

White Paper

From Reporting to Analytics: What Every Storage Administrator Needs to Know

By Wayne Pauley, Senior Analyst

August 2013

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

Contents

In Brief.....	3
Contemporary Storage Management Challenges	3
Shared Resources, Shared Problems.....	5
The Power of Analytics	5
EMC Storage Analytics and VMware vCenter Operations Management Suite.....	6
Management across Clouds, Converged Virtualized Servers, and Storage.....	7
Case Study – Major European Dairy Producer and Processor.....	8
The Bigger Truth	9

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.

In Brief

As IT organizations have moved from traditional siloed infrastructure to virtualization and now to cloud, the data center landscape has fundamentally changed. The ease of deploying resources is directly contributing to today's massive data growth. The challenge is that complexity in management cannot scale linearly with data growth; storage administrators need to continue to deliver high levels of service without a corresponding increase in operational costs. To do so, they must approach cloud management differently from the siloed management tools of the past and provide analytic solutions that support collaboration across storage, compute, and cloud organizations.

With the introduction of EMC Storage Analytics (ESA), [EMC](#) and [VMware](#) are approaching this problem proactively and intelligently with a new way to improve storage uptime with deep visibility into the dynamic dependencies between virtual machines, workloads, and storage arrays.

Contemporary Storage Management Challenges

The way IT is run and delivered has undergone tremendous changes and progress, resulting in boosts to productivity and business value. Virtualization and cloud computing are transforming traditional IT infrastructure foundations where the management of large and growing storage environments is fundamental to the success of the business.

For today's IT organizations, the storage management issue is far more than the "do more with less" mantra that is often bandied around. While cost reduction is a real requirement, the difficulties in today's IT come from doing new things in a highly dynamic world. Among the realities facing IT leaders are:

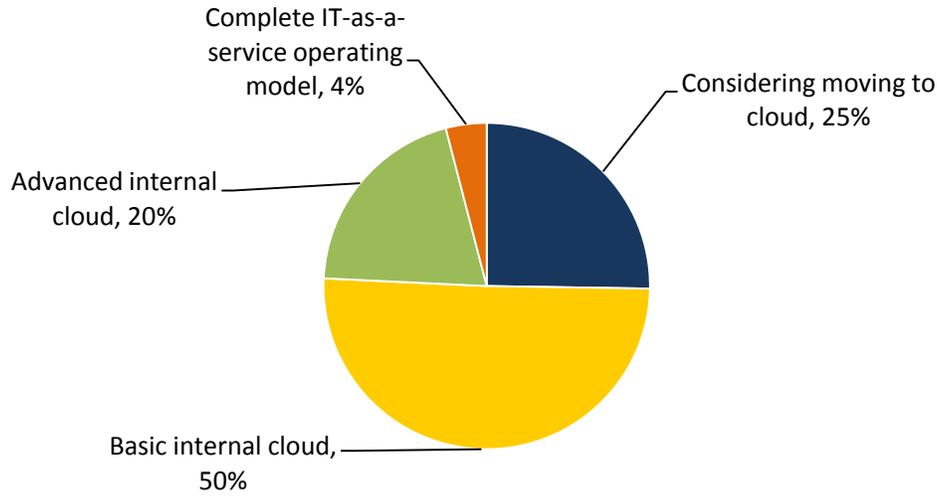
- **Ever-changing business climate:** Globalization, faster and less predictable economic cycles, and growing competitive pressures require increased agility, yet with reduced costs and improved efficiency.
- **IT transformation:** Operational and business model requirements drive IT to embrace trends such as cloud, virtualization, consumerization, and big data. Each of these solves certain issues and offers tremendous upside, yet each also comes at the cost of infrastructure sprawl. Multiple management toolsets and even more stringent SLAs make it harder to manage infrastructure without precise knowledge of the underlying resources and associated performance and availability requirements.
- **Complex storage:** Heterogeneous storage (different system types, vendors, and/or from cloud providers) and applications can obscure the visibility of issues, often forcing problems to be addressed after they occur. And self-service, end-user provisioning provides "ability without responsibility," where IT just has to "make it work."
- **Massive data growth:** Data volumes, types, and locations are growing faster than budgets and resources can meet.

