




Beyond
the
buzz

By Sarah C. Close





IT'S QUITE POSSIBLY THE MOTHER OF ALL BUZZWORDS—
BUT DOES ANYONE KNOW WHAT SCALABILITY REALLY MEANS?
MORE IMPORTANT, DO THEY KNOW HOW TO ACHIEVE IT?
HERE'S A QUICK ROAD MAP TO THE WHAT, WHY, AND HOW
OF BECOMING A SCALABLE ENTERPRISE IN TODAY'S ECONOMY

The past decade has been rough on the grammar police. Technology has spawned a new language, where proper names become verbs (“Google it!”), verbs become nouns (“We need partner buy-in before launching this service.”), and nouns gain new meanings almost monthly (“We don’t have the bandwidth for that project right now.”). In the business world, enterprises have even brought technical jargon to the boardroom—a sure sign of our dependence on IT for boosting company-wide profitability and productivity.

But amidst the buzzword frenzy, one word continues to stand out, both in frequency of use and urgency of definition. That word is *scalability*, and although it has a fairly firm meaning in terms of hardware and software, its strategic application in the data center might be its most advantageous and important connotation to date.

Say what?


To most engineers, scalability simply refers to the ability of a piece of hardware or software to maintain desired performance levels as system demands increase. For example, a server is considered scalable if it can successfully accommodate the addition of processors in order to withstand a larger or more complex workload. Likewise, a scalable application is one that can handle an influx of users or tasks without sacrificing functionality or speed.

But to many C-level executives, there is a hidden implication to this more architectural definition. True scalability means more than simply accommodating growth; it also demands that the process of scaling be both agile and cost-effective. Here’s why: If a company is feeling the need to scale its IT infrastructure to meet growing requirements, then it is likely achieving or anticipating some measure of financial success. Wasting valuable resources in extending the infrastructure might negate possible return on investment. Furthermore, lack of responsiveness could jeopardize customer service or corporate reputation. Scalability, therefore, means achieving flexibility from both technical and business perspectives.

According to a handful of leading IT vendors such as Dell, as well as industry analysts such as Frank Gens of the Massachusetts-based group IDC, today’s truly scalable enterprise is most often one whose IT architecture is based on industry-standard building blocks. “When we talk about ‘standard’ components,” says Gens, “we are referring to hardware or software built on or supporting widely accepted market-standard platforms, such as Microsoft® Windows® or Intel® x86. In particular, standardization should now be a fundamental principle in today’s new high-performance servers, storage, and network systems.”

On first mention, it might seem risky to base the core pieces of your infrastructure on such basic technology. But consider this: Proprietary systems can be like designer suits—attractive, expensive, and unadaptable. Standardized hardware, on the other hand, is like the white T-shirt of the data center: it’s relatively inexpensive, it’s easy to change in and out of, and it can pretty much be worn with anything.

“It’s that simple,” Gens says. “By leveraging standard, component-based IT assets, organizations can achieve easier



integration and operation. The end results include streamlined costs, increased reusability, and simplified support—and these are the things that ultimately make the enterprise data center architecture scalable. You need to optimize IT so that it can evolve efficiently and keep pace with the business.”

Why standards make cents

Engineers like Albert Chidiac, director of IT for seismic technology provider Input/Output, Inc., appreciate this modular scalability because it often provides a robustness, an agility, and an incremental growth pattern seldom available with costly proprietary configurations. “Industry standards are leading technology toward becoming a commodity business, where we no longer have to reinvent the wheel for every new application,” he explains. “Instead, we are using ‘cookie cutter’ pieces to more quickly implement applications or deploy hardware.”

A standards-based scalability strategy can also strike a harmonious chord with executive managers because it makes adapting to changes a less-expensive process, in terms of purchasing costs as well as operations. “Everyone in our business can get behind this approach,” says Chidiac. “Not only can standardized technology require much less up-front capital, but it can allow us the flexibility to pay as we grow. And more and more pieces of our infrastructure can work together in a heterogeneous environment, which makes us operate much more cost-effectively.”

Indeed, almost every facet of the enterprise can reap specific advantages from improved IT scalability. With the ability to step up performance and availability as needed, for example, computing infrastructures can deliver the uptime, data integrity, and system performance benefits required for running key business applications more responsively. Resources can be allocated where they are needed at the times they are needed, which helps balance critical workloads and improve server utilization. Additionally, employees have the resources they need to function more productively, while customers have more reliable access to valuable company services—no matter how demand might shift.

