White Paper

The Scale-out Storage World, and Intel Architecture's Place Within It

By Mark Peters, Senior Analyst

January, 2012

This ESG White Paper was commissioned by Intel and is distributed under license from ESG.

Contents

Introduction: Scale-out Storage and Intel Architecture

You would have to live under a pretty large rock to be completely unaware of the "Intel Inside" marketing and branding campaign. But the breadth of Intel's intent and capability to be "inside" might come as something of a surprise. After all, this paper has both scale-out storage and Intel Architecture (IA) in its title. Really? To most people, "Intel Inside" intuitively applies to the company's processors being found inside computing devices such as laptops and servers.

When the word "storage" is uttered, the minds of few people—even people in IT—leap immediately to think of Intel. However, that is something that Intel looks to be about to change, not only because Intel is a pragmatic and aggressive commercial organization, but also because the necessity for advanced, intelligent, and powerful computing in storage—especially today's increasingly dominant *scale-out* storage approach—means that storage is becoming more processor dependent.

Thus, IA has a key role and opportunity in the storage market. As storage infrastructures become more heavily based upon industry-standard X86 servers, so IA and storage become increasingly interlinked. "Intel Inside Scale-out Storage" may not yet be a phrase that trips off the tongue. But compelling reasons exist for users to be aware of the important relevance of processing in today's increasingly complex storage world.

Moreover, Intel wants to be seen as more than just a supporting cast member in the storage world; it wants to create the same sort of integrated value that it has achieved with its other, more traditional processing platforms. As the company states:

> *"The bottom line: The world needs storage, and it needs it in new ways. Intel® architecture-based storage solutions can support scale-out storage solutions that address your immediate need for efficiency and prepare you to meet tomorrow's challenges. Intel® Xeon® processors can help by providing more energy-efficient storage performance, delivering added data protection, and enabling new usage models and greater scalability."*[1]

This paper covers Intel's new focus from two perspectives. It contains two sections, each readable as a standalone piece.

- **Section A** examines and explains scale-out storage architecture, focusing on its emergence as the leading approach to contemporary storage and on the IT and business values it can deliver.
- **Section B** looks in depth at the specific IA capabilities that can address scale-out storage needs, emphasizing their operational and economic value.

Why does this matter? First, because modern storage generally, and scale-out storage in particular, put new demands upon the processing capabilities that support them, so it's important to know what underpins a solution. Second, the adoption and growth of scale-out storage as the modern architecture of choice means that the majority of users are impacted, and they'll need to understand all this so they can make good planning and purchasing decisions.

Indeed, ESG has delivered a scale-out storage market forecast[2] through the year 2015, at which point we expect the transition to scale-out architectures to be almost complete, with scale-out storage architectures making up 80% of all net-new external networked storage shipments on a revenue basis. From 2010 to 2015, scale-out storage continues its march into commercial IT and away from just the niche markets that were the first to require the bandwidth and performance it provides.

In that same time frame, our market forecast for scale-out external networked storage shipments in commercial and government sectors increases from 4,189 petabytes (PB) to 62,834 PB. That is a 72% compound annual growth rate (CAGR) that far exceeds the overall external networked storage growth rate. This significant growth is based on a near-wholesale transition from scale-up to scale-out architectures across the vendor landscape—and *that* makes understanding scale-out, together with IA's role within it, a matter of broad importance across IT.

[1] Source: Extract from the Intel Solution Brief, *Gain a Business Advantage with Scale-Out Storage*, 2011.
[2] Source: ESG Research, *Scale-out Storage Market Forecast 2010-2015*, March 2011.

Section A: The Evolution and Nature of Scale-out Storage

Market Background

The demands placed on today's storage systems are extremely different from those placed on storage two or three decades ago. Yet that was the era of the birth of the architectures, code, and supporting processors for many "modern" storage platforms. This isn't really an issue of whether those older platforms are capable of dealing with today's demands. Clearly many still get the job done. The issue is at what cost and with what level of ease, efficiency, and flexibility can they do the job? A new world requires a new way to do things.

The new demands are underpinned by the exponential growth in data driven by the Internet, by new users, by new applications (especially ones tied to science, entertainment, weather, genomics, and communications), by mobility, and by rising expectations. But there's more than sheer data volume in play. All that stored data must be budgeted for, managed, protected, secured, and be made increasingly available across multiple platforms and interfaces.

To address these demands, the traditional monolithic, discrete storage system is rapidly giving way to a more flexible, pragmatic, and converged storage approach called "scale-out." So what is it all about? And what was wrong with the old methods?

Scale-up Versus Scale-out Storage

The Architectural Concept

Scale-out storage certainly sounds like a promising cure for the multiple and growing storage challenges in both "regular" and server-virtualized data centers. But what exactly is it? Conventional solutions have been based on "scale-up" architectures (sometimes called monolithic or frame-based architectures) that have design roots reaching back to early development by IBM in the 1960s. Although they are still very common in many commercial environments, those architectures were designed to meet the needs of the mainframe world.

Scale-up storage sits behind a dedicated controller. Upgrades and capacity expansions to meet changing business conditions or keep pace with growth can present challenges, especially in maintaining acceptable system performance as the changes are made. Increasing capacity (and, therefore, the IO workload) without making a corresponding increase in bandwidth and storage processing power can decrease response times and service levels. [And it's also important to note that a malleable scale-out approach can help to better manage the cost challenges that a less flexible storage approach imposes (both in sense of the timing and amounts of expenditure). We'll revisit this aspect in more detail later]. A tactical workaround is to over-provision, but that is obviously a very wasteful approach. Furthermore, server virtualization exacerbates problems with scale-up architectures. Disparate "islands" of storage; poor utilization; complexity; and inefficient, inflexible, disruptive scalability are often the painful side-effects of scale-up architectures. They may also require specialized and complex management and encourage vendor lock-in. And "scale up" must happen when any one of the key components (disk, processor, network, etc.) is maxed out. It is the IT equivalent of traffic moving at the speed of the slowest vehicle.

