

New D-Link[®] SAN Makes Low Cost Storage Virtualization a Reality

White Paper
October, 2008

Abstract

Storage Area Networks (SANs) and virtualization software have traditionally belonged to the realm of large enterprises. Up until now, the technologies and professional services required to support them have been prohibitively expensive for small- to medium-sized companies. However, as data requirements, application demands, and storage and backup requirements escalate for all types of companies – especially those in information-intensive fields – more efficient ways of managing data and recovery are needed.

This white paper explores the cost and management challenges inherent to enterprise storage virtualization systems and describes the new solutions available to businesses. We conclude with a brief look at D-Link's low-cost, high-value storage virtualization solutions. D-Link's line of iSCSI SAN arrays, in particular, are affordable, flexible, easy to use and manage, and ideal for small-to-medium businesses.

Why Centralized Storage is Important

As more information is archived over networks and accessed from every conceivable IP-connected computer and device, centralized storage and remote backup becomes essential. Information-intensive industries like the legal, health care, insurance, banking, and financial services professions have known this for a long time. Not only do regulatory and compliance initiatives (both internal and external) require redundant archives of sensitive information, but day-to-day operations demand 24X7 access to the information that makes business run. Natural disasters as well as accidental and intentional data losses make the situation even more critical. Traditionally, large companies utilize data centers to manage complex information architectures and protect data. These data centers feature centrally managed disk farms that provide centralized storage and applications to widely flung client computers and dumb terminals. Duplicates of files are backed up at remote offices. In order to work, the data center model requires efficiently managed network storage. Files need to be efficiently 'served' to end users. Data needs to be backed up.

As hardware and software technologies continue to become more affordable, small businesses are now headed toward this scenario. The model provides optimal risk reduction and disaster recovery benefits. Eventually every size company will utilize the data center model – whether via a remote service or on the premises.

Direct Attached Storage – Risky Option

Traditionally, smaller companies have placed all their data on local storage. That works until disk space runs out or additional servers are required to accommodate the overflow of data. These days, even 50 to 100-person businesses are deploying multiple servers to handle email, applications, files, printing, media archiving, transactional databases and more. Recovery becomes a serious problem, even when a single server goes down. If a machine goes down, it becomes necessary to grab backup files from a remote location, which can take up to hours or days. Gigabit Ethernet connections are fine for day-to-day activity, but they are excruciatingly slow when it is necessary to repopulate a database after a significant server failure. For example, with up to 16 of 1TB drives in a 3U shelf there will be up to 16TB of storage that may contain data, but with only two gigabit uplink ports on the enclosure, getting data on and off a direct attached storage system can't be done fast enough. Most businesses simply cannot afford that downtime.

Storage Virtualization

Storage virtualization – the common methodology for managing SANs – consolidates multiple storage devices into a logical view so they can be managed from a single console. This allows multiple

storage devices to be accessed the same way no matter what their type or location. Applications can then read from and write to a single pool of storage rather than individual disks, tapes and optical devices. Resources can be shared in a fluid, efficient manner, and migrating data from one location to another (for backup or business continuity reasons) is transparent to applications. Storage virtualization helps administrators efficiently perform backup, archiving and recovery by disguising the complexity of the SAN.

This is different from traditional backup schemes, where administrators deploy systems completely identical to their production network at another remote location and then rehearse backup and restore processes.

Popular server virtualization software, like VMware, Microsoft Hyper-V and Citrix Xen, has ignited interest in storage virtualization. General virtualization uses Virtual Machines (or VMs) to simplify the management of multiple servers. The advantages of virtualization are well documented (server consolidation, stability and security, development flexibility, migration and replication, desktop virtualization) and are beyond the scope of this paper. Many companies – large and small – are using virtualization to reduce IT costs and simplify management.

Storage management in a VM setting does have its challenges, though. For example, in the event of hardware failure, it could take more than an hour to migrate a mere 30GB VM to another platform – if you can still reach the server that crashed. With a TB of storage it could take days to move the data from one server to another. Again, gigabit networks can only go so fast.

Cost Issues

Up until now, storage virtualization has been a costly proposition. Organizations with storage needs beyond the 3TB range were often forced to explore \$30,000 to \$70,000 SAN solutions from vendors like NetApp and EMC. Those solutions are designed for large enterprises with sophisticated backup and recovery schemes that require data de-duplication, live migration of LUNs (Logical Unit Numbers), snapshots and other high availability features. They are more suited to companies that need to have their systems online 24X7. These companies usually have large transactional data volumes or real-time data demands. For them, downtime is measured in tens of thousands of dollars per hour – even hundreds of thousands of dollars per hour.

