



By Kevin Wittmer

BUSINESS CONTINUITY FOR SMBs WITH EQUALLOGIC PS SERIES iSCSI STORAGE ARRAYS

Robust business continuity has historically been out of reach for small and medium businesses (SMBs) because it was often too costly and overly complicated. EqualLogic® PS Series Internet SCSI (iSCSI) storage arrays enable organizations of all sizes to deploy an enterprise-level storage infrastructure offering comprehensive data and application protection.

Costs and complexity have typically placed robust business continuity (BC) out of reach for small and medium businesses (SMBs). While large companies could afford expensive Fibre Channel storage area networks (SANs), the highly trained staff to manage them, duplicate data centers, channel extenders, and expensive replication software, SMBs were often restricted to making backup tapes and storing them off-site.

Today that has changed, and organizations of all sizes can cost-effectively build an infrastructure that offers robust protection without exorbitant costs or specially trained staff. Advanced functionality that was first developed for mainframe environments has been enhanced over time, making it available to Microsoft® Windows®, Linux®, and UNIX® environments. As a result, company size and platform selection no longer determine functionality level. This is due to multiple technological advances—in particular intelligent storage, Internet SCSI (iSCSI) connectivity, server virtualization, and wide area network (WAN) acceleration.

These advances carry significant advantages for SMBs, because data is just as critical for these organizations as it is for large enterprises—and downtime for SMBs can be much more damaging than it is for large enterprises. This article examines

how organizations of all sizes can create a cost-effective infrastructure that delivers comprehensive protection—from simple data protection to disaster recovery (DR) to BC.

THE STAGGERING COST OF DOWNTIME

In today's highly competitive business environment, data center outages can be devastating. Regardless of the cause—hurricane, fire, accident, hacker attack, or some other type of disaster—production downtime is not only costly, but in some cases can be ruinous. With a mobile workforce, global customers wanting to do business around the clock, and continually greater dependence on technology, companies need to not only protect data, but continue business operations virtually uninterrupted. The cost of downtime, depending on the industry, can range from thousands to millions of dollars per hour—costs due not only to DR expenses, but also to lost sales, customer defection, and lack of productivity. Add to that a damaged reputation in the marketplace and diminished shareholder confidence, and the cost of downtime can be staggering.

However, organizations that can continue business operations through outages, large or small, can gain competitive advantage—and sometimes even take market share from competitors.

Related Categories:

EqualLogic

Internet SCSI (iSCSI)

Storage

Visit DELL.COM/PowerSolutions for the complete category index.

TAPE: GOOD FOR ARCHIVING, BAD FOR AVAILABILITY

Tape backup has typically been the most common protection solution for SMBs, and remains a good long-term archiving method for most large and small organizations. However, there are several reasons that necessitate that organizations complement their tape solutions with other BC solutions. First, backup windows are shrinking because of the huge growth in data volumes, requirements for long retention and fast access, and generally high reliance on data and technology. Second, because tape backup can be time-consuming, many organizations cannot back up often enough to adequately protect themselves.

Instead, organizations need a continuum of protection schemes that include tape, storage array-based data protection, remote replication for recovery after a failure or disaster, and BC during outages and common IT maintenance procedures.

DATA PROTECTION

The first line of defense is to protect data where it is stored, and that means keeping storage units running as well as providing safety features. EqualLogic PS Series iSCSI arrays are fundamentally built to protect data with enterprise-level availability and reliability. Redundant, hot-swappable components—including disk drives, controllers, network interfaces, power supplies, and cooling fans—mean that component failure need not result in downtime or data loss. The system monitors disk drives in the background to help detect and correct problems before they occur; in addition, should a drive fail, a spare is automatically placed into service and configured. Disks and arrays “learn” the configuration from those already installed, so no administrator intervention is necessary. Controller caches are mirrored and battery backed, and the arrays support multiple RAID levels.

Of course, backup is a basic part of data protection—simply backing up to tape can help protect against complete data

loss, although restoring from tape can be time-consuming. Disk-based snapshot functionality is a big improvement—snapshots are created quickly and easily without disrupting operations, and can provide multiple restore points to minimize lost work. Many organizations take snapshots and then run tape backups from them, helping keep data available by virtually eliminating backup windows, speeding restore operations, and minimizing errors.

