

White Paper

EMC ProSphere v1.5

End-to-end Storage Management for Cloud Environments

By Bob Laliberte

March, 2012

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



Contents

The Transition to the Private Cloud	3
Storage Management Challenges in Highly Virtualized Environments	5
EMC’s Solution: ProSphere v1.5	6
The Bigger Truth	11

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 Transition to the Private Cloud

Heavy business dependence on IT, in conjunction with the constant change occurring in global markets, is increasing demand for infrastructure that is more responsive and better aligned with business need. As a result, organizations of all sizes are rapidly transforming their IT environments to meet new challenges and requirements. But this transformation goes beyond simple technology refresh cycles and consolidation efforts; organizations are making significant investments to create new, more agile IT environments commonly known as “private clouds.”

Private clouds are attractive for more than the consolidation benefits they bring (although these are widely documented and can result in significant savings). The real benefit, however, comes from the private cloud’s ability to improve business processes and IT responsiveness. Indeed, business process improvement topped the list of the most important criteria when justifying IT projects to the business, in the ESG Research 2012 IT Spending Intentions Survey¹. If an organization can spin up new business services in two hours instead of two weeks, it will clearly have a competitive advantage in the market. Additionally, many new technologies utilized in the private cloud enable IT to deliver higher levels of availability and mitigate risk with more cost-effective solutions. Ultimately, with greater integration and automation, private cloud infrastructures will enable organizations to reduce operational and capital costs.

A number of technology advances are enabling this evolution toward the private cloud, but the best recognized is server virtualization. ESG research indicates that server virtualization is ubiquitous and growing in enterprise environments; in a survey largely focused on the enterprise, 93% of respondents either had or were planning to deploy the technology. This is reinforced by the fact that respondents to ESG’s annual IT spending survey have reported that “increasing the use of server virtualization technology” has been a number one initiative for the last three years in a row (see Figure 1).²

Figure 1. Top Ten 2012 IT Initiatives



Source: Enterprise Strategy Group, 2012.


¹ Source: ESG Research Report, [2012 IT Spending Intentions Survey](#), January 2012.

² Ibid.

And for the first time, the research points to the building out of a private cloud as a top ten IT priority for respondents. ESG views this march toward private clouds as organizations maturing their rapidly growing virtualized environments to include more automation and orchestration. However, in order to get there, organizations must first have full visibility into these environments from end to end. This is relevant because these highly virtualized environments have strong interdependencies between the server, network, and storage domains. As a result, organizations are seeing significant growth in networked storage environments, which is also directly related to the increased use of server virtualization. In order to take advantage of all of virtualization’s mobility and high availability options, a networked storage environment is required. When ESG asked storage administrators about the number one impact server virtualization had on their storage environments, they responded that they had increased their use of SANs, so it shouldn’t be surprising that the number four IT initiative for 2012 was managing data growth. It is also worth noting that the increasing density of VMs per physical server will also drive bandwidth requirements for networked storage environments.

From a management perspective, increased use of server virtualization is having a significant impact on the tools and processes used to manage the server, storage, and network. While one would anticipate that the server team would report the most impact, the storage and network teams are not very far behind. ESG research reveals just how significant that impact is: an impressive majority of survey respondents reported that server virtualization is having either some or a significant impact on existing management tools and processes (see Figure 2).³

Figure 2. The Impact of Virtualization Technology on Management Processes and Tools



Source: Enterprise Strategy Group, 2010.

As virtualization technology continues to mature and IT moves further down the path to the private cloud, challenges are bound to arise. Organizations are reporting the need to make significant changes to management tools and processes. These changes need to take place in order to allow organizations to evolve their virtualized environments into private clouds. The focus of this paper will be to outline those challenges specifically related to the storage environment and demonstrate how end-to-end visibility with [EMC’s](#) latest storage management technology, ProSphere, can ease those challenges and help accelerate the adoption of private clouds.

