

INFOSTOR™

news analysis & trends

USER CASE STUDIES

DataCore counters performance criticism

■ BY DAVE SIMPSON

No longer a start-up after more than five years in business, feisty DataCore Software doesn't take criticism lying down. Responding to long-standing charges (mostly from large disk array vendors) that its in-band "storage area network (SAN) appliances" cause bottlenecks in storage networks, the company earlier this year ran tests based on the Storage Performance Council's SPC-1 benchmark, an industry-standard suite of tests designed to measure performance of disk subsystems.

Although benchmark tests only approximate real-life application performance, DataCore did manage to achieve 50,000 I/Os per second in the SPC-1 test—the highest number recorded at the time of testing. (The SPC tests are used by most of the large disk array vendors, with the exception of EMC.) DataCore achieved those results with off-the-shelf Intel servers in a configuration that was priced significantly below other SPC-tested array configurations.

DataCore's test configuration consisted of three Dell Wintel servers running the company's SANsymphony™ software, which was one of the first storage virtualization engines. The test configuration also included Seagate Fibre Channel disk drives in a JMR JBOD enclosure, QLogic host bus adapters (HBAs) and switches, mirroring across nodes, and software striping.

For a complete description of the SPC test configuration and details of the benchmark results, visit www.storageperformance.org.

However, user testimonials speak louder than benchmarks, and DataCore has lined up a number of customers that say that, not only does the SAN appliance not create a bottleneck, it actually increases performance in some cases (assuming the servers running SANsymphony, called Storage Domain Servers, are configured with sufficient cache).

AN UNEXPECTED BENEFIT

For example, Synovus Financial, a financial services company in Columbus, GA, implemented SANsymphony primarily for three reasons, according to Kevin Gowen, director of network services: to improve capacity utilization, for disaster recovery and business continuance, and to simplify storage management.

Synovus originally had about 50TB of direct-attached storage, but capacity utilization was only 40% to 50% of available capacity. The company migrated to a SAN that consists primarily of midrange disk arrays from IBM and JetStor RAID arrays from Advanced Computer & Network Corp. (AC&NC).

At first, the primary SAN application was a credit risk analysis program with a data warehouse, but Synovus has since migrated other applications to the SAN, including document imaging, large databases, and centralized backup.

"We went with SANsymphony for the flexibility of retaining existing storage arrays and the ability to put any type of array on the SAN back-end," says Gowen. "It also improves scalability." Synovus has two

Storage Domain Servers in its primary data center and one at its remote disaster-recovery site and uses DataCore's replication software to migrate data between facilities.

Gowen says that it's difficult to break down how much of the performance improvements are due to the SAN itself and how much are attributable to SANsymphony, but he reports that calculations that took a day or more in the direct-attached storage environment now take less than one hour in some cases. Gowen also ran comparative tests of the SAN with and without SANsymphony and says that performance was noticeably better with the SANsymphony appliances.

"We didn't choose SANsymphony for performance reasons," says Gowen. "In fact, we were surprised by the performance benefits."

INEXPENSIVE REPLICATION

At Burnett Staffing, a staffing, direct hire and training firm headquartered in Houston, the primary reason for installing SANsymphony was inexpensive data replication, according to Scott Hoffman, IT director at Burnett, which manages financial and payroll operations for thousands of temporary employees. The company originally had an EMC Symmetrix disk array, "and we were looking for an affordable way to do data replication for disaster recovery," Hoffman explains. "With EMC that would have meant buying another Symmetrix and SRDF software, which was prohibitively expensive."

Burnett got rid of the Symmetrix (primarily to save maintenance costs) and installed HP-Compaq disk arrays attached to SANsymphony servers in a storage network configuration. The all-Windows primary SAN includes three Storage Domain Servers (with seven application servers attached to them), and the replication site SAN has one Storage Domain Server.

Each of the servers has about 1TB of storage. The company replicates over a WAN via DataCore's Asynchronous IP Mirroring (AIM) software.

As with Synovus Financial, Burnett Staffing did not install the SANsymphony servers for performance reasons, but the company did realize performance benefits. "After installing the DataCore appliances our backup times dropped about 20%," says Hoffman, "and the reports for our main front-office application are noticeably faster."

In addition to storage virtualization, which is viewed as an enabling technology, SAN appliances are usually available with software for a variety of functions, such as auto-provisioning, volume management, remote IP-based disaster recovery (replication, mirroring, and snapshots), and iSCSI connectivity.

REPLICATION + PERFORMANCE

Community Health Network, a group of

five hospitals, a surgery center, and medical pavilions in Indianapolis, opted for SANsymphony for similar reasons. According to CTO Rick Copple, Community Health wanted to consolidate AIX and Windows servers in a SAN, and they needed affordable mirroring (asynchronous and synchronous).

The organization started with one IBM Shark array that supported 11 AIX servers running Oracle databases. For business continuity, instead of installing another Shark array and IBM's PPRC replication software (an expensive option), Community Health purchased two IBM FAStT disk arrays and DataCore's AIM replication software. The primary SAN includes the AIX servers running Oracle, as well as Windows servers running SQL Server. The two primary FAStT disk arrays are synchronously mirrored, and both are mirrored asynchronously to a third FAStT array at a remote site data center. (Total FAStT array capacity is about 20TB.)

On the performance front, Copple says that his requirement was that it "had to be as fast as the Shark array." After installing the FAStT disk arrays and SANsymphony appliances, Copple says that some database processes that used to take 30 minutes now take 10 minutes. On the Windows side the company realized a 3x performance improvement.

"Radiology applications that used to take 25 minutes to complete 2,200 transactions now take only seven minutes," says Copple, who attributes most of the performance gains to the Storage Domain Servers' cache utilization and cache algorithms. (Each of Community Health's SDS servers is configured with 4GB of cache.)

In addition to DataCore's replication and snapshot software, Community Health uses DataCore's Network Managed Volumes software for capacity-efficient storage provisioning, says Copple.

According to Augie Gonzalez, DataCore's director of product marketing, the primary reasons why users implement SANsymphony are better capacity utilization, easier management of distributed resources, and decreased downtime (planned and unplanned).

SAN appliances such as DataCore's enable users to create "tiered storage," where high-end (expensive) disk arrays are retained as primary arrays but lower-cost RAID or JBOD arrays are used as secondary storage.

As for the performance criticism leveled at in-band appliances, Gonzalez responds: "Performance criticism comes mainly from large vendors that fear losing sales of high-priced arrays. The installed disk suppliers have created a lot of FUD [fear, uncertainty, and doubt]." □



www.datacore.com
877.780.5111 (toll free)