Modular, Monolithic, and Scale-out

Traditional storage architectures have primarily been classified in two ways: as modular or monolithic. Both architectures are fixed-scale systems. Essentially, a single- or dual-controller system is loaded with disk capacity (either HDD spindles or solid-state drives) until maximum capacity is met. At that point, a new discrete system is added to increase capacity further.

- Monolithic systems typically support open systems and mainframe environments. They are massive, enclosed in a proprietary data center rack, and feature a built-in single or dual controller that fills with disks over time.
- Modular systems have one or two controller modules and self-contained shelves of disks that are housed in, and added to, a standard data center rack to increase capacity.

Scale-out Storage Defined

Although it can be elusive to define scale-out storage in a way that is acceptable to all industry participants, ESG uses the following definition. We believe it encompasses most of the core concepts of scale-out storage architectures:

"Scale-out" storage solutions increase performance, capacity, or throughput by adding resources (e.g., processors, memory, host interfaces) as a loosely coupled system composed of nodes that work side-by-side, in parallel. These systems differ from traditional "scale-up" storage solutions, which increase performance, capacity, or throughput by adding resources within a tightly coupled system that shares a common pool of resources that work in tandem.

Figure 1 offers a visual representation of the difference between scale-up and scale-out storage.

Figure 1. Scale-up vs. Scale-out Storage

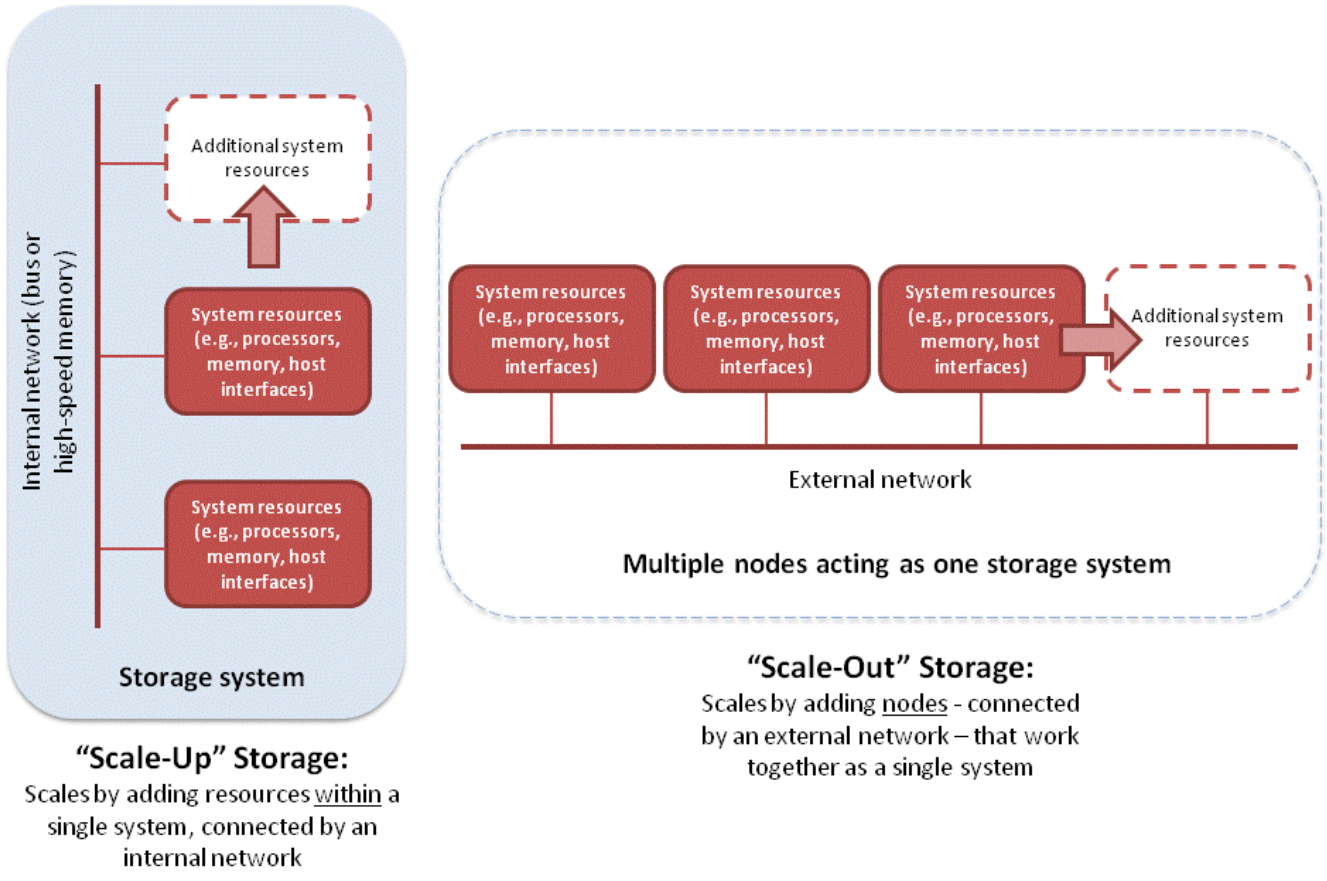

Source: Enterprise Strategy Group, 2011.

Scale-out Storage Advantages

Scale-out storage systems overcome the physical boundaries of racks and modules. They are virtualized systems whose front-end processing power and back-end capacity grow by adding new processor or capacity nodes, while they still function as a single system. They often don't need the complex individual physical disk and data-layout management, or the performance tuning, required by traditional monolithic and modular systems. Scale-out systems offer economic advantages as well, not just through better use of system resources, but also via their use of (sometimes) less-expensive hardware and standardized components. And very importantly, implementation time can be dramatically shorter as well, resulting in faster storage capacity deployment and, ultimately, faster application deployment. Naturally, all this integration, flexibility, and functionality demands much greater processing power as well – this is something that will be covered in greater depth in part B of this paper.