DISASTER RECOVERY

The focus on DR has been greatly magnified in recent years. Increasing virus and hacker attacks, hurricanes, and electrical brownouts serve as reminders of how vulnerable data is.

The keywords for DR are *copies* and *distance*. For DR, organizations must have more than one data copy, and they must store those copies some distance away from the primary data center on different physical systems—remote replication is the standard for DR. Today, remote replication is available for iSCSI-based SANs, and does not use the specialized equipment that Fibre Channel networks require. iSCSI enables organizations to replicate data across standard IP networks and provide multiple replicas from various points in time. Should a virus corrupt data, administrators can almost immediately roll back to a previous replica to help prevent data loss. Deciding how and where to replicate depends on an organization’s needs and available locations. Some organizations may replicate from the primary data center to one remote location; others may replicate the same data to multiple locations. Organizations with branch offices often replicate from each branch to a central DR site, and then back up data from there.

Configuring a DR implementation depends on two important factors that each organization must identify: recovery time objective (RTO) and recovery point objective (RPO). RTO indicates how quickly the organization can restore data—typically minutes or hours, and in some cases days. Some operations and

data types may only tolerate a very short RTO, while others can survive longer delays. RPO indicates how much data loss the organization can tolerate, and that determines how often it replicates data—every hour, three times per day, and so on. Many organizations define different RTOs and RPOs across the enterprise—uniformity is not important as long as they can easily and cost-effectively match data types to protection levels.

BUSINESS CONTINUITY

BC differs from data protection and DR because it describes not only a level of protection that helps speed recovery, but also a strategy that helps reduce the importance of recovery speed. For BC, the focus is not on how long it takes to get back in operation—the focus is on *staying in operation* regardless of failure, outage, attack, or corruption. For example, if an organization has a duplicate data center at another location with data copies, it can quickly bring operations online at that location.

Another way to continue business is with built-in storage features that help minimize disruption during standard IT maintenance tasks. Because of the patent-pending EqualLogic page-based volume management, data on PS Series arrays can be automatically moved *while it is in use*. This is an advancement of revolutionary proportions because it enables data movement without downtime. This capability means that organizations can add or move capacity among storage tiers, and automatically balance loads across disks and arrays, without interrupting users—to help IT staff get their job done and enhance performance while business continues.

Common performance improvement tasks such as load balancing across disks and arrays are done automatically by PS Series arrays—helping both avoid interruptions to business and optimize performance. Automatic disk sparing and multipath I/O enable the array to automatically replace a failed disk or manage network throughput without intervention or downtime. These features can add tremendous value during

