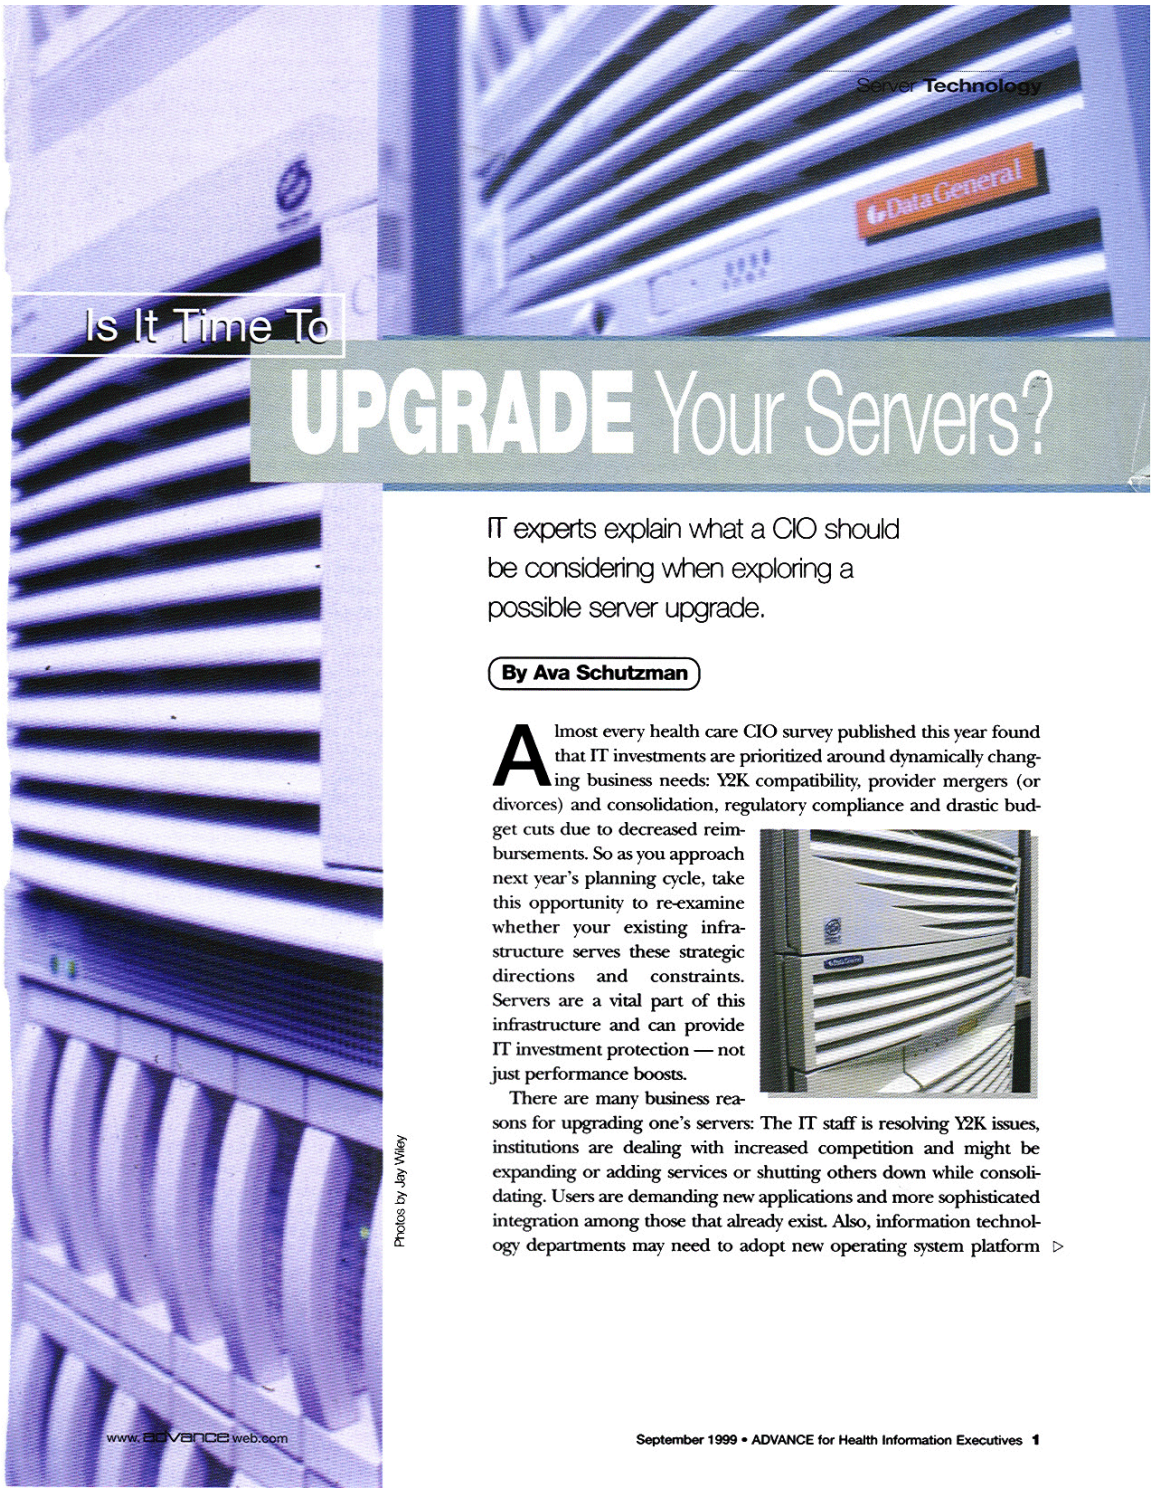

Is It Time To UPGRADE Your Servers?

IT experts explain what a CIO should be considering when exploring a possible server upgrade.

By Ava Schutzman

Almost every health care CIO survey published this year found that IT investments are prioritized around dynamically changing business needs: Y2K compatibility, provider mergers (or divorces) and consolidation, regulatory compliance and drastic budget cuts due to decreased reimbursements. So as you approach next year's planning cycle, take this opportunity to re-examine whether your existing infrastructure serves these strategic directions and constraints. Servers are a vital part of this infrastructure and can provide IT investment protection — not just performance boosts.

There are many business reasons for upgrading one's servers: The IT staff is resolving Y2K issues, institutions are dealing with increased competition and might be expanding or adding services or shutting others down while consolidating. Users are demanding new applications and more sophisticated integration among those that already exist. Also, information technology departments may need to adopt new operating system platform ▷

Photos by Jay Wiley

to support strategic applications such as CPR and clinical decision support systems. To respond to HIPAA requirements, institutions will have to add server and storage capacity to save patient records and documents for a much longer time, and address security concerns.

All types of software components within the IT infrastructure — application databases, middleware and operating systems — play a role in determining which computing and storage systems are most appropriate. Figure 1 summarizes the "layers" of hardware and software that comprise a typical enterprise solution. Server upgrades should address the organization's business needs by providing high availability and accommodating growth.

Let's examine each level of the enterprise solution and how it impacts the type of improvements that can benefit server performance.

Planned server upgrades or additions are often anticipated due to application- and vendor-specific requirements. Unplanned indicators are usually related to performance, but the explanation for a performance decrease is not always obvious. System performance can degrade at any component level indicated in Figure 1, or when the business demands on the servers have surpassed the capacity of the current operating environment to process system throughput in an acceptable timeframe. Because of this, vendors usually offer a variety of server upgrades that can address customer needs in a modular fashion.

