** [pcmag\\_info@vnu.co.uk](mailto:pcmag_info@vnu.co.uk)

**Editor** Tony Westbrook **Art Editor** Hazel Bennington **Production editor** Mark Hamilton **VNU Labs** Kelvyn Taylor, Paul Monckton, Simon Crisp, Alex Arias

# Dell enters PDA market with budget contender



◀ Dell's Axim X5 is bulky and doesn't offer any new features, but its low cost will sway some hesitant buyers

## POCKET PC

A perennial complaint about Pocket PCs is their high price compared to Palm OS-based PDAs. Starting at £169 ex. VAT for the entry-level model, Dell's new Axim X5 looks set to change all that. The Standard Model is powered by a 300MHz Intel XScale CPU with 32MB RAM/32MB ROM, while the Performance Model reviewed here offers 64MB RAM/48MB ROM, a 400MHz XScale CPU and an enhanced accessory bundle.

The X5 is larger and bulkier than most of the current crop of PDAs at just under 200g. Some of the extra width is due to the rubber surround Dell has built into the sides for added protection. For expansion there's an SD card slot in the side of the unit while a CompactFlash Type II slot is located at the top. These don't have covers, but use plastic dummy inserts, which will no doubt quickly get lost. The infrared port and headphone sockets are also located at the top of the unit.

The Performance Model features a stylish mirror-finish docking cradle with a recharging slot for a spare battery. The docking and undocking action isn't particularly smooth, however. A blue Dell logo illuminates when the PDA is in the cradle and there's a red/green light in the Axim's power button to show charge status. The Standard

Model comes with only a USB synchronisation cable, slip case and a power adapter. Both models use a 1,440mAh Li-ion battery, but a high capacity 3,400mAh battery is available for £69 (ex. VAT). In our informal tests using the standard battery, we achieved 6.5 hours of continuous MP3 playback on 50 per cent volume and the backlight switched off. With the backlight at 50 per cent brightness this reduced to 5 hours. Running just a timer utility with the backlight set at 50 per cent gave us well over 9 hours. Recharging with the device turned off took around three hours.

Navigation buttons follow the standard Pocket PC layout with four quick-launch keys and a central five-way paddle. There's also a jog switch on the left, above a recessed (and awkward) recording button. Audio quality's not too impressive.

The Axim X5 has a 90.3mm (diagonal) transreflective TFT screen that displays up to 65,536 colours and has a resolution of 320 by 240 pixels. Its brightness and clarity are good, and it's just about usable under office lighting without the backlight. The solid-metal, flat stylus is robust and the touch-screen's sensitivity is good.

There's nothing really revolutionary about the Axim X5, but it has no serious flaws and a price that should make it a popular choice with Pocket PC aficionados. *SIMON CRISP AND KELVYN TAYLOR*

## DELL AXIM X5

**Verdict** It's taken a long time, but finally there's an affordable Pocket PC

- ⊕ Good screen; useful battery life
- ⊖ Bulky; audio quality could be improved

### Fact file

<b>Dimensions</b>	18 x 81.5 x 128mm
<b>Weight</b>	196g
<b>Processor</b>	400MHz Intel XScale
<b>Memory</b>	64MB RAM/48MB ROM
<b>Display type</b>	Transflective colour TFT
<b>Resolution</b>	320 by 240 pixels
<b>Colours</b>	65,536
<b>Expansion</b>	1 x Type II CF card slot, 1 x SD card slot
<b>OS</b>	Microsoft Pocket PC 2002 Premium
<b>Other software</b>	Dell Switcher, Dell Data Backup/Restore,
<b>Accessories</b>	Slip case, USB docking cradle, AC power adapter
<b>Warranty</b>	One year, next business day advanced exchange
<b>Contact</b>	Dell UK • 0870 152 4699 www.euro.dell.com
<b>Price</b>	£229 (ex. VAT) for Performance Model

# Top quality at high speed



◀ The Phaser 7300's new printer drums and toner formulation combine to produce impressive output quality

## A3 POSTSCRIPT COLOUR PRINTER

Xerox's Phaser 7300 series is its latest range of workgroup A3 colour printers aimed at high-demand office use. Using single-pass technology to deliver up to 30ppm in colour and 37ppm in monochrome, the high-end Phaser 7300DT reviewed here powers through large, multi-page documents with impressive speed and quality.

The 7300 differs from the flagship Phaser 7700 in that it uses LED rather than laser technology. With previous printers, such as the Phaser 2135, this has meant a substantial drop in image quality, but this isn't the case with the Phaser 7300. New imaging drums and toner formulation combined with improved colour tables deliver impressive quality that only the fussiest of users would be able to distinguish from the output of the Phaser 7700. If you're not doing highly demanding pre-press work and you don't require the advanced profiling facilities of the Phaser 7700, then the Phaser 7300 is likely to deliver all you need and more.

Although fitted with a 500MHz G3 processor rather than a 500MHz G4, the Phaser 7700 easily packs enough power to keep up with the engine when printing from any office-type application.

For a printer so crammed with features, it's remarkably easy to use. The new driver software is now easier to install and shows a clear graphical representation of your selected settings as well as automatically detecting the page settings you may have configured inside Microsoft Office applications. Network installation has also been made easier—start the software installation then, when asked, press a button on the printer's front panel.

Your network settings will then be automatically configured to connect to the right printer. The printer controls themselves are of the improved kind found on all recent Xerox printers, offering a 6-line illuminated graphical LCD control panel complete with online help and a selection of printable documentation and demos.

The Phaser 7300 can cope with a variety of paper thicknesses including card stock, business cards and transparencies. New automatic thickness and transparency sensors remove the need to reconfigure the printer when changing stock and minimise jams and other problems that could arise from making the wrong selection.

Available as a range of four models, the Phaser 7300B starts at £2,325 (ex. VAT) for a non-networked 600 by 600dpi version with parallel and USB connectivity and 128MB of RAM. At £2,545 (ex. VAT), the Phaser 7300N adds 10/100Base-T Ethernet, oversize A3 capability (450 by 320mm), an extra 64MB of RAM and all of the usual Xerox networking software, including a built-in Web server. The maximum resolution is also boosted to a true 600 by 2,400 dpi. The DT model (£3,405 ex. VAT) adds automatic duplex and banner (328 by 914mm) printing, a 20GB hard drive and an additional paper tray while bringing the total memory to 256MB. The flagship DX model (£4,055 ex. VAT) offers 384MB RAM and yet another paper tray, bringing the total paper capacity to 2,300 sheets. All models come with true Adobe PostScript 3 and PCL 5c as standard.

The Xerox Phaser 7300 is an astonishingly fast printer, which makes no compromises on quality or ease of use. **PAUL MONCKTON**

## XEROX PHASER 7300DT

**Verdict** Perfect for high-demand office use with some overlap into the graphic arts sector

- Very fast, high-quality output with flexible media handling
- Bulky and heavy base models, although larger configurations have wheels

### Fact file

**OS** Windows 9x, ME, NT, 2000, XP, MacOS 8/9/X, NetWare 3 to 6x, Unix

**Engine technology** LED

**Maximum print resolution**  
600 x 2,400dpi  
(Phaser 7300B 600 by 600dpi)

**Print processor**  
500MHz PowerPC G3

**Memory** 256MB

**Hard disk size** 20GB

**Duty cycle** 83,000 pages/month

**Interface** Parallel, USB, 10/100Base-T

**Contact** Xerox • (0870) 241 3245  
www.xerox.co.uk

**Price** £3,405 (ex. VAT)



# Easy-to-use storage solution lacks redundancy

## NETWORK ATTACHED STORAGE DEVICE

Linksys, more commonly known for its wireless equipment, has now entered the Network Attached Storage (NAS) market with the Instant GigaDrive EFG80. Designed primarily for the small/home office markets it provides a simple-to-manage, quickly deployable storage solution.

The Linksys GigaDrive EFG80 is supplied with a single Maxtor 80GB Ultra ATA/100 drive installed in one of the two lockable drive bays. The additional drive bay can be fitted with any E-IDE drive up to 120GB (typical price £140 ex. VAT, but not sold by Linksys). This would allow you to set up one drive as a backup, giving you some measure of redundancy.

A cut-down version of Linux is the base operating system, although you'd never know it as all the administration is done via a Web browser. A setup wizard guides you quickly through the initial setup process and you'll find the Web interface simple to navigate.

It's possible to set up a variety of user groups with individual storage quota levels. Additionally there's a parallel port connection for the built-in Print server. This is controlled through the administration interface, although the options for managing the print queue are limited. There are a number of utilities to maintain the file system structure and performance and you can schedule SMART analysis on the drives to monitor any potential hardware problems.

Support for more advanced features, such as Internet Attached Storage (IAS) and Internet Printing Protocol (IPP), allows remote users to retrieve and print files over the Internet via their Web browser. However, while these are undoubtedly useful features, they can also introduce additional infrastructure and security issues unless used within a secure intranet.



▼ Nice in blue, but the GigaDrive EFG80 can't be striped, affecting its ultimate performance

Unfortunately the twin drives can't be striped (RAID 0), which reduces the ultimate performance, although the benchmark results show they would be capable of supporting a small office environment without a problem in any case.

As with all NAS products the main advantage over cheaper entry-level servers is the lower overall administration costs and the Instant GigaDrive certainly provides a simple and clearly explained administration interface.

There's plenty of competition in this field and, although slightly cheaper than other similar models, there are fewer redundancy options than some alternatives. That said, the EFG80 ease of use will appeal to those with limited time for administration or technical knowledge. *ALEX ARIAS*

### LINKSYS INSTANT GIGADRIVE EFG80

**Verdict** Rapid deployment will appeal to those who require a basic file and print server for quick additional storage

- ⊕ Inexpensive; ease to use
- ⊖ Basic services; no RAID available

#### Fact file

**Platform support** Wizard runs on Windows 98, ME, NT4, 2000 and XP

**Protocols support** SMB/CIFS over TCP/IP

**Maximum capacity** 240GB (drives purchased separately)

**Additional services** DHCP server, Print Server, FTP server


**Warranty** 2 years RTB

**Contact** D-Link • 020-8731 5555  
www.linksys.co.uk

**Price** £519 (ex. VAT)

## VNU Labs tests

### Linksys Instant GigaDrive EFG80



The lack of RAID level 0 affects the GigaDrive's performance, as the Snap 2200 demonstrates with its higher throughput levels. However, the performance is adequate for a small office environment of up to 20 users

# Guard against power surges

## UNINTERRUPTIBLE POWER SUPPLY

Uninterruptible Power Supplies (UPS) are generally perceived to be worthy, but ultimately dull, pieces of IT equipment. The BlackoutBuster B10E2 line interactive UPS from Power Kinetics does little to dispel this perception, but within its field it's a competent unit.

A small UPS for home/office use is a relatively low-cost way of protecting your system against power failures or surges. In many areas of the country they're indispensable and the surge pro-

tection element could prevent several hundred pounds' worth of damage.

The 500VA version of the BlackoutBuster series reviewed here is a very compact model. It provides a modem surge protection port and a single power outlet with a circuit breaker for overload protection. At the front are four warning lights and an on/off switch. Unusually the mains input cable is captive. The power outlet's connected to a Y-splitter mains cable to allow connection of two items—not an elegant solution but adequate for most needs. There's no pass-through surge protection port for high-powered devices. A serial cable sends status information to the PC for monitoring by the MonUPS software. The battery itself has a one-year guarantee and is user replaceable.