The combination of these factors leads to some interesting challenges for traditional IT functional boundaries. For example, virtualization impacts the "normal" delineation between system and storage administrators; in the meantime, the self-service and on-demand nature of cloud computing changes the IT budget process from an allocation model to a consumption model. This is validated by ESG research which reveals that among a group of current and potential private cloud users, nearly three-quarters have implemented a basic internal cloud (i.e., compute, storage, network stack) or an advanced internal cloud (i.e., advanced service automation and orchestration). (see Figure 1). Indeed, the same ESG research shows that storage management is one of the top three IT management capabilities companies want to see as an outcome of going to a private cloud (see Figure 2).¹

¹ Source: ESG Research Report, [Trends for Protecting Highly Virtualized and Private Cloud Environments](#), June 2013.

Figure 1. Organizations' Transition to Private Cloud Computing

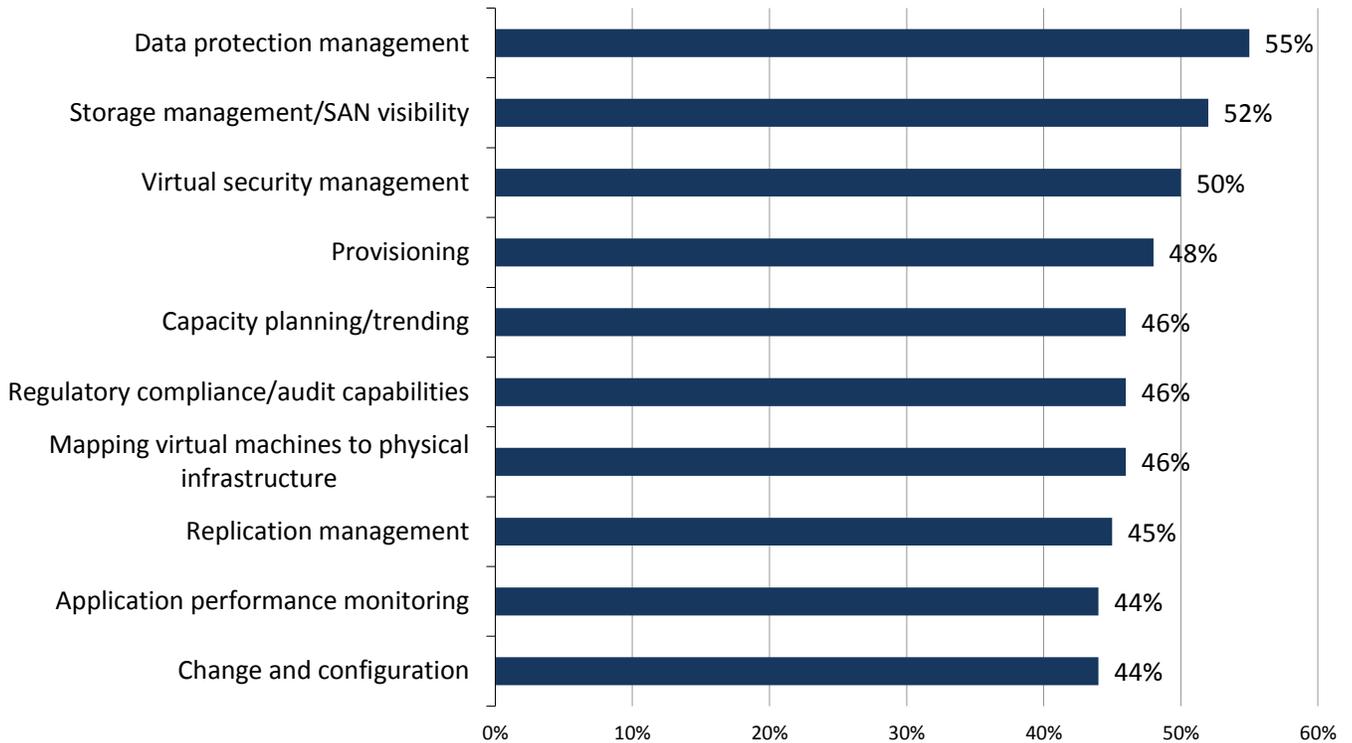
Currently, at what point would you say your organization is on its transition to private cloud computing? (Percent of respondents, N=290)



Source: Enterprise Strategy Group, 2013.

Figure 2. Important Management Capabilities for Private Cloud Environments

Which of the following management capabilities for private cloud environments are important to your organization? (Percent of respondents, N=151, multiple responses accepted)



Source: Enterprise Strategy Group, 2013.