Scale-out storage's ability to deliver increased performance *and* operational cost reductions has put these systems on the radars of an increasing number of enterprise IT buyers. Scale-out solutions can expand (granularly, of course) into the multi-petabyte range under a single system image. That means they provide an ideal platform for storage consolidation. They also help IT organizations reduce management costs, data center floor space footprints, and power and cooling costs. And of course, increasing storage consolidation onto a shared resource also increases utilization rates, giving users more "bang" for their storage buck.

A key budgetary benefit of scale-out storage is its cost-related predictability. IT managers know that whenever they want "X" TB and "Y" IOPS, they will pay "Z" dollars. ESG conducted an in-depth study of scale-out architectures,[3] and a subset of the findings provides insight into the IT and business values that scale-out storage can provide:

- **Adoption is growing fast.** As of late 2010, 18% of survey respondents reported they were using scale-out storage. But a further 40% were planning to deploy it by the second half of 2012.
- **Use of scale-out storage is—interestingly—correlated to users who have a more strategic view of storage.** The use of (or the plan to use) scale-out storage runs at 64% among those who view storage as "strategic." It runs at only 31% among those who view storage as "an afterthought."
- **Scale-out adoption in enterprise IT is driven by improved scalability and performance** together with lower infrastructure and management costs and faster provisioning.
- **Implementation time for a new storage system (from "need" to "use") is under three months for 38% of scale-out storage users.** Only 13% of users not using scale-out enjoy such a short implementation period. Conversely, 21% of scale-out storage users report implementation times of more than six months, compared with 61% of those not using scale-out storage (see Figure 2). The fact that 61% of traditional storage users endure implementation times exceeding six months underscores how dramatically IT projects such as server virtualization can be slowed by inflexible storage.

Figure 2. Length of Implementation Cycle for New Storage Systems, by Scale-out Storage Usage

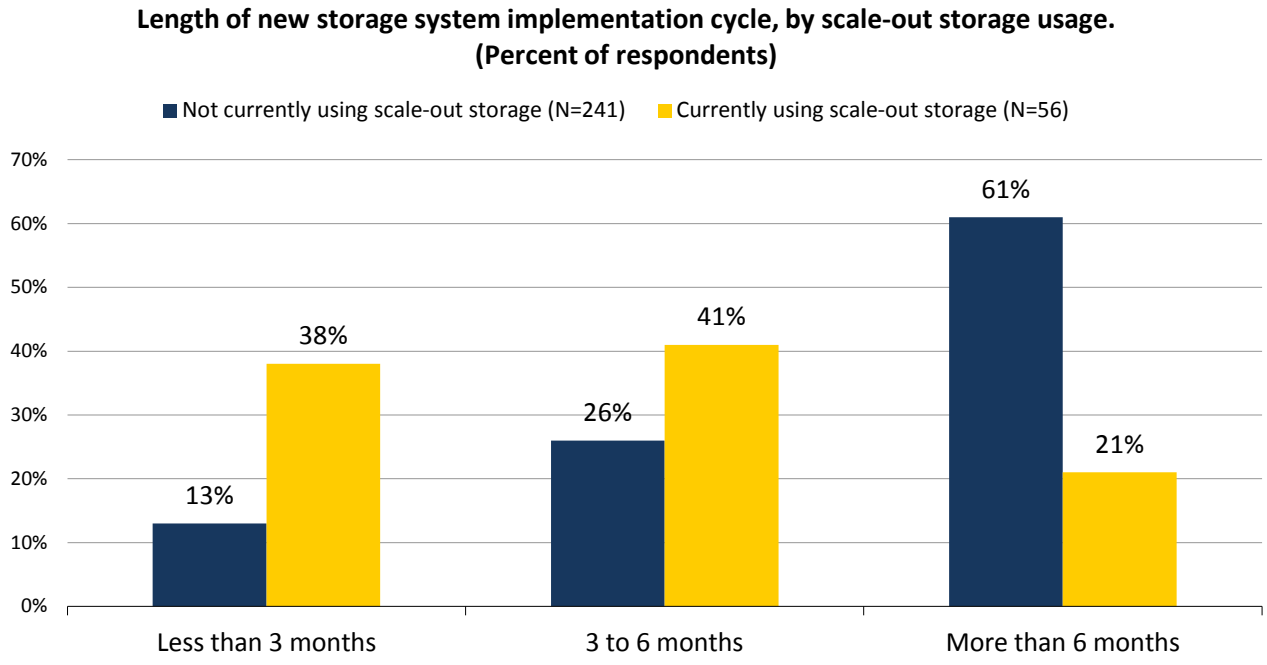

Source: Enterprise Strategy Group, 2011.

[3] Source: ESG Research Report, *Scale-out Storage Market Trends*, December 2010.

Finally, there is a common misconception (probably based on common examples and practices found in the market) that scale-up is a SAN-only approach, and that scale-out is synonymous with NAS environments. Neither is true. Scale-out is an overall architectural approach to delivering storage capabilities and is separate from the protocols employed. Scale-out architectures are flexible and scalable in performance and capacity, which means they can handle aggressive capacity growth, unpredictable workloads, and varying bandwidth requirements. And relatively simpler management means that, as storage grows to meet business requirements, addition of IT staff is not a requirement. More can be done on a lower budget.

The Value of Scale-out Storage in the Contemporary IT World

Scale-out Storage and Virtualization

The increased use of server virtualization is a major driver of scale-out adoption. A mismatch exists between the flexible, mobile capabilities of virtual server environments and the hard-mapped, stove-piped nature of traditional scale-up storage.

Scale-out storage has performance, flexibility, and efficiency attributes that address the challenges of traditional *and* server-virtualized environments. It helps users easily respond to rapidly changing business requirements with transparent scalability and simplified management—something that is impossible with traditional storage approaches.

