

Architecting Resilient Private Clouds

A New Approach to Transform Your Existing Data Center into Your Private Cloud

VR Satish, CTO, Storage and Availability Management Group

Architecting Resilient Private Clouds

A New Approach to Transform Your Existing Data Center into Your Private Cloud

Contents

Executive Summary	1
Private Compute Clouds	1
Private Storage Clouds	5
An Ideal Architecture: Servers and Storage Together	9
Building a Resilient Private Cloud	9
Conclusion	11

Executive Summary

This white paper proposes an approach to building resilient private clouds in an efficient and cost-effective manner. This idealized architecture incorporates solutions that address cloud server and storage needs and reconciles them with the realities of the data center such as heterogeneity and existing infrastructure. This paper also discusses how these solutions help data centers achieve the operational flexibility of a cloud-like model without compromising the service-level agreement (SLA) requirements.

Private Compute Clouds

From Virtualization to the Cloud

Virtualization of the server is the first step to creating an agile infrastructure that enables data centers to operate with the scale and economy of clouds. However, there is still a long way to go before Tier 1 or even Tier 2 applications can be moved into a virtual environment with confidence.

Discussions about virtualization in the IT industry inevitably focus on x86 servers with VMware®, XEN® or Kernel Virtual Machine; however, there are many Tier 1 applications that rely on traditional big iron UNIX® servers. For many of these applications, major UNIX vendors are offering virtualization in their own platforms. Solaris® Logical Domains (LDMs) are maturing at a rapid pace. While Solaris® Zones are already being used as a default in many environments, IBM Power® machines operate very efficiently with a large number of Logical Partitions (LPARs). Applications such as databases are moving into x86 environments, though not as fast as many want. The existence and use of multiple virtualization technologies means that discussions about the cloud need to look beyond x86 and consider how a heterogeneous cloud can be built to accommodate the workloads of a Tier 1 application.

Existing data center investment considerations should take into account that real-world cloud architectures will be heterogeneous with a mix of virtual and physical platforms with the applications that run on them. By design, these cloud architectures will also be able to accommodate larger scale than existing data centers. These realities dictate the following needs for servers in the cloud:

- **Scalable architectures**—Creating architectures and distributed availability components that allow efficient and easy scaling.
- **Holistic availability**—Availability across physical and virtual environments spanning heterogeneous vendors and applications; providing availability from the application layer to the database and coordinating it across the entire service to ensure actual availability.
- **Automation**—The scale and availability needs of the cloud mandate that you automate wherever possible to ensure predictable results and efficiency.

Scalable Architectures

Clouds establish operational efficiency (that is, services are delivered on-demand and IT resources are pooled and shared) to achieve agility and economies of scale. Virtualization enables the ability to pack more operating environments into a single physical machine and at the same time ensure isolation. While many applications or platforms include some degree of built-in resiliency, the missing piece in many of the solutions is availability that coordinates with the rest of the system. Symantec partnered with VMware to provide Symantec™ ApplicationHA, which provides application visibility and control in virtual environments. With the vast portfolio of application-specific monitoring software that was developed for Veritas™ Cluster Server, Symantec now solves the problem that limited the migration of business-critical applications to virtual environments. In virtual environments, most infrastructure availability solutions are unable to monitor the health of applications and automate their recovery in the event of a failure. The combination of Symantec ApplicationHA and VMware™ HA helps IT organizations keep critical applications up and running.

In essence, the key technological advancement is the splitting of infrastructure HA from the application HA. This is similar to what virtualization does; that is, split a physical machine from the operating system and eventually the application. Specifically, in virtual environments such as VMware or IBM PowerVM™ the control operating system is subject more and more to lock down. The platform, with its built-in availability software, ensures that the platform itself is available.

One of the key benefits of Symantec ApplicationHA is that it provides autonomous and distributed availability in the virtual environments while coordinating with infrastructure availability, whether it is the native tool or Veritas Cluster Server. By disassociating the levels of availability and making them autonomous, a foundation of distributed availability is created. This provides the building blocks for multi-tiered applications and business services that define a cloud.