³ Source: ESG Research Report, [The Evolution of Server Virtualization](#), November 2010.

Storage Management Challenges in Highly Virtualized Environments

Making the transition to highly virtualized environments and, eventually, to private clouds can create significant but not insurmountable challenges for the storage team—especially if they have the right tools and processes. The more prominent challenges include:

- **The creation of yet another siloed domain.** Managing storage is difficult enough in a fairly static environment where an application remains tied to the physical server it resides upon. Tracking problems to that application is easy because you know where it is. In a virtualized environment, the application could be on any server in the resource pool. With many virtualization tools limited to the virtual domain, it is difficult to get a complete picture. Now, instead of just dealing with physical servers and storage, the virtual server domain needs to be accounted for as well.
- **Visibility issues.** This new ability to abstract information also creates a challenge: in order to manage something, one first needs to see it. In the case of server virtualization, there are very strong interdependencies between the storage arrays and supporting network. In order to manage these environments successfully, management tools need to have visibility into adjacent domains up to the VM level. This is especially true in highly dynamic environments where applications frequently move around.

Another aspect of the visibility challenge in the storage domain is related to efficiency technologies like thin provisioning and automated storage tiering that create issues for administrators. In these cases, the ability to have complete visibility into physical, logical, and tiered environments is essential to ensure optimized performance and minimize risk. For example, without detailed information of a thinly provisioned environment, an application could run out of physical storage.

- **Performance issues.** As virtualized server environments mature and organizations begin to move production applications to this environment, performance information becomes more relevant. ESG research indicates that enterprises believe that 58% of their production applications will be run on virtualized server environments within two years. Given this fact, it is critical to ensure that end-to-end performance data be collected prior to a move and then used to benchmark applications in virtualized environments. Equally important is the ability to ensure that service levels are maintained even when dynamically reallocating physical resources. Also, it will be imperative that organizations have the ability to align service levels to the appropriate class of infrastructure to optimize efficiency.
- **Federation of information.** Organizations have always struggled to provide a holistic view of the entire enterprise's storage consumption. Even efforts to consolidate data centers and move more applications from remote locations to central locations have done little to enable a comprehensive view of storage in the data center. Worse is that the manual efforts often result in erroneous information or they become obsolete shortly after the data is collected. As environments grow more complex, the need to rapidly isolate and manage problems (performance, configuration, etc.) across multiple data centers will be in greater demand.
- **The ability to scale rapidly.** Server virtualization technology enables organizations to rapidly add new VMs and turn up new applications or services. However, this rapid scaling creates storage management challenges, with the first being the need to rapidly scale storage environments. ESG research supports this observation as the number one reported challenge for storage administrators is the ability to scale storage in virtualized server environments.⁴ The ability to tally the capacity of storage available will be critical to enable rapid scale and proper planning. Another issue pertaining to scale is agent-based discovery: as the environment rapidly grows, agent management becomes time-consuming and cumbersome. Administrators should spend their time managing the *environment*, not agents.

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

- **Process issues.** Almost any time a new technology is introduced, new tools and management processes follow. This is certainly the case for storage in a virtualized server environment. Change in any organization is difficult, but in IT, change is inevitable. What’s needed is adequate training and education.
- **Integration with existing management tools.** Highly virtualized or private cloud environments have highly interdependent technology domains, so it is critical that management tools be tightly integrated with adjacent domains. For the storage domain, this requires visibility into the storage network and up to the virtual machines. After all, it is difficult to manage a storage environment when applications are playing hide and seek. By collecting data from tools like vCenter or MSSC, storage administrators can determine which applications are affected by a storage outage and prioritize problem resolution.