Indeed, if we look at the aspects of storage infrastructure that users see as prerequisites for their broader use of server virtualization (see Figure 3), then a scale-out approach is highly attractive, especially in terms of allowing faster storage provisioning and, obviously, a scalable infrastructure.[4]

Figure 3. Top Five Storage Infrastructure Developments Required for More Widespread Server Virtualization

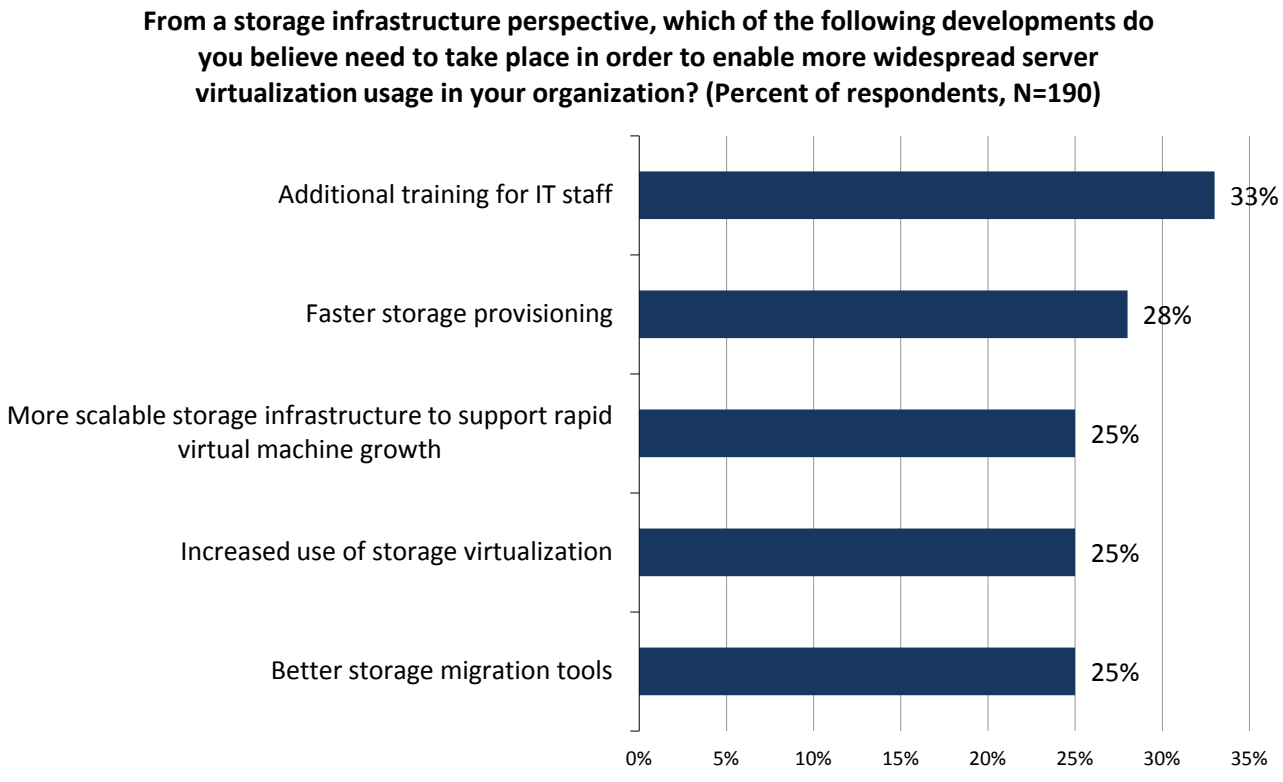

Source: Enterprise Strategy Group, 2011.

[4] Source: ESG Research Report, *The Evolution of Server Virtualization*, November 2010.

Server virtualization puts considerable performance and capacity-related stress upon storage systems. Because scale-out's adoption and extended use are such significant trends right now, it is worth delving a little deeper into this topic. Server virtualization software is now mature, and its adoption is widespread across organizations, as ESG research has confirmed.[5] When they were asked, "Is your organization currently utilizing x86 server virtualization technology?" 61% of 1,602 respondents said "yes," either for production or production and test/development. A further 13% said "yes" to test/development only, and yet another 12% were evaluating or piloting the technology.

While that information shows that the vast majority of users have at least some level of server virtualization implemented, it does not indicate the *extent* of its use within these organizations. ESG research (Figure 4) addressed this question by looking at the percentage of x86 servers that have actually been virtualized in organizations. The curves are general trend lines that demonstrate the current generally lower percentage levels of virtualization (blue arc) and users' expectations for an increase in higher-percentage levels of virtualization in the near future (yellow curve).

