

# Market Report

An orange rectangular graphic with a background of faint circuit board traces and a white line graph showing an upward trend. The text is overlaid on this graphic.

## The Future of Storage in a Virtualized Data Center

*By Steve Duplessie and Mark Peters*

January 2011



## Contents

The Tsunami of Server Virtualization .....	3
The Impact of Virtualization on Storage.....	3
The Overall (Storage and More) Problem .....	4
The Bigger Truth .....	7

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change from time to time. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of the Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at (508) 482-0188.

## The Tsunami of Server Virtualization

If anyone has any lingering doubts that the data center as we've known it is in the midst of total "virtual" upheaval, let us dispel them—it's over. The data center of yesterday is gone (or at least going); the data center of tomorrow is virtualized. Server virtualization, albeit still nascent in its overall capabilities, is not a fad. It is here to stay. It is not simply good marketing. It is logical and compelling, both operationally and financially. It simply makes business sense.

So, gradually and ever-faster, the world of IT has begun to virtualize its server infrastructure. Depending on how sophisticated the use of virtualization technologies and how much server virtualization has been deployed throughout a particular infrastructure will determine the overall value derived from such activities. What can be said unequivocally is that the more advanced the organization in terms of virtualization deployments, the greater level of benefit and value the organization can expect.

Therefore, we can conclude that the great virtualization experiments of the last ten years have been successes—and they will continue and accelerate in order for participating organizations to garner more and more value. The congruent growth of both data and user expectations only serves to make the change more necessary and more valuable.

## The Impact of Virtualization on Storage

Since we are now intellectualizing that server virtualization ubiquity is a foregone conclusion, the next logical areas to investigate are those which hinder its progress. ***Storage, as we have known it, is the dragon that must be slain in order for our IT society to truly progress out of the dark ages.***

ESG recently performed extensive research in order to categorize IT organizations' "maturity" which allowed the overall market to be segmented based on current server virtualization sophistication. Further, this Server Virtualization Maturity Model was used to identify specific challenges, concerns, and requirements in each maturity segment. Not surprisingly, there are significant obstacles to the advancement of server (and thus overall infrastructure) virtualization efforts caused by current—and often arcane—storage architectures, implementations, and management techniques. Put simply, many aspects of today's enterprise storage implementations grew out of monolithic mainframe-era designs—and they are showing their age.

***It is ESG's contention that IT organizations will be increasingly motivated to overhaul their network and storage infrastructures to keep up with radical changes at the virtualized compute layer. The network layer is rapidly shifting from Fibre Channel to Ethernet. The storage layer of the data center will increasingly look like the server layer: scale-out, virtualized, and flexible, with commodity hardware economics. The increasing political—and business—clout of virtualization advocates will only accelerate these trends.***