In Level 1, hardware-related upgrades are available for different system components. Performance issues can stem from a variety of sources that may be "maxed out" or otherwise stressed. For example, CPU, system bus, memory size and speed, and peripheral access time all interact to determine the performance impact. Bottlenecks are caused by three different sources: the CPU or cache, memory, and I/O (whether to use SCSI or fiber channel controllers that connect to disk drives, network cards, modems or graphics).

There are also issues of hardware to consider when upgrading (to resolve performance and scalability problems):

- processor speed and cache architecture;
- platform architecture and scalability;
- peripheral storage and I/O; and
- server form factor, flexibility and scalability.

Level 5: Enterprise

Level 4: Applications

Level 3: Middleware

Level 2: Operating System

Level 1: Server Hardware

Figure 1: Solutions components to consider when upgrading servers

Processor speed, cache architecture

Processor "speed bumps" are successively faster chip speeds (measured in MHz). These speed bumps can extend the life of a system, and give the system small performance increases, as the latest technology becomes available. They can also protect the peripheral system investment, the chassis and the motherboard — any of which may have to be replaced to be compatible with new applications or systems. However, while faster, speed bumps do not always translate to perceptible performance boosts.

The size, speed and position of the cache relative to the CPU chip — as well as the architecture of the cache — whether industry-standard or vendor-optimized for particular applications, also affect overall performance.

Upgrading servers by using multiple CPUs in the system to process more workload is a planned strategy. Computing systems have an upper limit to the number of CPUs they can accommodate. Most applications today are designed to take advantage of symmetrical multiprocessing (SMP), a capability that enables a customer to purchase a four-way server with two CPUs, and add one to two CPUs when needed.