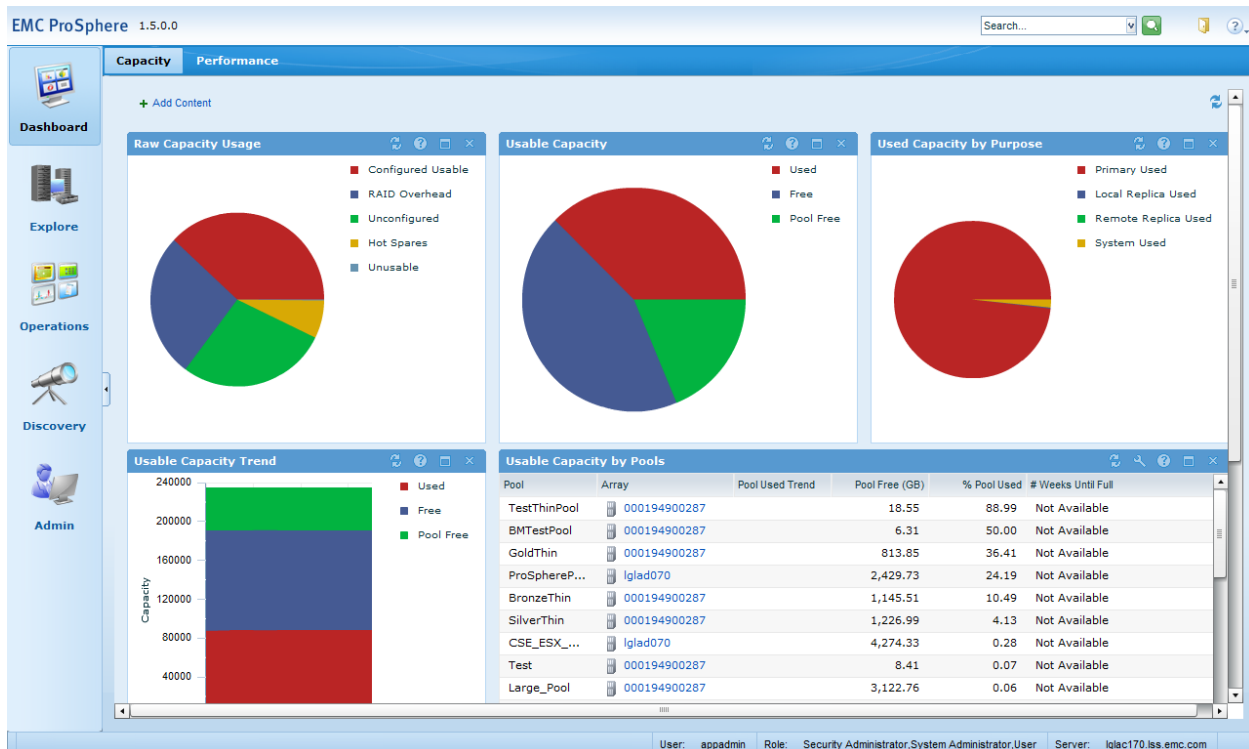
While these challenges seem daunting, they can be overcome. Organizations need to seek out new tools that better address the needs of these emerging virtualized and cloud computing environments. Solutions that address these challenges and help accelerate the adoption of cloud computing paradigms are available today. EMC has one such solution for cloud storage management: EMC ProSphere.

EMC’s Solution: ProSphere v1.5

To help organizations optimize their storage environments, drive enhanced ROIs and ultimately accelerate adoption of private clouds, EMC has re-architected its storage resource management software with the cloud in mind. The result of this multi-year effort is an entirely new, completely redesigned software solution that offers ease of deployment, and value extraction. EMC calls it ProSphere and has just released ProSphere v1.5, which will enable organizations to better address virtualization challenges and help them transition from either legacy or immature virtualization environments to platforms that are more agile and flexible. Specifically, EMC ProSphere v1.5 will deliver:

- **Capacity Reporting and Analysis.** The ProSphere capacity module provides a comprehensive enterprise-wide dashboard view of all storage (see Figure 3).