The BlackoutBuster does its job without fuss, but it's pricier than competing products with fewer features. Users with several peripherals to protect may prefer to look at the clunkier (but cheaper) B4E 400VA model. *KELVYN TAYLOR*

## POWER KINETICS BLACKOUTBUSTER B10E2

**Verdict Compact and functional, but not particularly well-featured for the price**

- ➕ Compact; easy to use software
- ➖ Single power output; no pass-through socket

### Fact file

**Dimensions** (DxWxH)

21.5 x 9.6 x 21.7

**Weight** 5.5kg

**UPS type** Line interactive

**Power output** 500VA, 2.2A

**Back-up time** 35 minutes (manufacturer's estimate)

**Software requirements**

Windows 9x, NT, 2000, XP

**Accessories supplied** Mains output splitter, modem cable, serial data cable.

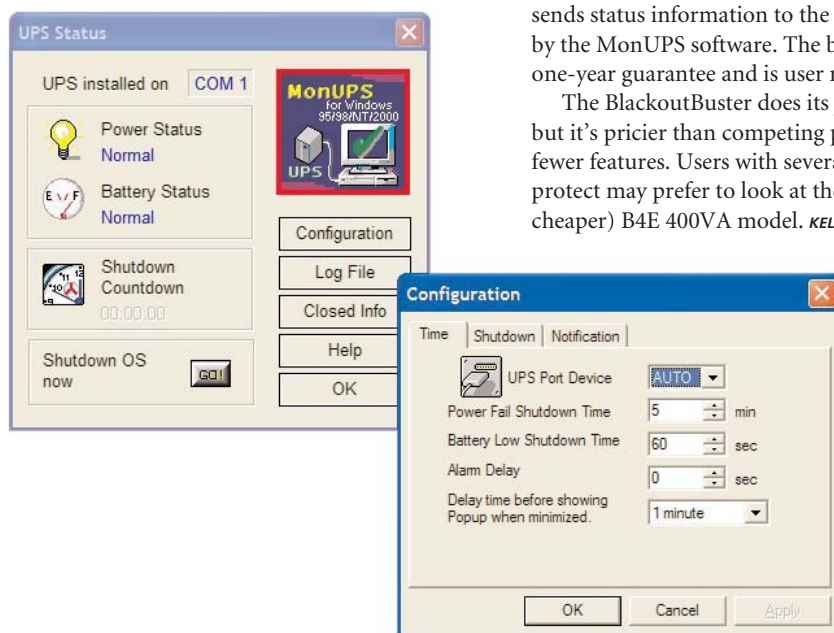
**Warranty** Three years RTB (battery—one year)

**Contact** Power Kinetics

(01634) 739188

[www.powerkinetics-emea.com](http://www.powerkinetics-emea.com)

**Price** £99 (ex. VAT)



◀ MonUPS is a simple monitoring and configuration utility that checks the UPS status via the serial port

# Latest driver boosts performance

## CHIPSET DRIVER PACKAGE

VIA's 4in1 chipset driver package has been the mainstay driver download for VIA-based systems for many years. The latest version of the package—version 4.45—also sees the launch of a new name, the Hyperion 4in1. This version is optimised for all the latest versions of VIA's chipset family but particularly the VIA Apollo KT400 and VIA Apollo P4X400 chipsets.

The most important parts of the the Hyperion 4in1 package are the IDE Filter driver 1.20a, which enables bus mastering on ATA drives, an INF driver, which identifies your VIA chipset and an updated AGP driver.

When tested on our KT400-based test-bed system this new driver update showed a significant performance increase over the previous driver package. *SIMON CRISP*

## VNU Labs tests

### VIA Hyperion 4in1 (4.45)



## System performance

Multimedia Content Creation Winstone 2003



SYSmark 2002



Reference—VIA 4in1 (4.43)

VIA's new Hyperion 4in1 unified drivers offer a 19.7 per cent performance increase in Multimedia Content Creation Winstone 2003 and an 8.5 per cent increase when benchmarked using SYSmark 2002

## VIA HYPERION 4IN1

**Verdict A useful driver update from VIA that gives a reasonable performance increase**

- ➕ Improved performance resolves some AGP problems
- ➖ It's still annoying that VIA needs to update its drivers so frequently

### Fact file

**Download** [www.viaarena.com](http://www.viaarena.com)

**File size** 1.29MB

**Supported chipsets** Any VIA chipsets except systems with hybrid northbridge/southbridge chipsets

**OS** Windows 98, ME, 2000, NT, XP

**Contact** VIA • [www.via.com.tw](http://www.via.com.tw)

**Price** Free





# What's next for the Office Suite?

We've only just started getting used to Office XP, but Microsoft's already well into the beta program for the next version of its ubiquitous office suite. *Geoff Eimon* investigates what the future of office productivity applications holds in store for us

## HIGHLIGHTS

- Expected to ship towards the end of the year
- Better support for collaborative working
- Core support for industry-standard XML
- Managing email in Outlook has been transformed by new automated features
- A new personal productivity tool, One Note, marks a new departure for Microsoft

Office 11 is the code name for the next version of Microsoft's Office Suite. Currently the Office 11 Beta 1 version is in limited release with selected testers. When it ships towards the end of this year, the most likely name is the unsurprising Office 2004. If, like most Office users, you're still struggling with the new features in Office 2000 or XP, then you can relax. Office 2004 will bring with it a couple of tempting new features that you'll appreciate—most notably a long-awaited reworking of Outlook. There's also a new Office family application called One Note (more of this later) that some of us will 'die for'. But the real thrust of Office 2004 is aimed at the corporate strategist rather than the day-to-day user and, because it's not an 'out of the box' strategy and deeply dependent on an organisation rethinking its information management strategy, it's unlikely to be implemented in the short term. According to Simon Marks, Redmond's product manager for Office 2004, its major goal is to create a 'rich technical environment which gives organisations and individuals tools for connecting business processes, integrating collaboration and giving better control over the management of information.'

This new information management strategy that Microsoft is offering organisations through Office 2004 is based on two developments. The first of these is core support for the Extensible Markup Language (XML), which is currently seen to be the best technical solution for the exchange, processing and analysis of information. While Word, Excel and Access previously had had only a limited capability to use and work with XML tagged information, in their 2004 incarnations, XML tagged data is a native form—and the onus will be on an organisation to structure its business data to be usable by Office 2004. The second development comes in the form of upgrades to Microsoft's SharePoint Team Service (STS) Collaboration server. At the time that Office 2004 is shipped, a new version of STS (version 1.1)

will also be available and it will give Office 2004 users significantly better facilities for collaboration and document management.

## XML developments

The freshly engineered support for XML allows Word, Excel and Access to use XML tagged files as native file formats. Essentially an XML file consists of plain text—where text and numeric data and document structure is indicated by XML tags. According to Microsoft, a benefit of moving to XML file formats is that the islands of information that are created by the arcane file formats of .xls, .doc and .mdb files will disappear allowing for free exchange of information between Word, Excel and Access, promoting easy re-use of information.

Word 2004, Excel 2004 and Access 2004 now also support XML schemas that allow the applications to tap in to any data source where the data is tagged using industry standard, in-house or customer or supplier defined XML tagging schema. Using Web Services an XML tag specifying the Sales Total for Widget in Quarter 3, for example, can go directly to a data source to

retrieve the current 'value' of item. To make this new information access system work as it's intended, an organisation must first create libraries of available tagged data items. Word 2004, Excel 2004 and Access 2004 then access these libraries through new Task panes that allow users to drag and drop in terms to documents or tables.

Built-in support for XML schemas and XML Web Services will then allow Office 2004 developers to more easily create documents and applications that tap in to an organisation's business processes and its data, and to extract specified information from external data providers. To allow for organisation-wide support for using XML-formatted documents, Word 2004 supports XML templates that can then be created to structure specific documents, such as reports, letters and memos, without the document creator needing to have any knowledge of XML. Similar support is provided for Excel. Build-



**The real thrust of Office 2004 is aimed at the corporate strategist rather than the day-to-day user**



## One Note

Personal productivity tools

## OFFICE 2004 PREVIEW

INTRO | XML | COLLABORATION | OUTLOOK | ONE NOTE

XML tagged fields is then automatically fed back in to the appropriate back-end databases—at least that's the theory behind smart documents, ensuring easy re-use of information.

### Developments in collaboration

At the time that Office 2004 is released, version 1.1 of SharePoint Team Services (STS)—Microsoft's collaboration portal creator—will also be released. Through STS 1.1, Office 2004 users will be able to create and access collaborative workspaces from within Word and Excel to support their work teams. In version 1.1, STS's document versioning and check-in/check-out control has been enhanced and individual sections of documents can be locked and modified by specified individuals only. Individual team members can 'subscribe' to documents to obtain email notification when any changes are made to that document.

These STS enhancements will enable team members to work much more securely on Office documents either individually or collaboratively. Later in 2003 or early 2004, a new collaboration server, currently code-named Greenwich, will provide 'presence' information to Word or Excel. Team 'buddy lists' will be supported, which will indicate what other members of a work group are currently online and available for discussion using embedded Instant Messenger services.

### Outlook enhancements

While features aimed at supporting the organisation predominate in Office 2004, individual users haven't been neglected. Regular Outlook users will love the new mailbox display options—these are good enough to warrant upgrading to Office 2004 in their own right. The basic order-by-time-of-arrival view has been transformed by the ability to group email by Today, Yesterday, Tuesday, Last Week, Two Weeks Ago and so on, and it lets you very easily 'eye-ball' your inbox. You can also view your inbox by groupings, such as email account, conversation, form, size, attachments and importance. Very usefully, in addition to your inbox and your archive folders, you can now create virtual folders that automatically display email that meets specific search criteria. You can set up a virtual folder to display all email from a specific sender—to display unread email or email flagged for follow up. Outlook 2004 then continuously monitors all its mail folders to locate items to display in the specified search folders. To support manual categorisation of email messages, Outlook 2004 implements a process called Quick Flagging. With a single mouse click you can add a coloured flag and the flagged message then automatically appears in the Follow-up virtual folder to help you better organise and manage your work.

Outlook's Preview pane also gets a work-over using technologies developed for Microsoft's e-reader. Now called the Reading pane, it implements ClearType, sub-pixel positioning, intelligent paragraph spacing and optimal line length determination to enhance the legibility of document preview—and the difference is impressive.

	Three Months Ended March 31		Nine Months Ended March 31	
	2001	2002	2001	2002
Revenue	\$6,403	\$9	\$18,719	\$9
Operating expenses:				
Cost of revenue	899	1,395	2,588	3,823
Research and development	1,069	1,066	3,015	3,123
Sales and marketing	1,193	1,240	3,526	3,864
General and administrative	239	246	621	1,266
Total operating expenses	3,405	3,947	9,770	12,076
Operating income	2,998	3,298	8,969	9,096
Losses on equity investees and other	(46)	(11)	(126)	(78)
Investment income	706	739	2,384	312
Income before income taxes	3,658	4,026	11,427	9,270
Provision for income taxes	1,207	1,288	3,771	2,966
Income before accounting change	2,451	2,738	7,656	6,304
Cumulative effect of accounting change	0	0	(375)	0
Net income	\$2,451	\$2,738	\$7,281	\$6,304
Earnings per share:				
Basic before accounting change	\$ 0.46	\$ 0.51	\$ 1.44	\$ 1.17
Diluted before accounting change	\$ 0.44	\$ 0.49	\$ 1.37	\$ 1.12
Basic	\$ 0.46	\$ 0.51	\$ 1.37	\$ 1.17

▲ An Excel 2004 Smart Document encapsulates a business process using information derived from its library of XML tagged data—shown in the XML structure pane at the left.

**Horizons Travel**  
Annual Report 2002

Prepared by  
John Test, Hernan Savastano, Anne Archambault, Raj Merchant, Ethan Bernstein, Roxanne Kenison, Greg Panter

**The Horizons Travel Mission**

- After a good trip, our customer will return home enriched with lasting impressions of beautiful places seen and interesting people encountered
- At Horizons Travel, we are committed to the idea that adventure, comfort, and good fun are essential ingredients in the making of an ideal travel experience

**Annual Revenue**

This table shows the annual revenue for the fiscal year 2002 (in the thousands of dollars). In comparison to fiscal year 2001, sales have risen 15% overall.

Gross Revenue (in thousands)	Europe	South America	Asia	Africa	Caribbean
Delta			\$528	\$64E	\$675
British Airways			\$345	\$73E	\$323

▲ To support collaborative working, Word 2004's Team pane displays team members currently on-line and available for contact via Instant Messenger services

Date	Description	Amount

▲ Excel 2004 and Word 2004 are specialised to work with libraries of XML tagged information to drag and drop into tables and documents to create their information structure

ing on this new capability, Word 2004 and Excel 2004 both support a new type of document, which Microsoft calls smart documents.