Multi-tiered Apps and Business Services

There needs to be a layer that can monitor the various components of a multi-tiered application in addition to the platforms on which they run on. This layer must provide resiliency for an end-to-end IT service. However, this can be challenging when there's a multi-tiered application that crosses platform boundaries because more than one technology must be managed and it must all be tied together manually. The layer of application HA on top of the platform technology needs to abstract the infrastructure HA and provide a uniform visibility of the health of the application component monitored by it. When tied together, these application components provide health to a complete business service.

With this in mind, Symantec is building application HA capabilities for many virtualization technologies that work with either the infrastructure HA provided by the vendor or Veritas Cluster Server. Symantec plans to roll out these application HA solutions in upcoming major product releases.

Clouds tend to commoditize the hardware and provide the benefits in software; resiliency is essentially provided by software rather than by hardware. This paradigm holds true even when mixing physical and virtual server environments.

Architecting Resilient Private Clouds

A New Approach to Transform Your Existing Data Center into Your Private Cloud

A pictographic representation of this concept is shown below:

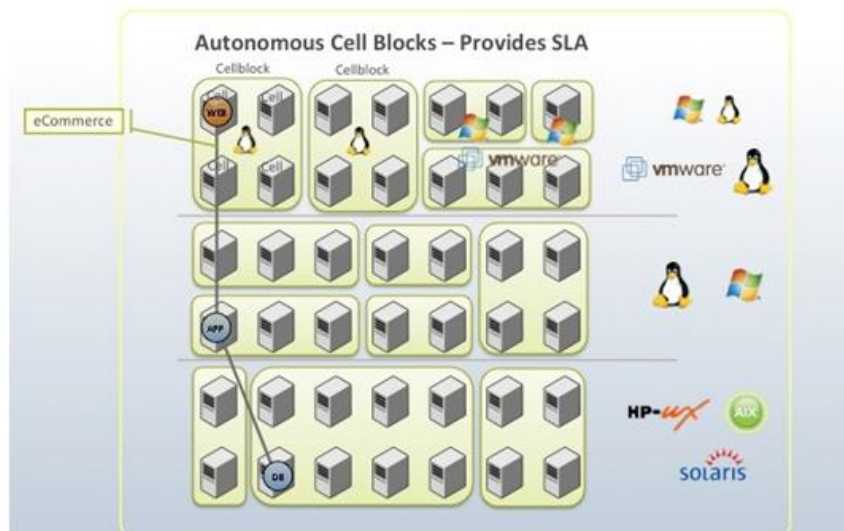


Figure 1: Typical multi-tiered application environment

Figure 1 shows a heterogeneous farm of servers. Some of the servers are x86 running VMware with Windows® or Linux® as the guest. Other physical servers are running Linux or Windows. The bottom layer consists of big iron UNIX servers. A business service called "eCommerce" is represented as a three-tier application. It consists of a Web server acting as the front end to the application, an application tier with an application server like JBOSS, and a database tier that runs on a Solaris machine.

In today's world, if this business service—eCommerce—were to be made highly available, a likely scenario would mix and match technologies for the various tiers and each one would act in a silo. There is not an easy way to uniformly start the eCommerce application or bring it down. There is no way to failover the entire eCommerce application from one site to another. Many models try to solve this using a centralized approach where a central server heartbeats with each of the physical/virtual machines and manages availability. These heartbeat solutions may have scalability problems, limiting their usefulness in a cloud-like environment where the benefits are derived from large scale pooling.

To solve this problem, we define a few fundamental concepts called "cells" and "cell-blocks." Conceptually a cell-block is similar to today's a clusters; it is homogeneous. The nodes in a cluster are like cells. Each cell-block provides availability in an autonomous fashion. If a cell were to fail in a cell-block, the cell-block is still capable of keeping the applications available because they failover from one cell to another cell. The virtual business service (VBS) called eCommerce is actually a virtual cluster that spans across multiple cell-blocks. This is analogous to a VLAN and LAN, or a VSAN and SAN. Any cell that is part of a VBS is capable of answering questions regarding the state of eCommerce.

Architecting Resilient Private Clouds A New Approach to Transform Your Existing Data Center into Your Private Cloud

Figure 2 below shows what happens in the event of a cell failure.

