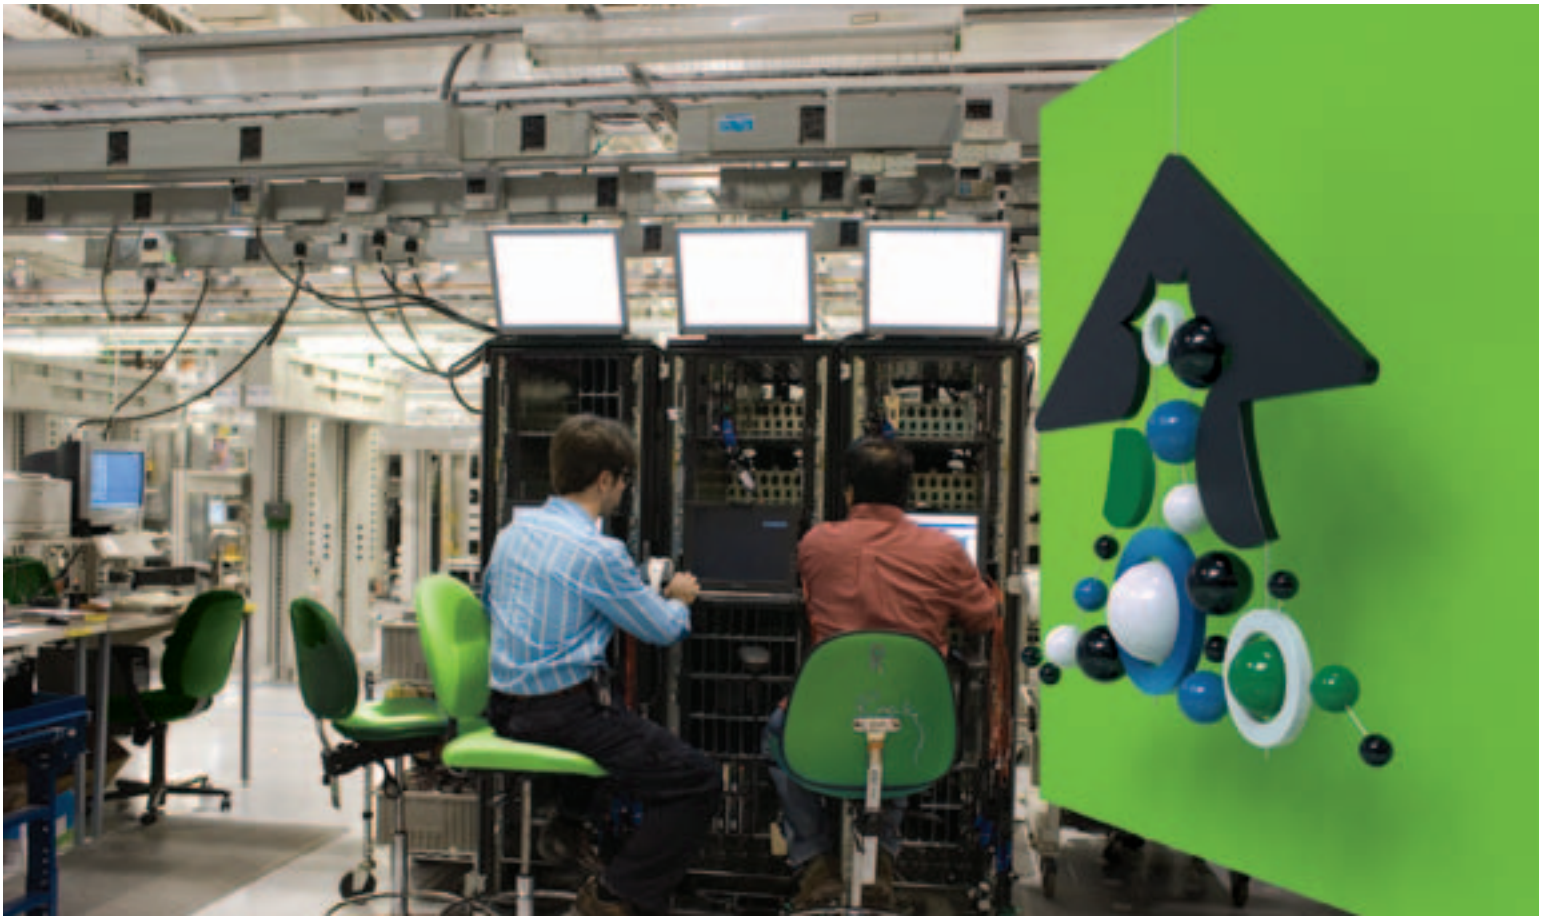

Rethinking server virtualization

Breaking performance and manageability barriers
White paper

Figure 1. Scale-out storage creates a pool of storage resources and delivers them to application servers as logical volumes.