At Electronic Arts, scalability is a way of life because it lets the company constantly move around servers when and where they are needed most. “We didn’t want to be married to UNIX®

“More and more pieces of our infrastructure can work together in a heterogeneous environment, which makes us operate much more cost-effectively”

hardware and the vendors’ road maps, which might not meet our needs in the future,” says Joe Kugler, vice president of Worldwide Technology Services and Operations at Electronic Arts. “We wanted a ubiquitous platform, like Intel processor-based servers, because it could drive reuse capabilities. Plus, we could get that hardware for a great price and reduce our ongoing maintenance costs.”

Furthermore, because standardized servers can be grouped in more efficient cluster configurations, procedures for deployment, upgrade, and management can become more universal and less complex. Standardized clusters are not only inherently more powerful and redundant, but they also enable faster reallocation of resources, thanks to simplified multinode cluster management tools.

“Technology is now on our side,” agrees Garry Egan, CTO at GBCblue, a premier provider of business center solutions for the hospitality industry. “When we work within this building-block framework, we have the tools available to create a fully managed and fully functional environment.”

Ease on down the road

Now comes the nitty-gritty. You know what scalability is, and you know how enabling it can significantly improve your ability to cut costs and boost agility. But how do you implement this standards-based strategy? Here are a few basic guidelines to send you on your way:

- **Scale down in size and complexity.** Don’t be fooled—scalability is not always about scaling up. In fact, most analysts these

“A key trend is to implement tools that can link, monitor, and manage the entire IT value chain”

days espouse methodologies for reducing, or *abstracting*, the infrastructure so that it can actually do more with less. IDC’s Gens refers to this idea as a vision of a “service-oriented architecture based on the abstraction of various key components in the computing environment, enabling greater levels of flexibility, optimization, and control.”

The use of industry-standard components facilitates this process by giving companies a set of increasingly smaller, yet more powerful, building blocks on which to quickly and easily extend the entire data center architecture. Consider starting with components such as standardized rack-dense servers. Dell recommends, for example, the Dell™ PowerEdge™ rack-optimized server lineup—which includes the PowerEdge 1850, 2850, 3250, 6650, and 7250 models—to help meet requirements for maximum performance optimization even in a limited amount of space. To further simplify the rapid deployment of these industry-standard systems, several of the new PowerEdge servers actually feature the same system board design, allowing customers to leverage a single system software image and speed up the actual deployment. Storage components should mirror this capability, offering starting points with room to expand as needed. Dell suggests the Dell/EMC CX series of networked storage arrays and Dell PowerVault™ tape libraries as core building blocks.

- **Improve resource utilization.** “Even though dynamic IT architectures need to grow the number of logical connections among IT resources and the business processes they support, they are also about eliminating—or hiding—as many physical interdependencies as possible to maximize flexibility,” says Gens. Two ways to demonstrate this concept are clustering and virtualization. Both options further enhance the processing and performance power of the standardized systems you deploy by distributing workloads for more efficient resource allocation and generating an overall improvement in bandwidth availability—potentially resulting in lower costs and greater productivity among users. Clustering may take the form of high-availability clusters with a handful of server nodes or high-performance computing (HPC) clusters with hundreds or even thousands of nodes. Virtualization typically includes deployment of special VMware® or Microsoft Virtual

Server 2005 software to enable deployment of multiple “virtual machines” across the server nodes.

- **Boost the I/O bandwidth.** Upgrading your network infrastructure from a mix of slower 10 Mbps and 100 Mbps connections to full Gigabit Ethernet connectivity makes sense, because many desktops, workstations, and notebooks include embedded Gigabit Ethernet capability—making it more efficient for them to communicate with application servers. Of course, you’ll also want to fortify the network with Gigabit switches to get the bits flowing at full speed. The Dell PowerConnect™ line of Gigabit Ethernet switches for the data center to the wiring closet are at the ready, and the supplied Dell OpenManage™ Network Manager software can help you keep them that way.
- **Manage everything.** Last, but certainly not least, you need the ability to manage and support your streamlined infrastructure centrally and easily in order to remain agile and scalable. “A key trend,” says Gens, “is to implement tools that can link, monitor, and manage the entire IT value chain. This refers to all IT operational elements that support a business activity—from hardware and system software to business applications, data, workflow, and business process.” The Dell OpenManage systems management framework and software components, for example, make it easy for network and system administrators to take control of and gain visibility into large server environments and multinode computing systems featuring PowerEdge servers. And the latest release, Dell OpenManage 4, further facilitates the incorporation of industry-standard components in the data center by including new, built-in Intelligent Platform Management Interface (IPMI)—compliant baseboard management controllers that allow administrators to monitor servers remotely using their choice of systems management software. As more vendors are turning to IPMI-compliant tools, the world of server management also grows into a standardized environment where administrators can be more productive and more proactive.