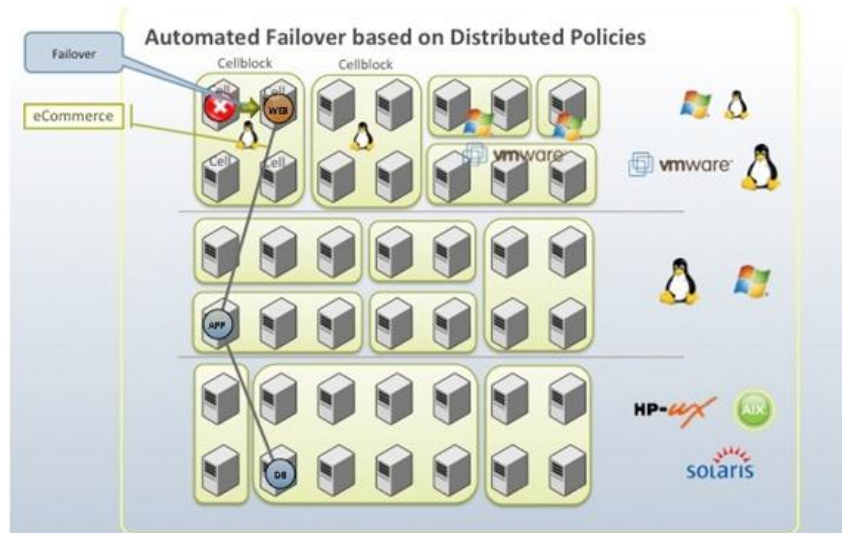


Figure 2: Automated failover in a multi-tiered application environment

By defining VBS with cells and cell-blocks as the foundation, a server farm can host applications that require a high level of SLA and manage applications holistically as IT services. The environment is now self-healing and self-monitoring, which are fundamental properties of the cloud. The vital element is a uniform layer that monitors applications and integrates with the infrastructure that's needed to build such an environment.

Let's consider three possible topologies: (1) a cell block with a typical non-virtualized two-node cluster, (2) a cell-block that's virtualized using ApplicationHA and that delegates infrastructure HA to a third-party provider (for example, VMware HA), and (3) a cell-block that's virtualized using ApplicationHA and that delegates infrastructure HA to Veritas Cluster Server running as a dedicated guest. In all three cases, there is a common underlying concept of service groups regardless of the technology. It is this common concept—along with resources interconnected by lanes across various cell-blocks—that creates a VBS.

Orchestration and Automation

To complete the picture, provisioning and an orchestration layer must be in place. This is what automates the environment. Automation requires a standard set of APIs, which are typically tied through a workflow engine that defines the workflow based on a data center's business practices. In the absence of a standard layer, you would need to hook into various discrete technologies and always be playing catch-up with the latest innovation.

The next step is to automate server provisioning. Organizations tend to stop at ease of deployment of a virtual machine and declare that they can now do on-demand provisioning. One of the key challenges is determining where to place the virtual machines required for the VBS.

Architecting Resilient Private Clouds

A New Approach to Transform Your Existing Data Center into Your Private Cloud

To address this issue, a heterogeneous monitoring framework as well as a data collection framework are required. As you assess key metrics, consider the following questions:

- What are the various OS instances that constitute the VBS?
- What kind of availability is required for the various components that constitute the VBS?
- What kind of storage is required (for example, replicated, tier 1, size, consistency group requirements etc.)?
- What kind of monitoring is required?
- What is the current state of the server farm? For example:
 - Central processing unit (CPU)
 - Memory
 - Resiliency

Once you have answered these questions you can designate a location for the virtual machines. The metrics listed above are by no means complete; however, they can be expanded based on data center topologies. Though several mechanisms exist to solve this min-max problem that is NP-complete in nature, in this case optimality is not a requirement; a heuristic approach is more than sufficient. Mechanisms for gathering all the metrics and making them centrally available for answering such questions will likely suffer from scalability issues.

Symantec is addressing these issues with a distributed map-reduce algorithm that determines the best fit for placement heuristically. It uses a self-learning and forecasting technique to find the available capacity of each host. Veritas™ Operations Manager, which provides a single pane for end-to-end visibility, will be the interface for this placement feature. Template-based storage provisioning is currently available in Veritas Operations Manager; similar capabilities for availability and server provisioning are expected in future releases.

