

# The Virtues of Virtualized Storage

Virtualization technology comes in several guises, but it's here today and can be a valuable part of your storage strategy

**A**S SENIOR DIRECTOR OF ENTERPRISE TECHNOLOGY operations at Corrections Corporation of America (CCA), a prison management firm that handles more than 60 facilities, Brad Wood faces several challenges. His group manages approximately 100TB of data — including inmate medical records, operational records, e-mail, and so forth — across four Hitachi Data Systems (HDS) storage arrays in two datacenters. Because of federal and state rules, much of the company's data is mirrored three or four times to keep it accessible in case of failure. Adding to the complexity, Wood buys his hardware based on current price and performance, so he has a mix of suppliers.

With costs escalating, Wood needed a way to slow deployment of new storage hardware and make better use of the disparate hardware he already had. The solution he chose was to implement storage virtualization.

The idea behind virtualization sounds deceptively simple. It aggregates storage systems (such as arrays) from multiple providers into a networked environment that can be managed as a single pool. In essence, "the storage controller is being disaggregated and spread around the network," says



Brian Garrett, lab technical director at Enterprise Strategy Group (ESG), a market research firm.

In Wood's case, his storage engineer can now manage the company's various hardware from a single console, using Symantec Veritas Storage Foundation and HDS HiCommand software.

Storage virtualization has several key IT benefits in addition to improved resource utilization. For one, it allows data to be moved to any storage device in the pool when subsystems fail or are replaced. Easier data migration also makes it possible to implement a tiered storage architecture cost-effectively, one in which data is moved to less expensive, lower performance systems as it becomes less business-critical over time.

Another benefit is easier replication, because virtualization removes the need for full redundancy. Without storage virtualization, IT typically must copy entire volumes from one storage type to another in order to ensure all related data is in one environment. With virtualization, it's easier to copy partial data — such as snapshots or delta files — and keep it linked to the entire data set across physical devices.

And, of course, maintenance costs are always a factor.

BY GALEN GRUMAN | ILLUSTRATION BY CAMPBELL LAIRD

## “Virtualization reduces the administration burden by getting storage folks out of the server business.”

— Rick Villars, IDC

“[Virtualization] also reduces the storage administration burden over time by getting storage folks out of the server business,” notes Rick Villars, IDC vice president of storage systems research.

If you listen to storage vendors, every product offers some form of virtualization. Since 2000, the term has been hyped consistently. Although skepticism is warranted, storage virtualization is now real enough that enterprises can begin at least exploring its use as an enabling technology to simplify storage management in heterogeneous environments. Many enterprises, like CCA, have found they can get real benefits today.

### Three Approaches to Virtualization

Vendors offer three approaches to true storage virtualization: host-client (via software), in-fabric (mainly through appliances but soon also through switches), and in-array (embedded functionality). While the vendors tout specific pros and cons of each approach (see the infographic on page 39), analysts agree that they all deliver in the end. The determining factor is usually how well a particular approach fits into your existing storage infrastructure, says Gartner research director Stan Zaffos.

The in-fabric approach is the most common method, offered in products such as DataCore Software’s SANsymphony, EMC’s InVista, FalconStor’s IPStor, IBM’s SVC (SAN Volume Controller), NetApp’s V-Series, and StoreAge’s SVM. These products, which have been on the market for just a few years, use dedicated appliances or software running on a server to discover storage resources and then build the metadata that lets IT manage them as a virtual pool. Of these, IBM and NetApp have the largest

installed bases (about 1,000 each), Zaffos notes.

Coming soon are switch-based products — often deployed as blades within fabric switches — that essentially do the same thing as a separate appliance. These will be from companies like Brocade, Cisco Systems, MaXXan Systems, McData, and QLogic. By putting the virtualization functionality in the switch, the theory is that operations are more efficient because data travels through one fewer device than if it also went through an appliance, notes ESG’s Garrett. He expects most of these ultimately to run a version of Symantec’s Veritas Storage Foundation host-client software, although Symantec says it has no immediate product plans to port its software to run on switches.

Storage Foundation has been around in various versions for a decade, running on file and application servers to detect storage resources and maintain the metadata used to manage them. Until recently, Veritas (now a division of Symantec) did not release its Unix and Windows versions in sync, so it was hard to use Storage Foundation in heterogeneous environments, Gartner’s Zaffos says. Still, he says, the technology is easy to use for many purposes, including data migration, load balancing, and flexible provisioning.