Figure 4. Percentage of x86 Servers That Have Been Virtualized

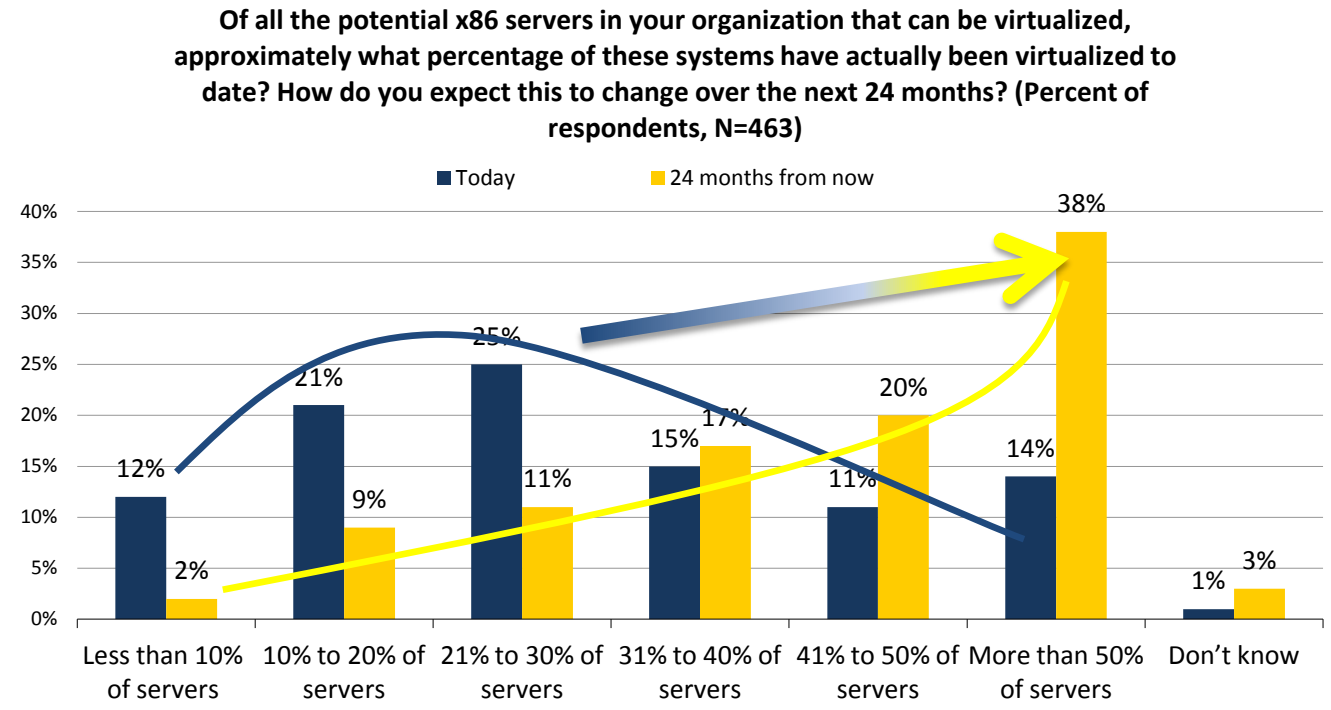

Source: Enterprise Strategy Group, 2011.

In other words, there is an awful lot of virtualization still to happen. As users move from "dipping their toes" in x86 server virtualization and progress to higher levels of usage, they should—as soon as possible in the process—be deciding upon and implementing an appropriate storage infrastructure.

This is crucial because the increasing adoption of server virtualization often comes with a concurrent realization that the limitations of traditional monolithic or modular storage architectures can significantly hamper a virtualization effort's success. In short, storage—as we have traditionally known it—is unable to support the performance or scaling needs of virtualized servers.

In fact, storage is becoming a throttle, or even an anchor, on the promise of virtualization.[6] Scale-out storage – supported by, if you will, scale-out storage-processing abilities—is a tool that can help address this challenge.

[5] Source: ESG Research Report, *The Evolution of Server Virtualization*, November 2010.

[6] For a much deeper investigation of this issue, see ESG's recent Market Report, *The Future of Storage in a Virtualized Data Center*, January 2011, which also details the value of Ethernet-based protocols for computing in general.

The Bigger Truth

Scale-out storage is something that, while complex to design and build well, is simple in concept and operation. The vast majority of the heavy lifting is borne by the vendors; the vast majority of the benefits are enjoyed by the IT users. It is possible to (albeit rather cynically) argue that traditional scale-up storage was, and is, just the reverse.

In both cases, the architectures can be effective. The big difference arises in the realm of efficiency. Scale-out architectures move storage in line with the contemporary IT world: It is flexible, is easy-to-use, and can drive both a TCO advantage (by requiring less overall equipment and less specialized management, for example) and an ROI improvement (providing more responsiveness and more speed to the organization). Of course, to support scale-out in an optimal fashion also needs (and here we can see the interest, and opportunity, for, Intel) considerable raw compute resources, as these are the resources that make a flexible, responsive and dynamic system, where parameters can change variably, actually work. A good scale-out storage system is indeed more akin to being a flexible server environment in many respects than it is simply a flexible box of storage components.

What do actual users say about scale-out storage? They confirm the value that scale-out storage can deliver:

- **Scale-out delivers on its promise of improved scalability and performance.** Scale-out early adopters have reduced their costs while improving scalability, IO performance, throughput performance, data availability, and hardware utilization.
- **Scale-out storage users are more likely to report having a sufficient number of storage administrators and significantly shorter implementation cycles.** Sixty-four percent of current scale-out storage users believe that they have sufficient staff to manage their storage environments, compared with 50% of organizations not using the technology. Seventy-nine percent of scale-out storage users have an average storage system implementation cycle of six months or less, versus only 39% of organizations that have not deployed scale-out storage.
- **Once users see scale-out storage in action, they quickly expand usage.** While users often initially employ scale-out for tier-2 data, they quickly embrace its capabilities, then they broaden and deepen their use of it.

With its suitability for a world of massive data growth (which demands an economically viable and easy-to-scale approach) and massive operational uncertainty (which demands a flexible and easily managed storage resource pool), scale-out has become the storage architecture of choice, delivering logic, simplicity, and value to an IT world crying out for all three.

Section B of this paper looks specifically at the impact of Intel Architecture on scale-out storage.

Section B: Intel Architecture in Scale-out Storage

Scale-out Storage

Section A of this white paper evaluated the relevance and importance of scale-out storage. It concluded that scale-out is suited for a world experiencing massive data growth (which demands an economically viable and easy-to-scale approach) and massive operational uncertainty (which demands a flexible and easily managed storage resource pool). Scale-out has become the storage architecture of choice, delivering logic, simplicity, and value to an IT world crying out for all three.

End-users have investigated and increasingly adopted scale-out storage, and vendors have responded accordingly. As emerging scale-out vendors such as Isilon, DataDirect, 3PAR, and BlueArc benefitted from the end-users' interest, major storage incumbents such as EMC, HDS, and IBM either acquired those firms, evolved their own offerings, or introduced new scale-out solutions.