### The Research

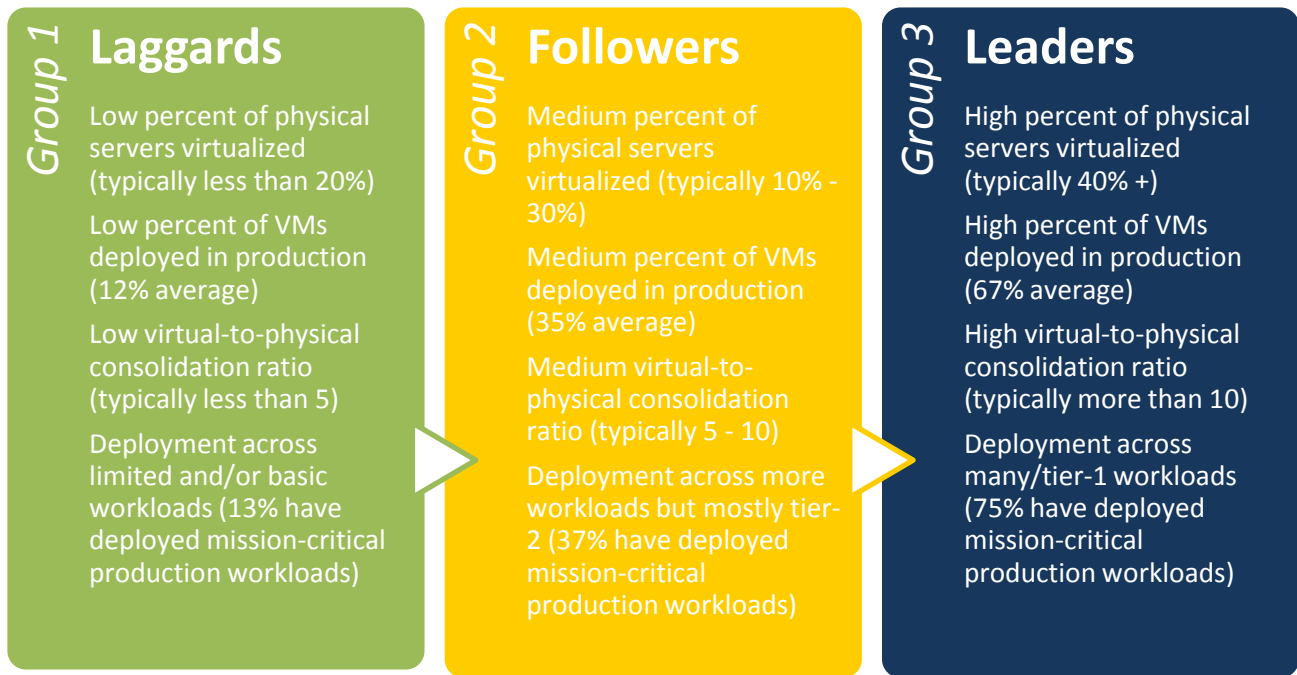
There are four dimensions to ESG's Server Virtualization Maturity Model:

1. Scope of Deployment: the percentage of physical servers that has been virtualized.
2. Virtual Production Ratio: the percentage of virtual machines in production.
3. Consolidation Ratio: the average number of virtual machines per physical machine.
4. Workload Penetration: amount and types of workloads deployed.

### Market Segmentation

The segmentation based on the research is shown in Figure 1. For purposes of clarity, the three segments—Laggard, Follower, and Leader—do not imply "worth" or (more importantly) a lack thereof; they are simply a means of categorizing the current state of the market by where IT organizations are in their current virtualization maturity. For completeness of information, analysis showed that some 22% of current virtualization users can be categorized as laggards, 53% as followers, and 25% as leaders.

Figure 1. ESG Server Virtualization Maturity Model Segmentation



Source: Enterprise Strategy Group, 2011.

Two things are abundantly clear from the research:

1. The benefits organizations derive as they progress in maturity are dramatic.
2. Storage challenges continue to be a major impediment to attaining improved virtualization maturity.

### Storage Problems

The data shows that the implementation of server virtualization inevitably and invariably requires substantial changes and upgrades to existing storage infrastructures. Specifically, storage groups find themselves adding new/additional networked storage and designing for higher IO and throughput densities in order to respond to the stresses and demands that server virtualization places upon them. There is, however, a little-discussed “dirty little secret” behind the technology curtain: in many cases, the cost of upgrading to faster storage systems can negate a large portion of the savings enabled by virtualization.

In addition to new equipment needs, virtualization also usually drives new requirements for business processes and planning. For example, storage organizations typically find they need to develop updated DR strategies, implement new backup policies, and increase their collaboration with other functional IT groups.

Virtualization education, which is significantly lacking across almost all IT areas, is also a major problem. Once educated, storage administrators report concerns in the areas of costs (capital and operational), storage performance, storage security, rapid provisioning, and interoperability. In short, things that a storage administrator has historically rarely had to worry about are now front and center in the new world of virtualization—and without some significant change, these challenges are only going to worsen as server virtualization becomes universal.