Figure 3. Capacity Dashboard

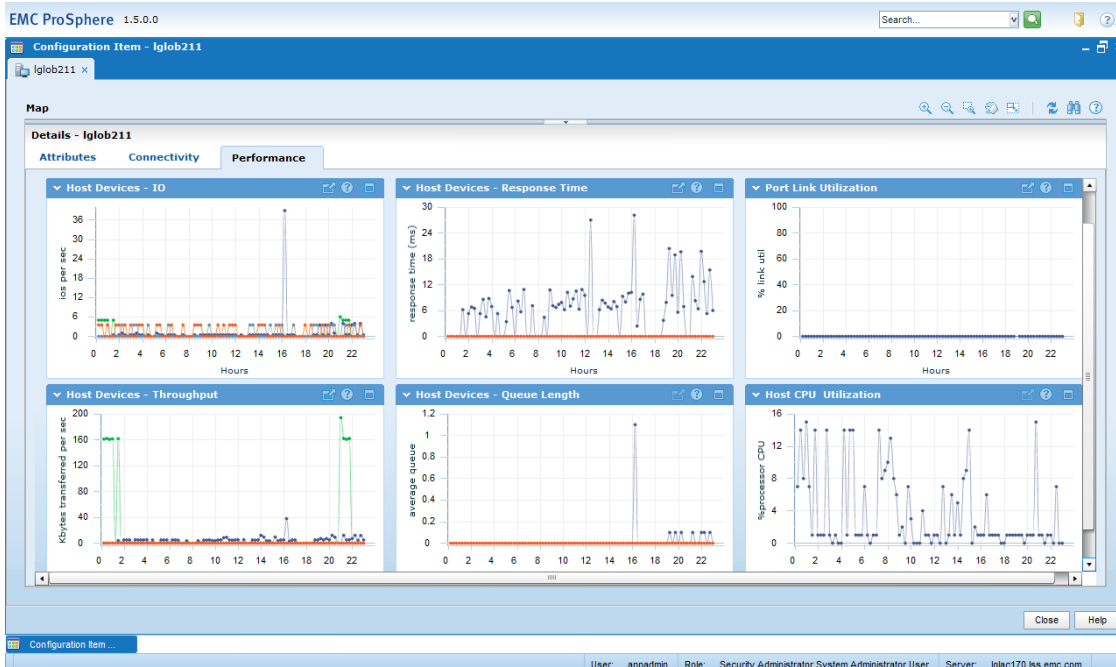


Source: EMC, 2010.

The ProSphere capacity module will provide the capability to drill down to a data center or even individual array level. Information such as raw capacity, configured usable capacity, RAID overhead, hot spares, and other unusable capacity are all available. From an analysis perspective, organizations can view capacity trends, tiered storage, primary versus secondary storage, and even storage by service level. This is all enabled by:

- **Definable service levels.** Users can now create or define a service level based on the disk type, RAID level, or other metrics. Then users can view and report on data or trends by each level—say bronze, silver, or gold.
- **An explore function.** Users can sort, filter, and identify storage by a number of areas, including amount of available capacity, RAID type, use (primary, local or remote replica), and whether it is thick or thinly provisioned, and generate reports based on the inquiry. This functionality covers arrays, LUNs, replication environments and pools of storage. Within each of these areas users can drill down further and get more specific data. For example, in the array tab, users can sort and filter by array name, type, total raw capacity, configured capacity with RAID, and the configured usable and RAID overhead. The Replication view shows detailed capacity information about each replicated LUN and the type of protection (mirror, snapshot, clone) each LUN has along with general array and LUN information.
- **Performance monitoring.** ProSphere's performance module is designed to monitor storage performance across the whole virtual infrastructure. Thresholds can be set to generate alerts in order to enable more proactive management of highly dynamic environments. End-to-end paths, illustrating all dependencies, can be easily viewed. This includes the network switches and ports on the host as increasing VM densities can shift the bottleneck to the network. Users can leverage trend charts of key performance indicators in order to better optimize performance. ESG research indicates that performance problems are an inhibitor of future growth of virtualized environments and private clouds. Thus the ability to rapidly identify and diagnose performance issues will enable more widespread adoption. Figure 4 illustrates the types of reports ProSphere can deliver based on performance criteria.