Shared Resources, Shared Problems

It might seem counterintuitive that streamlining IT—via technologies such as virtualization and cloud that are so valuable in a business sense—could actually make storage management tougher. Unfortunately, the realities speak otherwise:

1. Virtualization and cloud increase the demands on storage—there’s often increasing heterogeneity, while the frequent movement of VMs and data (to optimize and automate server workloads) creates a complex, dynamic storage environment. Users expect complete fluidity irrespective of the foundational underpinnings, making problem determination and resolution a nightmare: Is the problem in the VM, the guest OS, the network, or the storage? Such challenges can be multiplied further within cloud environments—whether public, private, or hybrid—where management of resources becomes a greater challenge to manage.
2. Data management innovations that should simplify tasks such as automated tiering, provisioning, or utilization-optimizing tools often makes managing resources more complex. As storage systems become more heterogeneous and siloed (by application, location, or tier), the more diverse the management tools and processes become. This can often force storage administrators to opt for manual control via the CLI (command line interface) and actually dilutes their ability to find and solve problems in a timely manner.

There is a commonality across these facts that sheds light on what the *real* challenges and needs are, and which suggests that we may have been concerning ourselves more with the symptoms rather than the cause of storage management challenges.

All the resource sharing that is happening in IT—via clouds or via converged infrastructures—(also known as pooling) is designed for better business outcomes and demands integrated resources. Yet, with pooling, the waters get even murkier for with the overwhelming amount of seemingly disconnected information that inundates today’s storage administrators. It requires a new approach and new tools with built-in intelligence and analytics that are in-step with today’s virtualized and cloud-based IT environments, including the underlying storage infrastructures that support them.

The Power of Analytics

If such information and analytics were available, the implications could be dramatic for storage and IT in terms of both effectiveness (i.e., getting the job done) *and* efficiency (i.e., using as few resources as possible). Management visibility into and control of pooled storage resources could make meeting SLAs far more likely, and could also help to optimally leverage the capabilities of virtualization and clouds by removing an anchor on their success.

The tool would need to augment historical reporting systems. Generically, it would need to be integrated within a converged infrastructure, and be as responsive and flexible as the systems it is designed to support. In terms of specifics, such a state-of-the-art management tool would need to encompass the following attributes and abilities in order to successfully manage highly virtualized and cloud environments:

- **Context Is King:** more than just UI consolidation, this should cover the physical/virtual compute, network, and storage layers, and must allow a drill-down capability while being application-context aware. The old management paradigm of each layer’s administrator receiving alerts or angry calls from users for just that layer has to be replaced by an integrated management platform that embraces private, public, and hybrid clouds, from both virtual and physical environments. In addition to analyzing each individual layer (system, network, and storage) of the infrastructure (capacity, performance, etc.), the tool must be capable of analysis across layers, while also taking into account the behavior of the application, the interdependencies, and its users.
- **Predictive Management:** Another “that’s-just-the-way-it-is” management standard that has to change is to move from merely reacting to problems to avoidance. This has been a long-time goal of management and the technologies are now maturing to support predictive analysis. What analytics should be able to do is to learn from past behaviors and detect anomalies rather than simply report on deviations from static

thresholds or alert boundaries. This need not be a rigid implementation; for instance, continuous monitoring (of, for example, performance and capacity trends) allows for normal baseline behavior to be adjusted and learned.

- **Support for Existing Technology Footprint:** Heterogeneity is so common now that the tool must be able to manage multiple technology types (storage can be file, block, object, or even a combination; servers can be white-box or from a myriad of vendors; and networks can be traditional hardware or software-defined technology). An integrated analytic and management tool must clearly encompass all the storage resources as well as other cloud components, where administrators can monitor, provision, define policies and thresholds, and manage the full lifecycle of all the resources. Having a separate set of tools, nomenclatures, and user-interfaces for clouds is inefficient for both the administrators (having to be proficient on yet another platform) and the use of the resources (as separate pools can rarely be utilized as well, or prove as flexible, as one pool).
- **Service Centric:** It is not the cost alone of the service, but the cost/quality trade off. IT executives and business stakeholders are no longer satisfied with storage-only availability metrics—they want to see the entire service availability. This ultimately demonstrates the value of internal IT departments over buy-by-the-credit-card public cloud services. Managing to SLAs is likely insufficient to deliver the desired cost-effective services; and ultimately the business must also see the bar met or exceeded when it comes to quality of service as well.