Customers also find that processor scalability is more extensive with today's technology. (Pentium III processors run up to 600 MHz.) Also, older processors may not be upgraded beyond a certain point.

Platform architecture and scalability

Generally, the more memory, the better the performance. However, the performance of 2GB of memory on one system can be very different from that on another system. Memory speed and the number of devices affect performance.

The design of the motherboard and the speed of the system bus that communicates with the rest of the system can significantly impact performance. Vendor designs can vary widely.

According to Adrianna Swartout, Data General's engineering manager for departmental servers, "A new motherboard with enhanced design and faster system buses can be a clue that overall performance and scalability will be favorably impacted by a technology-oriented server upgrade. You can take two servers, both using two-way Pentium III processors with the same amount of memory,

Customer-Friendly APPROACH TO SERVER UPGRADES

For more than a year, Data General has offered its NT customers the option to purchase robust 4-way servers as a cost-effective way to upgrade to AViiON AV8900 8-way Profusion-based NT platform. When the AV8900 begins shipping later this year, existing AV8600 and AV8700 customers can participate in an upgrade program.

The upgrade includes delivery and installation of a new chassis and preservation of many existing components without degrading performance. The upgrade is priced at $20,000 and includes an upgrade to an 8-way capable AV8900 four-processor configuration with 1GB of memory. The system will be configured and tested to reduce downtime during installation. The total time needed for an upgrade is less than six hours. Customers have 45 days to return unused parts from their older systems and will receive a three-year on-site hardware warranty on new components. Older parts will maintain existing warranty terms.

In comparison, some vendors are promoting an in-box upgrade to their existing four-way servers. While some in-chassis upgrades may be necessary, upgrading to a Profusion-based, eight-way server is a complex undertaking because several system boards need to be replaced, causing unacceptable downtime risk. The field service technician is basically manufacturing a new computer at the customer's location. The complexity doubles when the field technician upgrades in an NT cluster configuration.

Customers should also ensure that the upgrade pricing policy for the program is clearly spelled out and unambiguous. Make sure the upgrade program is simple and to the point — similar to the installation process.

Protecting your investment

Here are the AV8700 components that can be used with the new AV8900 system:

1. Pentium III Xeon processors;
2. Level 2 cache;
3. Most PCI cards;
4. CLARiiON disk arrays;
5. Tape backup devices;
6. Rack-mount cabinets; and
7. Monitor, keyboard, mouse and printers.

and get a big performance boost with successive-generation motherboards that use faster and wider system buses. The chip is only part of the equation. If speed bumps are implemented occasionally or if the motherboard changes, you'll enjoy a significant boost in performance."

Peripheral storage and I/O

The I/O architecture, which includes the number, speed and width of the system's PCI buses, combined with the total number of PCI slots, determines the maximum rate of access to disks and networks. Vendor systems vary widely in this area, although specifications may appear to be similar. As with system buses, I/O bus speed and width can be a major differentiating factor contributing to overall system performance.

Total peripheral storage isn't an issue as more capacity can be added by purchasing external arrays. However, application performance can be affected by the number of drives independent of the total amount of storage, as a single drive can only support a limited access rate. Disk drives come in different heights, and the highest drive yields the best possible capacity; however, there is a trade-off in environments that require compact design to conserve floor space or high-rate data access.

The type of storage controller on the server (e.g., SCSI vs. UltraWide or Ultra2) determines which storage devices can be added. At some point, an IT department will want to consider the higher performance and improved reliability offered by fiber storage technology.

Flexibility and scalability

The purchase of rack-mounted, rather than pedestal-based servers is a strategy that provides flexibility to install, upgrade and swap modular components. In the rack space you can easily accommodate different types of processors integrated into the same chassis.

If your vendor offers scalable systems within a family of departmental to large enterprise servers that can be upgraded, then you can assume that hardware components are interchangeable or can be re-deployed on other models, offering investment protection.

Bruce Reirden, vice president and CIO of Care New England in Providence, R.I., recently upgraded his older servers from other vendors and earlier Data General products to newer Intel processors. He now runs four AViiON 3650 servers at one site, and an AViiON 6600 server at another. He points to a recent statistic that showed 100 percent uptime, which means that his end user/customers were online and productive. The systems support QuadraMed/Compucare clinicals and MEDITECH's lab system at one site, and Keane's HIS at two sites. Reirden's upgrades took advantage of improved technology, while his investment yielded higher system availability, reliability, and performance.