Private Storage Clouds

To make any cloud solution viable, the storage infrastructure needs to deliver on the Storage as a Service paradigm, which adheres to the following requirements:

- On-demand storage with elasticity
- Controlled storage costs
- SLA driven
- Chargeback enabled
- Security at the source

On-Demand Elastic Storage

Clouds are based on the fundamental premise of operational efficiency through the pooling and sharing of resources. This requires discipline from the applications and infrastructure to use only what is required and then return resources that are not used to the pool. In the context of on-demand storage, new techniques are being adopted to meet this need by providing the ability to provision storage on the fly in an efficient manner. One of the technologies that enables this vision is thin provisioning. By provisioning storage on an as-needed basis and utilizing the benefits of pooling, the amount of

storage consumed by applications can be significantly reduced and storage can be delivered in a timely manner. A logical unit number (LUN), for example, 10 GB LUN or 100GB LUN, etc., will become obsolete. A LUN should be thought of as a conduit to a pool of storage and in particular to one that satisfies a specific SLA requirement. With thin provisioning, wide striping is done in the array, so there is no need to import hundreds or thousands of LUNs on a host for performance. Instead, a LUN represents a specific tier of storage. This reduces operational expenses and increases utilization substantially, and its allocation on-demand model provides the functionality of Storage as a Service.

Veritas Storage Foundation™ from Symantec improves upon storage hardware's thin provisioning feature to deliver an on-demand model of storage by providing the storage pool and providing the ability for applications to consume storage as needed. Thin provisioning by itself provides some of the benefits we discussed. However, the ability to return unused storage for use by other applications requires intelligence that can only be provided at the application host server. Thin reclamation, a feature of Veritas Storage Foundation, provides the ability for data centers to reclaim storage once the space is no longer used by any application, helping maintain operational efficiencies on an ongoing basis.

Controlling Storage Costs

One of the common practices that existed in the many legacy data centers was the allocation of top tier storage for all applications, regardless of the requirements of the application. This was usually done to ensure that the minimum SLAs and performance were met for any application. It also usually reduced administrative overhead. However, in the private cloud, this strategy doesn't make sense. Storage vendors now routinely include multiple storage tiers within a single array. For that reason, it does make sense to put data on the appropriate class of storage based on a set of criteria that can be easily configured. A software solution such as Storage Foundation has the ability to combine multiple tiers of storage across physical arrays and dynamically move data between those tiers based on business value as well as SLA or performance requirements. This ensures that the storage cost profile of data matches the business value.

SLA Driven Storage Provisioning

Just as availability on the server side is critical to the success of the cloud initiatives, storage SLAs will play a significant role in the success of the cloud. It is not enough to be able to create storage on demand; the right kind of storage based on the needs of services and applications is also necessary. For instance, some business-critical applications may need gold storage with enhanced RAID support and extremely fast solid-state drive (SSD) storage, while others can make do with slower Serial ATA (SATA) storage. Given a non-critical application, the highest SLA storage would be excessive and could result in increased costs without significant benefits.

Chargeback Enabled

Long-term viability of the cloud solutions depends on having a sustainable model wherein services and resources are tracked and charged back to the consumers based on their usage. Chargeback also helps control the storage sprawl by creating incentives on the demand side to reduce storage usage. Veritas Operations Manager provides storage chargeback capability across heterogeneous environments, including commodity storage, giving visibility into storage usage and enabling chargebacks.

Security at the Source

It's important to use the network-attached storage (NAS) gateway as an offload engine for performing antivirus security scans, enforcing compliance of virtual images, and for data classification. This allows for a single point of management for these security-oriented services. When such a device works in tandem with a corresponding piece of software running resident in the guests, the overall efficiency of the environment improves. Offloading storage in this way is essential because that's where the data is resident. This allows scanning subsystems to take advantage of the layout and functionality of file systems such as shared extents and deduplication.

In addition to the requirements listed above, a storage cloud should be designed to handle the workloads expected for the applications hosted and provide security at the source.