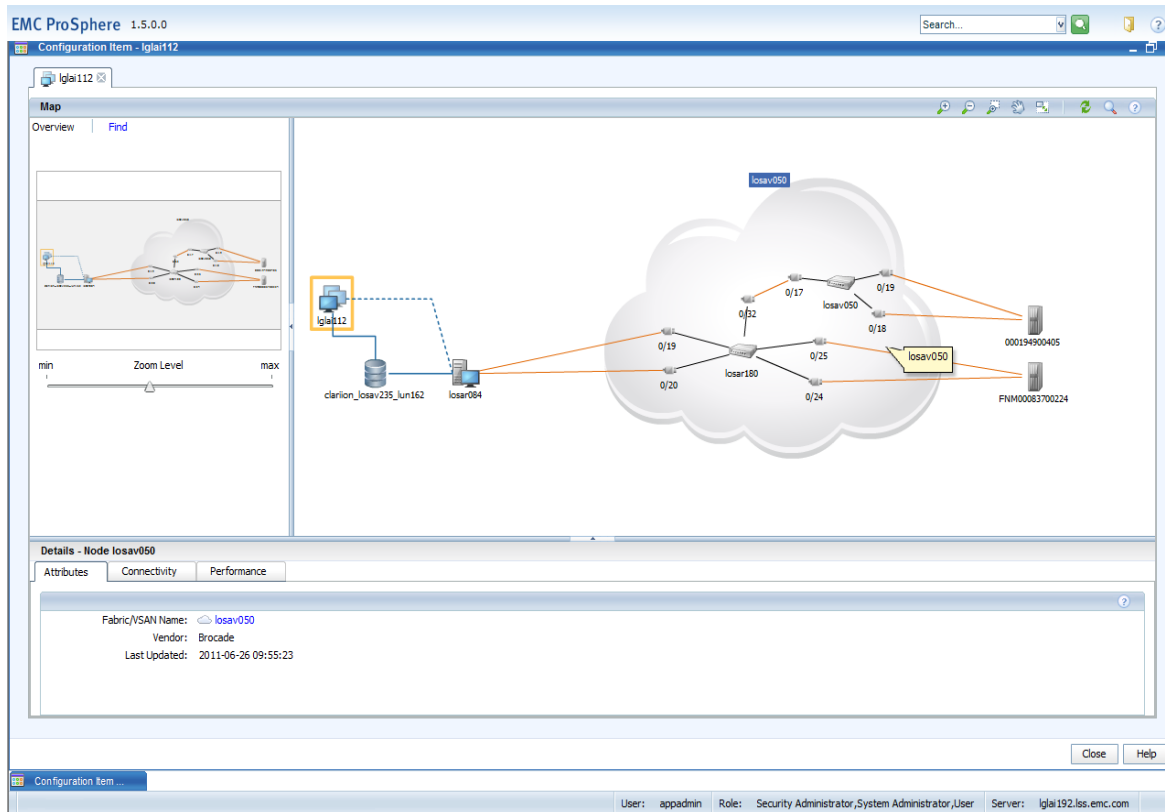
Figure 4. Performance Data



Source: EMC, 2012.

- **A complete, end-to-end view of virtualized infrastructure.** Understanding that one must have visibility in order to effectively manage an environment, ProSphere is capable of providing a view that extends from the VMs through the Fibre Channel switches to the storage and includes innovative EMC storage technologies like Virtual Provisioning and Fully Automated Storage Tiering (FAST). Because highly virtualized environments are so interdependent and dynamic, understanding how everything is aligned at any point in time will be critical for effective management. This screenshot (Figure 5) includes a detailed view of the storage network environment as well.