Operating system-dependent upgrade

If an institution's IT architecture is standardized around preferred operating systems, the CIO may want to limit the choice of applications and hardware to those that run on those operating systems. Server upgrades therefore depend on this factor. Many server vendors specify a minimum level for each operating system. Selecting a de facto standard system, such as Microsoft NT or UNIX can increase the number of options available.

When Windows 2000 is released, Microsoft will require minimum hardware configurations, and will no longer support some options that had been supported under Microsoft Windows NT Server 4.0. And, to use some newer technologies (e.g., digital cameras and scanners that use the IEEE 13 94 interface ["fire-wire"]), it may be necessary to install Windows 2000 on a dedicated PC or server.

Selecting systems that can handle multiple generations of processors is key; this strategy allows for future innovation.

Combining hardware and software can be a powerful incentive to develop an upgrade strategy based on a flexible architecture, such as Non-Uniform Memory Access (NUMA). NUMA systems are extensions to SMP architectures, allowing more processors, memory and I/O to be used in a single system.

Another consideration is the ability to adapt new technologies and applications to the computing environment. New applications should be prototyped and tested on the hardware on which it will be deployed without disturbing the production system. For example, new applications can be developed and activated (for Y2K testing or quality assurance testing) in a partition (part of one system), then integrated into the production environment.

Many analysts are estimating that by the year 2000, more than 90 percent of

all enterprise installations will include both UNIX and Windows NT. The ability to integrate NT applications within the same system is beneficial to a major system implementation effort. This allows a cost-effective way to integrate best-of-class applications and to phase in new systems in an evolutionary rather than revolutionary manner.

Middleware-dependent upgrades

Database and middleware vendors typically have version compatibility restrictions that may impact a customer's selection of server vendor or model when upgrading. When software capabilities are significantly enhanced or expanded, system performance on the same server may not be as good as on the previous release because of the additional workload. Leveraging the new features may require additional memory, disk or CPU power.

Application-dependent upgrades

As the first level in Figure 1 relating to actual business needs, the application will be a primary determinant of the upgrade strategy.

Typically, application vendors impose restrictions on operating systems and versions supported, and database vendor and version. Less obvious factors may include whether the application was written for a specific hardware architecture, or the amount of marketing and technical cooperation between the vendors.

The network and technical operations staff at Cape Cod Healthcare (CCHC) in Hyannis, Mass., successfully leveraged application enhancements to increase the capacity of its physician transcription database. CCHC used MEDITECH's MAGIC "big block" feature that increased maximum database size from 8GB to 32GB per server.

To complement the application upgrade, CCHC upgraded its Data General AViiON servers from a Motorola-based to an Intel high-density configuration, which MEDITECH has benchmarked at a 10x improvement in CPU performance.

A long-time customer, John Kilroy, vice president and CIO at CCHC, ensured that overall system throughput would realize the benefit from the new processor technology by becoming one of the first customers to adopt full-fiber channel storage arrays. Certain applications, such as billing and laboratory, are extremely I/O-intensive, and Cape Cod's investment in storage technology ensured a comparable performance increase on the I/O side.

Another example of balancing the importance of systems performance and protecting the hardware investment when deciding to upgrade comes from St. Agnes HealthCare (SAHC) in Baltimore. Its CIO, William Greskovich, and his IT staff had to address this and other challenges when they needed to upgrade their client/server hardware platform earlier this year. The sun, moon and stars seemed to converge as Greskovich's team addressed the following challenges at once:

St. Agnes HealthCare in Baltimore is an example of an organization that balanced system performance and protected its server investments.

- *Preparing for the electronic medical record.* SAHC tripled its storage capacity for the enterprise,
- *Improving data storage capabilities for recently deployed clinical systems.* SAHC satisfied increased user demand for storage resulting from the move to electronic signature throughout the enterprise and the growth of its recently deployed clinical systems.
- *Validating Y2K readiness for core information systems.* SAHC took advantage of the downtime to apply Y2K hardware code updates, test the Y2K information systems contingency plan and evaluate all downtime procedures.

Before moving to an NT-based client/server environment, SAHC operated a system that required daily downtime and a four-hour shutdown every weekend. On SAHC's new platform, there is no scheduled downtime, and users are now accustomed to 100 percent system availability. As the user workload increased and SAHC established baseline data to determine storage requirements, Greskovich found it necessary to accelerate his upgrade plans to accommodate new clinical and electronic signature applications and the introduction of the electronic medical record.