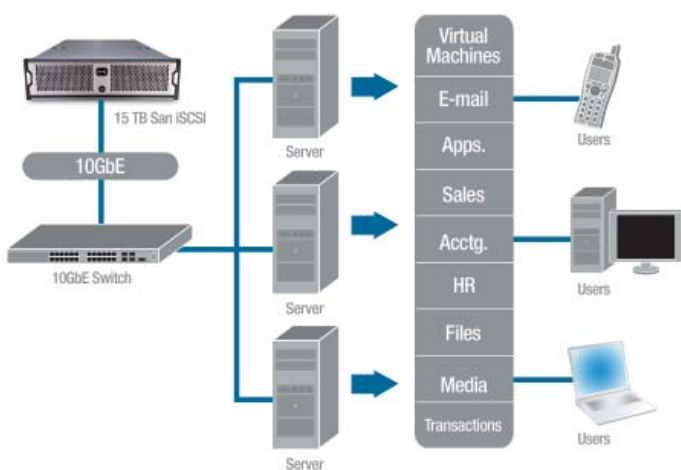
Most mid sized companies don't need that kind of robust feature set. And, they are certainly not prepared to spend \$150,000 on hardware, implementation and professional services to add several terabytes of storage.

So, what's available for the company that is interested in virtualization (or already has VM solutions deployed) and recognizes the risks involved with direct attached storage and traditional backup scenarios?

Low Cost VM-SAN Solution

Organizations that cannot afford a high-end SAN storage solution and do not have the expertise to maintain and support one now have an alternative. By connecting an iSCSI SAN to hypervisor-managed servers, an IT infrastructure can benefit from most of the benefits of storage virtualization without sacrificing big chunks of the IT budget.

If, for example, there are VMs running on Hyper-V, Xen or VMware, it



D-Link iSCSI SANs and D-Link switches work seamlessly with standard server hardware to enable storage virtualization.

is possible to quickly zip up, duplicate and transfer exact machine configurations to SAN storage devices at any location. Since the VM file is a flat file, it can reside anywhere, be backed up on any host, and it works just like the original machine. The capability is unique to virtualization, and the file transfer times required are nowhere near those of a traditional replication scenario. This is especially true when implementing 10GB Ethernet hardware at the SAN and switch locations. If something happens to one of the virtual servers, then simply restart a new virtual machine on another host in just a couple of minutes.

With D-Link's iSCSI SAN solutions, per-terabyte costs approach \$1,000 as opposed to comparably equipped NetApp or EMC equipment that can easily run \$10,000 per terabyte. D-Link DSN-3400, DSN-3200 and DSN-2100 SAN arrays all allow IT to quickly duplicate the virtualization environment at other locations throughout the enterprise. The systems balance storage capacity, enterprise features (like redundant power supplies) and connectivity

options (8X1GB Ethernet and 1X10GB Ethernet) so that smaller companies can have all the advantages of big corporate enterprises.

Easy Manageability

Typical IT staffs at small- to medium-sized companies can easily manage a D-Link storage virtualization solution. Advanced engineering capabilities are not required, and implementation requires very little support. This is not the case with the more sophisticated SAN/VM solutions out there. Unlike complex Fibre Channel SANs that require specialized skills, the iSCSI technology that D-Link SANs employ ties directly into standard Ethernet frameworks that most IT staff are comfortable with. The D-Link solution also allows you to procure your own drives. Most of the other SAN solutions on the market require the customer to buy their recommended drives at a ridiculous mark up.

Heterogeneous Virtualization on One SAN

With D-Link's storage virtualization solutions it is possible to run VMware alongside Citrix XenServer and Microsoft Hyper V – all on a single SAN. The hybrid environment allows combining heterogeneous VM environments with heterogeneous OS environments and back it all up quickly and easily. This provides the flexibility and software that users demand while sleeping easier because of enhanced business continuity and backup capabilities.

Advanced Virtualization Features

With a D-Link SAN, D-Link switch and virtualized servers, it is possible to also unlock advanced virtualization features, like VMware's VMotion or XenServer's XenMotion. This is not possible without centralized storage. VMware VMotion and Citrix XenServer XenMotion allow movement of a running virtual machine from one server to another. This offers companies zero-downtime for server maintenance since VMs can be moved from server to server without service interruption. It is also possible to move running application workloads to take advantage of available compute power. That's unprecedented flexibility.

What Size Company Needs Storage Virtualization?

Companies that run more than one serious business application on separate servers are likely candidates for storage virtualization solutions. It is not the right fit for a single-server IT shop.