A third type of storage virtualization is exemplified by the TagmaStore network controller, from Hitachi Data Systems, which lets HDS’s management software work with multiple vendors’ storage systems as if they were one pool. Approximately 45 percent of the roughly 1,700 current TagmaStore customers implement its virtualization technology, says Claus Mikkelsen, HDS’s chief scientist. Its key benefit, according to Zaffos, is that “you’re

not adding another element in the I/O path, so you’re not buying another asset.” Because it’s usually cheaper to replace storage arrays than to pay for their annual maintenance, Zaffos expects TagmaStores to be used mainly to ease migration from old arrays.

Pricing for all three strategies is fairly equivalent, though that’s not immediately evident when comparing, say, a TagmaStore controller with a NetApp appliance or a Symantec software license. “The pricing variables are driven more by scale,” says IDC’s Villars. “For example, a TagmaStore is more expensive than IBM’s SVC, but you’re buying more.”

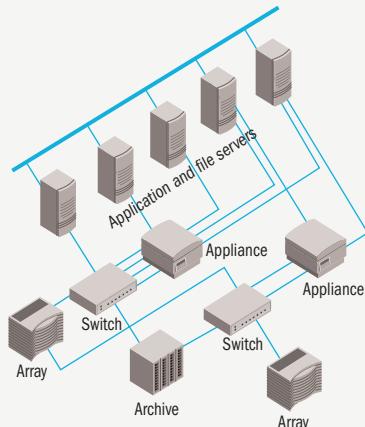
One potential gotcha is that while virtualization promotes the idea of cross-vendor storage utilization, all three strategies also enforce vendor lock-in. In-array products obviously lock you into a specific vendor’s array hardware, says Mark Lewis, EMC’s chief development officer, but in-fabric and host-client products lock you into the virtualization software or the appliance that embeds that software.

### Where Virtualization Works Today

ESG’s Garrett says his research shows that storage virtualization applied to storage environments with at least six storage fabrics reduces costs in several areas: Hardware costs drop 23.8 percent, on average; software costs drop 16.2 percent; and administration costs drop 19.3 percent.

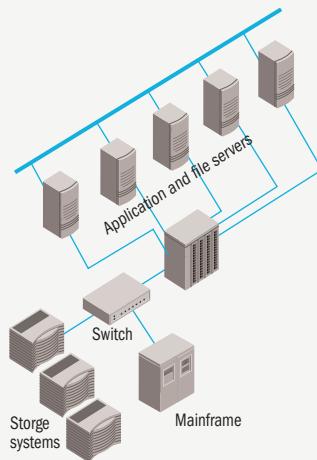
Once an enterprise has deployed storage virtualization, the technology is “relatively easy to use,” Garrett says. The real effort lies in getting up and running. So Garrett recommends that IT focus on a specific tactical issue, such as getting nondisruptive data migration in place. If you apply storage virtualization to that specific

## In-Fabric (appliance)



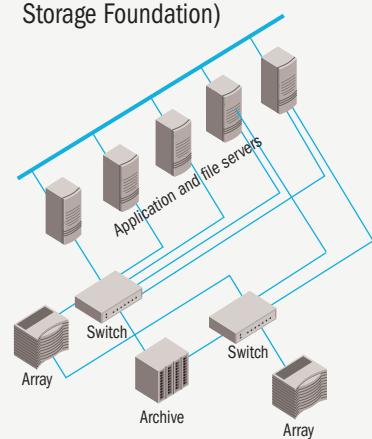
An appliance monitors the SAN for all available storage and builds metadata about stored data and its location. **Pros:** Cross-platform; allows multiple logical pools; fairly low cost; can scale incrementally. **Cons:** Uses additional network resources; adds additional devices; metadata can be scattered across multiple appliances.

## In-Array (HDS TagmaStore)



A network controller pools all SAN-attached Fibre Channel storage and presents it as a single pool. **Pros:** Handles all processing off the network; maintains metadata in internal storage; centralized management. **Cons:** Designed for large-scale deployments; single point of failure.

## Host-Client (Symantec Veritas Storage Foundation)



Software at the file and application servers monitors data traffic and storage, building metadata on the fly. **Pros:** Can be deployed and scaled incrementally; mature technology. **Cons:** Server software not always in sync; large operations can be difficult to manage; metadata-exchange traffic can be burdensome.

issue, he says, “then you can extend into the other stuff as you get more experienced.”