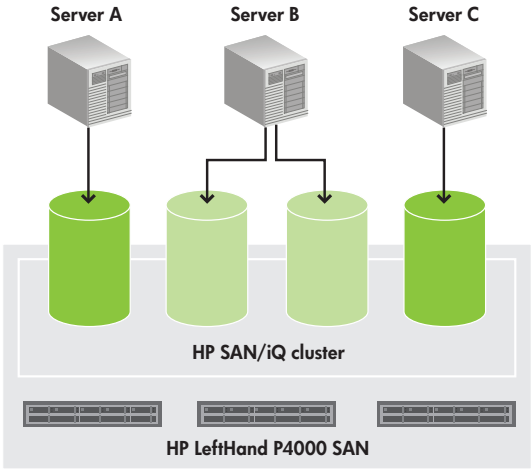

Introduction

Server virtualization is an ever more important tool for reducing cost, increasing availability, and enhancing business agility. But for many organizations, the savings that come from server consolidation are the primary reason for server virtualization.

Many of virtualization's additional benefits—including built-in, cost-effective high availability (HA) and disaster recovery (DR)—require external shared storage with a comprehensive feature set to support HA and DR. Unfortunately, it is all too easy to spend the savings from server virtualization on inefficient storage systems. Because they overcome many of the cost and management limitations of traditional SANs, iSCSI SANs are becoming the preferred choice to support virtualized environments.

Advantages of scale-out iSCSI SANs

HP LeftHand P4000 SAN Solutions take iSCSI SANs one step further because they are based on *storage clustering*—a form of scale-out storage that creates a scalable storage pool by aggregating the critical components of a number of storage systems into a single pool of resources, or storage cluster (Figure 1). The cluster accepts and responds to iSCSI requests as a single system. All physical capacity is aggregated and is available to the volumes created on the SAN. When more storage is needed, additional storage systems can be added to the cluster; the cluster seamlessly, non-disruptively reorganizes its storage to incorporate the new system into the cluster.

HP LeftHand SANs provide low-cost, built-in support for HA and DR implementations; superior, scalable performance; and straightforward management that any server administrator can understand and put into use. This white paper describes the advantages of scale-out iSCSI storage in four areas—cost, high availability and disaster recovery, performance, and management—and illustrate how iSCSI SANs, particularly the HP LeftHand SANs, provide better overall support for virtualized environments.

Reduce cost with efficient capacity utilization

Today's economic conditions present a challenging set of trade-offs for IT organizations. Reducing costs is paramount. But a focus solely on cost reduction can become a threat if the business loses the agility to respond to the next economic upturn. Server virtualization, however, makes it easy to reduce cost through server consolidation, which increases ROI by running multiple workloads on a single server. It also increases business agility by making it easy to deploy new applications or to scale them up or down.

The difficult realities of cost savings

It is all too easy to negate the cost savings that come with server virtualization by purchasing a SAN solution that is unable to properly support the advanced requirements of virtualized environments. The following issues related to storage systems can drive up the cost of virtualization projects:

- External shared storage can be very inefficient because it often requires more storage to be purchased up front. For example, a 100 GB volume that is only half full still requires 100 GB of storage—even though its full capacity may not be used for years. The problem of excess capacity becomes even more evident when temporary volumes to support test and development environments still consume 100 GB—although they will never be filled and will only be needed for days or, at most, weeks during testing.
- More virtual machines equates to more logical volumes—and, thus, more snapshots for backups. Traditional SANs require a 100 percent reserve, meaning that the same half-full 100 GB volume requires an additional 100 GB for each snapshot.
- Implementing high availability across sites requires an additional storage system and additional synchronous replication software to keep data continuously available in the event of a single SAN failure. This is not only expensive; it is also complex.
- Disaster recovery using remote replication can drive up storage costs. Many storage replication solutions require multiple copies of data and massive capacity reserves that can increase storage inefficiency by a factor of three or more (Figure 2):
  - A snapshot of the source volume must be created, requiring a 100 percent reserve.
  - The remote copy requires a 100 percent reserve on the remote system.
  - A third copy with 100 percent reserve space must be created before a server at the remote location can mount the copy for use.

  In other words, the remote copy requires not just the 50 GB that is actually needed, but 300 GB. In this example, it drives up the cost of storage by a factor of six.
- The reality of today's economy means that it is difficult to budget for future storage needs. Yet the deployment model for traditional SANs calls for purchasing a storage system with headroom for future expansion—difficult to do when today's budgets are being slashed.
- Budgets for small and medium-sized businesses often dictate the purchase of a low-cost SAN solution. But low-cost SANs often lack the features required to support synchronous replication for HA, asynchronous replication for DR, and enterprise-class storage management features. Even when these features are available, they are often expensive add-on software options that drive up the cost of storage.

Figure 2. On a traditional SAN, disaster recovery through remote replication requires three copies of each volume on a traditional SAN. This compounds storage inefficiency because each copy must be fully provisioned.