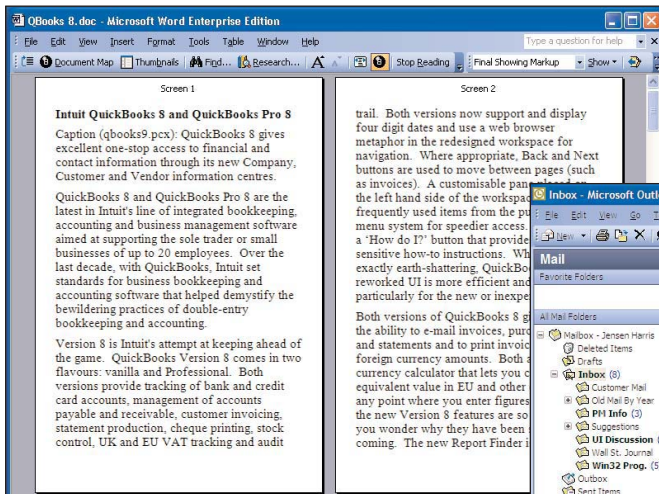
Microsoft describes a smart document as a business process encapsulated within a document and the idea is that all core business processes are supported through libraries of smart documents. A simple example could be an invoice in a Word or Excel document. Created by an Office developer using XML tags to define data fields and the sources of the data, this invoice smart document would automatically draw down prices and customer-related discounts for each item as it's entered in to the smart document invoice template. On completion of the invoice, the new information held in its

## Office Tutor

How to transfer  
from Excel to Access

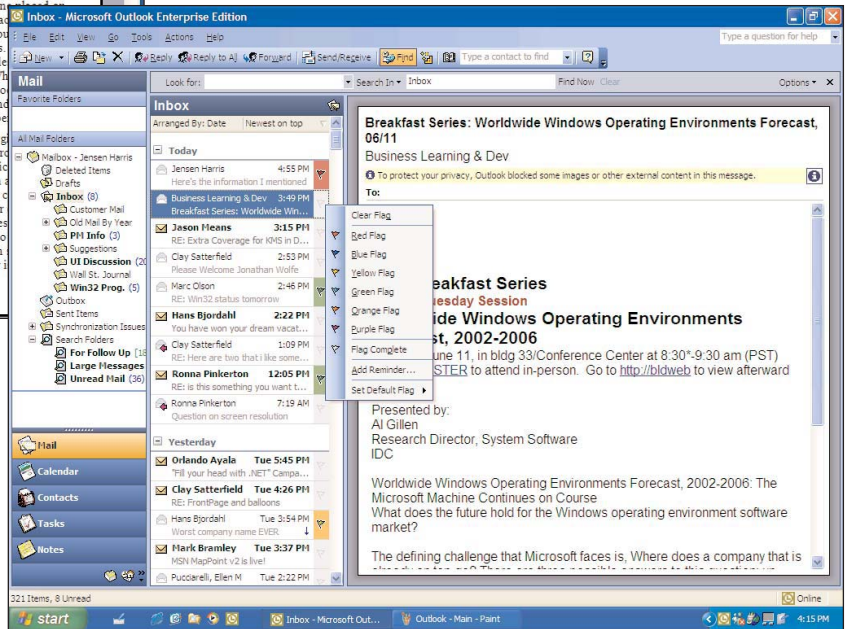
## OFFICE 2004 PREVIEW

INTRO | XML | COLLABORATION | OUTLOOK | ONE NOTE



Word 2004 gains ClearType technology to making document more legible on-screen

Outlook 2004's Inbox workspace has been dramatically re-vamped to make email management more effective



To support mobile mail-use better, Outlook now detects when the user is online or offline and when online can be set to automatically cache a local copy of the mailbox to allow synchronisation of any offline work on the next connection.

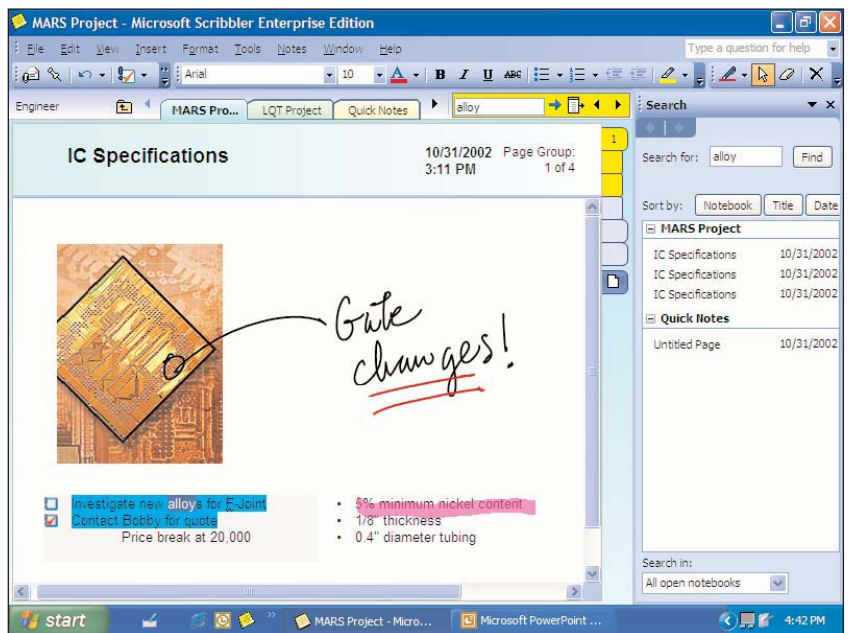
### One Note

Another development in the Office 2004 family is the entirely new One Note application. Part outliner, part notebook, part audio recorder, part sketchpad, One Note is a personal productivity tool that marks a new venture for Microsoft.

One Note provides a tabbed workspace giving access to multiple note pages. Its best feature is that you can enter information anywhere on a note page just as you would on a piece of paper—using keyboard, pen stylus, microphone or by drag and drop from graphic files and source documents, such as HTML pages. In a note pad page you can create a collage of items—that can be just as 'messy' as when you take hand-written notes in a meeting. One Note is, however, intelligent enough to treat each as a separate item within the page. Hand-written notes—generated through a tablet PC or a pen/tablet device—are stored using Microsoft's Ink technology, which displays the hand-written text but also converts and stores it as if it were typed. One of the advantages of Ink is that you can write a list of items (say) and later re-order them and also number or bullet the items as you would with typed text. One Note also has a container for audio notes—so that they're stored in context. While taking meeting notes, you could record part of the discussion, which will be stored along with your manual notes.

The highly innovative One Note promises to be an absolutely must-have program. For Microsoft, its importance is that it's likely to kick start movement to Tablet PCs and develop a revenue stream that is currently dormant. At present it's not clear whether One Note will be available only with Office 2004 or will also be available unbundled.

One Note and the Outlook enhancements are sufficient 'bait' to lure the individual user into upgrading to Office 2004. For the organisation, the technical implications of Office 2004 deployment will need very careful consideration. To use Office



2004 in the way that Microsoft designed it for will require an organisation to rethink its information management strategy. It will need require a consistent use of XML tags to represent information, such as stock items, prices and sales returns in its back-end databases and the infrastructure to implement access through Web services. This is no small undertaking. Will it be a price worth paying to achieve the levels of business process automation that Microsoft promises will accrue from its XML-based strategy for Office 2004? It's going to take a lot of time—well beyond the release date—for the answers to questions such as this. ■

Microsoft's new One Note lets you position typed and penned text, graphics and audio anywhere within a page—just like you're working on paper



➔ Memory  
1GB at 1GHz

## NVIDIA GEFORCE FX

INTRO | CLOCK SPEED | TEXTURES AND VERTICES | COOLING | ANTI-ALIASING | DIRECTX 9 | EXTRAS

# Graphics giants go head-to-head

With a long-delayed launch, does nVidia's FX chipset still have what it takes to see off the established competition from ATI's Radeon 9700 Pro? *PC Magazine* examines the evidence...

Programmable pixel shaders like those in nVidia's GeForce FX will be used to create realistic lighting and modelling effects in games such as the long-awaited Doom III

## HIGHLIGHTS

- Clock speed hits a **staggering 500MHz**, with a lower-cost 400MHz option available
- The high-speed processor core means a **special ducted cooling system** takes the place of a second PCI slot
- nVidia claims that the GeForce FX **exceeds DirectX 9.0** in its pixel and vertex shader specifications
- **Eight texture units** give a peak rate of 4Gtexel0/s

**W**e've just got our hands on the first sample hardware of nVidia's long-awaited FX graphics product and, as you can see from our performance tests, it appears that nVidia's just been able to match ATI's Radeon 9700. Here we take an in-depth look at how the GeForce FX works.

With the GeForce FX—or NV30 as it has been known until now—it is clear that nVidia is planning to reclaim its position as the manufacturer of the fastest graphics chip in the market. This is a title currently held by ATI, who stole the show with last year's groundbreaking Radeon 9700 Pro. Taking a quick look over nVidia's releases, however, it's clear that the US graphics giant has pulled out all the

stops to deliver a chip that's unlikely to fall short of the coveted Number 1 slot.

Unlike ATI, who chose the safer and simpler 0.15-micron option with the Radeon 9700 Pro, nVidia is the first to apply a 0.13-micron die size to a graphics chip. Although this move was rumoured to have caused the production delays that stopped nVidia from launching the GeForce FX in November last year, it does carry several advantages, not least of which are reduced power dissipation and size, as well as increased speed.

It's a flip-chip design using a total of 125 million transistors—nearly twice that of the previous GeForce4 Ti—and the core clock speed of the GeForce FX chipset will be a staggering 500MHz (a more economical 400MHz version will also be avail-

## NVIDIA GEFORCE FX

INTRO | CLOCK SPEED | TEXTURES AND VERTICES | COOLING | ANTI-ALIASING | DIRECTX 9 | EXTRAS



▲ NVidia's 'Dawn' demo shows the flexibility of the vertex and pixel shaders found in the new GeForce FX chipset, providing a realistic range of facial expressions and movements

▼ The Time Machine demo demonstrates how effective a custom pixel shader can be by creating the effects of rust and ageing on the pickup truck's painted surface



able as part of the product range), compared to the Radeon 9700 Pro's 325MHz. It, too, is equipped with eight pixel pipelines, giving it an outstanding peak fill rate of 4Gpixel/s—that's 1.4Gpixel/s more than the Radeon 9700 Pro.

#### Clock speed

Raw processing power is worthless unless the rest of the graphics card can support the workload, so the memory clock will also be set to 500MHz. The GeForce FX is designed to address up to 1GB of DDR2 RAM, so this provides an effective memory speed of 1GHz. Like conventional DDR RAM, DDR2 RAM times data transfers to the rising and falling edges of the clock, but uses a 4-bit prefetch rather than 2-bit in these transfers as well as



a lower access time of as little as 1.8ns. To make best use of this feature, the GeForce FX has been specifically designed to handle these longer data bursts, so it benefits more from this new memory type than a Radeon 9700 Pro would, even though the latter is still capable of addressing DDR2.

Although a 1GB DDR2 memory load would be an attractive prospect in a high-end graphics card, the first products will be fitted with more affordable 128MB and 256MB payloads. The GeForce FX uses a similar memory interface to the earlier Lightspeed Memory Architecture (LMA) II of the GeForce4 Ti series. This means that it has a crossbar interface comprised of four 32-bit independent memory controllers providing a total bandwidth of 15GB/s ( $1\text{GHz} \times 128\text{bit/s} \div 8 \div 1,024 \div 1,024 \div 1,024$ ). Although this is impressive, it should be noted that the Radeon 9700 Pro uses a 256-bit interface and 310MHz DDR1 RAM, giving it a peak memory bandwidth of 18.5GB/s ( $620\text{MHz} \times 256\text{bit/s} \div 8 \div 1,024 \div 1,024 \div 1,024$ ), something that could well tip the balance in ATI's favour during memory-intensive operations.

However, in addition to LMA II's bandwidth optimisations, such as Z-data compression, fast Z-clears and visibility subroutines that cull 'invisible' objects before they clog up the renderer, some new routines have been added that include lossless compression of colour data, which comes in particularly handy when anti-aliasing is being applied (but more of this a little later).