That’s exactly the approach taken at the Baylor College of Medicine, in Houston, Texas. Two years ago, the college decided to integrate dozens of file servers and ERP stores attached to Unix and Windows servers in order to reduce unused storage capacity and lower administration costs. Despite the initial expense, Baylor decided to replace its storage devices with a single FC (Fibre Channel) storage fabric and a set of HDS arrays, recalls Mike Layton, director of IT for enterprise services and mainframe systems. Not having a heterogeneous environment to support — “a luxury,” he says — made the decision to deploy storage virtualization fairly safe.

Today, the Baylor system manages 200TB of data, including patient records and university operations data. HDS hadn’t yet released its TagmaStore array, so Layton deployed NetApp V-Series appliances instead. Baylor’s use of storage virtualization is mainly to pool storage resources,

although the college is also considering how to use the technology to implement data lifecycle management, where patient data can be highly available during treatment but later moved to lesser systems for analysis, auditing, or other needs.

Dallas-Fort Worth International Airport had a different problem. It stored flight data (such as passenger lists, arrival times, baggage tracking, and gate information) in two SANs using Oracle RAC (Real Application Clusters). Oracle RAC could treat one storage target as the primary target and then replicate to secondary systems, but this process simply took too long, recalls John Parrish, associate vice president of terminal technology. If one terminal’s SAN goes down, the other SAN has to step in immediately so flight boarding and baggage handling isn’t delayed. DataCore’s SAN-symphony appliance made Oracle RAC think it was working with just one SAN, and Parrish has seen no latency issues crop up in this deployment.

Replication issues were also a problem for Freeze.com. The online retailer

needed to keep its 400GB Microsoft SQL Server transaction databases in sync with its reporting databases, but SQL’s resource requirements prevented the reporting tools from working on the same database as the transaction management, recalls Freeze.com IT director Kyle Ohme. He would mirror the database periodically, but replication took so long that the reporting database was hours behind, preventing the kind of analysis needed to manage supplies properly. Ohme deployed tools from FalconStor to pool the storage into a virtual volume so both sets of applications can access it in real time. That way, he could send snapshots of the transaction database to the reporting tools, rather than replicate the entire thing.

### A Long-Term Effort

Although the storage virtualization promises touted in 2000 turned out to be premature, today’s technology does deliver at least the first step toward an automated, self-managing storage infrastructure that functions more as an IT utility, notes Gartner’s Zaffos.

## The large enterprises that benefit most from storage virtualization are the most conservative in deploying new technology.

But that version of storage virtualization is many years away.

One reason is economics. The storage vendors can't afford suddenly to lose the profits from their hardware businesses if storage hardware becomes a commodity managed by software tools. As a result, they're only likely to take measured steps to support the independent standards needed by third-party management tools, Zaffos says.

Operating system vendors may help force the issue if they adopt some of the Storage Networking Industry Association standards, such as Volume Shadow Copy Service, and start moving storage virtualization from network/storage middleware to the OS, says George Teixeira, CEO of DataCore. In fact, IDC's Villars says that in

five years Microsoft Windows could implement storage virtualization for mid-tier enterprise deployments.

Another impediment to the grander storage-virtualization vision is that the large enterprises that benefit most from storage virtualization are the most conservative in deploying new technology, since their risk of failure is greater. "We're not seeing the deployments pushing to the technical extreme," Villars says.

The third reason is that the tools are still immature. Current tools focus on giving IT a common console for managing storage. Over time, vendors will begin to add automation and policy-based intelligence for provisioning storage and managing data migration and replication. But because of the

complexity of the infrastructure that such tools must manage, "I don't see it for years and years," says Tom Clark, director of solutions and technologies at McData and author of "Storage Virtualization" (Addison-Wesley, 2005). In the interim, IT should expect to see point solutions such as heterogeneous replication and snapshots, he says.