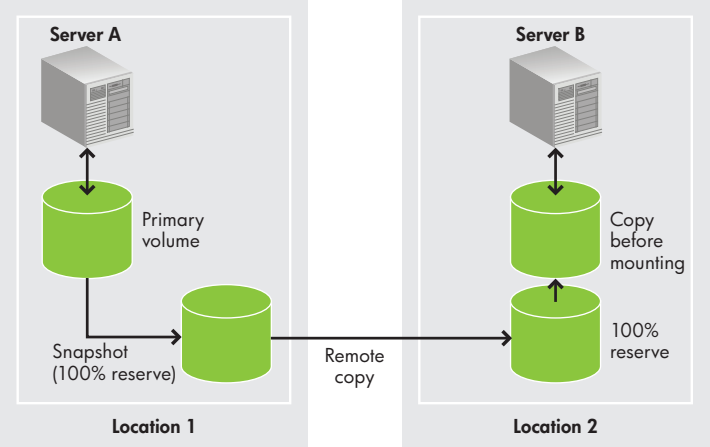

Built-in storage efficiency with HP LeftHand P4000 SAN Solutions

The combination of storage clustering and thin provisioning in HP LeftHand P4000 SANs dramatically improves overall storage efficiency, delivering a solution that increases ROI by increasing capacity utilization and by matching the needs of server virtualization.

Thin provisioning is a built-in feature that underlies every P4000 SAN. The storage cluster manages all of the storage allocation underneath a logical volume; thin provisioning allocates space only as data is actually written to the volume. Thus, thin provisioning allows you to purchase only the storage you need today, adding more as application data grows. This not only increases storage utilization and ROI; it also helps you spread out your capital expenditures, a significant business advantage in times of tight budgets.

Consider how storage clustering and thin provisioning can minimize the cost of storage as it increases the cost-effectiveness of storage virtualization:

- By allocating disk blocks only as they are needed, thin provisioning uses storage more efficiently. For example, a 100 GB storage volume that is half full uses only 50 GB of storage, allowing you to defer storage purchases until more disk space is actually needed.
- Snapshots are space-efficient because they are also thin provisioned. A snapshot of a 100 GB storage volume requires *no reserve*; it allocates space only as the contents of the primary volume and the snapshot diverge.
- Multi-site high availability can be implemented at no additional cost. Simply split a storage cluster between sites, and built-in synchronous replication handles block replication, failover, and failback automatically. Your volumes are continuously available and ready to support failover of virtual machines—even in the event of dual disk drive failures, storage system failures, rack and power failures, and site disasters.
- Disaster-recovery solutions now require only two copies of a volume's data. First, a space-efficient snapshot of the volume's data requires no space reserve. Second, the remote copy can be used at the remote site without creating an additional copy for mounting (Figure 3). Using the example of a half-full 100 GB volume, a traditional SAN requires 300 GB of storage while the P4000 SAN solution requires only 50 GB—a six-fold improvement.
- HP LeftHand P4000 SANs make copying existing volumes—or "golden master" images—as easy and cost-effective as creating the new virtual machines that use them. HP SAN/iQ® Software with SmartClone™ Technology makes space-efficient volume copies that allocate space only as the volume contents are changed by the virtual machine using them.

Figure 3. Disaster recovery solutions using HP LeftHand SANs require only two copies of volume data. Each copy uses only the storage actually allocated.

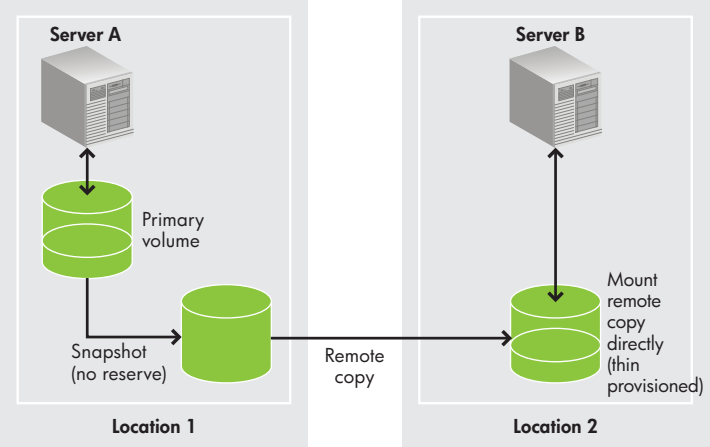

Simple, cost-effective HA and DR

Server virtualization makes high availability and disaster recovery more straightforward and cost-effective than ever. Without requiring any application-level programming, virtualization software can recover from server failures by automatically restarting virtual machines on alternate servers, making practically any application highly available. If disaster strikes and takes an entire geographic location out of service, virtualization allows virtual machines to be restarted at an alternate site from remote copies of their logical volumes. Unlike the painstaking setup required with physical servers, virtual machines can be started at the remote site easily. Even the assignment of virtual machines to servers can be changed in order to keep costs under control.

High availability and disaster recovery are among the key reasons that customers deploy virtualization solutions. But in order for those solutions to work effectively, SAN storage must have equivalent virtualization features:

- High availability—Data must stay online and available during site failures; virtual machines must be able to access their logical volumes after a failure.
- Disaster recovery—Remote copies must be kept up to date at alternate locations in order to enable data recovery after a disaster.