Ultimately, the design of the cloud must be tailored and molded to its needs. It is important to consider the kind of workload that will be addressed before embarking on a private cloud. The workload determines the access protocols, the amount of flexibility needed with respect to connectivity to server, the scale of storage, and the backup strategy associated with it. Organizations sometimes want to move data centers completely to a virtualized model; however, certain workloads and availability assurances prevent that from occurring. A hybrid approach that trades flexibility with performance is usually preferred and more effective.

Demands of Virtualization on Storage

As virtual machines move (as in relocating a virtual machine to another physical host), connectivity must be preset so that the new host can see the storage required by the virtual machine it is going to host. Network File System (NFS) as an access protocol has received a lot of attention, particularly in a virtualized environment, due to its simplicity in configuration and connectivity management. iSCSI is another protocol that is used in the virtual environment. However, there is an insistence on Fibre Channel protocol for Tier 1/Tier 2 workloads, and in many cases rightly so. The SLA assurances of a SAN fabric have been high when compared to IP networks, and moving storage requires a net new fabric that is designed for storage and bridging IP to FC SANs.

The private cloud must be architected so that it can satisfy the requirements for the targeted workloads. At the same time, storage utilization must improve, and operational and capital costs must decrease. Let's consider boot images. NFS is becoming predominant for boot images, particularly for virtual desktop infrastructure (VDI). The reason for this is ease of management and performance. Data center administrators find that with Fibre Channel-based storage solutions, deploying large scale VDI is difficult. Concerns such as zoning storage, live mobility (such as VMware vMotion™), storage cost, and performance all become challenging in a Fibre Channel world. NAS, however, provides a welcome abstraction because anywhere there is network connectivity, virtual desktops can be launched. Furthermore, NAS solutions are file-based, making it easier to deduplicate common data between desktop images as well as provide various performance accelerators.

As virtualization proceeds at its current pace, between 60 and 70 percent of all servers are likely to be virtualized by 2015.¹ A good scale-out and scale-up NAS solution is required to support such large-scale deployments of virtualization. The recently launched Symantec™ VirtualStore product is suitable for these virtual environments and provides IP network

1-Gartner Data Center Conference 2010: Phillip Dawson and Ray Paquet, The Virtualization scenario: 2010 to 2015

connectivity while leveraging Storage Area Network (SAN) storage. Traditional NAS solutions, however, have one drawback; they typically consist of a prepackaged set of controllers and storage. This coupled bundle of hardware does not usually scale well because the controllers saturate long before the NAS filer reaches its maximum capacity. To solve this problem, vendors are looking to creative software-based clustered NAS solutions.

The requirements for a truly high performance, large scale VDI environment include:

- Independently scalable controllers and storage
- Fast boot image provisioning powered by an innovative caching technique
- Some form of deduplication or compression to eliminate redundant data across virtual images
- Integration with the management consoles used to provision new virtual machines
- Ability to use commodity servers and storage with the performance of high-end storage
- All of the NAS heads should be able to serve up any image for efficient load balancing

NAS filers satisfy many of the requirements above, except typically they cannot decouple the controller from the storage and often times require expensive SSD modules for performance. Software tools solve this problem by using inexpensive blades or in some cases virtual appliances running on the hypervisor to accelerate performance. While a NAS filer has controllers that are embedded, software NAS controllers can be placed almost anywhere in a VDI deployment. The closer the NAS controller is to the hypervisor, the better performance characteristics one can expect. With a traditional filer, accelerators in the NAS controller itself are embedded in the filer and can only provide acceleration after data has travelled from the hypervisor, over the network, and into the NAS filer. With a software virtualization layer running on inexpensive x86 hardware in front of your storage array, input/output (IO) will flow from the hypervisor over the network to the software NAS controller where many of the requests will be served directly out of cache and never hit the underlying storage array or disks. If we move the software NAS controller into a top-of-rack configuration, for example, within a Cisco™ Unified Computing System (UCS) block, data requests are now served over a high speed backplane within the rack. In this configuration, much of the IO requests for VDI can be served out of cache within the rack, and the requests never actually leave over the network. The best performance can be achieved when the software controller resides in a virtual appliance running on the hypervisor itself. In this case, much of the data requests can be served directly out of cache on the same host as the virtual machines are hosted, and therefore never even leave the physical server.

