AUTOMATIC LOAD BALANCING RELIEVES ADMINISTRATIVE BURDEN

Load balancing across your SAN helps to optimize performance, but in most environments is a manual, iterative process that significantly disrupts operations.

Storage consolidation is an accepted best practice because it enables organizations to improve operational efficiency, reduce costs, and simplify management. An obvious benefit of consolidation is better storage capacity utilization. A less commonly discussed topic is how to ensure that consolidation does not cause performance issues: underutilization or overutilization of hardware resources, especially as workloads change and grow over time. To avoid performance utilization problems administrators must provision storage so that loads are balanced across resources evenly.

STORAGE PROVISIONING AND LOAD BALANCING ARE DISRUPTIVE

Change and growth in workloads cause organizations to regularly re-spread workloads across hardware resources – load balancing – to avoid running out of space or performance. Unfortunately, this ongoing re-provisioning is a manual, iterative task that is time-consuming and costly.

Before beginning, the administrator must determine the necessary capacity, performance, and service levels (for peaks and backups as well as normal operations) for each server, application, and data volume. The effort continues with manually planning the data layout to load balance capacity and performance across disks, controllers, cache, and network ports, taking care to prevent conflicts between applications. Finally, the administrator implements the new layout plan by shutting down applications and servers, backing up the data, provisioning additional storage, copying the data, reconfiguring applications, and restarting. The time, effort, and disruption are significant.

But what happens when workloads vary, as they do for most environments? For example, when I/O rates double near the end of a quarter, can you take the time (and downtime) to go through the whole process again?

THE PROMISE OF SANs

SANs make storage consolidation possible and extend scalability and connectivity. They also enable better data protection and simpler management – when well implemented they can provide higher application service levels. However, just implementing a SAN does not deliver these enhancements – improvement depends on the deployment choices administrators make. The problem is that administrators cannot predict all the variables needed to solve the problem. Honest mistakes caused by the complexity result in administrators spending nights and weekends moving data and reconfiguring servers.

Unfortunately, many SAN-based storage arrays are designed to be statically installed and run – yet application environments are naturally dynamic and variable. To properly grow and load balance a SAN, two key features are required: 1) Virtualization, so that when components are added to the environment application servers don't need reconfiguration and downtime; and 2) Automatic data placement and optimization in the SAN to ensure maximum utilization of both capacity and performance – without placing impossible planning demands on administrators.

SCALE UP VS. SCALE OUT

There are two ways to scale storage so that all storage resources take advantage of virtualization and automation. *Scale up* means replacing the current storage array with a larger one. The benefit: the storage remains a single device to manage with controllers, memory, disks, and networking. However, scaling up has disadvantages: a) upgrades are difficult because of planned downtime to copy data from the old to the new; b) there is eventually a hardware growth limit; and c) obsolescence occurs for all resources simultaneously. In addition, scaling up late in a product's lifecycle is expensive.

The other option is *scale out* – to grow by adding more arrays, not replacement with larger ones. With this method additions occur incrementally without impacting installed units, and there are no hardware limits imposed by a single controller/head design. Web server farms are an example of this method – when "largewebsite.com" goes from 99 web servers to 100, the first 99 are not impacted. Adding storage in this manner is conceptually simpler, but requires the ability to physically divide workloads among the arrays and to manage multiple devices efficiently.

PS SERIES DELIVERS "SCALE OUT" DESIGN WITH "SCALE UP" MANAGEMENT

Unlike most storage arrays, EqualLogic's PS Series offers the best features of both methods. Administrators can add disks or arrays online without disruption, and data are automatically load balanced across disks, controllers, cache, and network ports. At the same time the SAN retains the single device view – not only to servers accessing data, but also to administrators, who have a single point of management regardless of the number of arrays. Multiple arrays join together (transparently to applications) to act as a single device. When arrays are added, the PS Series automatically and continuously distributes and balances workloads across the environment to optimize performance. A PS Series SAN is a single-image distributed system that can grow incrementally and can be comprised of different units that fully interoperate as a single pool or multiple tiers of storage.

Load balancing is important for optimal storage performance, but when done manually it can be extremely disruptive and labor intensive. With EqualLogic's PS Series, automatic load balancing is a standard feature making storage consolidation, growth, and high performance easy to achieve and maintain.

To view other Coffee Break Bulletins or to learn more about EqualLogic, visit us at **www.equallogic.com**.

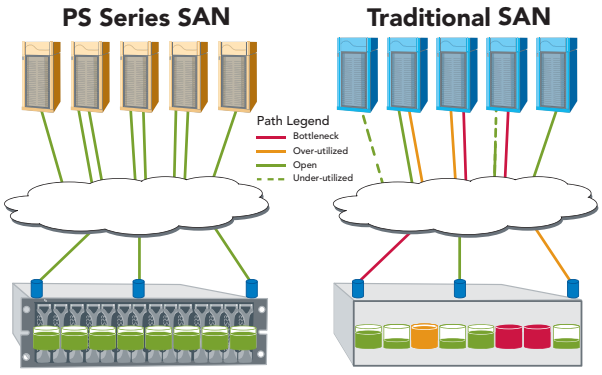

A standard feature of PS Series arrays is automatic, continuous load balancing across your SAN – storage can scale without disruption and have data automatically placed for optimal performance.

EQUALLOGIC

110 Spit Brook Road, Building ZKO2, Nashua, NH 03062
Tel 603.579.9762 / Fax 603.579.6910 / www.equallogic.com

Copyright © EqualLogic, Inc. All rights reserved. EqualLogic, PeerStorage, and Simplifying Networked Storage are trademarks or registered trademarks of EqualLogic, Inc. All other brand or product names mentioned are the trademarks or registered trademarks owned by their respective companies or organizations. CB121_USA_011807–012308