Storage pitfalls with most systems

In most storage systems, data is not distributed—so a failure that takes a data center offline can easily take the SAN offline as well. The high-availability solution for traditional SANs uses redundant components. While this approach protects against a hardware failure, it does not protect against data center failures ranging from power and cooling problems to simple human error. To protect against a comprehensive set of failures with most storage systems, you must purchase a second storage system—and optional synchronous replication software—and then maintain two copies of data in two SANs at two locations. In that case, if one site fails, a backup copy of the data is ready to support virtual machine failover.

But failover of a logical volume requires human intervention, which negates the benefits of automatic virtual machine failover. Because the failover volume's logical unit numbers are different, systems must be reconfigured to access an alternate storage system in order to accomplish a failover. Failback is a complex and time-consuming operation that requires manual synchronization of volume copies.

Implementing a disaster-recovery solution with traditional SANs is both expensive and complicated. Remote-copy software is usually a costly add-on option, and failover and failback of both virtual machines and logical volumes is a complex operation. One mistake can result in a catastrophic loss of business-critical data.

Integrated HA and DR with HP LeftHand P4000 SANs

Both synchronous and asynchronous replication features are integrated—at no additional cost—with every P4000 SAN. With synchronous replication to support high availability as well as asynchronous replication to support disaster recovery, the P4000 SAN enables you to implement cost-effective HA and DR solutions. Because management for all features is integrated into the centralized management console (CMC), administrators can use the same storage interface to provision and manage both storage and replication. No third-party or add-on software is required.

Best of all, the implementation of both synchronous and asynchronous replication provides a simpler, more cost-effective way to support server virtualization with HA and DR features in storage.

High availability through built-in synchronous replication

P4000 SANs incorporate high availability to provide data availability above and beyond what hardware RAID can provide. While hardware RAID protects against disk drive failures, P4000 SANs protect against multiple component failures—including complete storage system and even site failures. Their ability to provide continuous availability despite failures makes P4000 SANs an ideal match to the needs of the HA solutions provided by virtualization software.

High availability is implemented with network RAID, a synchronous replication technique that is completely transparent to application servers. Failover and failback are automatic and transparent, with the cluster maintaining a consistent logical volume state at all times. Gone are the worries concerning resynchronization of primary and secondary volumes. Also gone is the need to manually change LUNs in the event of a failure.

Figure 4. HP LeftHand P4000 SANs use high availability to maintain copies of each logical volume's blocks in two locations, providing continuous availability.