#### Textures and vertices

If you're curious as to the texture handling of the GeForce FX, you should know that nVidia has followed a similar path to ATI and has only fitted eight texture units—one for each processing pipeline. To fit two per pipe would be extremely costly in terms of transistors, as well as increasing the possibility of bottlenecking at the memory interface. Even so, eight texture units still provides a peak rate of 4Gtexel/s, which is still far more than any other chipset we've seen and should be more than enough to deal with both current and upcoming software releases.

For any software that utilises the vertex engine of a graphics chip rather than the CPU, polygon processing power is absolutely vital. Previously, you could simply calculate the figure by examining the number of shader units in a chipset, but the programmability and efficiency of the GeForce FX floating-point array means that it's not as simple as that. According to nVidia, however, the GeForce FX is capable of processing 350 million polygons per second, with the Radeon 9700 Pro left trailing behind with 325Mpoly/s.

While the Radeon 9700 Pro was the first on the market to provide 128-bit floating point accuracy in its pipelines, any pixel shader operations it does are limited to 96 bits, or 24 bits for each channel. This is still impressive, but not quite to the same extent as the GeForce FX, which provides 128-bit control



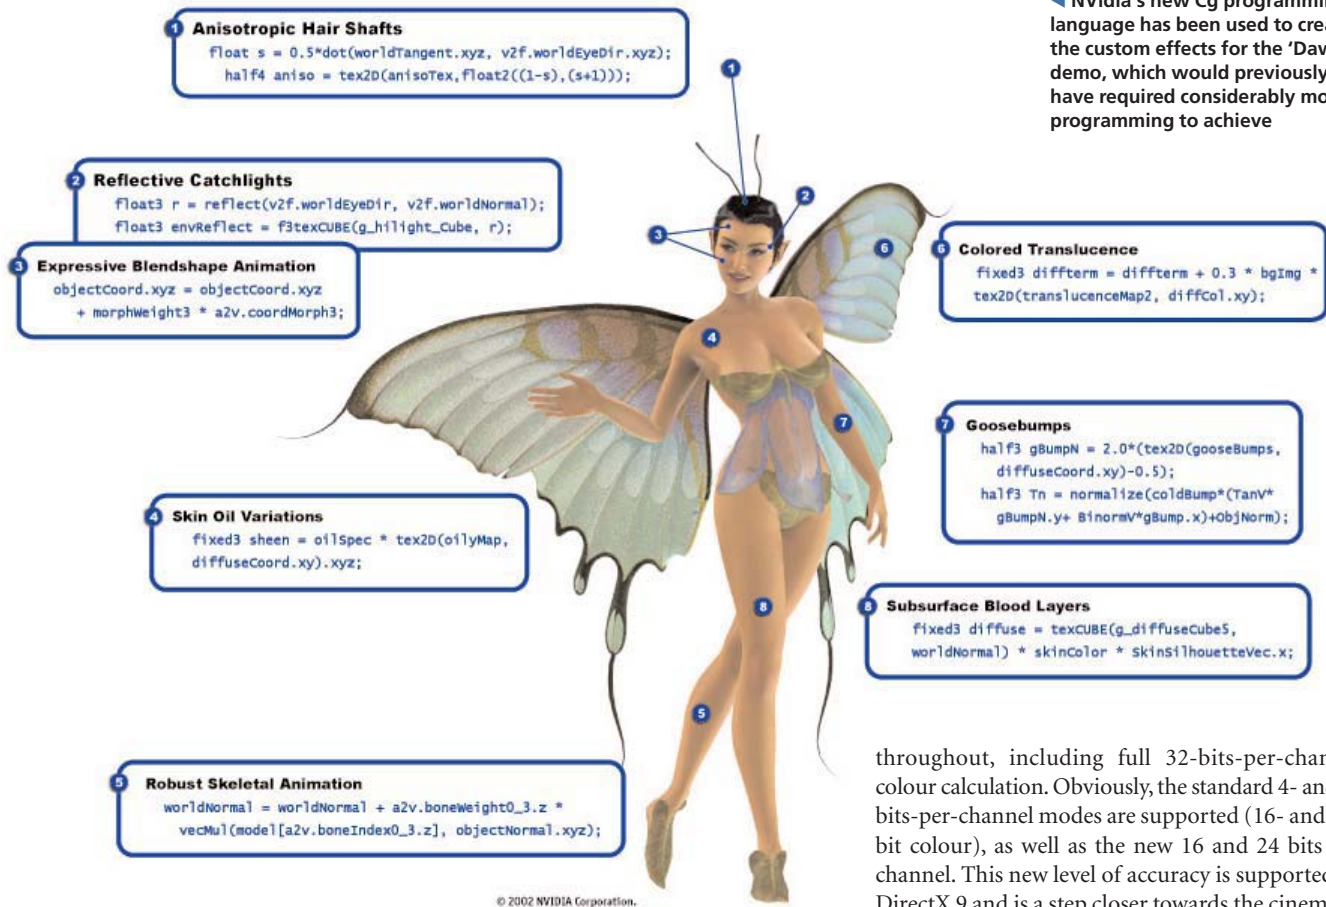
## DirectX 9

Exceeding the standard

INTRO | CLOCK SPEED | TEXTURES AND VERTICES | COOLING | ANTI-ALIASING | DIRECTX 9 | EXTRAS

## NVIDIA GEFORCE FX

◀ NVIDIA's new Cg programming language has been used to create the custom effects for the 'Dawn' demo, which would previously have required considerably more programming to achieve



throughout, including full 32-bits-per-channel colour calculation. Obviously, the standard 4- and 8-bits-per-channel modes are supported (16- and 32-bit colour), as well as the new 16 and 24 bits per channel. This new level of accuracy is supported by DirectX 9 and is a step closer towards the cinematic realism that nVidia is chasing, particularly when it comes to lighting and colour calculations.

## What is Cg?

### Nvidia's high-level programming language

One of the most common issues affecting the release of new graphics hardware is the lack of support for any new features in current or immediate software releases. NVIDIA's answer is Cg, a high-level programming language similar to C++, but with an emphasis on graphics instructions.

According to nVidia, Cg allows developers to create complex and customised routines, like the custom shaders that the GeForce FX is so good at, without the need for such complex coding before compiling. It works with both DirectX and OpenGL calls, so it's not hardware-specific and allows developers to work across platforms, such as consoles and PCs, without having to rewrite any of their custom shaders or effects. It's also the only way that they can access features that may be outside

the DirectX or OpenGL specification, such as Pixel and Vertex Shader 2.0+.

At present, Cg is a new tool and is currently only optimised for nVidia GPUs, but it has been made open source, so other hardware manufacturers could release their own Cg optimisations should they want to. It's already integrated into SoftImage XSI, with plug-ins for 3DSMax and Maya, so cinematic effects built on these platforms can be scaled down to suit the desktop market. Ironically it will take time for the developer community to fully adopt Cg (a move which may rely on the arrival of optimisations from other manufacturers) before we'll begin to see the proposed benefits of Cg. Until that point, we'll still have to wait for a year before we see software that utilises the hardware that's out today.

### Keeping cool

Despite the 0.13-micron die size, the 500MHz core speed of the GeForce FX chip is obviously pushing the silicon to its limits and the 125 million transistors need substantial cooling, which is why nVidia has introduced nFlow. It's a seriously large active cooling system that piggybacks the chipset and occupies the space normally reserved for your PCI Slot 1. It consists of a copper heatsink, a U-shaped air duct with intake and exhaust ports at the rear, and there's also a fast fan and passive cooling fins on the underside of the card.

The obvious benefit of this design is that it keeps the chipset within operating temperatures without increasing the heat inside your PC. However, early test samples have been criticised for their noise when running at full throttle and you might not want to lose a PCI slot for the sake of a cooling system. It should be pointed out that manufacturers are completely free to apply whatever cooling system they want, as long as it meets nVidia's specification, so we may see quieter alternatives.

### Anti-aliasing


It wouldn't be a full-blooded new graphics chipset release without a new spin on the must-have feature of all current graphics cards—full-scene anti-aliasing (FSAA). In addition to 2X, 4X, 6X and 6XS, the GeForce FX offers an 8X FSAA mode thanks to the

## NVIDIA GEFORCE FX

INTRO | CLOCK SPEED | TEXTURES AND VERTICES | COOLING | ANTI-ALIASING | DIRECTX 9 | EXTRAS

## A developer's view

John Carmack, the technical director of ID software, has discussed the two graphics chip heavyweights in his recent plan outline ([www.webdog.org/plans/1/](http://www.webdog.org/plans/1/)). Using performance on the eagerly awaited Doom III game as his yardstick, he appears to place the GeForce FX ahead of the Radeon 9700 Pro, but adds that it's a complicated issue 'because of the different ways the cards can choose to run the game'. He, too, mentioned the noise of the nFlow cooling system on his sample board.



◀ The GeForce FX has nearly twice the number of transistors as the GeForce4 Ti series and uses the increasingly common FlipChip package

chip's new colour compression optimisation. Where pixels in the scene have allocated colour data (rather than calculated colour data), this data is compressed to a quarter of the original size both losslessly and in real time, similarly to the Z-buffer data compression. This doesn't apply to the anti-aliasing calculations that occur at the edges of polygons within the scene, but does mean that there's a greater amount of resources available to perform the anti-aliasing calculations, otherwise an 8X sample would be totally impractical.

### Beyond DirectX 9.0

When ATI released the Radeon 9700 Pro, it proudly proclaimed compliance with all DirectX 9.0 graphics features—even though the final release of the API didn't happen for another four months. Unsurprisingly, nVidia has gone a step further, announcing that the GeForce FX exceeds the DirectX 9.0 graphics specification in several ways, notably in its vertex and pixel shaders—the heart of programmable graphics technology. By labelling its pixel and vertex

shader versions 2.0+, nVidia is keen to underline the increased capabilities that the GeForce FX has over the DirectX 9.0 specification.

As far as Vertex Shader 2.0+ is concerned, this means that developers can build more complex light routines, skeletal animations, palette skins and full-flow control, including branching and looping for repeated or interruptible subroutines. nVidia also claims that Pixel Shader 2.0+ will allow up to 16 texture maps in a single pixel shader routine with up to 1,024 texture fetch instructions. To put it in simpler terms, this allows for custom effects to be built that include soft shadows, depth of field, blurs, cartoon-like finished and haloing—as well as more complex texture sampling routines, such as bicubic filtering. Whether developers choose to take advantage of these proprietary abilities and what the results will bring to future software releases remains to be seen. A lot depends on the success of Cg (see our sidebar).

### Conclusion

It's obvious that the GeForce FX is going to be one of the fastest graphics card in circulation when it's released in the UK with fully developed drivers, but it's unlikely to cause the same stir as its major rival because the Radeon 9700 Pro got there first with eight pipelines, DirectX 9.0 support and features like adaptive filtering and tessellation. Even though the advanced Pixel and Vertex shaders are impressive, developers generally go for the mainstream option—unless they're paid to do otherwise—meaning that any custom effects using pixel or vertex shaders are likely to fall within the DirectX 9.0 specification, rather than the Pixel and Vertex Shader 2.0+ specification of the GeForce FX.

So, as always, it comes down to an issue of raw performance and it's clear from this chipset's specification that it's got what it takes to keep ATI at bay for a while at least. Assuming that you're prepared to sacrifice a PCI slot and perhaps put up with some extra fan noise, the GeForce FX looks like the graphics chip to watch in 2003. ■

### PC Magazine Lab tests

**3D Mark 2001**  
(1,024 x 768 32-bit Triple Buffering)

#### nVidia GeForce FX

Driver Version 6.14.01.4263

14,194

#### Sapphire 9700 Pro

Driver Version Catalyst 3.0

14,461

The nVidia GeForce FX we tested was an early example of the board. It also used an early example of the drivers, whereas the ATI card was tested with a more mature version of ATI's Catalyst drivers



## USER TO USER

KELVYN TAYLOR

# Creating PDFs using GSView

Adobe's Portable Document Format (PDF) is one of the most popular cross-platform file formats today. This walkthrough was submitted by *PC Magazine* reader Steve Hurley and shows you how to create PDF files by using simple freeware tools

### KEYPOINTS

- What you'll need
- Installing the printer driver
- Installing the software
- Converting documents to PDFs

Using Ghostgum's free GSview software you can convert your standard Microsoft Word, plain text, or Star Writer files into Adobe PDF files. The advantage of this is that Adobe PDF files can be read by any system using the free multi-platform software Adobe Acrobat Reader—just as you're viewing this issue of *PC Magazine*. PDF files can be uploaded to the Internet and can be downloaded to and read by any system as long as Adobe's free Acrobat Reader is installed. Also, in the case of many files (such as Word), the equivalent PDF files are much smaller than the original ones. PDF files can contain graphics, tables, photos and text.