But none of these factors prevent enterprises from benefiting from storage virtualization today. In the immediate term, companies should apply virtualization to solve specific problems, such as easing data migration or pooling data across SANs. As time goes on, IT can incrementally broaden its use of the technology, taking advantage of the continued improvements vendors will make over the coming years. ↻

### The iSCSI Issue

AS MORE AND MORE PRODUCTS ENTER THE MARKET, iSCSI IS becoming an increasingly attractive alternative to FC (Fibre Channel) SAN technology ([infoworld.com/3347](http://infoworld.com/3347)). Not only is iSCSI cheaper than Fibre Channel, but the technology is less complex to implement. Because it uses the familiar IP network protocols, it simplifies the IT skill set needed to maintain the SAN. Thus, though it's not as fast and has a lower maximum capacity than FC systems, iSCSI meets the needs of many small businesses and non-mission-critical enterprise storage applications, such as departmental file sharing and near-line data storage.

If you're considering iSCSI and virtualization is on your storage road map, however, you may want to think again — at least for the time being. The management tools that take advantage of storage virtualization to aggregate storage subsystems across a network are not yet mature enough to handle iSCSI SANs, says Brian Garrett, lab director at the Enterprise Strategy Group, a market research firm.

That's why Corrections Corporation of America has held back deploying cheap iSCSI SANs to handle lower-priority storage, according to Brad Wood, the company's senior director of enterprise technology operations. Wood says today's storage-management software can't handle iSCSI SANs well. For

example, iSCSI doesn't have a single facility to provide global names, as Fibre Channel does.

Kyle Ohme, IT director at online retailer Freeze.com, also struggled to deploy storage virtualization on iSCSI SANs. Even with help from vendors BlueArc and FalconStor, the difficult effort has led him to scale back his iSCSI plans.

One reason for the lack of tools is iSCSI's comparatively small market share. Garrett estimates that there are maybe 10,000 iSCSI SANs in place, compared with hundreds of thousands of FC SANs. Plus, he says, most iSCSI deployments are in smaller enterprises that don't use the tiered storage or storage lifecycle management techniques that these tools address.

"iSCSI SANs are not yet well supported," says Mark Lewis, chief development officer at EMC, but he expects that to change as IP-based storage gains market share. The Storage Networking Industry Association's SMI-S standard may help bring to the iSCSI and IP world the same storage-management capabilities that Fibre Channel and SCSI enjoy — and a basis on which to translate among them, notes Roger Wofford, IBM's storage virtualization marketing manager. Until then, even in environments where iSCSI and FC SANs are interconnected, the management tools will remain separate.

— G.G.

Apply online at: <http://subscribe.infoworld.com>

**PRIORITY CODE: WW5PDF**

I wish to receive a free subscription to **InfoWorld**.  
 **Yes**    **No**

SIGNATURE \_\_\_\_\_ DATE \_\_\_\_\_

**A. MAILING ADDRESS**

Publisher reserves the right to limit the number of complimentary subscriptions. Free subscriptions available in the U.S. (including APO and FPO) and Canada

NAME \_\_\_\_\_

TITLE \_\_\_\_\_

COMPANY NAME \_\_\_\_\_

DIVISION / DEPT. / MAIL STOP \_\_\_\_\_

MAILING ADDRESS \_\_\_\_\_

CITY / STATE / ZIP / POSTAL CODE \_\_\_\_\_

Is the above address a home address?    1. Yes    0. No

E-MAIL ADDRESS \_\_\_\_\_

BUSINESS PHONE (INCLUDING AREA CODE) \_\_\_\_\_ BUSINESS FAX NO. (INCLUDING AREA CODE) \_\_\_\_\_

**1. Over the course of one year, do you buy, specify, recommend, or approve the purchase of the following products or services worth:**

Please include amounts for all locations of your organization. Consultants: please include what you recommend for your clients as well as what you buy for your own business.

- |                                  |                                |                            |
|----------------------------------|--------------------------------|----------------------------|
| 01. \$100 million or more        | 06. \$5,000,000 to \$9,999,999 | 11. \$100,000 to \$399,999 |
| 02. \$50,000,000 to \$99,999,999 | 07. \$2,500,000 to \$4,999,999 | 12. \$50,000 to \$99,999   |
| 03. \$30,000,000 to \$49,999,999 | 08. \$1,000,000 to \$2,499,999 | 13. Less than \$49,999     |
| 04. \$20,000,000 to \$29,999,999 | 09. \$600,000 to \$999,999     | 14. None                   |
| 05. \$10,000,000 to \$19,999,999 | 10. \$400,000 to \$599,999     |                            |

| Product category                         | Write code in box    |
|--|----------------------|
| Large systems                            | <input type="text"/> |
| Client computers                         | <input type="text"/> |
| Networking / Telecom (including servers) | <input type="text"/> |
| Wireless                                 | <input type="text"/> |
| Internet / Intranet / Extranet           | <input type="text"/> |
| Security                                 | <input type="text"/> |
| Storage                                  | <input type="text"/> |
| Peripheral equipment                     | <input type="text"/> |
| Software                                 | <input type="text"/> |
| Service/Support / Outsourcing            | <input type="text"/> |