Storage Clouds

The emergence of mobile computing platforms with cloud backends has seen the birth of storage clouds that can handle scale-out workloads and huge unstructured data. Scale-out workloads and unstructured data repositories such as storage for analytical processing and content management systems need a different type of access. These are read-intensive and can access storage using HTTPS—a new technology designed for modern Web-based technologies. A lot of content analysis can be done because access is file-oriented. Typically, large quantities of data are present in these kinds of workloads. A new breed of middleware-oriented storage layer is a requirement—a layer that federates amongst various file systems and provides a uniform namespace through HTTPS as the access mechanism. Several large public clouds implement such a system. One of the predominant use cases for such a system is media archiving and backup. These use cases are read-intensive, with many reads and a few writes. These types of storage are unfit for traditional workloads and are typically associated with “storage clouds.” Retrofitting a file system-based access semantics is unnecessary and can

Architecting Resilient Private Clouds

A New Approach to Transform Your Existing Data Center into Your Private Cloud

cause complexity and scalability implications on the architecture. Because storage consumption resulting from unstructured data is increasing at an alarming rate, it is necessary to reduce the cost of storage hardware and move towards commodity storage rather than building it out using Tier 1 storage subsystems.

An Ideal Architecture: Servers and Storage Together

A typical private cloud buildout should take into account both storage and servers. To summarize, here is an example of an ideal architecture:

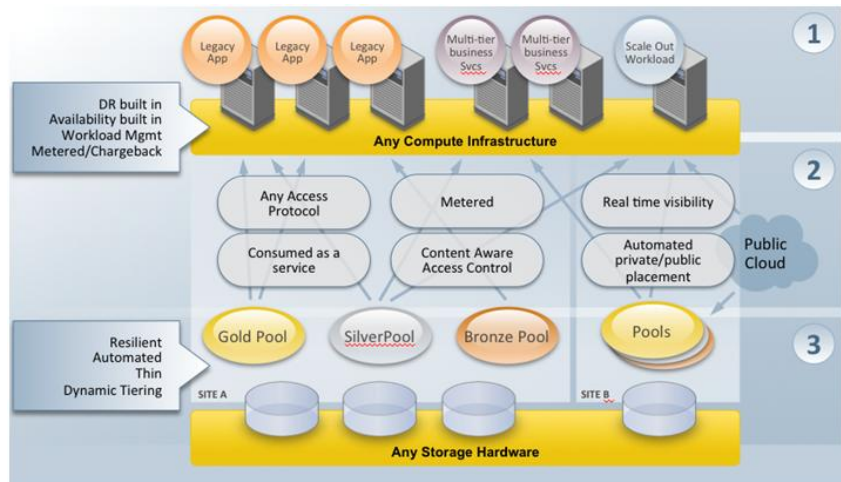


Figure 3: An idealized cloud architecture

In figure 3, layer 1 includes servers that are built to be resilient and can provide the SLA required for multi-tiered applications. This environment also has requirements for managing traditional workload applications. The architecture should support scale-out workloads for analytical processing. These applications should be capable of accessing storage in an agile and secure mechanism using the protocol that is most suitable for the applications. An SLA-oriented storage service that is capable of moving data from one tier to another is based on the context of data and performance requirements. Any type of storage based on an application requirement should be capable of being used in this environment.

Ultimately, the private cloud will be bridged to the public cloud through an access gateway. In short, the whole environment should be heterogeneous from the server to storage, and cost should be one of the major driving factors. All of this is managed by an orchestration subsystem that is capable of provisioning virtual machines and analyzing the environment along with tools that ensure the SLA of the application.

Building a Resilient Private Cloud

Symantec's existing portfolio of products already provides the building blocks to enable a resilient private cloud. In order to build an effective orchestration mechanism, a true layer that bridges the server and storage subsystems is required. This layer needs to have an interface to allow integration of existing storage and server solutions. Symantec is hardware

Architecting Resilient Private Clouds

A New Approach to Transform Your Existing Data Center into Your Private Cloud

and platform agnostic, and is heterogeneous in regard to physical and virtual server environments for example server, availability, virtualization, and storage, positioning it well to solve this problem.