Figure 5. Detailed End-to-end View



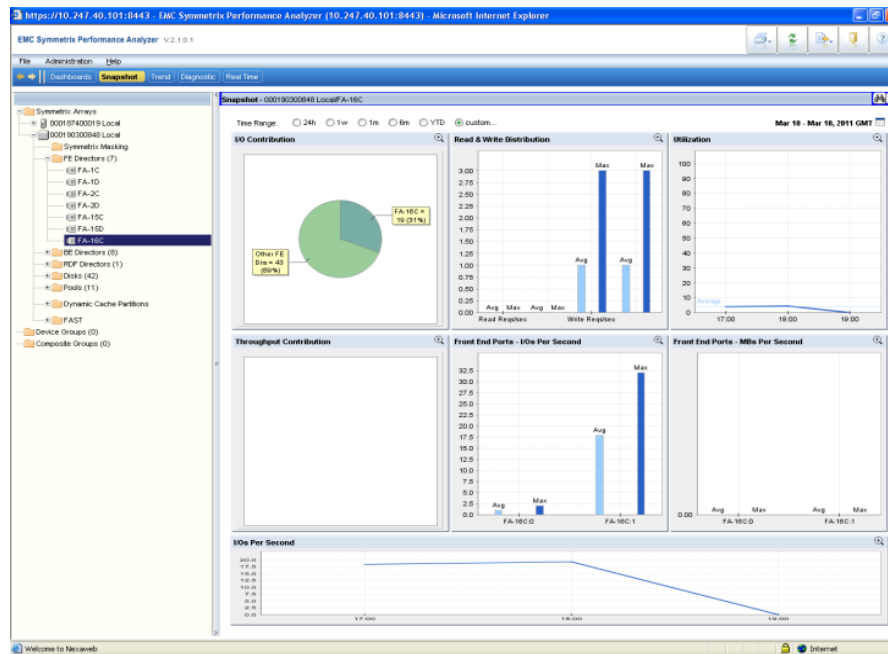
Source: EMC, 2012.

- **Ease of deployment, value extraction, and management.** As mentioned, ProSphere is neither a rebranding effort nor an upgrade to existing technology. It is a new architecture designed from the ground up to handle cloud architectures. As part of that redesign, EMC placed a heavy emphasis on ease of deployment and rapid time to value. The result is software that discovers the environment without an agent. In fact, EMC leverages an innovative discovery model that allows organizations to find not only known servers, but servers that were deployed but not documented. It does this by searching and auto-discovering an IP range. This software is deployed as a virtual application which can be done in less than 60 minutes. The ability to rapidly deploy will enable organizations to recognize value from the software on the very first day and, in highly dynamic environments, every day after that. EMC also estimates that organizations could save as much as 75% on the infrastructure required to host the ProSphere software in this new model.
- **Common look and feel.** EMC ProSphere is part of an effort to deliver a common interface and flow across all of EMC's software products. Thinking along the lines of the usability of Microsoft Office, this could be referred to as the "principle of least astonishment." Basically, organizations that are familiar with any new EMC software should be able to transfer that knowledge base to leveraging ProSphere. This should save time and money on training. It also means users should be able to generate reports more easily and begin

leveraging the software to make intelligent business decisions faster than before. For example, an organization currently using Storage Configuration Advisor will find ProSphere to be very familiar.

- Tight integration with other management and domain tools.** Recognizing that highly virtualized and private cloud environments are highly interdependent, ProSphere is tightly integrated with virtual management consoles like VMware’s vCenter to extract detailed VM information to provide end-to-end insight. Version 1.5 tightens the integration with storage array software like EMC’s Symmetrix Performance Analyzer (SPA), shown in Figure 6, and Symmetrix Management Console (SMC), allowing launch in context and single sign on as well as integration with FAST VP to provide detailed array level information and enable service level policies. Organizations can also launch the Brocade Network Advisor in context to drill down into greater detail on the FC fabric.