### What you'll need

Besides the Acrobat Reader (downloadable from [www.adobe.com](http://www.adobe.com)) you'll need:

- 1 A driver for any printer that supports the Postscript printer standard. Note that you don't need to have the actual printer—just the printer driver. In this example we've chosen the HP Laserjet 4/4M Postscript driver supplied with Windows 98, but you can use almost any printer that supports Postscript. You'll find the drivers on your Windows installation CD-ROM if they aren't already installed. You can also find the HP Laserjet drivers we used at [www.hp.com](http://www.hp.com) if you can't locate them for any reason.
- 2 GSview 4.3 (<ftp://mirror.cs.wisc.edu/pub/mirrors/ghost/ghostgum/gsv43w32.exe>)
- 3 Ghostscript 7.04 (<http://belnet.dl.sourceforge.net/sourceforge/ghostscript/gst704w32.exe>)

For the purposes of this walkthrough, you'll need to create a folder in your Adobe folder called MyPDF and then open a sub-folder

called Downloads. Download the two programs listed above to this sub-folder.

### How it's done

There will be slight variations in the installation procedures depending on your version of Windows. We've used Windows 98 SE.

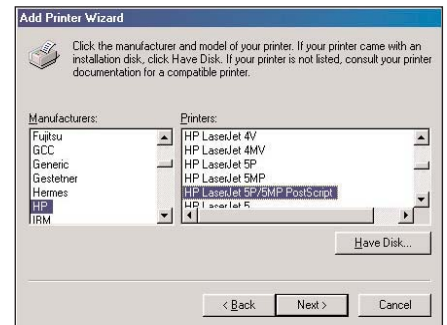
## STEP 1: Installing the printer driver

Click Start/Settings/Printers/Add Printers to launch the Add Printer wizard (see Figure 1). Click Next and choose the HP Laserjet 4L/4ML (or another Postscript printer if this isn't present) from the list of printers. Select the correct printer port (normally LPT1 for a parallel connection—see Figure 2) and a dialogue box will open, which asks if you want to make this your default printer, check No. Then choose Local printer, click on Next and choose No to the 'Do you want to print a test page?' question. Then click on Finish. If prompted for the Windows CD-ROM, insert the CD-ROM and click OK. Type the path to your CD-ROM drive and click OK again.

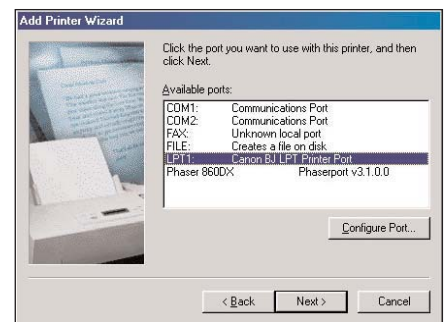
The printer driver will be installed and the printer installation will continue. During installation you may be asked whether you want to share your printer—you should select the 'Do not share this printer' button. On successful completion of the wizard, the new printer icon will appear in the Printers folder.

## STEP 2: Installing the software

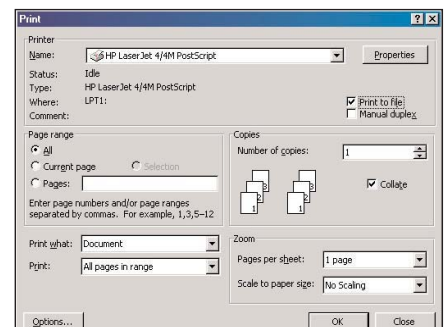
Install Ghostscript and GSView from the downloaded self-extracting executable files GS704W32.EXE and GSV43W32.EXE. Note that Ghostscript (GS704W32.EXE) must be installed before GSView (GSV43W32.EXE).



1 First you'll need to launch the Add Printers wizard from the start menu

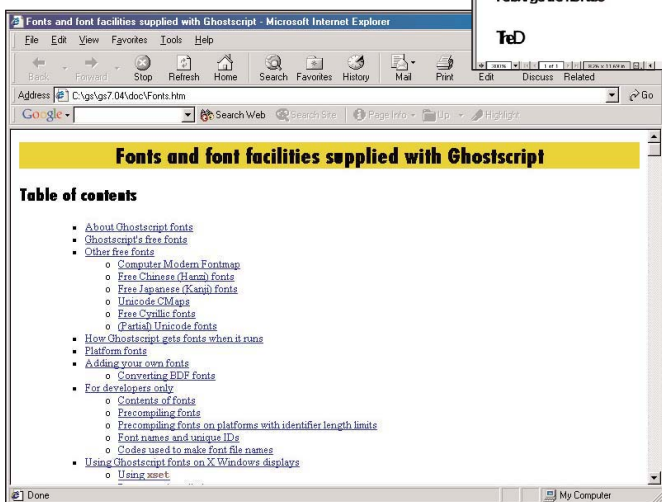
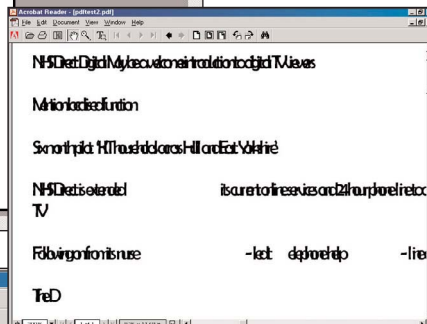
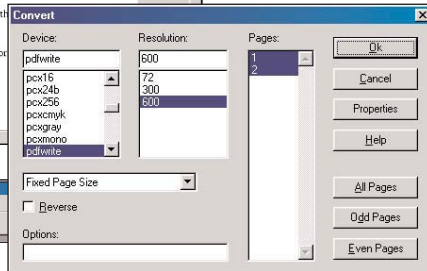
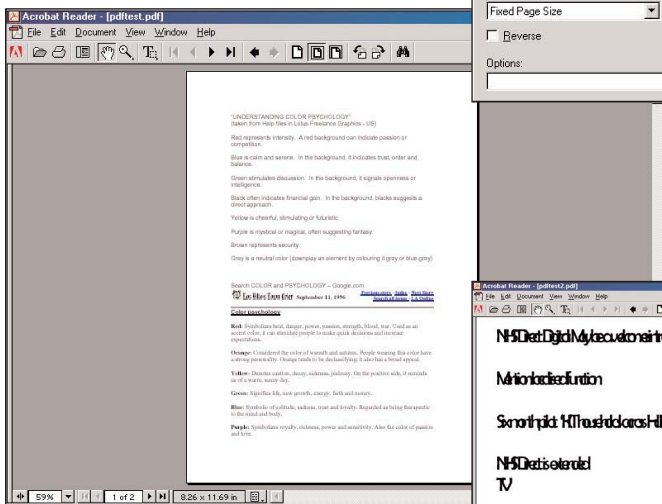
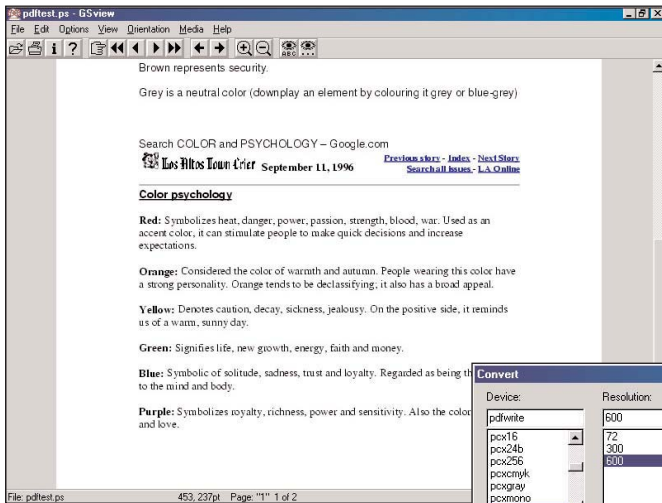


2 Having chosen a Postscript printer, select the correct printer port—usually LPT1



3 When you've chosen your document, use the print selection box to Print to File





Follow the installation commands but leave the options boxes alone until you've had time to read the help files. You need make no changes to the default installation to be able to convert your documents. However, there may be some problems with fonts .

### Step 3: Converting your documents to PDF format

Open the document you want to convert. Choose File | Print. In the printer selection box, choose the printer you've just installed (See Figure 3). Check the box labelled Print to File. Then click OK, which will bring up a Save dialog box. Navigate to the MyPDF directory then key in a name for the file, using the extension .PS. This saves the document in Postscript format. In the lower drop-down list choose All Files and then click Save. Now close your original document. Part one of the conversion is now complete.

Double-click on the .PS file you've just created and the program GSView should open the file—if not, launch GSView and open the file manually using File | Open from the menu bar (see Figure 4). When you've opened the Postscript file, choose File | Convert from the menu bar (see Figure 5). Scroll down the left-hand Device window and choose PDFWRITE. In the Resolution window choose your preferred resolution—the default is 600dpi. Choose Fixed Page Size from the drop-down box (see Figure 6). Click on All Pages then click OK. In the resulting dialog box, save your document with the extension .PDF.

You've now converted your original document to a PDF file. To check that the conversion's been successful just launch the new file using Acrobat Reader (see Figure 6).

#### Possible font problems

If, after converting your .PS file to a PDF file and re-opening it, there appear to be font corruption, overlapping words or other formatting problems (see Figure 7), they're probably caused by a non-compatible font. You won't see this in GSView as it uses your system fonts. To overcome this, create your initial document in a standard Windows font, such as Times New Roman, Ariel or Courier. GSView has a wide variety of free Postscript fonts, but as there are thousands available, problems can occur. One way to check many fonts at the same time for compatibility is to format every two lines of a text document in a different font. Convert the document into a .PDF file, open it in Acrobat Reader and check the results. For a more in-depth explanation of the font issue read the file Fonts.htm (shown in Figure 8), which can be found in the DOC subdirectory of the Ghostscript install directory (by default \gs\gs7.04).

#### SOURCES FOR PROGRAMS

Adobe Acrobat Reader  
www.adobe.com

Ghostscript 7.04  
http://belnet.dl.sourceforge.net/sourceforge/ghostscript/gsv704w32.exe

GSView 4.3  
ftp://mirror.cs.wisc.edu/pub/mirrors/ghost/ghostgum/gsv43w32.exe





**TUTOR**  
GEOFF EINON

# Moving from Excel to Access

**Excel's fine, as far as it goes, but it does have limitations when it comes to managing your data. Access offers more power and specialised tools, but how do you transfer all that data?**

**KEYPOINTS**

- Normalisation
- Eliminate repeated columns
- Split data into separate tables
- Create related tables

Excel, although it is primarily a spreadsheet, is also designed to double as a database and works very well for storing routine business data—up to a point, that is. Inevitably there comes a time when you need the power and more specialised tools of Access to help you out. This month we concentrate on a hands-on experience in converting flat-file Excel data into related tables that you can import into Access to take advantage of its relational capabilities. A follow-up article next month will show you how to import those tables into Access and create a working Access application.

Open the Excel file HOMar03.XLS, which you'll find on the *PC Magazine* Web site ([www.pcmag.co.uk/pcm/solutions.jsp](http://www.pcmag.co.uk/pcm/solutions.jsp)). Look at the SourceData worksheet—shown in **Figure 1** opposite. This contains quarterly sales data from five sales people who are located in four different regions. In constructing this Excel database, the person who designed it first created a staff contact list and then went on to add columns at the far right to hold quarterly sales information related to each individual. At the end of each quarter a new column is added to the table, as is the individuals' sales totals for that quarter and new individual and quarter totals calculated.

