

Storage Virtualization a Work in Progress

Already reaping utilization and cost benefits from virtualizing storage, Cisco IT awaits accomplishing even more.

Storage at Cisco® will be considered virtualized when all the storage in a data center, or among data centers, is in one pool completely abstracted from the applications it serves. This is a milestone in data center virtualization that Cisco IT sees as crucial—and is eagerly waiting to achieve.

So far, the IT group has virtualized storage to the level of creating monolithic 300- to 500-terabyte storage arrays containing up to 960 disk drives. Within these arrays, logical units are created as a single physical device, part of a physical device, or across many devices. This abstraction from physical to logical has been used in storage arrays for many years, to increase reliability (e.g., mirrored or parity

Increasing utilization of storage devices by 1 percent saves Cisco about US\$1 million annually.

protection) or increase the performance over a single physical device (e.g., striped META devices). New array software advances are taking logical abstraction to the next level, says Rich Harper, senior storage manager at Cisco, and allowing for technologies such as data compression, thin provisioning,

oversubscription, and data de-duplication, which present a greater amount of usable storage to end users than might actually be physically present in the array.

Cisco IT's use of storage virtualization goes a bit further. Virtual SAN (VSAN) technology, supported by the Cisco MDS 9500 Series Multilayer Director Switch, partitions a single physical SAN into multiple VSANs, allowing different business functions and requirements to share a common physical infrastructure.

"The application host can only interface with the logical devices presented to it on its own virtual fabric," explains Harper. "There is a logical separation even though multiple applications and drives are managed by the same switch." In addition, he says, many of Cisco IT's internal clients now view storage as an IT-provided service rather than a resource owned and managed by the application users.

The ability to pool storage across storage arrays and entire data centers is something Cisco IT must wait for. Why? The cross-array virtualization software available today, says Harper, does not have all of the capabilities needed to move from virtualized storage arrays to virtualized storage across the entire data center.

First, Pooled Arrays

In the 1990s, Cisco, like many other companies, used direct-attached storage (DAS), but DAS was severely limited in scalability and flexibility. For example, Cisco's enterprise resource planning (ERP) software required more than 400 ports into storage facilities, and that, in turn, required a complex, difficult-to-manage architecture of DAS subsystems.

In 1998, Cisco IT started moving the ERP applications to Brocade and McDATA departmental SAN switches, which made it easier to share storage among multiple hosts. This approach was an

improvement over DAS, but it still left some storage arrays “stranded” in SAN islands. “Available storage could not be used by applications on a host that was inconveniently located in a separate SAN island,” Harper says.

In January 2003, Cisco IT put the Cisco MDS 9509 Multilayer Director into production. With up to 224 Fibre Channel ports in this single switch chassis, storage and hosts on individual SAN islands could be collapsed into a single physical infrastructure that was easier to manage and more cost effective than DAS and small, often underutilized SAN islands. So, the IT group began moving other applications to storage managed by the Cisco MDS 9500 Series. In 2003, IT also began using the Cisco MDS 9500 to further consolidate isolated SANs—which might be adjacent to each other in a data center—into larger ones.

Cisco IT now employs more than 100 of the MDS Series switches and maintains many large consolidated SANs. Included among them is a 2500-port SAN in Research Triangle Park, North Carolina, which spans multiple data centers via Fibre Channel and interconnects with Cisco data centers in San Jose, California, and Lawrenceville, Georgia, over the Internet on the company’s WAN.

Using a plug-in adjunct to the switch, the Cisco MDS Family Storage Services Module, Cisco IT is investigating multiple vendors’ specialized storage services within the SAN. The module incorporates application-specific integrated circuits (ASICs) and software that interfaces with various code, and can be used for a variety of storage services including SAN-based storage virtualization and continuous data protection (CDP) applications.

Second, Virtual SANs

The Cisco MDS 9500 Series offered Cisco IT two ways of consolidating SANs: simply linking the existing SAN infrastructure together through physical connections into a larger SAN, or VSAN technology to carve out logical SANs from a large MDS-managed physical SAN. Cisco IT used both techniques, the latter mostly when it wanted to partition storage and

dedicate the partitions to a specific storage function, e.g., separating development from production while providing each with what appears to be its own dedicated storage infrastructure.

The Cisco MDS 9500 Series helped form the foundation for broader virtualization, according to Harper, by not only separating the application from its storage, even though both link to the switch, but also by enabling VSANs. VSANs operating through the switch can reach across and between floors and even among buildings. With storage requirements doubling every year, Cisco IT has found that a common physical SAN infrastructure and logical VSANs are essential to reducing costs and keeping storage manageable.

VSANs also allow the IT group to provide different quality of service (QoS) levels and access speeds for different applications, ensure separation of one application and its data from another, and even provide tailored levels of services such as security.

Working with the fabric of a storage array, Cisco IT increases utilization by employing techniques such as oversubscription—intentionally overbooking within parameters that still ensure QoS—and also migrates data from one device to another transparently to enhance resiliency.