Figure 6. Symmetrix Performance Analyzer Launch in Context for Detailed Array Analysis



Source: EMC, 2012.

- Higher levels of control at cloud scale.** The ability to retain complete control of all storage assets while scaling rapidly will be critical for organizations to successfully transition to a private cloud. EMC’s ProSphere was designed to accelerate the transition from legacy environments to the cloud and as such is capable of scaling to accommodate very large environments. It is also able to federate data from multiple locations with a simple synchronization method that is easy to set up and enables a holistic enterprise-wide view. From a scale perspective, the first version of EMC ProSphere is capable of managing 18,000 hosts, 36,000 SAN ports, and over 1,200,000 volumes. EMC anticipates future releases will scale to even higher numbers. However, in order to make all that data more easily consumable, devices can be quickly grouped into smaller segments, such as by application, business unit, or even project, by users.
- The ability to scale the number, and type, of devices.** Recognizing that the pace of change in IT is also accelerating in the cloud computing era, ProSphere has been designed to accommodate new technologies more quickly using two methods: first, by integrating with a device’s element manager and leveraging launch in-context for detailed analysis. Second, it leverages industry-standard protocols like SNMP and SMI-S and is augmented where necessary with vendor APIs.
- Improved reporting capabilities.** ProSphere will enable organizations to have greater insight and more flexibility, including the ability to generate reports more easily. This includes a simple and intuitive dashboard that can be easily customized. Also by leveraging the explore function, users can create detailed

reports on arrays, LUNs, replication and storage pools. Figure 7 is an example of a report based on storage pools.

Figure 7. ProSphere Reporting Leverage the Explore Function

Pool	Pool Type	Serial Number	Pool Usable (GB)	Pool Free (GB)	Pool Used (GB)	% Pool Used	Pool Subscribed (GB)	% Pool Subscribed
BMTTestPool...	All	287 000194900287	12.63	6.31	6.31	50.00	11.23	88.94
BronzeThin	All	287 000194900287	1,279.90	1,145.51	134.38	10.49	1,783.28	139.32
CSE_ESX...	Local Replica Mirror Pool	APM00113101619	4,286.35	4,274.33	12.02	0.28	4,096.00	98.17
GoldThin	RAID Group	287 000194900287	1,279.90	1,087.18	192.72	15.05	1,805.93	141.09
Large_Pool	Snap Pool	287 000194900287	3,124.87	3,122.76	2.11	0.06	6.49	0.20
Large_Pool	Storage Pool	854 000194900854	0.00	0.00	0.00	0.00	0.00	0.00
Local mirror...	Thin Pool	APM00113101619	0.00	0.00	0.00	0.00	0.00	0.00
Local mirror...	Thin Pool	APM00113101619	0.00	0.00	0.00	0.00	0.00	0.00
ProSphereP...	Thin Pool	APM00113101619	3,205.24	2,432.74	772.50	24.10	3,072.00	98.46
RAID GROU...	RAID Group	FNM00093200345	3,668.59	668.59	3,000.00	81.77	0.00	0.00
RAID GROU...	RAID Group	Iglad070	534.78	54.78	360.00	67.31	0.00	0.00
RAID GROU...	RAID Group	FNM00093200345	7,337.19	0.19	7,337.00	99.99	0.00	0.00
RAID GROU...	RAID Group	Iglad070	2,141.33	1,181.33	600.00	28.01	0.00	0.00
RAID GROU...	RAID Group	FNM00093200345	3,668.59	3,600.59	68.00	1.85	0.00	0.00
RAID GROU...	RAID Group	FNM00093200345	7,337.19	803.19	6,534.00	89.05	0.00	0.00
RAID GROU...	RAID Group	FNM00093200345	0.00	0.00	0.00	0.00	0.00	0.00
RAID GROU...	RAID Group	FNM00093200345	3,668.59	368.59	3,300.00	89.95	0.00	0.00
RAID GROU...	RAID Group	FNM00093200345	3,668.59	598.59	3,070.00	83.68	0.00	0.00

Source: EMC, 2012.