The data structure shown in **Figure 1** looks a logical way of recording this information—but what's wrong with it? One answer is that it's difficult to work efficiently with this data structure, which contains three different information categories—staff, regions and sales—in the same table with the sales data located in three separate columns. In the terms of database theory and language, this data table is 'de-normalised'. If you're familiar with Excel's PivotTable analysis, try to create a pivot table from the SourceData data. You'll find that the data structure precludes the cre-

1	LastName	FirstName	DOB	Region	StartDate	Phone	Address1	Address2	Q102Sale	Q202Sale	Q302Sale	Total
2	Smith	Mike	21/07/1984	London	01/11/2000	12345	Tower Bul	Grange st	100	75	88	26
3	Jones	Barbara	23/08/1969	Manchest	05/03/2001	24567	Archway F	Wollfender	78	129	45	25
4	Lewis	John	30/07/1962	Cardiff	16/09/1999	67890	Prestige F	Allan St	136	108	67	31
5	Evans	Donna	12/04/1976	Edinburgh	04/03/2001	19877	Fairmile T	Grampian	99	86	78	26
6	Jenkins	Megan	23/01/1980	Cardiff	07/09/2000	67890	Prestige F	Allan St	67	99	120	28
7									480	497	398	
8												
9												
10												

**1** The SourceData table that is to be normalised

1	LastName	FirstName	DOB	Region	StartDate	Phone	Address1	Address2	Qtr1	Qtr2	Qtr3	Sales
2	Smith	Mike	21/07/1984	London	01/11/2000	12345	Tower Bul	Grange st	021			100
3	Smith	Mike	21/07/1984	London	01/11/2000	12345	Tower Bul	Grange st	022			75
4	Smith	Mike	21/07/1984	London	01/11/2000	12345	Tower Bul	Grange st	023			88
5												
6												
7												

**2** An intermediate stage in the process of eliminating repeating columns of data

ation of a meaningful pivot table—for example, of sales by person by quarter.

Our task in this issue is to normalise this data. The task for the next issue will be to import into Access the normalised data and create a working application that will let us add new data and report on it.

**Normalisation**

To normalise the table we have to do three things. First of these is to eliminate repeating columns of data. Second is to split the data into separate tables with one table for each significant data category. In the case of the SourceData table this means that, we'll create three tables from it containing information about staff, sales totals and regions. The third task will be to create relationships between the tables. When this is accomplished, in database-speak, the data will be in 'third normal form'. As you'll see, this will allow you to work much more efficiently with the data.

**Deleting repeating data columns**

In the DataSource worksheet table (Figure 1) we need to rework columns containing similar data I, J and K and replace them with two columns: one specifying the Quarter (021, 022 or 023) and the second the transaction amount. At this stage, we'll forget the Totals information in column L.

To start the elimination process we'll use the blank Elimination worksheet. From the SourceData worksheet, copy the range A1:K2 and paste it in to the same position in the Elimination worksheet. Now that you've got the row of data labels as well as the Smith information, our immediate goal is to record each of Smith's quarterly sales totals in a separate row. To do this, in the Elimination worksheet, first copy the range A2:H2 and first paste it in to A3 and then into A4. This creates three identical sets of information about Smith and his Region address.

	A	B	C	D	E	F	G	H	I	J
1	LastName	FirstName	DOB	Region	StartDate	Extn	Address1	Address2	Qtrr	Sales
2	Smith	Mike	21/07/1984	London	01/11/2000	12345	Tower Buil Grange st	Q021		100
3	Smith	Mike	21/07/1984	London	01/11/2000	12346	Tower Buil Grange st	Q022		75
4	Smith	Mike	21/07/1984	London	01/11/2000	12347	Tower Buil Grange st	Q023		88
5	Jones	Barbara	23/08/1969	Manchester	05/03/2001	24567	Archway F-Wolfenden	Q021		78
6	Jones	Barbara	23/08/1969	Manchester	05/03/2001	24568	Archway F-Wolfenden	Q022		129
7	Jones	Barbara	23/08/1969	Manchester	05/03/2001	24569	Archway F-Wolfenden	Q023		45
8	Lewis	John	30/07/1962	Cardiff	16/09/1999	67895	Prestige H AllAn St	Q021		136
9	Lewis	John	30/07/1962	Cardiff	16/09/1999	67896	Prestige H AllAn St	Q022		108
10	Lewis	John	30/07/1962	Cardiff	16/09/1999	67897	Prestige H AllAn St	Q023		67
11	Evans	Donna	12/04/1976	Edinburgh	04/03/2001	19877	Fairmile Tc Grampian	Q021		99
12	Evans	Donna	12/04/1976	Edinburgh	04/03/2001	19878	Fairmile Tc Grampian	Q022		86
13	Evans	Donna	12/04/1976	Edinburgh	04/03/2001	19879	Fairmile Tc Grampian	Q023		78
14	Jenkins	Megan	23/01/1980	Cardiff	07/09/2000	67890	Prestige H AllAn St	Q021		67
15	Jenkins	Megan	23/01/1980	Cardiff	07/09/2000	67890	Prestige H AllAn St	Q022		99
16	Jenkins	Megan	23/01/1980	Cardiff	07/09/2000	67890	Prestige H AllAn St	Q023		120

3 The Eliminated worksheet contains the results of the Elimination process

1	LastName	FirstNa	DOB	Region	StartDate	Extn	Address1	Address2	Qtrr	Sales
2	Smith	Mike	21/07/1984	London	01/11/2000	12345	Tower Buil Grange st	Q021		100
3	Smith	Mike	21/07/1984	London	01/11/2000	12346	Tower Buil Grange st	Q022		75
4	Smith	Mike	21/07/1984	London	01/11/2000	12347	Tower Buil Grange st	Q023		88
5	Jones	Barbara	23/08/1969	Manche	05/03/2001	24567	Archway F-Wolfenden	Q021		78
6	Jones	Barbara	23/08/1969	Manche	05/03/2001	24568	Archway F-Wolfenden	Q022		129
7	Jones	Barbara	23/08/1969	Manche	05/03/2001	24569	Archway F-Wolfenden	Q023		45
8	Lewis	John	30/07/1962	Cardiff	16/09/1999	67895	Prestige H AllAn St	Q021		136
9	Lewis	John	30/07/1962	Cardiff	16/09/1999	67896	Prestige H AllAn St	Q022		108
10	Lewis	John	30/07/1962	Cardiff	16/09/1999	67897	Prestige H AllAn St	Q023		67
11	Evans	Donna	12/04/1976	Edinbur	04/03/2001	19877	Fairmile Tc Grampian	Q021		99
12	Evans	Donna	12/04/1976	Edinbur	04/03/2001	19878	Fairmile Tc Grampian	Q022		86
13	Evans	Donna	12/04/1976	Edinbur	04/03/2001	19879	Fairmile Tc Grampian	Q023		78
14	Jenkins	Megan	23/01/1980	Cardiff	07/09/2000	67890	Prestige H AllAn St	Q021		67
15	Jenkins	Megan	23/01/1980	Cardiff	07/09/2000	67890	Prestige H AllAn St	Q022		99
16	Jenkins	Megan	23/01/1980	Cardiff	07/09/2000	67890	Prestige H AllAn St	Q023		120

19	Sum of Sales	Qtrr		
20	LastName	Q021	Q022	Q023
21	Evans	99	86	78
22	Jenkins	67	99	120
23	Jones	78	129	45
24	Lewis	136	108	67
25	Smith	100	75	88
26	Grand Total	480	497	398
				1375

4 The Eliminated data structure is suitable for PivotTable analysis

You need to transpose Smith's sales totals figures so that they're held in cells J2 to J4. So, in cell J2 type '100'; in J3 type '75' and in J4 type '88'. To identify the quarter in which the sales occurred, in cell I2 type the label '021' (year 2002, quarter 1); in cell I3 type the label '022' and in cell I4 type the label '023'. Don't forget that to insert these numeric codes as labels you should precede each with the single quote character to indicate to Excel that it's a label and not a value. Finally, in cell I1 type the label 'Qtrr' and in cell J1 'Sales', and the Elimination worksheet should look like **Figure 2**.

For the SourceData data the basic process of elimination entails creating three records

for each individual and populating each with a sales total for a specified quarter. To complete the elimination process, you should now, one at a time, copy the rows for Jones, Lewis, Evans and Jenkins into the Elimination worksheet and process each as we have just done for Smith—so that there are three quarterly sales records for each individual. If you don't want to take on this task, the Eliminated worksheet (see **Figure 3** above) contains this data structure. You can see the benefits of this re-structuring of the data, which is now in what is described as 'first normal form', by carrying out a PivotTable analysis—such as that shown in **Figure 4** above. Looking at Figure

4, you can see that a 'problem' associated with this data structure is that every new record requires the entry of a lot of information. The way that this information is handled in Access is to split it in to separate tables—which we'll do a little later. First, though, we have to ensure that relationships are in place so that related data in the separate tables can be linked. To do this we set up 'keys' that uniquely identify items in the Staff and Region columns and also each Sales item (rows) in the Eliminated worksheet.

### Creating relationships

The next stage in this process of structuring the data is to uniquely identify the items within the data categories in which we are interested—people, regions and transactions. To do this, insert a new worksheet and name it 'Relations'. Copy the entire data set from the Eliminated worksheet into Relations. Start by identifying the regional information. To do this, in the Relations worksheet, select column D (Regions) and insert a column before it (Insert | Columns). Label this column-RegionNo. Work your way down the column to insert a number corresponding to each of the regions. Identify London as '1'; Manchester as '2'; Cardiff as '3' and Edinburgh as '4'—as shown in **Figure 5**.

Next, we'll uniquely identify each of the Sales people. To do this, insert a new column before column A and label it StaffNo. In this new column B—as shown in **Figure 3**—for each entry for Smith enter the numeral '1'; for each Jones enter '2'; for Lewis enter '3'; for Evans enter 4; and for Jenkins enter '5'.

To uniquely identify transactions, in the Relations worksheet, add a column before the new column A. To do this select column A and label it 'SaleNo' and beneath it number each transaction sequentially starting at '1'. You should have 15 entries as shown in **Figure 5**.

### Creating the separate data tables

The next stage in restructuring the data for use in Access is to transform the flat-file data structure we've worked on so far into three separate tables with Sales, Staff and Region data. To do this, insert three new worksheets after the Relations worksheet. Label these: Staff, Sales and Regions. From the Relations worksheet, copy the range A1:M16 (the complete data set and column labels) and paste it in to each of the three new worksheets.

Starting from the right-hand side in the Staff worksheet, delete columns M to J inclusive, and columns G and column A. Finally, delete duplicate rows so that you're left with one entry for each of the five people. When you've done this the Staff sheet will look like **Figure 6**. Notice that the table includes the RegionNo column, which uniquely specifies the person's region. By using this relational



reference we've managed to avoid the duplication of the Region address data that was created in the Elimination table.

Starting from the right-hand side in the Sales worksheet, delete columns K to C inclusive. When you've done this, the Transactions worksheet will look like **Figure 7** below. Note that each row in this table includes, the StaffNo code that identifies the sales person responsible for the sales

Starting from the right-hand side in the Regions worksheet, delete columns M to L inclusive, I to H inclusive and E to A inclusive. The resulting data set will look like **Figure 8** below. Note that there's no relational information in this table.