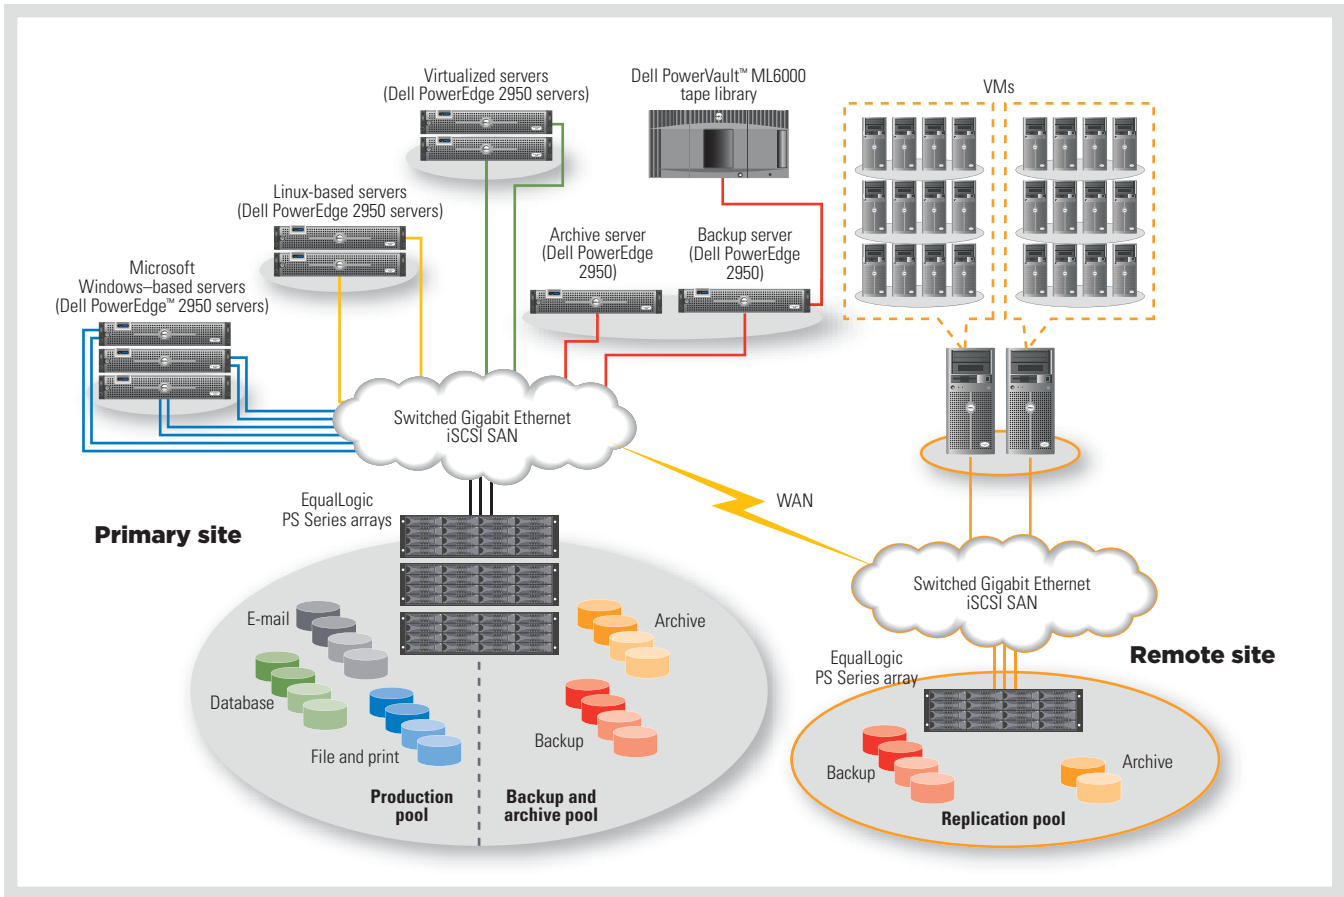


Figure 1. EqualLogic PS Series iSCSI arrays enable business continuity for organizations of all sizes

regular business operations, and can actually help keep problems and accidents from becoming disasters.

BUSINESS CONTINUITY STRATEGIES AND TECHNOLOGIES

SMBs need the same comprehensive protection as large organizations—but, in addition, they need it to be easy to use, nondisruptive, and cost-effective.

Data center equipment costs have dropped sufficiently to make *duplicate data centers* possible for many organizations—for example, EqualLogic PS Series iSCSI arrays make remote replication much more cost-effective than it has been in the past because the arrays are not cost-prohibitive and they include base and advanced software functionality (see Figure 1). In addition, while physical server costs have dropped, the development of *server virtualization* has dramatically

altered the landscape. Products from companies such as VMware enable multiple virtual machines (VMs) to reside on a single physical host—as a result, organizations can build a second data center without duplicating hardware. For example, an organization’s production data center may have 100 physical servers running business applications, with data shared on a SAN. At its secondary data center, it can run those 100 servers as VMs using only 10 physical servers, with each physical server running 10 VMs. This approach enables the business to maintain a duplicate data center cost-effectively to help remain operational following a failure or other outage. Some organizations simply could not have BC without using virtualized servers.

Server virtualization also lets organizations replicate multiple applications to the same target that normally would not work

well on the same system. For example, administrators might replicate servers running Microsoft SQL Server™, Microsoft Exchange, and Oracle® application software to different VMs on the same replication target, which typically costs much less than buying individual physical servers for each application. Different versions of the same application can be treated the same way—Microsoft Exchange Server 2003 and Exchange Server 2007 can both be replicated to the same target server, but to different VMs.

Virtualization also enables organizations to share data from a SAN and move VMs between physical servers to enhance application management—and administrators can make these modifications without users even knowing it. For primary or remote data centers, this capability helps provide BC during maintenance and performance optimization.

STORAGE TIERING

Another key to keeping BC cost-effective is to not treat all data the same way. Instead of continuously replicating all data to a remote location on high-speed disks so that any piece of data is almost instantly recoverable, most organizations set up storage tiers in which some data is replicated less often, to higher-capacity (less expensive) disks with slower performance—such as Serial ATA (SATA)—than the organization's most critical data. Some data should simply be archived to tape—available if needed, but at a lower cost than disk storage.

