

Dana Farber Cancer Institute Demands Scalable, Reliable High Performance Storage Infrastructure



When Dana Farber Cancer Institute, a leading Boston-based medical research hospital, changed the name of the Dept. of Biostatistics to the Dept. of Biostatistics and Computational Biology it was more than a departmental name change. With the new name the focus of the department shifted and the organization expanded, which increased the demands on its system and storage infrastructure.

The department already had several clusters and SANs as part of its systems infrastructure, “but this was a new project that we were building from scratch,” says Nikos George, director of scientific computing. The challenge facing George’s team was to put together the best systems and storage infrastructure for the demands of their new, expanded mission and do it quickly at the best price.

Challenge: A Green Field Opportunity

The name change actually represented the formation of a completely new group within the department. Although closely related to the work it had been doing, the new group brought in new people and set about doing new and different things in new ways. This called for a new server

and storage infrastructure. They weren’t bound by any previous technology Dana Farber may have had.

The Department’s new computational biology group goal is to develop software, databases, and bioinformatics techniques that will allow the development of new diagnostics and a more complete understanding of the cellular networks that are mechanistically responsible for diseases. Specifically, it would be developing research and analytical tools, databases, and methodologies and making them freely available to the research community via the Internet to enable research beyond its own. This is an ambitious mission, and it fell to George to assemble a system and storage infrastructure to meet the

Dana Farber had used Winchester Systems storage technology in the past and found its reliability, performance, ease of management unmatched.

needs of this expanded organization and enable it to achieve its goal. He turned to Winchester Systems, a long-time Dana Farber technology provider, to assemble a high performance, highly reliable and scalable storage infrastructure and do it at the right cost.

Since the newly reorganized Department was considered a green field opportunity—meaning that it was starting out fresh—it could seek

the right servers and storage and design its storage infrastructure to meet the needs of its new mission. “We wanted to put in a new SAN for performance, high availability, and scalability,” says George.

Solution: Familiar but New

The new organization hired more researchers with the intent of ramping up operations fast. To get the infrastructure up and running quickly, George’s team turned to familiar technology that it already understood and had used in the past. “The backbone would be a high performance computer cluster, a SAN and an Oracle RAC (real application cluster) database,” says George. Forty compute nodes would connect to the SAN, five would handle the Oracle RAC database and five would act as Web servers.

The compute nodes consisted of Dell server blades connected via Fibre Channel (FC). All the compute nodes would connect to about 8 TB of shared disk storage on the SAN. The five Oracle servers ran Oracle’s clustering file system, which allowed them to see the same space on the SAN and write simultaneously. The clustering file system protects the integrity of the stored data in the face of simultaneous writes.

For the SAN storage Dana Farber turned to Winchester Systems’ FX-608e RAID head and disk drive

enclosures holding up to 16 drives each. DCFI had used Winchester Systems storage technology in the past and found its reliability, performance, ease of management unmatched. The flexibility of the Winchester Systems technology proved particularly appealing in this case, allowing George's team to mix high performance FC drives with less costly SATA drives, which offered higher capacity although the performance was lower. In addition, the team opted for four QLogic FC switches to connect all the nodes to the SAN. "We looked at other SANs, even iSCSI, but we really wanted Fibre Channel for the high performance," says George. With Winchester Systems Dana Farber could have both high performance Fibre Channel and low cost SATA in one RAID system.

In addition, the SAN can easily scale well beyond what the Department would likely need just by adding more Winchester Systems expansion enclosures. "We can scale simply by adding more drive shelves. We can

"Once, an issue came up at 3pm on a Friday afternoon. I sent Winchester an email and had a reply and a solution by 4pm. The big vendors don't do that." But Winchester Systems did."

Nikos George,
Director Scientific Computing

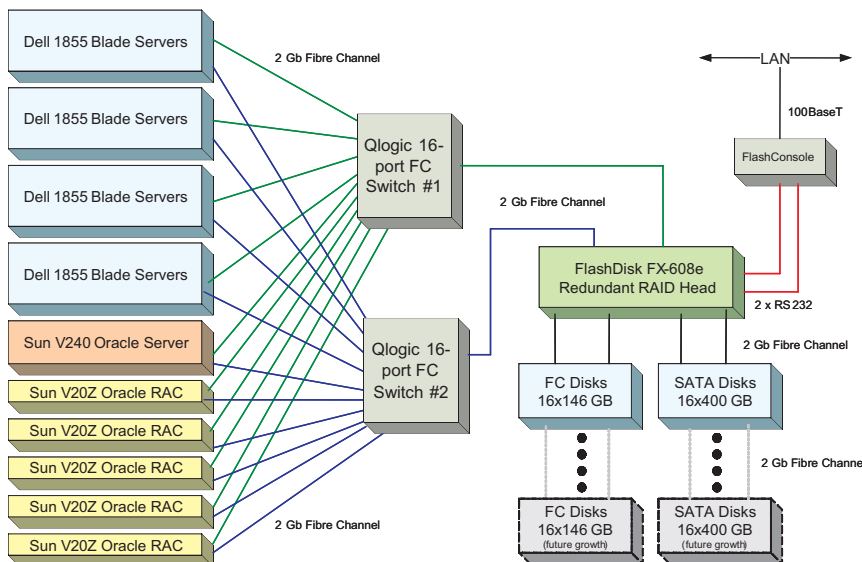
also add another RAID head," says George. Currently, the Department is running two 4Gb/s FC connections. "We could double the number of 4Gb/s connections if we ever needed more, although we are not yet near to saturating the link," he notes.

Results: Performance and Scalability at the Right Price

The Department is still forming up, and the system has not yet gone into full production. Early indications, however, confirm that the Winchester Systems SAN is exceeding George's expectations. "We have been able to avoid the kind of performance hits that others had experienced, and the price is right," says George. Based on their testing and experience to date George expects no problems

when the Department gears up to full production. Looking ahead, George is starting to plan for more backup and disaster recovery. The database presents the biggest backup challenge and currently it is backed up disk-to-disk to another filer, from which it is backed up to tape. Eventually, George hopes to replicate the SAN in another data center and mirror everything to that.

As a green field operation, George could have gone with any storage vendor, including the big name vendors that are constantly calling on Dana Farber. However, he found that Winchester Systems gives him the storage performance, scalability, and flexibility the Department needs and does so at an attractive cost. "Price is always a consideration," he notes. Winchester Systems' responsive support is an added plus. Says George: "Once, an issue came up at 3pm on a Friday afternoon. I sent Winchester an email and had a reply and a solution by 4pm. The big vendors don't do that." But Winchester Systems did."



The flexibility of the Winchester Systems technology proved particularly appealing in this case, allowing George's team to mix high performance FC drives with less costly SATA drives, which offered higher capacity although the performance was lower. In addition, the team opted for two QLogic FC switches.

WINCHESTERSYSTEMS®
Storage Without Complexity

149 Middlesex Turnpike, Burlington, MA 01803 • 800-325-3700 • 781-265-0200 • fax: 781-265-0201 • www.winsys.com