**2. What is your primary job title? (PLEASE CHECK ONE ONLY)**

- |   |   |
|---|---|
| <b>IT / Technology Management</b>                                     | <input type="checkbox"/> 10. IT Staff                                   |
| <input type="checkbox"/> 01. CTO, CIO, CSO, Vice President            | <input type="checkbox"/> 11. Other IT Professional                      |
| <input type="checkbox"/> 02. Director                                 | <b>Corporate / Business Management</b>                                  |
| <input type="checkbox"/> 03. Manager / Supervisor                     | <input type="checkbox"/> 12. CEO, COO, President, Owner, Vice President |
| <input type="checkbox"/> 04. Network Manager / Director               | <input type="checkbox"/> 13. CFO, Controller, Treasurer                 |
| <input type="checkbox"/> 05. Engineer                                 | <input type="checkbox"/> 14. Director                                   |
| <input type="checkbox"/> 06. Systems Analyst / Programmer / Architect | <input type="checkbox"/> 15. Manager / Supervisor                       |
| <input type="checkbox"/> 07. Other IT Management                      | <input type="checkbox"/> 16. Other Business Management Title            |
| <b>IT / Technology Professional</b>                                   |   |
| <input type="checkbox"/> 08. Consultant / Integrator                  | <input type="checkbox"/> 98. Other Title                                |
| <input type="checkbox"/> 09. Developer                                | (specify) _____   |

**3. Please indicate your job function(s)? (PLEASE CHECK ALL THAT APPLY):**

- |  |  |
|--|--|
| <b>IT / Technology Functions</b>   | <b>Corporate / Business Functions</b>                          |
| <input type="checkbox"/> 01. Executive   | <input type="checkbox"/> 09. Executive                         |
| <input type="checkbox"/> 02. Department Management - IT                                  | <input type="checkbox"/> 10. Department Management - Business  |
| <input type="checkbox"/> 03. Networks / Systems Management                               | <input type="checkbox"/> 11. Financial / Accounting Management |
| <input type="checkbox"/> 04. Applications Development                                    | <input type="checkbox"/> 12. Research / Development Management |
| <input type="checkbox"/> 05. Management of Enterprise Applications (CRM, ERP, SCM, etc.) | <input type="checkbox"/> 13. Sales / Marketing Management      |
| <input type="checkbox"/> 06. Research / Development Management                           | <input type="checkbox"/> 14. Other Business Functions          |
| <input type="checkbox"/> 07. Consultant / Integrator                                     | <input type="checkbox"/> 98. Other Functions                   |
| <input type="checkbox"/> 08. Other IT Functions  | (specify) _____  |

**4. Are you involved in buying, specifying, recommending or approving the following IT products / services?**

(PLEASE CHECK ALL THAT APPLY):

- |   |  |
|---|--|
| <b>Software / Products / Technologies</b>   | <input type="checkbox"/> 18. Web / Video Conferencing                |
| <input type="checkbox"/> 01. Customer Relationship Management                       | <input type="checkbox"/> 19. Storage                                 |
| <input type="checkbox"/> 02. Enterprise Resource Planning                           | <input type="checkbox"/> 20. Disaster Recovery                       |
| <input type="checkbox"/> 03. Business Process Management / Outsourcing              | <input type="checkbox"/> 21. Security                                |
| <input type="checkbox"/> 04. Business Intelligence / Data Mining / Data Warehousing | <input type="checkbox"/> 22. Anti-Virus / Content Filtering          |
| <input type="checkbox"/> 05. Portals  | <input type="checkbox"/> 23. Firewall                                |
| <input type="checkbox"/> 06. Financials / Payroll / Billing                         | <input type="checkbox"/> 24. VPN                                     |
| <input type="checkbox"/> 07. Performance / Application Management                   | <input type="checkbox"/> 25. Identity Management                     |
| <input type="checkbox"/> 08. .NET   | <input type="checkbox"/> 26. Authentication / Authorization          |
| <input type="checkbox"/> 09. Other Software   | <input type="checkbox"/> 27. Intrusion Detection & Prevention        |
| <input type="checkbox"/> 10. Networking   | <input type="checkbox"/> 28. Encryption                              |
| <input type="checkbox"/> 11. Web Services   | <input type="checkbox"/> 29. Other IT Products / Technologies        |
| <input type="checkbox"/> 12. Content Delivery Networks                              | <b>Hardware / Peripherals</b>  |
| <input type="checkbox"/> 13. Network and Systems Management                         | <input type="checkbox"/> 30. Servers                                 |
| <input type="checkbox"/> 14. VoIP (Voice Over IP)                                   | <input type="checkbox"/> 31. Notebooks / Laptops                     |
| <input type="checkbox"/> 15. Telecommunications                                     | <input type="checkbox"/> 32. PDAs / Handhelds / Pocket PC / Wireless |
| <input type="checkbox"/> 16. Wireless   | <input type="checkbox"/> 33. Printers                                |
| <input type="checkbox"/> 17. Remote Access  | <input type="checkbox"/> 34. Other Hardware / Peripherals            |