In a key partnership with VMware, Symantec ApplicationHA software works with VMware HA (infrastructure HA) and ensures that when software failure happens, the infrastructure is notified and appropriate action is taken. For environments that do not use VMware, Veritas Cluster Server is the market-leading high availability software that provides infrastructure HA services. This allows for a common mechanism to represent business services and its dependencies, regardless of the underlying HA technology, and it facilitates orchestration.

Veritas™ File System from Symantec is the only file system that runs on multiple operating systems that can perform thin reclamation. Reclamation is a fundamental aspect of leveraging thin provisioning to provide storage as a service. It is important that file system is thin-friendly, context aware and is relevant to the storage subsystem so that reclamation can happen. For example, if a file is deleted in a file system, only the file system can tell the storage subsystem that the space can be reclaimed. It is the file system that has to understand the underlying reclamation unit size and properly lay out its data so that free space can be produced by automatically defragmenting. One of the most important functions of Storage Foundation is to be a conduit and provide context and intelligence to the storage subsystem so that the storage subsystem can operate effectively and efficiently.

The Symantec™ VirtualStore is a product that is built as a solution using Cluster File System technology that can allow for serving a large number of boot images as well as providing storage over iSCSI and NFS. File level snapshots is a feature that allows for efficiently taking snapshots of virtual machine images to create more virtual machines that are space-optimized (only the deltas are captured). Data centers that are thinking of providing a VDI-like offering should look at this to scale both horizontally (more throughput and less latency) and vertically (more storage) independently. VirtualStore is a software-based solution that can be used on any commodity server along with any type of storage to provide a scalable cost-effective solution for serving boot images.

In addition to traditional workloads and boot images, Symantec also envisions building a middleware layer that federates instances of VirtualStore to provide a scale-out, object-based data repository. This middleware layer will be resilient by providing built-in object replication and versioning thereby allowing redundant copies for redundancy purposes. The middleware layer will provide hooks to enforce content analysis and security enforcement to maintain privacy and multi-tenancy.

Dynamic multi-pathing is another key enabler for Symantec's vision of enabling the private cloud. It is the market-leading multi-pathing solution that works with almost every storage subsystem and most operating systems. It is this layer that acts as a bridge between the file system and storage subsystem.

To tie all these pieces together, a centralized management solution called Veritas Operations Manager was developed. This allows for end-to-end visualization and a solution-oriented set of operations for increasing operational efficiency. One of the key features of Veritas Operations Manager is the ability to create and provision file systems and volumes based on templates. For example, you can have a template for a multi-tier application with three sub templates—one for each tier—and allow for careful picking of appropriate storage based on the Tier's requirement. This feature, when coupled with virtual machine provisioning and automatic set up of application monitoring and availability, is the key to orchestrating a private cloud through a simple and standard software layer.

Conclusion

Symantec understands that many organizations are struggling to visualize how their specific environments can benefit from the cloud paradigm. The primary challenges organizations face include:

- How do we figure out the public vs. private cloud models and which one is right for us?
- What kind of architectures can help us get to the private cloud?
- How should I think about my existing infrastructure and how can I leverage it?

Symantec understands these challenges. Our leadership position in the data center software solutions space provides insight into the needs of our customers and data centers in general. Our solutions will help you to leverage your existing infrastructure and realize the benefits of the cloud.

For more information about Symantec's solutions for the private cloud, visit <http://go.symantec.com/cloudavailability>.

About Symantec

Symantec is a global leader in providing security, storage, and systems management solutions to help consumers and organizations secure and manage their information-driven world. Our software and services protect against more risks at more points, more completely and efficiently, enabling confidence wherever information is used or stored. Headquartered in Mountain View, Calif., Symantec has operations in 40 countries. More information is available at www.symantec.com.

For specific country offices and contact numbers, please visit our website.

Symantec World Headquarters
350 Ellis St.
Mountain View, CA 94043 USA
+1 (650) 527 8000
1 (800) 721 3934
www.symantec.com

Symantec helps organizations secure and manage their information-driven world with storage management, email archiving, and backup and recovery solutions.

Copyright © 2011 Symantec Corporation. All rights reserved. Symantec, the Symantec Logo, and the Checkmark Logo are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.
5/2011 21193123