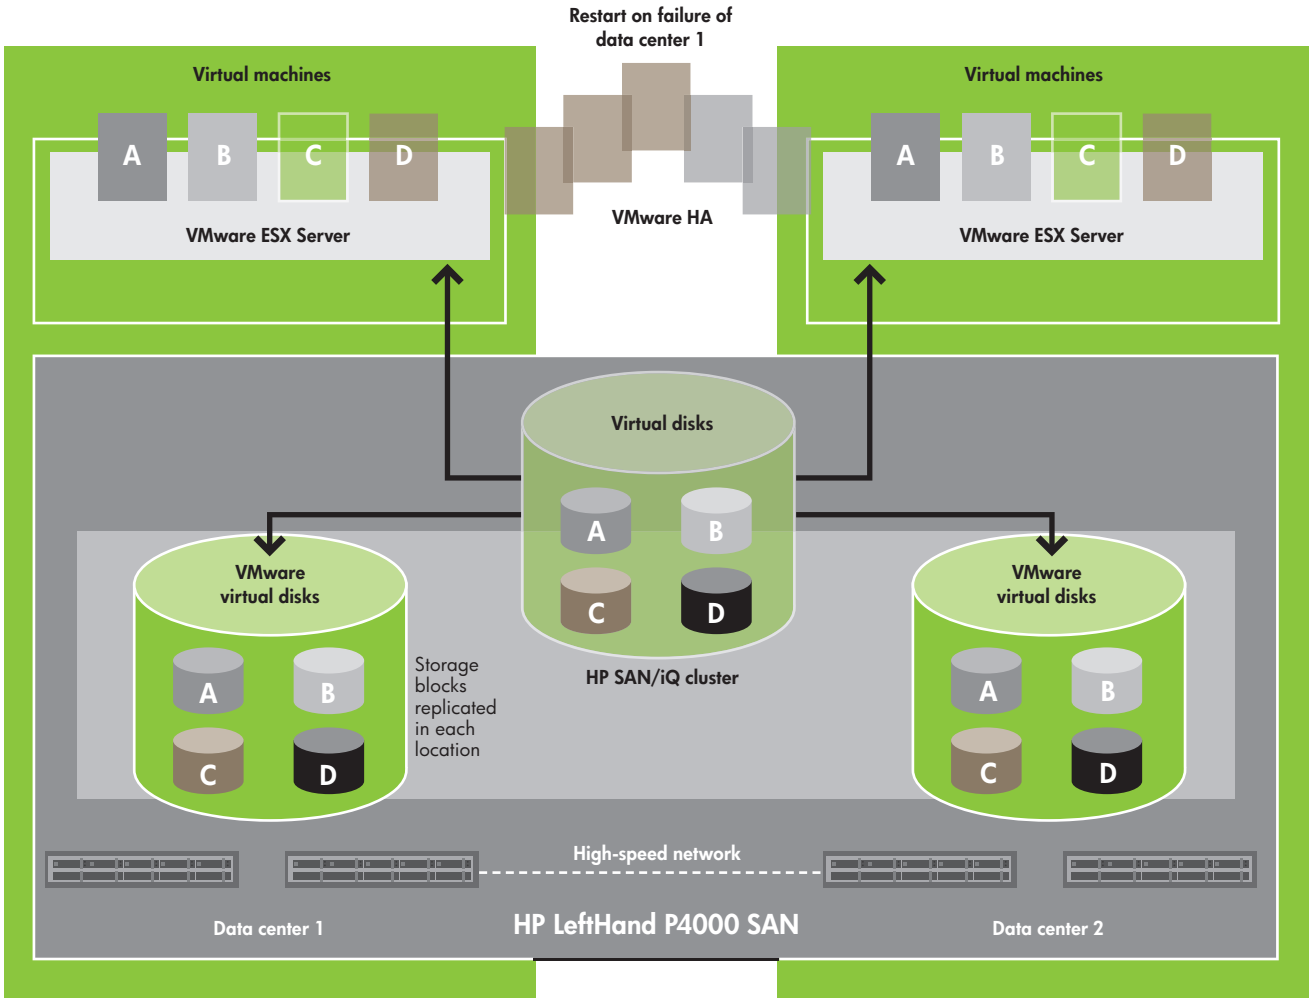

Figure 4 illustrates how this built-in HA works in conjunction with virtual machine failover. The figure illustrates a storage cluster that stretches across data center 1 and data center 2, and is interconnected through a high-speed network. The cluster supports a logical volume that contains four virtual disk files, one for each of virtual machines A–D. Network RAID maintains two copies of the volume's blocks, one in each of the two data centers. The cluster presents only one volume to application servers, maintaining the copies internally to the cluster. In the event of a failure in data center 1, virtual machines A and B can be restarted in data center 2 using the exact same volume that they had used in data center 1. The failover takes place automatically, without human intervention.

Disaster recovery through Remote Office Solution Pack

In the event of a failure at a primary site, P4000 SANs handle disaster recovery through space-efficient remote copies that can be used at an alternate site.

Because remote copy is integrated into the HP SAN/iQ® Software that runs on each storage system, managing remote copies is simple and straightforward through the CMC. In the event of a failure, the CMC's failover/failback wizard takes the guesswork out of geographic failover and failback, as well as reducing the chance of error.

For backup and recovery, a best practice for companies with multiple sites is to have one central site maintain copies of every remote site's data. HP SAN/iQ Remote Office Solution Pack is a replication client that can work in almost any remote office using existing server resources (Figure 5). This solution is cost-effective because no SAN hardware is required at the remote sites. It is also simple to manage because the replication clients are managed through the CMC. There is no additional management software to install and learn.

Figure 5. HP SAN/iQ Remote Office Solution Pack helps remote offices maintain copies of their data at a single, central site for backup and disaster-recovery purposes.