With the pull-and-push factors at play, ESG anticipates that by 2015, scale-out storage architectures will comprise 80% of worldwide storage shipments.[7] We are witnessing a near-wholesale transition from scale-up to scale-out architectures across the storage landscape, which makes close examination of the prerequisites for good, effective, optimized scale-out storage important. It also makes understanding the impact of Intel Architecture (IA) upon this paradigm shift simultaneously intriguing and necessary.

Why the Underlying Infrastructure Matters

To set the foundation for what IA can contribute to scale-out storage, it makes sense to start by examining why the raw foundations of a storage platform itself matters. Scale-out storage architectures use a networked collection of components in a peer architecture. The components collaborate as a pooled resource, delivering maximum performance, advanced utilization, economic efficiency, non-disruptive scalability, and robust data protection.

The underlying elements are modular, but that modularity is transparent to users: The system always has one logical management interface. Simplicity and elimination of management complexity are signatures of scale-out storage architectures. But, like a duck appearing to glide effortlessly across a pond yet paddling away furiously under the surface, scale-out works hard to deliver a simple, unified, efficient user experience.

This is where processing power *and* sophistication come into play; and it is why IA has relevance in a scale-out storage discussion. The tie-ins relate to:

- **Raw processing headroom:** The general move of storage controllers away from specialized hardware toward standard hardware is manifested in the move from application-specific integrated circuits (ASICs) to software that runs on servers that are pretty much standard. While the move from purpose-built processing chips to general-purpose chips improves economics, adds flexibility, and eases upgrades, it typically comes with a need for more power and flexibility which are things that you don't always get - either often or enough - when a storage implementation is tied to customized hardware.
- **Advanced storage functions:** Users are demanding ever-more advanced storage-optimization features. However, these typically sophisticated capabilities—such as thin provisioning and deduplication—will not be implemented if they hurt performance. They represent compute-intensive workloads *within* the storage infrastructure that demand additional processing power. A host of technologies fall into this category: thin provisioning, data-reduction technologies such as deduplication and compression, automated failover, erasure coding (which distributes private data across multiple data centers and/or public cloud storage for added protection), encryption, auto-tiering, extended metadata, and caching.

[7] Source: ESG Research Report, *Scale-out Storage Market Forecast 2010-2015*, March 2011.

All of them are compute intensive. Moreover, all are very popular, as the broad-based ESG research data in Figure 5 shows. These desired features require significant compute power, but all are either absolutely required or strongly preferred as part of any purchase.

Figure 5. Importance of Various Features/Functions When Considering Scale-out Storage Platforms

Please rate the importance of the following features/functionality in considering a scale-out storage platform. (Percent of respondents, N=178)