This list embraces elements that have traditionally been separate—operations management, storage analytics, and cloud automation—and so it can genuinely be seen as the criteria for a converged management tool. The good news is that the list can also serve as the description for the key abilities of VMware's vCenter Operations Management Suite in concert with EMC Storage Analytics.

EMC Storage Analytics and VMware vCenter Operations Management Suite

True IT added-value can be claimed when improving processes that allow new technical and business benefits to be realized. EMC has developed a predictive analytics tool to support the changing needs of storage administrators, built on the VMware vCenter Operations Management Suite engine. This new offering, EMC Storage Analytics, can deliver value standalone or can be used in conjunction with VMware vCenter Operations Management Suite to improve application uptime and collaboration across storage and virtualization teams.

EMC Storage Analytics combines and enhances several key elements, such as Unisphere from EMC, with plug-in APIs from VMware. The essential ability of ESA is to provide the storage admin true end-to-end visibility in heterogeneous storage environments that span physical, virtual, and cloud. ESA's strength lies in its extensive analytics and management capabilities, of which some key components are:

- A patented analytics engine based on VMware vCenter Operations Management Suite scans and compiles extensive metrics, which uses baseline data correlated to the current workload to predict problems and bottlenecks as well as determine what the optimized behavior/fix is.
- Dashboards provide dependency mappings between storage and VM workloads, and allow storage and virtualization administrators to collaborate and improve application uptime.
- Alerts can detect potential problems before they occur and make adjustments as the system learns from normal business behaviors. Additionally, the tool predicts potential degradation of application performance, so making adjustments and streamlining processes makes the system even more efficient.
- Intelligent Automation uses automated root cause analysis to allow for simplified and optimized workload rebalancing across VMs and storage resources, automatically.
- Optimization of storage uncovers application hot spots and over/under provisioned capacity in order to meet SLAs.

- Currently supports multiple platforms, including EMC's VMAX, VNX, CLARiiON, and NS offerings, with expansion to VPLEX and Isilon coming soon.

vCenter Operations Management Suite provides incremental functionality to customers who are using ESA to improve uptime of their virtual environments and storage and tighten collaboration between teams, including:

- Prebuilt and configurable dashboards for real-time performance, capacity, and configuration management across cloud, virtual, and physical environments help gain visibility into infrastructure and applications (in addition to the storage-centric dashboards provided in ESA).
- Automated root cause analysis, including correlation of health, performance, and change events at the infrastructure and guest OS-level.
- Capacity optimization to right-size and reclaim overprovisioned capacity and increase virtual machine consolidation ratios.
- Automatic detection and monitoring of metrics across 120 web technologies to find, fix, and prevent performance problems in custom web applications.
- Automated configuration management across virtual and physical servers, workstations, and desktops to avoid configuration drift and maintain continuous compliance.
- Automatic discovery and dependency mapping of applications running on virtualized compute, storage, and network resources.
- Prebuilt and custom integration with third-party cloud, monitoring, storage, and network systems to apply vCenter Operations Management analytics to the full environment.

vCenter Operations Management Suite and ESA are extremely useful and valuable independently; combined, however, they add greater management value for the ever increasing converged, cloud-centric IT world. The more complex, dynamic, and distributed the environment, the greater the need for analytics to apply insight to the vast amounts of data collected.

Management across Clouds, Converged Virtualized Servers, and Storage

Clearly having a real-time "at-a-glance" view of everything that is going on means that hot spots and problem areas can be addressed before they create operational issues or negative business impacts. Administration becomes something that is done from an integrated applications viewpoint, as compared with the granular fiefdoms that are more traditionally found, because there's just one management tool with an end-to-end view of all the layers that can configure, provision, monitor, and automate virtual machines and storage. Once the analytics engine in the tool identifies, for instance, a performance problem anywhere in the system, any given administrator can rebalance the workloads by reconfiguring any of the layers (compute, network, or storage) or add resources as required.

Comprehensive management has to include the cloud, which is further enhanced by ESA's ability to enable an administrator to add resources for the service catalog for all layers. In addition, vCenter Operations Management Suite works with VMware vCloud Automation Center to deliver automated provisioning of services, including storage dependencies and ongoing lifecycle management (i.e., scale up, scale out, retirement).

The combination of vCenter Operations Management Suite and ESA provides a powerful streamlining mechanism that simply makes optimizing performance and resources much easier across an entire application delivery stack. It can remove role specialization and in doing so it makes each administrator more powerful and capable than before: After all, there's no need to argue who gets to translate when you're all speaking the same language!