Storage tiers should be linked to the business value of the data involved, the ability to re-create it, and the requirement for speedy access. Organizations can first identify the required RTOs and RPOs for all data and set up tiers of storage and services that help maximize utilization while lowering total cost of ownership. Also, platforms like Microsoft Windows Storage Server 2003 and software like Microsoft Exchange Server 2007 include single-instance storage, enabling organizations to store only a single copy of identical data, such as e-mail attachments that go to multiple mailboxes. This approach can help streamline data movement and minimize bandwidth requirements for DR and BC.

WAN OPTIMIZATION

Bandwidth is expensive, and accounts for a significant portion of the cost of both DR and BC; in addition, bandwidth latency creates challenges for recovery and restore processes. These issues can affect the selection of remote sites, the amount of data that can be replicated, RTOs, and RPOs. However, today organizations like F5, Riverbed, and Citrix offer WAN acceleration and optimization technologies that enable organizations to replicate a large amount of data while increasing efficiency in a cost-effective way. Using various compression, de-duplication, and optimization techniques, these solutions accelerate WAN traffic in ways that can be significantly less costly than purchasing more bandwidth.

CLUSTERING, ISCSI SAN BOOT, AND THIN PROVISIONING

Clustering servers is a way to create BC, because applications can be distributed across multiple nodes for performance optimization as well as failover. This approach is aided by SAN implementations, because data can be shared among servers. Booting servers from the array itself, another standard feature of EqualLogic PS Series arrays, also enables BC; if a server fails, another can be deployed almost immediately using the boot volume on the SAN. This feature also enables central provisioning and management of VMs, as well as implementation of diskless servers such as blades, making a robust infrastructure more cost-effective than it has been in the past.

Thin provisioning, included as a standard feature in PS Series arrays, is another strategy designed to prevent downtime by enabling applications to grow nondisruptively—administrators can add capacity on demand up to preset limits. For systems that do not have online expansion, IT staff can allocate virtual disk capacity up front but not logically provision it. As a result, applications can grow when needed without downtime by allocating increments of actual capacity on demand from a free pool. This strategy does require some diligence, however, because real and perceived capacity limits may differ.

BUSINESS CONTINUITY FOR ORGANIZATIONS OF ALL SIZES

EqualLogic PS Series iSCSI arrays enable organizations of all sizes to create an enterprise-level storage infrastructure that provides data protection, DR, and BC in a cost-effective way. iSCSI connectivity means the SAN uses standard Ethernet instead of a complicated and expensive Fibre Channel network. This approach helps make the array itself much more cost-effective than it would be otherwise, so organizations can apply the savings to create robust DR and BC implementations. iSCSI also helps eliminate the need for IT staff to be specially trained in a new network protocol—the organization can leverage the existing skills of its staff.

Equally important to BC, load balancing and other management tasks are handled *by the array*, not by administrators. Continuously monitoring themselves, PS Series arrays allocate disk space—along with connectivity, security, and performance—dynamically for every application as needed. This proactive management helps prevent downtime and keep business applications running. Remote replication can be done over a WAN without expensive add-on components like channel extenders. Organizations can replicate data between PS Series arrays using Ethernet without buying additional software and licenses for various servers, because the arrays come with the necessary functionality already included: all of the management and protection features described in this article come standard with all PS Series arrays.

Because EqualLogic PS Series arrays include this advanced functionality, upgrading from direct attach storage can provide organizations with data protection, DR, and BC capabilities at the same time. They do not have to think of building a BC infrastructure over time—when they install a PS Series SAN, they get all the capabilities in one cost-effective package.

There is no software to add on to use the enterprise-class data management and protection features of the arrays, and as new capabilities are developed, they are delivered at no additional cost for systems under warranty or service plan. This is a fundamental difference from many storage vendors who charge for add-on software to provide snapshot, replication, and other features. Features such as these help make BC a reality for organizations of all sizes, with pools of storage, server power, and network bandwidth designed to operate like an integrated business utility. A few years ago this kind of environment was typically reserved for the very largest, wealthiest organizations. Today, a robust BC infrastructure is available and cost-effective for organizations of all sizes. 

Kevin Wittmer is director of product marketing for the Dell™ EqualLogic product family.