The advantages to supporting storage with the Cisco MDS 9500 Series switches are compelling: higher availability (applications have multiple paths into their stored data) and utilization; lower costs; simpler management; faster provisioning; and greater performance. According to Harper, Cisco IT currently has about 4 petabytes of storage on SANs. The total cost of ownership (TCO) for storage has dropped from US\$0.21 per MB to US\$0.012 per MB over six years, and Cisco is saving US\$5.5 million annually in storage maintenance.

Cisco VFrame Data Center (VFrame DC), a new service orchestration and provisioning solution, should bring additional cost savings “by further abstracting applications from storage and streamlining the process of giving each application the storage it needs,” says Kumar Ramachandra-Rao, senior storage manager at Cisco. After storage administrators create logical unit numbers (LUNs) for a given storage array supporting a SAN, he explains, VFrame DC will allocate the LUNs to applications and provision them. For more on Cisco IT’s testing and use of VFrame DC, see the article “VFrame Data Center First Customer” (Further Reading box, page 4).

Cisco IT also provides network-attached storage (NAS) for many applications that require NAS for accessing data. For consolidating NAS, the group uses techniques similar to those used for SANs, and will employ VFrame DC to create quota trees and provision partitioned storage.

Third, the Giant Step toward Virtualization

The next step on Cisco IT’s path to data center virtualization is to create multi-petabyte storage pools. And the group has a pressing need to accomplish this.

“Like Cisco, most companies depreciate storage hardware over a 30-month cycle, which means that we are intensively involved with data migration every three years,” says Harper. “We are due for a hardware refresh, and it would be so much easier with the right data migration tools.” But this is one example of the capabilities Cisco IT has not seen yet in the industry.

Cisco IT is also waiting for software that has capabilities at the data center level which the group now uses at the storage array level, such as oversubscription and data duplication within storage pools and the ability of virtualized storage to communicate with unvirtualized storage.

“The goal is to do anything we want at the back end—moving data in and out of the virtual pool or

from one storage device to another, backing up data, changing out devices—all to increase storage utilization,” says Ramachandra-Rao.

For good reason: Cisco IT estimates that increasing utilization of storage devices by 1 percent saves the company about US\$1 million annually. With the tools it wants at the data center level, Cisco IT estimates that it can boost utilization by 50 or even 100 percent over where it is today. To this end (and because of an impending hardware refresh), Cisco IT is planning a pilot of virtualization at the data center level that includes use of data migration software. Slated to be included in the pilot are EMC Corporation’s Invista SAN virtualization product and the Cisco MDS Data Mobility Manager, which enables movement of blocks of data from a source device to a destination device.

To other enterprises undertaking storage virtualization, even at the array or SAN level, Ramachandra-Rao and Harper recommend doing so with an eye toward reliability and stability. No rushing in. Pooling SANs is a good place to start, as is looking at larger storage arrays and non-hardware-based techniques, such as oversubscription, which can be done effectively today at the array level.

And, they emphasize, always bring application users into the planning process to help ensure that IT understands all of the subtleties of an application. This will be sound practice for the data center-level virtualization that is sure to come. ■

For more Cisco IT featured content and case studies, visit <http://www.cisco.com/go/ciscoit>.

FURTHER READING

- Cisco on Cisco Featured Content: “VFrame Data Center First Customer”
http://www.cisco.com/web/about/ciscoitwork/featured_content/fc_july_aug_2007_article02.html
- Cisco on Cisco Featured Content: “Server Virtualization on Fast Track”
http://www.cisco.com/web/about/ciscoitwork/featured_content/fc_july_aug_2007_article01.html
- Cisco on Cisco Case Study: “How Cisco IT Virtualizes Data Center Application Servers”
http://www.cisco.com/web/about/ciscoitwork/case_studies/data_center_dl2.html
- Cisco MDS 9500 Series Multilayer Director
<http://www.cisco.com/en/US/products/ps5990/index.html>
- Design Zone for Data Centers: Cisco Validated Designs for Data Center Networking
http://www.cisco.com/en/US/netsol/ns743/networking_solutions_program_home.html
- Cisco Data Center Networks Blog
<http://blogs.cisco.com/datacenter/>



Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Europe Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: +31 0 800 020 0791
Fax: +31 0 20 357 1100

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

©2006 Cisco Systems, Inc. All rights reserved. CCVP, the Cisco logo, and the Cisco Square Bridge logo are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn is a service mark of Cisco Systems, Inc.; and Access Registrar, Aironet, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, Follow Me Browsing, FormShare, GigaDrive, GigaStack, HomeLink, Internet Quotient, IOS, IP/TV, IQ Expertise, the IQ logo, IQ Net Readiness Scorecard, iQuick Study, LightStream, Linksys, MeetingPlace, MGX, Networking Academy, Network Registrar, Packet, PIX, ProConnect, RateMUX, ScriptShare, SlideCast, SMARTnet, StackWise, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0609R)