Typical multi-server application scenarios include the following applications:

- Accounting (e.g. QuickBooks)
- Sales (e.g. Salesforce, CRM apps)
- Email/Messaging/Microsoft Exchange

- Web and e-commerce hosting
- Relational database
- Transaction processing
- Other Line of Business (LOB) applications

The scenario usually maps to companies with 50 to 100 employees or more. A 100 person company with \$10 million a year in revenues would have the types of computing and backup requirements that storage virtualization and virtualization in general solve. Information driven businesses – like legal, health care, insurance, banking and financial services firms – are particularly well suited to the solution. Storage virtualization works well for distributed companies, as well. Satellite offices across the U.S. often manage their own messaging and file/print servers, for example. While the headquarters office might employ a solution from NetApp or EMC, the remote offices could certainly gain advantages from a low-cost D-Link solution.

D-Link Solutions

D-Link storage virtualization solutions deliver big business storage functionality at a small business price. Our wide variety of SAN and switch solutions help IT regain control of data and reduce exposure to data loss disaster. They protect corporate intellectual property, enable instantaneous restores and ensure business continuity.

DSN-3400-10 1x10GbE iSCSI SAN Array, 15 Bays, 3U, w/o Drives, with Trays

Features:

- Integrated iSCSI System-on-a-Chip (SoC) Solution that can Handle over 80,000 I/Os per Second
- Single 10GbE Port
- 15TB Capacity with 1TB Hard Drives (Supports Higher Capacity Drives as They Are Introduced)
- Embedded Disk Controller Supporting RAID Levels 0, 1, 1+0, and 5

DSN-3200-10 8x1GbE iSCSI SAN Array, 15 Bays, 3U, w/o Drives, with Trays

Features:

- Integrated iSCSI System-on-a-Chip (SoC) Solution that can Handle over 80,000 I/Os per Second
- Eight 1GbE Ports
- 15TB Capacity with 1TB Hard Drives (Supports Higher Capacity Drives as They Are Introduced)
- Embedded Disk Controller Supporting RAID Levels 0, 1, 1+0, and 5

DSN-2100-10 4x1GbE iSCSI SAN Array, 8 Bays, 2U, w/o Drives, with Trays

Features:

- Integrated iSCSI System-on-a-Chip (SoC) Solution that can Handle over 80,000 I/Os per Second
- Eight 1GbE Ports

- 15TB Capacity with 1TB Hard Drives (Supports Higher Capacity Drives as They Are Introduced)
- Embedded Disk Controller Supporting RAID Levels 0, 1, 1+0, and 5

D-Link xStack® Switches - High switching capacity, high-speed ports, feature rich, high security, and versatile manageability

DXS-3250 - 48-Port Gigabit Wireless-Ready Switch + (4) Combo SFP Ports + (2) Optional 10-Gig Copper/Fiber Uplinks

Features:

- ‘Wireless-ready’ Layer 2+ Switching for Seamless Roaming and Centralized AP and Wired Network Management
- 48 10/100/1000 Ports with 4 Combo SFP in 1RU Chassis
- Optional Dual 10-Gigabit Stacking/Uplinks through Fiber XFP or Copper CX-4
- Preconfigured ACLs Protect Against Known Attacks

DXS-3350SR - 48-Port 10/100/1000 Switch+ 4 combo SFP + 2 10GbE ports

Features:

- 176Gbps Switching Capacity
- Recoverable Ring or Star Stacking Architecture
- Dual 10-Gig Stacking Ports and Optional 10-Gig Uplinks
- Stacks up to 8 Units per Stack*

* Stack up to 8 DXS-3350SR. Up to 12 units in a mixed stack with other xStack 3300 series switches.

For more information about D-Link storage virtualization solutions, please visit www.dlink.com and www.support.dlink.com or call 714-885-6000.

Definitions

Virtual Machine/Virtualization: (VM) One instance of an operating system along with one or more applications running in an isolated partition within a server. The strategy enables different operating systems to run in the same computer at the same time. VMs also prevent applications from interfering with each other. This differs from a dual-boot or multi-boot environment, in which the user has to choose only one OS at startup. All virtual machines run simultaneously and function as if each owned the entire server.

Hypervisor: A program that provides a virtual machine environment. The term came from the IBM mainframe world, which first introduced the virtual machine (virtualization) concept in the 1960's.

Storage Virtualization: Also called “file virtualization,” storage virtualization consolidates multiple storage devices into a logical view so they can be managed from a single console. This allows multiple storage devices to be accessed the same way no matter what their type or location.

iSCSI: A protocol that serializes SCSI commands and converts them to TCP/IP. This allows block-level storage data to be transported over widely used IP networks. Businesses can then manage storage networks from a central location by using the ubiquitous IP infrastructure.

SAN: (Storage Area Network) A network of storage disks that connect multiple servers to a centralized storage pool via iSCSI or Fibre Channel cabling (typically). As opposed to managing hundreds of servers, each with their own disks, SANs simplify resource administration. As a result, disk maintenance and routine backups are easier to schedule and control.

Source: TechWeb TechEncyclopedia