This upgrade resulted in a storage configuration of 1.2 terabytes as opposed to the original 450 gigabytes on six CLARiiON Model 3000s connected to 12 AViiON 5900 application servers. While the servers were preserved and look the same to a casual observer, the older 4GB disk drives were swapped out and replaced by 43 new 10,000 RPM 18GB drives. Everything was backed up at least twice to two separate locations then swapped out and replaced. Re-initialization of the high-availability RAID groups brought the systems back online. This proved to be the largest upgrade evolution of its kind, and in spite of the enormous number of files being backed up, copied and restored to six different storage arrays, the task was completed in 19 hours (over a weekend to minimize disruption).

Enterprise-level upgrades

The last category in Figure 1 is interoperability, and ties various system components together:

- network capacity, impact or management;
- new computing architectures;
- high availability and backup; and
- service concerns.

The network infrastructure also impacts the upgrade decision in a variety of ways. From an enterprise management perspective, one might choose to upgrade servers that aren't supported by your preferred network management tool or suite. Vendors write agents so the system management software can recognize and manage them as components of the enterprise network. A server vendor will provide a compatible suite, but you should check specifically when utilizing third-party tools that your platforms support.

Expanding the capacity of the network and implementing modular extensions to the IT architecture can also extend the life of the server investments.

A relatively new strategy that has successfully helped organizations preserve

legacy investments is the use of thin-client computing. Adding this to client/server architecture enables maintenance of older, even proprietary applications on a single server while continuing to use the existing PC and dumb terminal clients.

Another recent architectural innovation, the Storage Area Network (SAN), allows separate storage from the server by combining storage devices into one pool that services all the servers in the network. If the SAN architecture is fiber-based, the technology offers significant performance; it can be configured to provide scalability benefits for enterprise applications, and facilitates backup of systems in an extremely efficient manner.

High availability and backup capabilities contribute to the customer satisfaction and the business value of IT by providing higher system uptime and insurance against downtime. Server upgrades may be needed to take advantage of the benefits derived from high availability technologies.

Over time, the maintenance costs of older hardware become more expensive, and it may be less expensive to replace the servers with newer technologies. In less than three years, the cost of ownership can drop dramatically.

Kilroy of CCHC also suggests investing in a "hot spare" server to minimize downtime during peak usage hours.

Vendor programs cater to your needs

With rapid advancements in technology, more customers are becoming interested in upgrade programs that maintain a competitive advantage while still protecting their investment. In fact, some customers see upgrade programs as an inexpensive way to deliver necessary "headroom" when a particular department or division experiences rapid growth. However, not all upgrade programs are alike, and while some offer true customer value, others should be avoided. A vendor should offer upgrades or trade-in programs that cater to its customer's needs over time and contribute to investment protection.

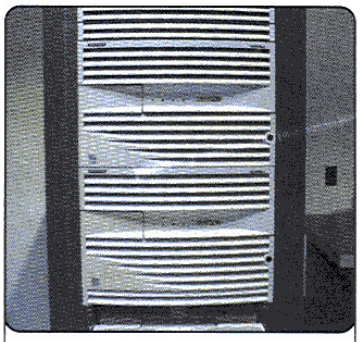

Technology upgrades alone don't always yield performance increases.

There are a few points to keep in mind when evaluating server upgrade programs. According to Frank Nicolo, product manager for Data General's AViiON Enterprise servers, "First, you should consider the complexity involved in taking your system from a relic to a roadster. The more complex the effort, the longer the downtime. Ask the supplier how long it expects the upgrade process to take and whether it can be completed over a weekend.

"Second, map out the cost of the upgrade to ensure that it does not exceed the cost of a new system. Ask the supplier to provide a list of components that must change along with the cost associated with each item," he said.

"Third, make sure you understand the performance implications of the upgrade. Will older parts in the upgraded machine provide sub-optimal performance? Be sure to consult with your supplier before upgrading.

"And last, be sure to review the terms and conditions of the upgrade program and understand the warranty. Forty-five days could allow a contingency plan through the migration," he said.

Server upgrades can be offered in different packages, most simply being an upgrade that only involves swapping out an older CPU or disk component for a newer model. Sometimes the system design won't be quite as modular, or the vendor will offer a "box swap" that provides a fair trade-in value for your old (refurbished and resellable) servers.

Keep context in mind

As we've seen, technology upgrades alone don't always yield performance increases. Carefully planning server upgrades within the context of the enterprise strategy can improve system performance, contribute to investment protection and improve customer satisfaction. ■

Ms. Schutzman is the health care product marketing manager at Data General Corp. in Westboro, Mass. She thanks numerous health care IT and Data General executives for their contributions to this article, and welcomes feedback from readers. She can be reached via e-mail at Ava_Schutzman@dg.com.