Putting the pieces together

Assembling the puzzle that is the enterprise data center can seem a daunting task, especially in today’s economic climate, where budgets are still tight and every item requires specific purchasing

Q.&A.

We asked IT professionals from GBCblue, a premier provider of solutions for the hospitality industry, and Input/Output, Inc., a seismic technology provider, for their perspectives on the meaning of scalability. Here's how they define the buzzword in the trenches.

Q. WHAT DOES SCALABILITY MEAN TO YOU?

A. Garry Egan, CTO at GBCblue:

“One word: *automation*. Most people think of scalability merely from a hardware and/or physical capabilities perspective. For example, bandwidth, CPU cycles, load balancing, redundancy, electrical power, and HVAC are common scalability topics in the IT industry. But IT solutions and management are a combination of both hardware and software—which means that, in order to scale, you need to automate and simplify as many processes as possible to maximize and maintain a stable IT environment. And I am not just talking about ‘5,000 nodes’ scalability; I am referring to 100,000+ nodes over a worldwide network where any type of manual process quickly becomes unmanageable. Process automation in an IT environment is not easy, but you can attempt to control your own destiny with a good quality assurance (QA) process and expert-level staff that can troubleshoot issues beyond the surface layer.”

A. Albert Chidiac, director of IT, Input/Output:

“Scalability has both a technical and an enterprise meaning. Technically speaking, my applications need to be scalable to meet my company's current and future needs. Therefore, my engineers need to architect a technical environment for each application that will grow as the business needs grow—and many great technologies exist in today's marketplace that offer my applications this scalability.

“But my IT department also has to be scalable within the enterprise. As my company grows or changes its strategy, my IT department needs to grow and change without massive retooling efforts. Any workable definition of scalability must encompass this requirement as well.”

justification. But scalability is so important that most CFOs would be hard-pressed to contest an argument for a standards-based strategy. Cost and performance benefits aside, interoperable standardized components are the new innovations of the technology world, demonstrating the latest and greatest engineering feats designed to address the latest and greatest business needs.

Q. HOW IMPORTANT IS IT TO BECOME A SCALABLE ENTERPRISE TODAY?

A. Egan:

“Most would rank scalability as number one in terms of importance in an IT solution. I would rank it second, with stability of the core system first. That being said, however, scalability is necessary for growth—especially on a global scale where bandwidth must traverse the trans-Atlantic and trans-Pacific Internet pipes. In 2000, much focus was placed on bandwidth as being the scale point. Now, bandwidth is cheap—so automation is the next scalability point.”

A. Chidiac:

“Scalability should be a top priority on any IT department's to-do list. To compete in today's market, companies need to grow and change business strategies, and their IT departments need to be integral to corporate growth and strategic development. CEOs cannot wait long for an IT department to reengineer its structure and processes after an acquisition. A change of business strategy is implemented quickly, and the IT department needs to support the strategy immediately to ensure corporate success.”

Q. WHAT'S THE BEST ADVICE YOU CAN OFFER A COMPANY IN TERMS OF ACHIEVING SCALABILITY?

A. Egan:

“Know code. Almost every piece of hardware today—from cars to toasters—has a computer chip. Behind every computer chip is code to operate it, and this is the basis of all technical systems. The Windows operating system is based on code. Routers and switches are based on code. Scalability is success. Scalability is dependent on automation, and automation is derived from—you guessed it—code.”

A. Chidiac:

“You will need to partner with a visionary vendor that has the resources and knowledge to support you in building your entire enterprise to meet today's and tomorrow's business challenges. This vendor needs to support all aspects of your IT department, not just one area. Input/Output partnered with Dell and has been successful in meeting its global business challenges.”

And admit it: There's more to some buzzwords than meets the ears. Once you understand what they mean, words like *scalability* actually play a tune worth listening to. **D**

¹ This term does not connote an actual operating speed of 1 Gbps. For high-speed transmission, connection to a Gigabit Ethernet server and network infrastructure is required.