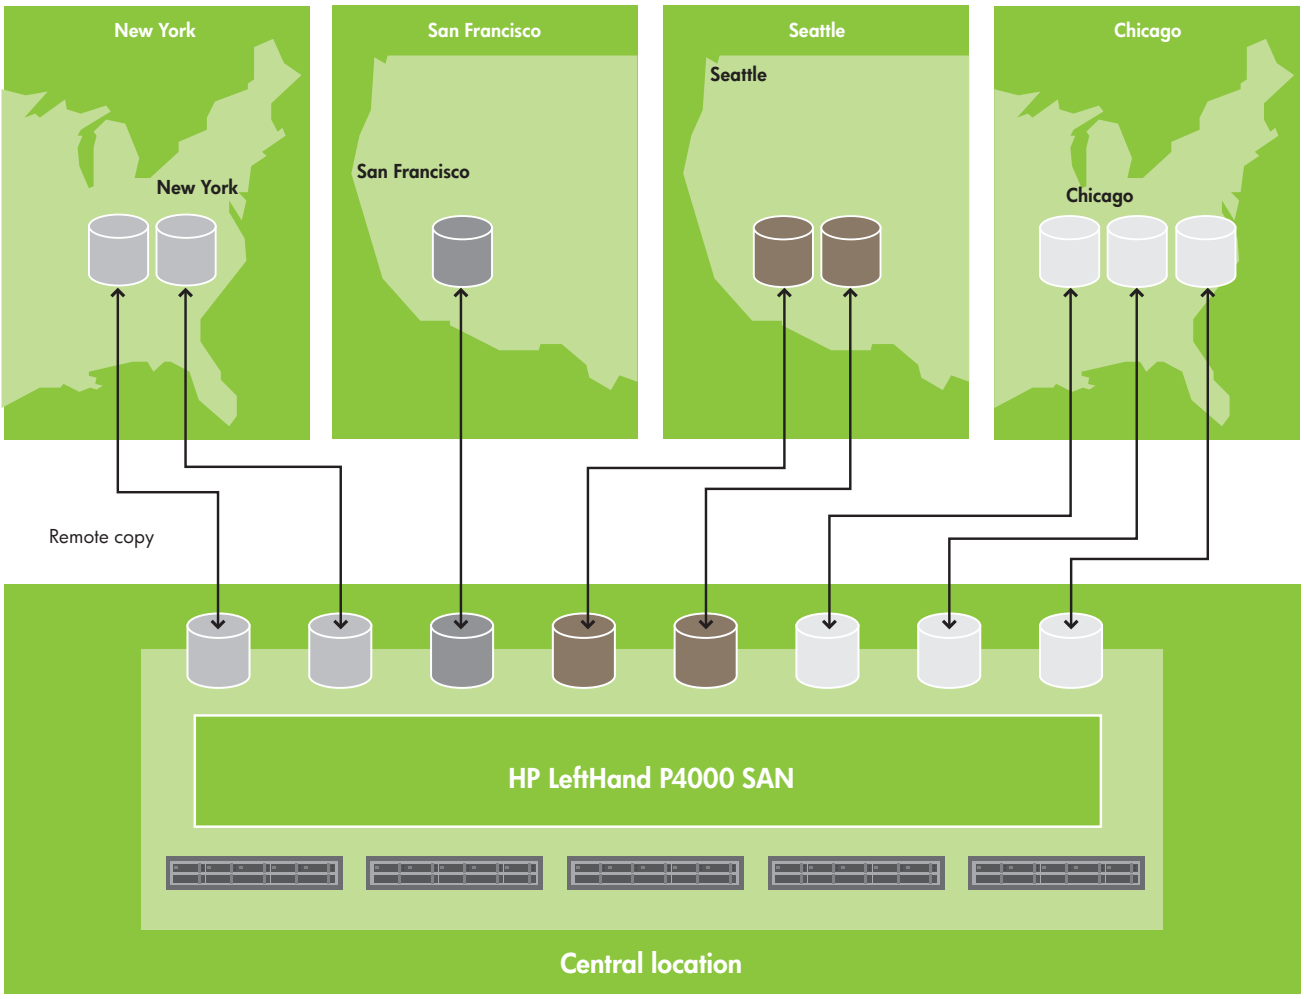

Avoid performance issues by using scale-out storage

Server virtualization helps to improve ROI by increasing server utilization. However, increased server utilization means a greater demand for I/O bandwidth. SAN storage must be up to the challenge by providing a higher level of scalable performance than is needed by traditional server environments.

Because of the layers of abstraction that can obscure a problem's root cause, virtualization can make it difficult to resolve performance issues. SAN storage must contribute to the ability to diagnose and resolve these problems.

How bottlenecks develop

Most storage systems come with built-in performance limitations. A pair of central controllers handles all incoming requests for storage and coordinates all access to disk storage, often acting as a bottleneck. Logical volumes are stored on RAID groups that have a fixed number of disks—with performance limited by the combination of disk drives and the central controller—and limited controller resources. Mapping logical volumes to resources is also fixed, resulting in performance that is limited by the resources in the RAID group.

Virtualization exposes these limitations by demanding more from SAN storage, including ports, zones, and RAID groups. For most storage offerings, the options to scale bandwidth are to reconfigure RAID groups and zones, to upgrade controllers, or to purchase

Figure 6. In contrast to traditional SANs whose performance is limited by fixed RAID groups and a centralized controller, HP LeftHand P4000 SANs use storage clustering where every storage node contributes to the performance of every logical volume.