Looking at the three tables, one advantage is clear—adding a new sales total or region or changing the region in which a sales person works is a simpler task involving changes in single tables. From your experience with Excel it's also clear is that this relational approach isn't suited for the Excel environment which doesn't provide the tools for working with normalised data in multiple tables in a coherent way. In the next article we'll import these three tables into Access and create an application to add new records and create reports to explore the way that it's specialised to work with relational data. ■

	A	B	C	D	E	F	G	H	I	J	K	L	M
1	SaleNo	StaffNo	LastNam	FirstNam	DOB	RegionNo	Region	StartDate	Extn	Address	Address2	Qtrr	Sales
2	1	1	Smith	Mike	21/07/1984	1	London	01/11/2000	12345	Tower Bt	Grange st	Q021	100
3	2	1	Smith	Mike	21/07/1984	1	London	01/11/2000	12347	Tower Bt	Grange st	Q023	88
4	3	1	Smith	Mike	21/07/1984	1	London	01/11/2000	12347	Tower Bt	Grange st	Q023	88
5	4	2	Jones	Barbara	23/08/1969	2	Mancheste	05/03/2001	24567	Archway	Wolfender	Q021	78
6	5	2	Jones	Barbara	23/08/1969	2	Mancheste	05/03/2001	24568	Archway	Wolfender	Q022	129
7	6	2	Jones	Barbara	23/08/1969	2	Mancheste	05/03/2001	24569	Archway	Wolfender	Q023	45
8	7	3	Lewis	John	30/07/1962	3	Cardiff	16/09/1999	67895	Prestige	Allan St	Q021	136
9	8	3	Lewis	John	30/07/1962	3	Cardiff	16/09/1999	67896	Prestige	Allan St	Q022	108
10	9	3	Lewis	John	30/07/1962	3	Cardiff	16/09/1999	67897	Prestige	Allan St	Q023	67
11	10	4	Evans	Donna	12/04/1976	4	Edinburgh	04/03/2001	19877	Fairmile	Grampian	Q021	99
12	11	4	Evans	Donna	12/04/1976	4	Edinburgh	04/03/2001	19878	Fairmile	Grampian	Q022	86
13	12	4	Evans	Donna	12/04/1976	4	Edinburgh	04/03/2001	19879	Fairmile	Grampian	Q023	78
14	13	5	Jenkins	Megan	23/01/1980	3	Cardiff	07/09/2000	67890	Prestige	Allan St	Q021	67
15	14	5	Jenkins	Megan	23/01/1980	3	Cardiff	07/09/2000	67890	Prestige	Allan St	Q022	99
16	15	5	Jenkins	Megan	23/01/1980	3	Cardiff	07/09/2000	67890	Prestige	Allan St	Q023	120

	A	B	C	D	E	F	G
1	StaffNo	LastName	FirstName	DOB	RegionNo	StartDate	Extn
2	1	Smith	Mike	21/07/1984	1	01/11/2000	12345
3	2	Jones	Barbara	23/08/1969	2	05/03/2001	24569
4	3	Lewis	John	30/07/1962	3	16/09/1999	67895
5	4	Evans	Donna	12/04/1976	4	04/03/2001	19879
6	5	Jenkins	Megan	23/01/1980	3	07/09/2000	67890

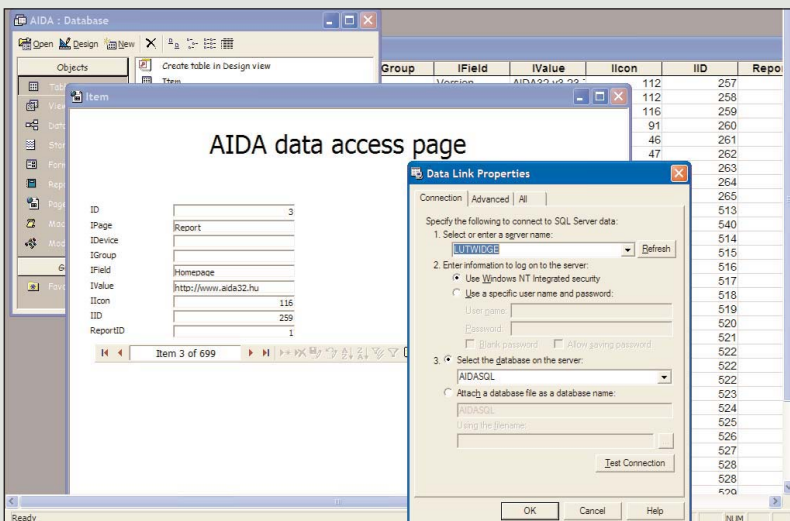
- 5 Each instance of the Sales, Staff and Region categories is uniquely identified.
- 6 The Staff table

## What is Access?

MICROSOFT'S ACCESS 2002 DATABASE is a powerful relational database management application that ships with the Professional and Developer versions of Office. It can also be bought as a standalone product, although there's little point as it costs approximately

the same as the full Office suite.

With Access you can either create standalone database applications, using data from direct input or external sources such as Excel, or even use Access to connect to an enterprise system based on SQL Server.



▲ Access is a powerful and flexible database management tool, which is capable of interfacing with enterprise-level SQL Server databases

	A	B	C	D	E
1	SaleNo	StaffNo	Qtrr	Sales	
2	1	1	Q021	100	
3	2	1	Q022	75	
4	3	1	Q023	88	
5	4	2	Q021	78	
6	5	2	Q022	129	
7	6	2	Q023	45	
8	7	3	Q021	136	
9	8	3	Q022	108	
10	9	3	Q023	67	
11	10	4	Q021	99	
12	11	4	Q022	86	
13	12	4	Q023	78	
14	13	5	Q021	67	
15	14	5	Q022	99	
16	15	5	Q023	120	

7 The Sales table

	A	B	C	D
1	RegionNo	Region	Address1	Address2
2	1	London	Tower Buildings	Grange st
3	2	Manchester	Archway House	Wolfenden Place
4	3	Cardiff	Prestige House	Allan St
5	4	Edinburgh	Fairmile Tower	Grampian Road

8 The Regions table



## MASTERCLASS

ALEX CRUICKSHANK

# Content management systems

In this last instalment of the Perl/CGI content management system series that we've spent eight months developing, it's finally time for your Web site to go live

### KEYPOINTS

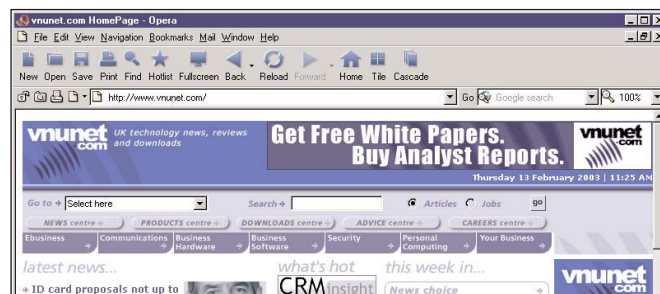
- What you need
- Converting existing HTML
- Going live
- Further information

Content management systems form the backbone of every major site on the Web. The benefit of having such a system is to separate content from design, so that anyone—even those with limited or non-existent knowledge of HTML—can update the content of the site and add new pages.

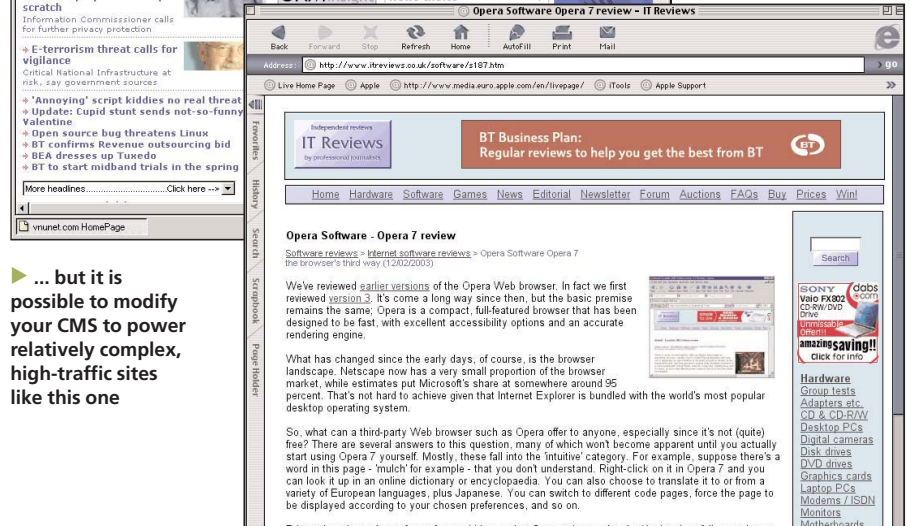
The front-end of our CMS gives a user-friendly screen where users can create and edit page text, add links to other pages, define the publication date and page category and so on. The CMS then takes this information, adds HTML formatting, embeds the result in an HTML template and serves it on the site.

The CMS we've been developing over the last eight months, which is now complete, is included in our file downloads available on our Web site at [www.pcmag.co.uk/pcm/solutions.jsp](http://www.pcmag.co.uk/pcm/solutions.jsp) along with all the necessary support files, as described over the previous months' issues. The CMS is free and open source, so if it doesn't entirely satisfy your requirements, you're welcome to modify the code as necessary. Nothing has changed since the file from the last issue. So if you've already installed those CMS files you can continue to use them. In fact, we won't be adding to the core CMS at all this month. Instead, we'll be looking at a complementary Perl script called `convert.pl`, which is also included on the Solutions download area this month.

As we've mentioned over the last few issues of the print version of the magazine and the e-book, creating the CMS is only part of the solution to the problem of managing online content. We've also had to deal with security issues and the problems associated with mapping old, static URLs onto the new dynamic ones. There's one further problem that we haven't yet considered, though: how are you going to transfer the contents of your existing HTML pages into the new CMS format?



◀ The CMS that powers PC Magazine's parent company's site is a little more powerful than the one we've described here, but works on similar principles...



▶ ... but it is possible to modify your CMS to work relatively complex, high-traffic sites like this one

For very small Web sites, this isn't likely to be a major problem. After all, it's not going to take you long to cut and paste a few paragraphs of existing text into the CMS editor screen and then put them live. But for larger sites, any transition is likely to be time-consuming. Although it's not possible to provide a script to seamlessly convert all the HTML pages on your site to the format supported by the CMS, the script we're looking at here will at least help in that potentially arduous process. With a little tweaking, it should be possible to modify it to suit your

site. Even in its basic form, it should at least reduce the amount of time you'll need to spend cutting, pasting and editing.

If your site is already run via a different content management system and you're planning to switch to this one for reasons of cost, simplicity or its open-source nature, you'll have to take one additional step before you can begin the conversion process—you'll have to export the dynamic CMS pages as static HTML. Some CMS systems will do this for you, otherwise there are utilities available on the Internet that will simply save user-



## What you need

THE MOST IMPORTANT REQUIREMENT of the content management system is a Perl interpreter. If your server's running on a Unix box (including Linux) you'll find it's almost impossible to install Apache or Zeus without also installing Perl. The same isn't true of Windows servers running IIS, but a free port of Perl for 32-bit Windows machines is available from [www.activestate.com](http://www.activestate.com). This comes with documentation and is easy to install. One problem is that file paths may be different compared with a Unix box. If in doubt, use absolute paths (for example, /site/www/data rather than ../data)

Perl scripts need to be set to 'executable' in order to run. Using your FTP client, you should be able to right-click on a script once you've uploaded it to your server and select 'File properties' or 'File attributes' or 'chmod'. Then you'll be asked to define the file's attributes. The shortcut for this is to type '755' into the 'chmod/manual' box if your FTP client allows it. If not, you need to set the permissions as follows:

owner: read, write and execute  
group: read and execute  
public: read and execute

Generally this is only necessary with Unix servers. With Windows servers you probably won't need to change any permissions. But you may have to tweak IIS to ensure that .pl and .cgi files in the cgi-bin or cgi-local folder are executable.

The scripts also use some additional Perl modules that may or may not be included as part of your standard configuration. The two main ones are 'Time::Local' and 'Image::Size'. If these aren't already installed on your server, you can download them from [www.cpan.org](http://www.cpan.org). See your Perl documentation for notes on installing new modules. Usually it requires nothing more than a single 'install' command.

### 1 Configuring variables

```
# the following describes which directories to search (plus root),
# separated by hashes, leading '/'. Others will be ignored.
$indirpaths = "/#/products#/news";
#
$base = "wwwroot";
# the root from which to search for files
#
$outputdir = "wwwtext";
# where to write new files (structure must be same as above -
# sub-directories are NOT automatically created)
```

### 2 Preparing files for processing

```
] @indir = split(/#/, $indirpaths);

opendir(ROOT, $path);
@files = readdir(ROOT);
closedir(ROOT);

foreach (@files) {
    next if /^\.|\.\.$/;
    if ($_ =~ m/_vti_cnf/i) { next; }
    $fullfilename = "$path/$_";

    $shortfilename = "$path/$_";
    $shortfilename =~ s/^$base///i;
    $shortfilename =~ s/\.htm$/;

    if (-d $fullfilename) {
        $tmp = $fullfilename;
        $tmp =~ s/$base//;
        foreach $tpath (@indir) {
            $ok2dosub = 0;
            if ($tmp eq $tpath) { $ok2dosub = 1; }
            if ($ok2dosub > 0) { checkFiles($fullfilename); }
        }
        next;
    }

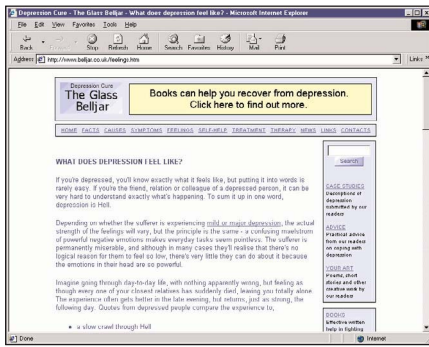
    # change following extensions if necessary
    if (($fullfilename =~ m/\.htm$/i) || ($fullfilename =~
m/\.html$/i)) {
```

specified dynamic pages as HTML. A quick search should reveal a selection of them.