**5. What is your organization's primary business activity at this location? (PLEASE CHECK ONE ONLY):**

- |  |  |
|--|--|
| <b>General Business Industries</b>   | <b>Technology Providers</b>  |
| <input type="checkbox"/> 01. Defense Contractor / Aerospace                            | <input type="checkbox"/> 12. Service Provider (MSP, BSP, ISP, ASP, etc.)                     |
| <input type="checkbox"/> 02. Retail / Wholesale / Distribution (non-computer)          | <input type="checkbox"/> 13. Computer / Network Consultant                                   |
| <input type="checkbox"/> 03. Pharmaceutical / Medical / Dental / Healthcare            | <input type="checkbox"/> 14. Systems / Network Integrator, VAR / VAD                         |
| <input type="checkbox"/> 04. Financial Services / Banking                              | <input type="checkbox"/> 15. Technology Manufacturer (hardware, software, peripherals, etc.) |
| <input type="checkbox"/> 05. Insurance / Real Estate / Legal                           | <input type="checkbox"/> 16. Technology - Related Retailer / Wholesaler / Distributor        |
| <input type="checkbox"/> 06. Transportation / Utilities                                | <b>Government / Education</b>  |
| <input type="checkbox"/> 07. Media (print / electronic)                                | <input type="checkbox"/> 17. Government: federal (including military)                        |
| <input type="checkbox"/> 08. Communication Carriers (telecomm, data comm., TV / cable) | <input type="checkbox"/> 18. Government: state or local                                      |
| <input type="checkbox"/> 09. Construction / Architecture / Engineering                 | <input type="checkbox"/> 19. Education   |
| <input type="checkbox"/> 10. Manufacturing & Process Industries (non-computer)         | <input type="checkbox"/> 98. Other   |
| <input type="checkbox"/> 11. Research / Development                                    | (specify) _____  |

**6. How many people are employed at this organization, including all of its branches, divisions and subsidiaries?**

(PLEASE CHECK ONE ONLY):

- |   |  |
|---|--|
| <input type="checkbox"/> 1. 20,000 or more  | <input type="checkbox"/> 5. 500 - 999    |
| <input type="checkbox"/> 2. 10,000 - 19,999 | <input type="checkbox"/> 6. 100 - 499    |
| <input type="checkbox"/> 3. 5,000 - 9,999   | <input type="checkbox"/> 7. 50 - 99      |
| <input type="checkbox"/> 4. 1,000 - 4,999   | <input type="checkbox"/> 8. Less than 49 |

**7. Which of the following operating systems are in use or planned for use at this location?**

(PLEASE CHECK ALL THAT APPLY):

- |  |   |
|--|---|
| <input type="checkbox"/> 01. Windows XP    | <input type="checkbox"/> 04. Linux / Unix / Solaris |
| <input type="checkbox"/> 02. Other Windows | <input type="checkbox"/> 05. Other                  |
| <input type="checkbox"/> 03. Mac           | (please specify) _____                              |

**B. CONTACT PREFERENCES**

You may receive a renewal reminder via e-mail. May we send other information about InfoWorld products, services, or research via e-mail?    1. Yes    0. No

We occasionally send our subscribers email messages with news about technology solutions and special offers from qualified third parties. Would you like to receive these messages?  
 1. Yes