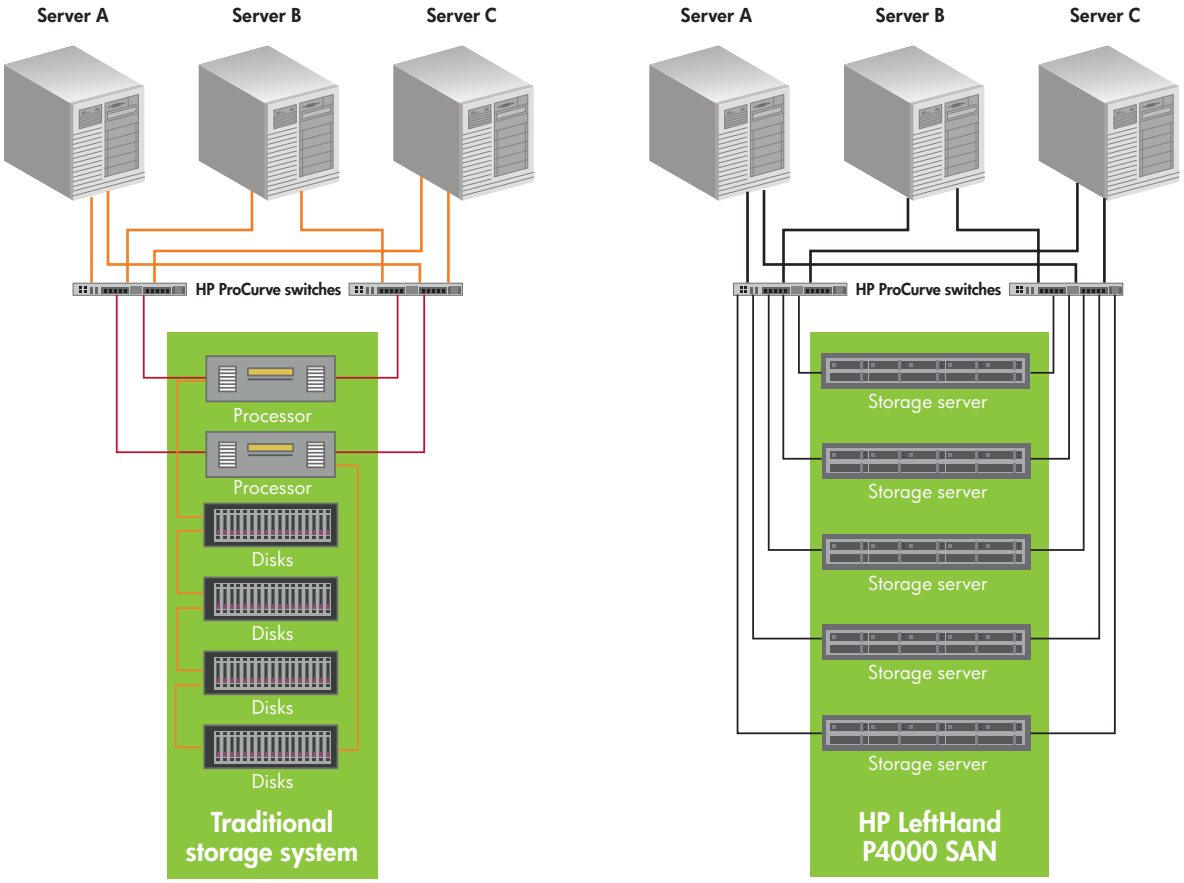

a new storage system—each of which results in application downtime. With traditional SAN storage, applications may not get the storage performance they require; virtualization makes it even more difficult and complicated to troubleshoot the problem.

Scalable performance with storage clustering

In contrast to traditional SANs, P4000 SANs use a scale-out architecture to maximize performance. With traditional SAN storage, performance is limited by centralized controllers (Figure 6). In a P4000 SAN solution, every storage node in a cluster contributes to performance because each node contributes resources—including network bandwidth, processing power, RAID controller, cache, and disk storage—to the cluster. Whereas adding disk trays to a traditional SAN improves performance only to the extent of a central controller's capabilities, adding storage nodes improves performance and capacity in a near-linear fashion because:

- Logical volumes are distributed across the entire cluster, with every storage node contributing to overall performance. This creates a flexible relationship between logical volumes and the resources that deliver performance.
- Storage clustering delivers maximum performance for every volume on the SAN because every storage node contributes to every logical volume's performance. The performance of a logical volume is the aggregate amount of bandwidth, disk drives, processors, RAID controllers, and cache in an entire cluster.
- Storage capacity and performance are scalable because there is no central bottleneck. When higher server utilization results in higher-than-expected storage demands, the storage cluster can meet those demands or be scaled to meet them. Simply add a storage node, and the cluster re-stripes itself to deliver higher performance with no disruption to application availability.

- HP LeftHand P4000 SANs distribute the logical volume blocks throughout the storage cluster, allowing hundreds of network ports and disks to support a single logical volume's performance. Compare this to replication in traditional SANs, which further reduces performance.

Better performance management for virtualized environments

The P4000 SAN's performance management system provides an easy, intuitive way to monitor SAN performance. Because it provides performance information for each application server/virtual machine, logical volume, snapshot, storage cluster, and storage node, it helps to manage both physical and virtual machine performance more effectively. It also provides an intuitive interface for quick and easy identification and diagnosis of any performance issues that come up.

The performance management system helps with trend analysis and decisions regarding the right time to add a new storage node in order to increase performance or capacity.

Simple management for storage flexibility

Server virtualization brings about unprecedented flexibility and agility in data centers. Virtual machines can be created or cloned with a few clicks of the mouse—helping to meet business objectives or create test and debug environments in far less time than physical servers require. To optimize performance and meet service-level objectives, the number of virtual machines can be increased or the relationship of virtual machines to physical servers changed.

The problems with static storage configurations

The dynamic nature of virtualized environments creates havoc with most storage systems. While it is easy to create, copy, move, or delete virtual machines, doing the same with logical volumes is antithetical to their "configure once, leave alone" design. In practice:

- Storage configurations are difficult to change because an administrator must first change storage system parameters, and then carefully modify intermediate switch and HBA settings to establish end-to-end connectivity.
- Provisioning is complex, requiring knowledge of storage systems and switches, including concepts such as volumes, RAID groups, zoning, LUN masking, and mapping LUNs to HBAs.
- Traditional SANs require specialized skills to administer. Administrators must be involved for every change, even though server administrators can quickly modify virtual server environments. Thus, changes to the SAN can lag far behind.
- Configuration changes tend to be disruptive, resulting in application downtime. Changing the RAID characteristics of a volume, for example, means bringing an application down while its data is moved to a new RAID group.