Assuming you now have a collection of HTML pages that you want to convert, we'll now go through the code in convert.pl, explaining how each section works. We'll cover this in some detail, because of all the scripts included in the CMS package, this is the one that you're most likely to want to modify to suit your particular purposes. Note that the convert.pl script isn't designed to be run from a Web server or accessed via a browser. Instead, it should be run locally. If you have Perl installed on your machine, you should be able to launch the script by running 'perl convert.pl' (but don't do so until you've read through the rest of this article and

made any necessary changes). Windows users may first need to download and install the ActivePerl package (see our 'What you need' sidebar on the left for details).

Start by configuring the variables, as shown in Listing 1 above. You should define one root for the original HTML files (any file ending in .htm or .html will be included) and another for the exported text files. These should both have the same directory structures, so if your Web site is located under 'wwwroot' and contains the directories 'products' and 'news', and you want to export the text content to 'wwwtext', then your configuration will look like the one here. You must create the empty 'products' and 'news' directories (or any others, depending on your



▲ This not-for-profit site is powered by the CMS described in this series of articles

site structure) under the 'wwwtext' directory before running the script.

The next part of the script looks through the directories you've specified in the configuration section and makes a list of the files they contain, before adding on the appropriate path and storing the short and long file names separately. If you want to read other files, not just those ending in .htm and .html, change the last line as appropriate. Generally, though, you shouldn't need to make any changes to the code in **Listing 2**.

The same applies to **Listing 3**, which simply opens one of the files in the list and read its contents into a variable. This part of the script places the contents of the relevant HTML file in a variable called '@lines'. The other significant variables are emptied prior to running through the loop, to ensure that no leftovers from the previously processed file are present to affect the way the current file is processed.

At this point, the script will work its way through the file contents on a line-by-line basis, extracting the relevant information, as shown in **Listing 4**. For example, anything in a <TITLE> tag is assumed to be the page title, while everything between the <BODY> and </BODY> tags is placed into the 'body' variable. At this point very little processing has taken place. Even so, you may want to consider adding extra conditions here to cope with, for example, meta tags or some other element that is likely to be found in each HTML page of your particular site.

Once the basic information has been extracted, the trial and error approach comes in to effect. The code in **Listing 5** works through the body text performing various search and replace operations to try to convert all HTML tags to the proprietary format that's used by the CMS.

Part of this procedure is the same as the 'process\_html\_tags' section found in the editor part of the CMS; it converts a basic range of HTML tags—mostly those associated with text formatting—into a proprietary format. This is done so that you can then choose to serve the data in any way that you see fit (WAP, XML, HTML and so on.) However, the

### 3 Reading in the contents of an HTML file

```
open(FILE, $fullFilename);
@lines = <FILE>;
close(FILE);

$filetitle="";
$body="";

$dowrite=0;
```

### 4 Extracting information from the HTML file

```
foreach $line (@lines) {

    # handle end-of-line spaces carefully
    if (($line =~ m/>$/) || ($line =~ m/^</)) {chomp($line);}
    else {
        $line =~ s/\n/ /g;
        $line = " " . $line;
    }

    if ($line =~ m|<TITLE>(.*?)</TITLE>|i) {
        $filetitle = $1;
        next;
    }

    if ($line =~ m|<body(.*?)>|i) {
        $dowrite = 1;
    }

    if ($line =~ m|</body>|i) {
        $dowrite = 0;
    }

    if ($dowrite == 1) {$body .= $line;}
}
```

### 5 Converting the HTML tags into proprietary format

```
# do some tidying
$body =- s|<body([^\>]*)>||i;
$body =- s|</body>||i;

$body =- s|</P>\s*<P>|\(\(!paragraph\)|ig;
$body =- s|<P>|\(\(!paragraph\)|ig;
$body =- s|</P>|\(\(!paragraph\)|ig;
$body =- s|<BR>|\(\(!newline\)|ig;

# convert HTML to proprietary tags
$body = &process_html_tags($body);

# do some more tidying
$body =- s/\(\(!newline\)\)\s?$/i;
$body =- s/\(\(!paragraph\)\)\s?$/i;
$body =- s/\(\(!newline\)\)\(\(!paragraph\)\)|\(\(!paragraph\)\)|ig;

$body =- s/ / /g;
$body =- s/\s*\(\(!paragraph\)\)\(\(!paragraph\)\)/ig;
$body =- s/\(\(!paragraph\)\)\s*\(\(!paragraph\)\)/ig;
$body =- s/\s*\(\(!newline\)\)\(\(!newline\)\)/ig;
$body =- s/\(\(!newline\)\)\s*\(\(!newline\)\)/ig;
$body =- s/\s*$/i;
```

## Time for a live performance

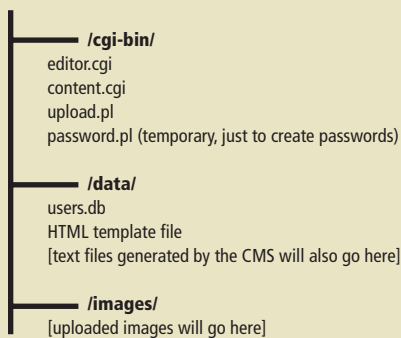
**MAKE SURE YOU USE THE EXISTING DIRECTORY structure for the script and data files, which you'll find at `www.pcmag/pcm/solutions.jsp`, unless you plan to change the relevant variables in the scripts. This is best left until you've tested the default configuration to check that everything's working properly. You should also remember that the `editor.cgi` script lets you create and edit text documents, while the `content.cgi` script displays those documents as HTML pages.**

The layout of the files on your server should be similar to that shown in Figure 2. Note that the editor script only allows you to edit the files in one directory. If you want to store your pages in several directories, you have two choices. You can either make duplicates of the scripts and change their path variables, or you can modify the main script to handle the directory structure as part of the query string. This latter is relatively easy and is left as an exercise for the reader.

Perl scripts can be opened and edited in any text editor, although a programmer's text editor is best because you'll then have colour-coded statements, variables and so on, which makes the whole thing easier to read and edit.

As is the case with all new software, be sure to test the CMS thoroughly on a local machine before you even consider running it on your live Web server. That way any customisations can be made and any new bugs can be eliminated.

### WWW Root



#### 2 File/directory tree

chances are you'll need to add some lines of your own here, as the existing conversions are unlikely to cover every tag found in your site. This is particularly true where tables are concerned. Remember that if you make any changes here you should also change the relevant sections of the editor and content-serving scripts in the same way.

The routines used to process the text and perform these search and replace operations are called regular expressions (or regexps). To find out more about them, search the Web for 'regexp' or 'regular expressions' and you'll find all you need to know.

Once this processing operation is complete, the variables generated so far are placed in a text file with the same name stem (but a .txt extension) as the original HTML file, by the code shown in Listing 6. Each entry is attributed to a variable name that's recognised by the CMS. Note that the 'description', 'keywords' and 'publication\_date' entries aren't specified, as we can't generalise about them. However, if you know where in your HTML files these values are stored (perhaps in meta tags, for example), you can add the appropriate lines to the script to extract them.

#### 6 Exporting the CMS content to a file

```

# write to file
$newfilename = $fullFilename;
$newfilename =~ s/$base/$outdir/i;
$newfilename =~ s/\.htm/\.txt/i;

open (FILE2, ">$newfilename");
print FILE2 "heading=$filetitle\n";
print FILE2 "pagename=$shortfilename\n";
print FILE2 "description=\n";
print FILE2 "keywords=\n";
print FILE2 "publication_date=\n";
print FILE2 "body=$body\n";
close (FILE2);
  
```

#### 7 A basic error check is performed on the file contents

```

print "$fullFilename done (to $newfilename)";
if (($body eq "") || ($body =~ m/</) || ($body =~ m/>/)) {
{print " - with errors";}
print "<BR>\n";
}
  
```

If the body text is empty or if it contains any tag markers, the file will still be written but the script will warn that there are errors. In practise, this is almost certain to occur unless you can find and eliminate all unwanted tags. If you're not keen on as much programming as this may entail, it's best to use this script as a starting point and then edit the text files it produces by hand—it will still save you a lot of time.

And that's it. You should now have everything you need to run—and, if necessary, modify—the CMS to serve and edit pages on your site. If you make any useful modifications, discover any bugs or improve the functionality of any of the scripts included in this package, please contact me with details ([alex@itreviews.co.uk](mailto:alex@itreviews.co.uk)). We've heard from several people who are already using the CMS to power their site, so feedback from you will help others too. ■

## Further information

If you plan to delve more deeply into the workings of the CMS, you'll want some sort of reference to guide you through it.

Books on Perl are plentiful, but two are especially recommended by the author. They're 'Programming Perl' and the 'Perl Cookbook', both published by O'Reilly. These contain pretty much everything you'll ever need to know about Perl programming and are useful both as references and as tutorial guides.

# Offers Direct

## PERSONAL COMPUTER WORLD CD-ROM NUMBER 3

**PCW on CD-Rom 3** contains the entire editorial content of the January to December 2002 issues in searchable, browsable and printable format.

Our latest addition to the ultimate *PCW* library features new, enhanced search tools enabling you to search for key words, phrases, products and authors, or simply browse each issue at your leisure. The CD also includes over 300MB of Windows shareware.

System requirements:  
**PC:** Windows 98, ME, NT, 2000, XP, 64MB of Ram, CD-Rom drive.

**Mac:** Powermac or compatible, Mac OS 9.1 up, 64MB of Ram.

You will also need Acrobat or Acrobat Reader 5 installed.

Please quote order code **MQ19**

**READER PRICE**      **PCW SUBSCRIBER PRICE**

£14.95

£13.45



**ONLY  
£14.95**  
(subs price  
£13.45)

## THE ULTIMATE PCW LIBRARY - CD-ROM 1, 2 & 3

Includes the entire editorial content of issues **August 2000 to December 2002** in searchable, browsable and printable format. Search for key words, phrases, products and authors or simply browse each issue at your leisure. The CD-Roms contain the following issues: CD-Rom no 1 includes August 2000 to July 2001, CD-Rom no 2 includes February 2001 to January 2002 and CD-Rom no 3 includes January 2002 to December 2002.

Please quote order code **MQ20**


**READER PRICE**      **PCW SUBSCRIBER PRICE**

£33.95 - save £8.90

£31.45 - save £8.90



## OFFERS >>



**ONLY  
£19.50**  
(subs price)

### ROBOT BUILDING FOR BEGINNERS

David Cook

Understand robot development and find out what tools and parts you need.

Please quote order code **MQ21**

**RECOMMENDED RETAIL PRICE**      £21.50

**READER PRICE**      £21

**PCW SUBSCRIBER PRICE**      **£19.50**

### 25 HIGH-SPEED CD-Rs

Copy or transfer your data fast with some of the first 56-speed CD-Rs to arrive in the UK! Save money when you buy 25 or 50 *PCW*-branded CD-Rs.

	25 CD-Rs	50 CD-Rs
<b>READER PRICE</b>	£9.95	£17.95
<b>PCW SUBSCRIBER PRICE</b>	<b>£9.25</b>	<b>£16.50</b>

Quote order codes:      **MQ22**      **MQ23**



**ONLY  
£9.25**  
(subs price)

## HOW TO ORDER

>> **ONLINE**

via our secure website

[www.offersdirect.pcw.co.uk](http://www.offersdirect.pcw.co.uk)

>> **PHONE**

on 01858 438 883

>> **FAX**

on 01858 468 969