- Consolidated alerting.** Enhancements to v1.5 include enhanced alerting capabilities which include the ability to set alerts on more than just performance issues, but also extend them to capacity levels, such as threshold alerts for thin Group. It also has the ability to aggregate capacity and availability alerts from Symmetrix and VNX arrays. All of these alerts can be forwarded to a management framework. An example of the consolidated alerting can be found in Figure 8.

Figure 8. ProSphere Alert Dashboard

Severity	State	CI Type	CI Name	Source	Message	Age
CRITICAL	ACTIVE	Array	000194900149	SMC(SMI-S)	Component state has changed to Online. Object is SB...	3 days 22 hrs
WARNING	ACTIVE	Array	https://lglac170.lss...	SMC(SMI-S)	Device configuration has changed.	4 days 5 hrs
WARNING	ACTIVE	Array	000194900287	SMC(SMI-S)	Component state has changed to Online. Object is SB...	5 days 2 hrs
WARNING	ACTIVE	Array	000194900287	SMC(SMI-S)	The configuration of a Virtual Provisioning Pool has ch...	5 days 8 hrs
WARNING	ACTIVE	Array	000194900287	SMC(SMI-S)	Device configuration has changed.	5 days 12 hrs
WARNING	ACTIVE	Array	https://lglac170.lss...	SMC(SMI-S)	Device configuration has changed.	6 days
WARNING	ACTIVE	Array	https://lglac170.lss...	SMC(SMI-S)	The configuration of a Virtual Provisioning Pool has ch...	6 days 1 hr
WARNING	ACTIVE	Array	https://lglac170.lss...	SMC(SMI-S)	Component state has changed to Online. Object is SB...	6 days 2 hrs
WARNING	ACTIVE	Array	FNM00093200345	ProSphere	Storage Processor dirty pages for SP-B is at 99%	6 days 10 hrs
WARNING	ACTIVE	Array	000194900854	SMC(SMI-S)	Component state has changed to Online. Object is SB...	6 days 15 hrs

Alert	CI
Message: Component state has changed to Online. Object is SB-1/ENCS-4/MM-B Category: Health Source: SMC(SMI-S) Severity: CRITICAL # Occurrences: 11	CI Type: Array CI Name: 000194900149

Source: EMC, 2012.

The Bigger Truth

As organizations transition to highly dynamic environments and build out private clouds, management will play an increasingly important role. While many are solely focused on server virtualization management, it will be equally important to have comprehensive management in place for the storage and the storage network. The ability to ensure service levels and demonstrate end-to-end performance will help to accelerate the transition to the cloud.

EMC purpose-built ProSphere to make it easy to manage the transition across physical, virtual, and storage cloud environments. With version 1.5, it provides an end-to-end view of the environment, capacity reporting, performance information, and tight integration with other domain solutions to enable a successful entrance into a new era of computing. Its architecture will help lower ownership costs and make it easier for IT to deliver value. ProSphere v1.5 will enable organizations to more effectively control their data growth and thereby rein in storage costs, improve their storage return on investment, and align resources with the appropriate tier of storage. Ultimately this will enable organizations building out private clouds the ability to provide showback or chargeback numbers to individual business units.

Maturing virtualized computing environments require solutions that can scale, provide end-to-end views, and easily integrate with other solutions. Only with a complete understanding of all of IT's domains will organizations be able to take advantage of the benefits cloud computing can provide. In short, EMC's ProSphere v1.5 delivers the requisite capabilities to effectively manage and optimize cloud storage environments for a new cloud computing era.



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