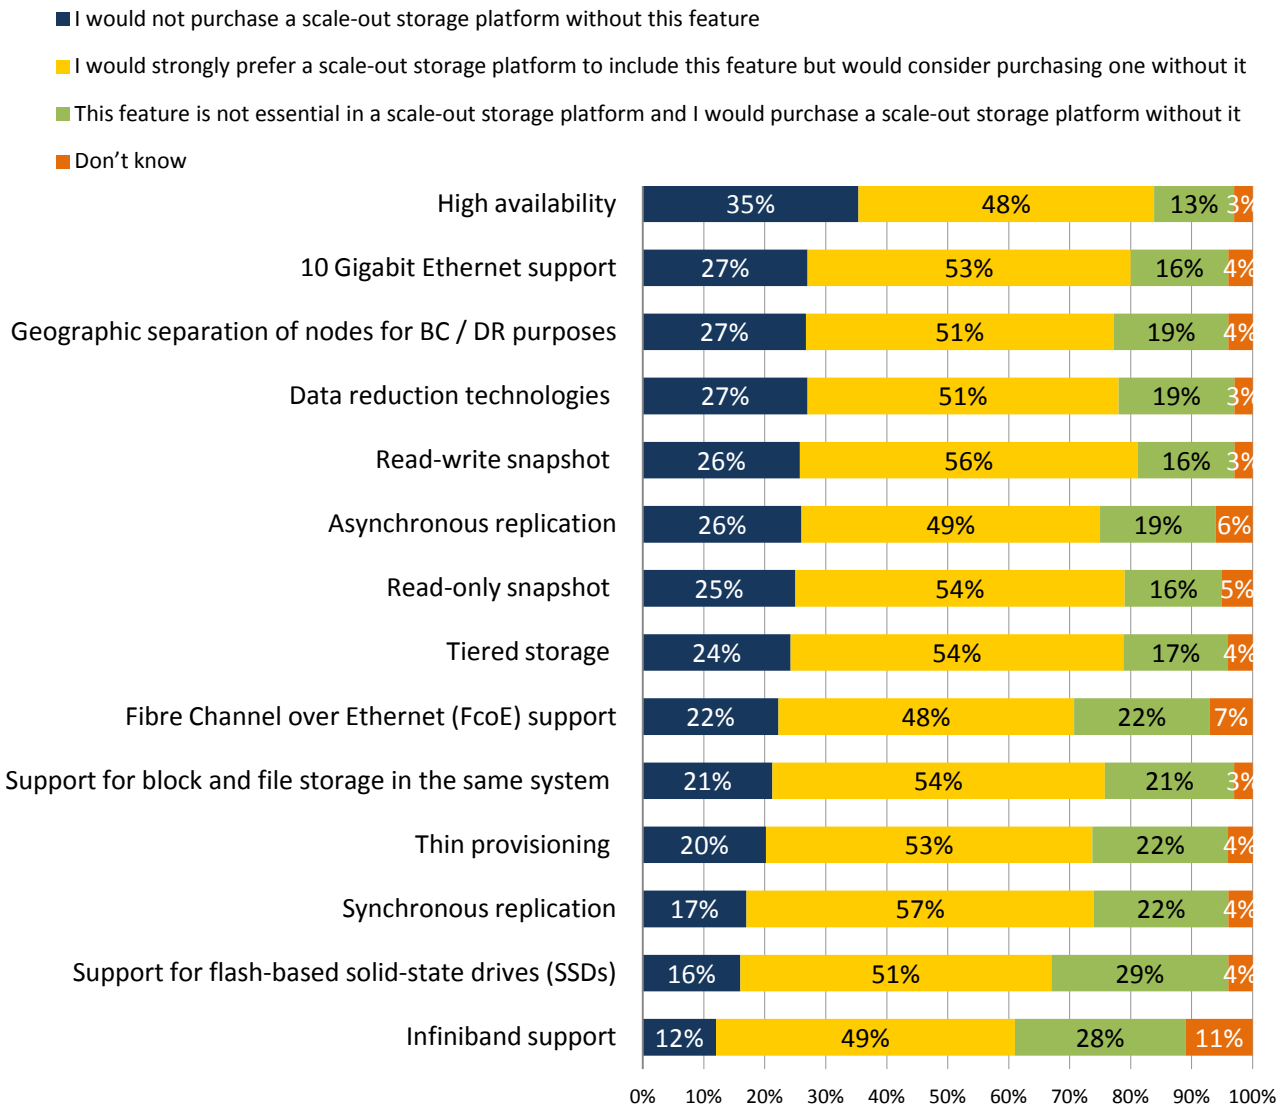

Source: Enterprise Strategy Group, 2011.

Other Needs and Variables

- Because of the random nature of much IO, the link between having an "X% better processor" to particular, predictable "Y% storage benefits" is not as direct or linear as it is with servers and applications. Again, this means more processing power is vital, both in sheer "horsepower" and in terms of flexible packaging and built-in functionality that is well integrated with server systems and user needs.
- Many users have electrical power limitations—or at least they aspire to reduce their electricity use. Being energy-frugal is required in order to move advanced storage functions into scale-out storage.

- Concerns about the security of corporate data are leading to the broader implementation of technologies such as encryption, driving a further need for higher levels of computing power in storage.

All of these needs are being compounded by the overall complexity and IO intensity of many now-common applications for data discovery, data mining, data warehousing/business intelligence, and analytics. Compute power and sophistication are needed, but they must also be economically suitable.

IA Delivers Value to the Scale-out Storage World

IA can support all those infrastructure-related needs, both generally for IT and specifically for scale-out storage. in regard to general IA capabilities, Intel has committed to (and so far has delivered against) what it calls the "tick-tock model" of computing power. At least every two years, it is developing faster cores with higher performance and more efficiency. The company's credibility and continued ability to deliver is impressive. Intel also is active in the networking arena and in specific storage media, notably non-volatile memory (NVM), which shows great promise for wider adoption as the economics become more attractive.

The convergence of items such as these—combined with standards-based implementation—boosts the relevance and value of IA in the scale-out storage world, especially as many in the industry (including Intel) expect a future where storage is more of a server-based service. We can already see signs of this happening with the general adoption of x86 *into* storage systems and in the emergence and growth of hypervisor-based storage control.

Looking more specifically at some of the things IA can deliver in the scale-out storage world, we see that:

- IA offers plenty of processing power for "raw headroom" needs and for supporting compute-intensive encryption, deduplication, and other tasks. This same power is valuable in virtualized server environments that put stress on storage subsystems and demand higher performance. As the "scale" of scale-out grows, so does the need for strong compute-power foundations.
- IA has attractive economic and energy-efficiency characteristics. For example, Intel Xeon processors can invariably provide the growing compute power required by scale-out storage—without increasing the watts consumed.
- IA has integrated advanced storage features directly into its chips to make data protection (accelerated encryption via AES-NI) and high availability (memory parity, Asynchronous DRAM Refresh (ADR), Non-Transparent Bridging (NTB)) simpler, without sacrificing performance.

It's a fairly simple equation: In order to drive-up and enjoy the capabilities of scale-out storage, there is a need for scalable computing power as well. This power helps to integrate storage features with the IO and memory so as to support the sophisticated usage models that enable storage efficiencies to be achieved.

Sample TCO Considerations of IA in Scale-out Storage

With economics underpinning everything in storage, it is important to try to gain at least some idea of how much TCO improvement IA can bring to a scale-out storage world. To do this, we will use proof rather than mere assertion by referring to ESG research findings[8] that inform us of real-world results. By adding some (open and declared) assumptions regarding cost, we get an indication of the potential TCO impact of scale-out with IA.

- **Management and efficiency:** 17% of organizations with more than 100 discrete storage systems have at least six full-time storage employees, compared with 3% of those with 25 or fewer systems. If IA/scale-out can reduce the number of discrete systems and management effort/complexity, then perhaps multiple employees can be redeployed. With burdened costs likely to exceed $100,000 per storage administrator, this could reduce TCO dramatically.

[8] Source: ESG Research Report, *Scale-out Storage Market Forecast 2010-2015*, March 2011.

- **Faster implementation cycles:** The essence of this consideration is shown in Section A, Figure 2. More than three-quarters (79%) of scale-out storage users have an average storage system implementation cycle of six months or less, as opposed to only 39% of organizations that do not have scale-out storage deployed. Clearly there is a TCO effect tied to personnel costs, management time, and opportunity costs. However, financially, this is as much about ROI impact as TCO. Faster implementations mean expenses are incurred more in line with business needs and/or that new applications and business endeavors can come online more quickly. The former might be worth a few (tens of) thousand dollars, depending on the size of the organization. The latter could be worth six or seven figures in time-sensitive competitive situations.
- **Scalability:** Respondents were asked: "When your organization purchases and deploys new storage systems to accommodate application data growth, what is the most common reason you have added those new systems?" Only 10% of respondents claimed to need improved performance. Only 8% needed availability. At 25%, the number one reason users added storage was: "Needing additional capacity and we can't expand existing systems." Assuming that this requires management time and effort, as well as (most likely) the effective waste of some existing resources (they need capacity, not performance), then the TCO impact for a medium-sized organization could easily be in the six-figure range. Scale-out would not only help to manage growth in a more granular and gradual fashion without "forklift upgrades," but it would also probably enable better overall resource utilization, resulting in a general infrastructure cost reduction.
- **Staff levels:** When looking at important TCO factors, most users stick to the basics of infrastructure-related and staff-related costs. That makes the information in Figure 6 notable.

