

SAN/iQ® Network RAID

Integrates Synchronous Replication with Automated Failover and Failback


NETWORK RAID AND SYNCHRONOUS REPLICATION

One of the key features that sets SAN/iQ® software apart from the competition is Network RAID. Network RAID dictates how a logical volume's blocks are laid out across the cluster, providing reliability that can be configured on a per-volume basis to best meet the needs of each application's data. Depending on a logical volume's Network RAID level, 1, 2, 3, or 4 copies of each block are striped across the cluster. You can change a volume's RAID level with no interruption in data availability.

BETTER THAN TRADITIONAL SANs

Traditional SANs make you define RAID groups ahead of time by allocating the disks to a particular RAID group. Changing volume's RAID level means taking its RAID group down for restriping, or moving the volume to a different RAID group — either way, all the applications that depend on it suffer down time during the move. Also, to support different RAID levels for different types of data, you must purchase and allocate physical disks for each, causing underutilization, creating storage islands, and forcing application downtime for copying.

With SAN/iQ clusters, Network RAID is an attribute of each logical volume, and you can change it on the fly. This allows you to closely match your storage characteristics to your business requirements. If you change the Network RAID level for a volume, the cluster re-adjusts how it stores the volume's blocks internally, and transparent to the applications that depend on it — all with no downtime. To accomplish this with traditional SANs means purchasing a second storage system and special asynchronous replication software.



KEY FEATURES AND BENEFITS

SAN/iQ clusters distribute each logical volume's blocks according to its Network RAID level, allowing administrators to closely match data availability and reliability to business requirements.

- Replicates and stripes up to four copies of each data block
- Configured on a per-volume basis
- Automatic restriping when Network RAID level is changed
- Synchronous replication with automatic failover and failback
- Non-disruptive volume migration
- Proactive self-healing reduces risk of block-level failures

INTEGRATED SYNCHRONOUS REPLICATION

Network RAID replicates blocks as they are written, providing integrated synchronous replication that has built-in failover and failback. It can be used to replicate data across a datacenter, building, campus, or even across geographies. Because this form of synchronous replication is inherent to each logical volume, it provides continuous availability by handling failover and failback automatically and transparently to application servers.

BETTER FOR BUSINESS

The bottom line is that SAN/iQ Network RAID and synchronous replication helps you to utilize your storage better. Rather than having to maintain a different storage system for each storage policy you need to support, you can support multiple policies in a single SAN/iQ cluster. Because all of your storage is pooled, you reduce fragmentation and the wasted space that multiple, isolated storage systems bring. With the SAN/iQ grow-on-demand model, you can expand each cluster dynamically as you need the storage. When you add nodes to the cluster, Network RAID re-stripes data and re-balances its workload internally, with no application downtime. No downtime means easier management: no intricate coordination between storage and application administrators, and no more middle-of-the-night struggles to restore volumes by dawn.

NETWORK RAID AT A GLANCE

SAN/iQ clusters are based on high-performance, highly reliable, enterprise-class x86-architecture servers. In the context of a cluster, these storage nodes protect against many single component failures

by storing data with hardware RAID. Each node is also equipped with features such as dual NICs and power supplies. Network RAID replicates and stripes blocks across the SAN/iQ cluster, improving further on the availability that each node provides. With Network RAID, your data is protected against multiple component failures that can cause a node to fail, multiple node failures, and even site failures.

SAN/iQ Network RAID protects against failures that the storage systems in traditional SANs cannot.

Protection against:	SAN/iQ Network RAID	Traditional SANs
Single disk failure	✓	✓
Double disk failure	✓	(RAID 6 only)
Storage system failure	✓	✗
Site failure	✓	✗

Network RAID stripes and replicates blocks on a per-volume basis, automatically retrieving block replicas in the event that a primary block is unavailable due to a failure. It automatically re-synchronizes a repaired or replaced node when it is brought online. Volumes can be configured with one of four Network RAID levels, giving organizations flexibility to closely match application and data availability requirements (Figure 1):

- *Network RAID Level 0* stripes data across the cluster and stores one copy of each block. Useful for temporary files and backup data, this RAID level can provide continuous availability across many single-component failures. If a failure makes a node unavailable, the volume becomes unavailable.
- *Network RAID Level 2* stores two copies of each volume's block, providing continuous data availability across any single node failure. This is the most popular Network RAID level with customers.
- *Network RAID Level 3* stores three copies of each block for mission-critical data that needs to be available despite any double node failure.
- *Network RAID Level 4* is used in situations where a cluster is divided between two locations, and the data must be continuously available in the event of both a site failure and a node failure at the alternate site, as described below.

DELIVERING UNPRECEDENTED AVAILABILITY

When volumes are stored using Network RAID Level 2 and above, clusters can be split into multiple locations so that data remains

available even if an entire site fails. If those locations are within the reach of a single subnet, a "Campus SAN" configuration can be used to split the cluster between closets in a building or across buildings on a campus. By placing even-numbered nodes in one location and odd-numbered ones in another, the cluster can provide continuous data availability despite an entire site failure.

For organizations with multiple sites, the SAN/iQ Multi-Site HA/DR Solution Pack can be used to manage SAN/iQ clusters spread across multiple subnets. With the Multi-Site HA/DR Solution Pack, SAN/iQ clusters become geographically aware, optimizing network bandwidth and performance by localizing data access.

SYNCHRONOUS REPLICATION — ONLY BETTER

In traditional SANs, synchronous replication is an add-on product that maintains two separate storage systems — and volumes on them — with the need for an elaborate, external failover and failback mechanism. SAN/iQ clusters manage synchronous replication, failover, and failback internally, on a per-volume basis. Because the volume's identity never changes even if a node or site fails, applications experience continuous data availability. The choice of whether to provide the application server a primary or backup block is made by the cluster itself. When a failed storage node is repaired or replaced, the cluster automatically re-synchronizes its data transparent to applications and without interrupting data availability.

NON-DISRUPTIVE VOLUME MIGRATION

Network RAID supports the ability to move volumes between clusters with no disruption in data availability. Just as Network RAID level is an attribute of a volume that can be changed at any time, so is the volume's membership in a cluster. If you implement tiered storage and wish to migrate a volume, you simply change which cluster owns the volume. Network RAID always knows where a volume's blocks reside, and it moves data blocks between clusters while continuing to make the volume available to applications.

PROACTIVE SELF-HEALING

SAN/iQ Network RAID keeps clusters healthy on an ongoing basis, with constant data scrubbing to identify and repair bad blocks before applications encounter them. It defragments storage without taking volumes offline, and it optimizes sequential read performance by using larger stripe sizes, reducing the number of seeks required to retrieve large amounts of data.



ABOUT LEFTHAND NETWORKS

At LeftHand Networks, we deliver physical and virtual SANs that are easy-to-install, easy-to-manage and designed to perform optimally in today's global data centers. LeftHand Networks pioneered IP-based SANs in 2001, and its innovative SAN products are engineered to deliver the highest availability and scalable performance, with integrated enterprise-class features.

Corporate Headquarters
2580 55th Street
Boulder, CO 80301
United States
303.449.4100

European Headquarters
10 Fenchurch Avenue
London, EC3M5BN
United Kingdom
+44 (0) 203.178.3904

www.lefthandnetworks.com

© 2008, LeftHand Networks. All rights reserved.