### The Overall (Storage and More) Problem

**How we got here.** There are clearly some immediate areas within storage that need to be addressed for it to be a compliant and valuable contributor to the “new IT.” In order to understand exactly why, a little history is needed.

Commercial computing took hold when one single infrastructure stack executed one specific application for one specific purpose. The original mainframe was a glorified calculator. Centralized computing was predictable and controllable, albeit expensive. But it could be managed: one processor system and one IO subsystem.

Decentralized (or distributed) computing was developed largely to try to solve the economic challenges of centralized computing (essentially CAPEX) and yielded low-cost, commodity servers—which we promptly plugged into proprietary, large, expensive, monolithic storage boxes. Servers became cheaper and more interoperable while storage has remained proprietary and expensive. In the old days, the server was the thing that cost all the money. You picked your server by your OS. You picked your OS by your application. Storage was a “peripheral.”

Today, servers are cheap and interoperable while storage is outlandishly expensive, complex, incompatible, and difficult; in many respects, it is the last bastion of IT awkwardness: the peripheral tail wagging the purposeful dog!

**Where to next?** If the primary objective is to create a data center built on virtualized assets (for all of the reasons you already know), then in order to achieve that objective, all of the virtual “layers” must coexist and support the same functional capabilities. These layers (or assets) must then “act/react” to changing conditions within their own infrastructural areas ... and to those around them.

Let’s take for granted that we want to virtualize in general because we can gain efficiencies in asset utilization, take advantage of the commoditization of hardware, leverage common infrastructures, provide seamless mobility options, etc. If we can do all this, then we are set up to drive the next (higher) level of value where we can then aspire to provide infrastructure that:

- **Self-optimizes:** boxes that tune/reconfigure themselves for the workloads that are presented—and change as those requirements change.
- **Self-heals:** infrastructure deals with fault-scenarios autonomously, remapping/rebuilding itself so that the application is not affected.
- **Scales dynamically:** up or down, in or out; infrastructure that extends—virtually—to whatever requirements the workload(s) presents.
- **Self-manages:** adapts to changing scenarios based on policy and enforces those policies via automation.

This set of values is nirvana when it comes to infrastructure today, but it is exactly what we are moving toward. Along the way, however, there are still several basic requirements that have to be addressed in order for it to occur:

- **Is everything connected to everything else?** Your servers all connect to Ethernet, but much of your storage exists (some would say “is trapped”) behind Fibre Channel networks that originated in the mainframe era to carry FICON and ESCON traffic. As 10-Gbit Ethernet prices continue to drop and as 40-Gbit and 100-Gbit emerge, it is inevitable that everything within the data center will be Ethernet connected. That particular war is over, even though the peace has yet to be fully implemented. Thus, Fibre Channel vendors are advocating FCoE as a way to preserve the existing protocol over the new network; the reality is that although it is a welcome interim step for many committed FC users, it is nonetheless an admittance of where the future lies. The good news is that FCoE runs on Ethernet switches, but the bad news is that it does little to address the complexity and cost challenges of legacy Fibre Channel. Meanwhile, iSCSI and NAS vendors are emphasizing their Ethernet stories (which will indeed play well in a number of environments) and other vendors are introducing new models for purer forms of “Ethernet storage,” such as ATA over Ethernet (AoE), which are capable of supporting—and designed to support—more flexible virtualized and cloud architectures.