Figure 7. Provisioning a volume or changing its attributes using HP SAN/iQ Software requires only a few clicks of a mouse. Illustrated are the Basic and Advanced tabs of the Create Volume dialog box.

Simple, non-disruptive management enables dynamic environments

HP LeftHand P4000 SANs are designed to support dynamic environments. They use standard iSCSI protocols and standard LAN technology—i.e., a standard storage node and a simple, intuitive GUI that every server administrator can understand. Because server virtualization makes data centers more dynamic than ever before, P4000 SANs are built to support these fast-changing environments:

- Just like the virtual machines they support, logical volumes in P4000 SANs are virtual, meaning that their data is striped across every storage node. The cluster, not a storage administrator, manages all of the details.
- P4000 SANs make it easy to create, clone, and manage volumes. Through the CMC, creating a new logical volume, cloning an existing one, or making a snapshot takes only a few clicks of the mouse (Figure 7).
- Making changes after the fact is easy. Most storage attributes—including network RAID level, thin or full provisioning, and even the cluster on which the volume is stored—are maintained on a per-volume basis. Changing a volume's attributes at any time is transparent and non-disruptive to applications.
- Because P4000 SANs are based on standard servers and IP networks, virtually everyone has the skills to install and support a P4000 SAN. And virtually every data center has the equipment to connect SANs with servers.
- Adding a new storage node to scale capacity and performance through the centralized management console takes only two steps: discovering the storage node; and adding it to a cluster. The cluster then incorporates the new storage node by restriping data with no impact on storage availability. Adding new clusters—and even moving volumes between clusters—is all initiated through a single management interface that accomplishes these tasks without taking data offline.

Conclusion

Virtualization opens up a new range of possibilities for data centers. These include higher server utilization with lower costs, the economies of server consolidation, and high availability and disaster recovery supported with built-in features. Additional benefits include performance management through dynamic workload balancing, as well as ease of management through the harnessing of all servers as a single, uniform pool of resources.

Virtualization is changing the IT landscape by enabling more efficient, adaptive IT infrastructures. HP LeftHand SANs promote this transformation, making them an ideal fit for the next-generation virtualized data center. HP LeftHand SANs help you drive down costs by increasing storage utilization. They help you work within budget constraints by allowing you to purchase storage only when you need it, not months or years in advance. With HP LeftHand SANs, every storage purchase adds capacity and performance, making the most of your storage investments and avoiding the bottlenecks of monolithic storage systems. Built-in HA and DR features enhance server virtualization capabilities to deliver better business continuity—without disruption. The ease of management that HP LeftHand SANs provide means that both server and storage administrators can manage the SAN, helping to make the most of limited administrator resources.

HP Education Services for HP LeftHand P4000 SANs

A well-trained IT staff helps make your HP LeftHand P4000 SAN Solutions even simpler to use and brings still more agility—and greater value—to your business. Educated end users experience improved solution reliability, fewer support requests, speedier support issue resolution, and faster project implementation.

To help you get the most from your HP LeftHand P4000 SAN Solutions, HP offers two levels of training:

- Basic training imparts the knowledge needed to understand, manage, and configure your P4000 SANs.
- Advanced training outlines best practices for P4000 SANs, based on field experience and applied industry knowledge.

Visit **www.hp.com/learn/storage** for more information.

HP Services

Put the strategic and technical know-how of HP Services experts to work for you:
When you buy HP LeftHand SAN Solutions, it's a good time to think about other levels of service and support you may need. You can trust the service professionals at HP to collaborate with you to make technology the difference in your business.

Recommended services

- **HP Support Plus 24 Service**—for around-the-clock, reactive onsite hardware support and over-the-phone software support
- **HP Installation and Startup for HP LeftHand SAN Solutions**—fast, reliable startup for enhanced server virtualization and business continuance with SAN solutions

Related services

- **HP Proactive 24 Service**—integrated proactive and reactive services for businesses looking to achieve better performance, higher availability, and greater stability
- **HP Proactive Select Service**—to improve IT performance and manageability for businesses looking for services flexible enough to cover the IT product lifecycle and adapt to changing needs

When technology works, business works. For more information, contact your HP sales representative or HP-authorized Channel Partner, or visit **www.hp.com/hps/storage**.

HP Financial Services

HP Financial Services provides innovative financing and financial asset management programs to help you cost-effectively acquire, manage, and ultimately retire your HP solutions. For more information on these services, please contact your HP representative or visit **www.hp.com/go/hpfinancialservices**.

For more information

For more information about HP LeftHand P4000 SAN Solutions, contact your HP representative or visit **www.hp.com/go/p4000**.

Technology for better business outcomes

To learn more, visit www.hp.com

© Copyright 2006–2009 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

4AA2-5423ENW, April 2009