The logic of all the above is inescapable, but ultimately the only thing that matters is whether this all means something useful for IT and business or not. To verify this, ESG conducted an end-user interview exclusively for this paper.

Case Study – Major European Dairy Producer and Processor

Demographics & Prior Management Methodology: ESG conducted an in-depth interview with a Belgian-based dairy foods company, operating ten locations in Belgium and one each in France and the Netherlands. Each year, the firm processes 1.1B liters of milk, as well as associated products such as ice cream, cheese, and other dairy products. The company's 2,000 employees are supported by an IT department that manages over 700 PCs, a fully virtualized server environment of 250 VMs, and some 25TB of storage, with all systems being mirrored to a matching DR site.

The interviewee is the systems administrator responsible for all the layers—storage, networks, and virtualization. vCenter Operations Suite is being used to manage primary applications (heavily used Oracle and SQLServer databases), with Unisphere as the storage management tool.

Why ESA? Using a separate tool for the application/VM and storage environments was proving suboptimal. Deploying ESA provided two key advantages:

- 1) The ability to manage all the layers of the stack (application, compute, network, and storage) both from end to end and via one interface.
- 2) Easy-to-use converged functionality to enable the proactive rebalancing of workloads.

Another benefit, for example, is that when LUNs move on the array, ESA notifies storage administrators, which allows them to verify the move and analyze its root cause, therefore being able to determine that the move had not been caused by a less directly related issue, such as WAN latency (something that was not possible pre-ESA without inordinate manual work). The customer also spoke highly of the ability to set SLA policies, which meant that when users called to complain about perceived performance problems, the customer could immediately prove that storage was not the issue and drill down to determine the root cause. Finally, this customer leverages the “Dashboard Community” and found it easy to import and deploy a cross-domain dashboard.

Specific Example: The customer related an example of positive ESA impact concerning the migration of a mission-critical database from its older storage platform to a newly purchased VNX. They had kept the same basic storage layout from the old array to the VNX, but were immediately alerted to hot spots as a result of ESA's ability to drill down all the way from the application to the VM and to the storage. Database files were quickly and easily rebalanced...and precluded the performance issues before they developed and had a chance to impact business productivity.

As the case study demonstrates, the ability to proactively respond to performance issues can minimize or even eliminate impact on users. Although it is something of a cliché, it is worth stating that new computing models (such as cloud, virtualization, converged infrastructure, etc.) cannot rely on old methods of analysis and management (i.e., spreadsheets, separate tools, best-guesses, and the like) if the optimum results are to be delivered. It's true that virtualization and cloud computing can help drive resource efficiencies via higher levels of utilization and agility in their inherently pooled environments. However, they can also present a conundrum for IT organizations that must manage an increasingly dynamic and diverse environment. Clearly then, the new IT environment requires new tools—such as ESA integrated with vCenter—that deliver a holistic capability.

The Bigger Truth

Modern business realities are forcing IT to adapt and adjust as quickly as—if not more quickly than—other functional groups within contemporary organizations. This adaptability must reflect the accelerated pace of business (and “time to value”) and of course, impact the storage administrators. Ultimately, these are the key drivers behind the evolving and increasing use of virtualization, cloud computing, and converged infrastructures. But without equal progress supporting the storage administrators in terms of analysis and management, a significant portion of the potential benefits can easily be lost. Just as a mechanic can’t afford to search through his toolbox for 15 minutes for a wrench, storage administrators can’t afford to solve problems by looking at the storage then handing it off to the network team who in turn hands it off to the server team. Not only must management tasks be more integrated and holistic, they must also enable a single language of control. In order to meet the challenge of dealing with information overload from the storage, networks, and servers themselves, IT organizations must have more advanced tools like predictive analytics to proactively prevent problems from manifesting themselves.

EMC Storage Analytics, which uses the capabilities of the analytics engine in VMware vCenter Operations Management Suite, provides a much bigger and clearer picture: a comprehensive view from virtual servers all the way through to the storage. Working together, EMC and VMware are delivering a tightly integrated (i.e., converged) management tool that provides application context and state-of-the-art predictive analytics across servers, networks, and storage. At a time when competitive businesses do not have time to deal with outages or performance degradation (which can cost money, customers, or both), EMC Storage Analytics provides cross-domain monitoring, measurement, and management capabilities that are truly synchronized with modern IT.



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