*Ethernet SANs are the future. Protocols are secondary.*

- **The data center will converge on a single “uber-network” as soon as it is conceivable.** Much as the human body is comprised of nerves that have differing functions, it is nonetheless true that those nerves share a common structure. It does you no good to be able to move virtual workloads all over the data center if the data doesn’t—or can’t—go with it.
- **Scale-out storage is a must, but it’s just starting.** Think of it this way: virtual servers can be thought of as scale-out servers: a plethora of commodity machines all connected to Ethernet with various capabilities that allow workloads to be almost instantly brought up anywhere. But this wondrous capability is then all-


too-often restricted by having to serve the requisite data from a “monolithic box” originally architected in the IT Stone Age. Today, you can buy any server you want, from anyone, plug it into your environment, and use it almost instantly—but if you’re running monolithic storage, you’re no better off than you were ten years ago. It is therefore inevitable that storage needs to become commodity-based, Ethernet-based, and scale-out-capable to match (or exceed) the capabilities of the server and Ethernet layers.

- We have to accept the commoditization principles of servers and networks in our storage environments.** We can no longer continue to create proprietary monolithic “boxes” of functionality where complexity is layered upon complexity to deal with the complexity that was put there last year in what might be described as “Ponzi storage.” After all, these monolithic systems are actually just boxes of disks and memory—cheap disks and cheap memory. As an industry, we need to acknowledge that our value cannot and will not be contained in a box of disks and memory. Any functionality of legitimate value has to be monetized independently of the commodity, not locked within it. Elsewhere in the IT world, this lesson has been learned, voluntarily or otherwise. HP/IBM/DELL etc. cannot (nor do they try to) charge 5X their competitors for a motherboard with the same Intel processors on it; so why is it that storage manufacturers can vary so greatly in their pricing for what is essentially the same box of disks?

ESG also specifically investigated the storage-specific developments that would enable more widespread usage of server virtualization.<sup>1</sup> The responses are summarized in Figure 2 and need no additional explanation other than to say that in a world that has elected to “go virtual,” we clearly need storage to become more simple and flexible as well.

*Figure 2. Storage Developments That Would Enable Wider Server Virtualization Usage*

**From a storage infrastructure perspective, which of the following developments do you believe need to take place in order to enable more widespread server virtualization usage in your organization? (Percent of respondents, N=190, multiple responses accep**



Source: Enterprise Strategy Group, 2011.

<sup>1</sup> See: ESG Research Report, [The Evolution of Server Virtualization](#), November 2010.

## The Bigger Truth

Today, we can provision a virtual machine (server) for a business in minutes; yet we still provision storage in days, weeks, and—often—many months. We can instantly move a virtual machine onto another server and give it more CPU power than entire countries had ten years ago; but if data isn't moved with it, it's a useless exercise.

These competing forces cannot continue to proceed. ESG's position is simple: storage will cease to be implemented as we know it and will instead become a virtualized complement to the server and network layers. It is only a matter of time. In the future, storage within IT will match (or exceed) the competency and functionality of other elements of the IT infrastructure. It will:

- Be entirely virtualized.
- Leverage commodity hardware components not only to the benefit of the manufacturer, but also of the buyer in the same way that buyers get more powerful servers for less money every year.
- Leverage Ethernet's ubiquity.
- Become self-managing, self-optimizing, and self-healing.
- Scale to what seems (today) like unimaginable levels.

The technologies to make this happen are not futuristic pipe dreams—they all exist right now. The industry, which makes a lot of money by continuing to propagate existing monolithic technology, will be forced to adapt to these new realities or, like the old monolithic server vendors, perish. Historically, storage vendors have competed by having a better "function" trapped within their boxes, but those functions are either going to be available differently from those vendors or they will be consumed by the virtualization layer.

Indeed, it could also be argued that the "storage administrator" function as we currently know it will not exist in five to ten years. The "virtual administrator" will be the thing—automatically, autonomically—setting performance, protection, and recovery policies that will be enforced at the virtualization orchestration layer. Even "set and forget" will be surpassed by "get and forget."

As an IT industry, we have opened Pandora's box by implementing server virtualization technologies. We cannot go back. We now have set expectations with our businesses and they now expect assets to be provisioned almost instantaneously; having to explain that we need to take months to make sure all the pieces behind the curtain are in place before we "really" turn it on is not going to be an option—unless we're interested in seeking other employment.



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