Figure 6. Ability to Support Storage (Scale-out vs. Non-scale-out) with Available Trained Staff

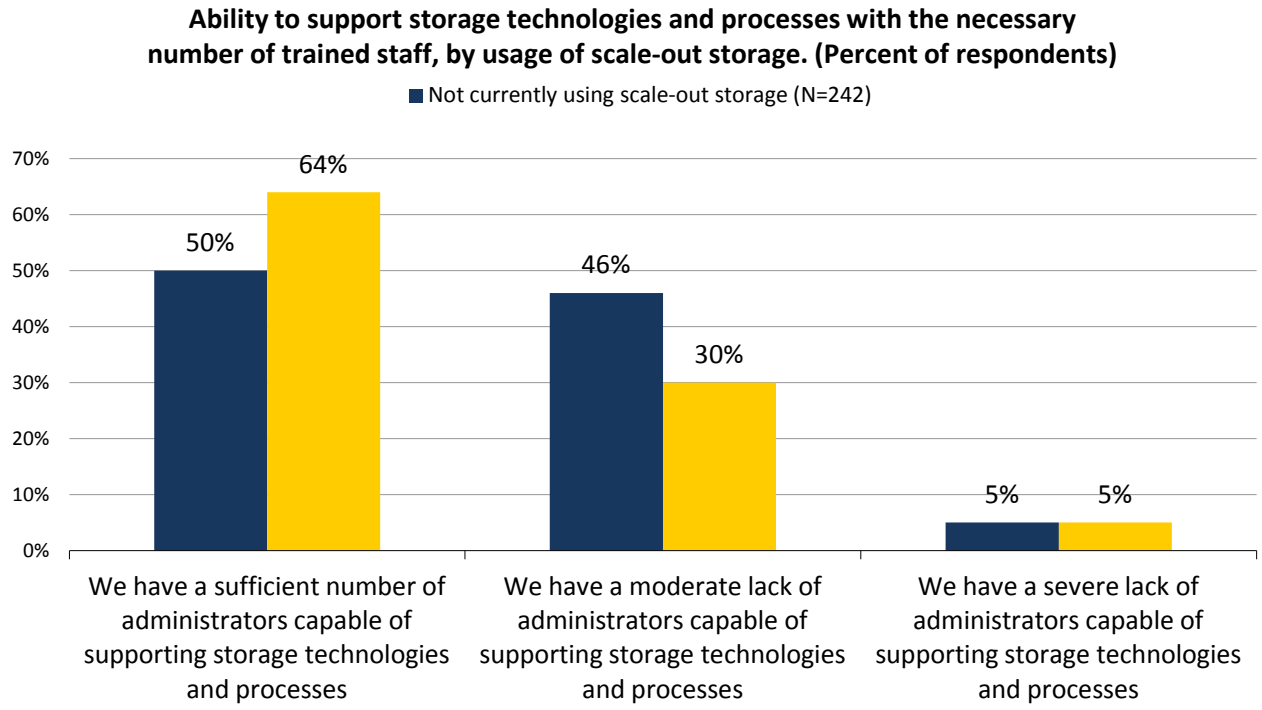

Source: Enterprise Strategy Group, 2011.

Users of scale-out storage are considerably more likely to feel they have a "sufficient" number of administrators and considerably less likely to feel they have a moderate lack of capable administrators. Whether this reflects a higher comfort level and less-complex training with scale-out systems or merely a lack of administrators in non-scale-out environments, the effect is probably to require more staff-related expenditures (raising TCO by $100,000 or so for each one).

These TCO ideas are just examples. This discussion intentionally keeps to the very basics and does not try to account for the value of the advanced features and functions that IA is also designed to enable in scale-out storage. What is clear, however, is that the positive impact of even these "vanilla" TCO aspects can be significant.

The Bigger Truth

The summary here is pointed and easy. Scale-out storage is a clear, attractive, and proven option for much of the user community. And IA adds specific value in terms of supplying the necessary compute power to support scale-out in general and key advanced storage functions in particular.

While Intel has been somewhat quiet in this area, it is now declaring itself. There appears to be several reasons for this:

- First, there's the **obvious potential commercial gain** in associating itself with a significantly large and growing market. Not only does Intel bring a solid brand reputation to bear, but its processor role in storage delivers advantages that can have measurable impacts on IT and on a business.
- Second, **Intel wants to enhance storage** (and support the storage market/providers) via increasingly specialized chips that add specific value (such as the security play)—and that ultimately could separate Intel from its competitors.
- Third, Intel **has a very definite view of the future** in which IT is run "holistically" on a unified pool of physical compute, network, and storage resources, with storage available as a service via server-based software. The idea of "converged storage" on standard server hardware is not unique to Intel of course. But by adding things such as its security, networking, and NVM capabilities, it is exceedingly well positioned to become a key player, not just a key participant.

Promoting the value of IA in scale-out storage is a very deliberate first step in Intel's broader ambitions. It is commendable and credible, both in-and-of itself, and as a smaller piece in the larger picture.

Enterprise Strategy Group | **Getting to the bigger truth.**

20 Asylum Street | Milford, MA 01757 | Tel: 508.482.0188 Fax: 508.482.0218 | www.enterprisestrategygroup.com