

Caltech Selects FlashDisk® To Provide Fast Data Access to Over 10,000 Users Worldwide

CERN, the European Organization for Nuclear Research, is well known throughout the scientific community not only as the world's largest particle physics center, but also as the place where the World Wide Web was invented. This invention was in response to the need for instantaneous information sharing of experiment data from the facility among physicists at universities and institutes throughout the world - institutions that used many different computer platforms and operating systems.

CERN is now in the process of building yet another grand scale data sharing environment for physicists who are involved in the new, high-energy physics facility, the Large Hadron Collider (LHC). While the LHC project will not be completed until 2006, the project is expected to produce several petabytes (one petabyte equals 1000 terabytes) of data in its first year of operation. In order for CERN to prepare to handle this massive amount of data, it is implementing a five-tier data distribution model. Development of the regional, or Tier 2, centers is being implemented at the California Institute of Technology and the University of California, San Diego. As Tier 2 centers, these regional centers will provide about half of all the


computing capability from the LHC experiments worldwide.

"The computing challenge is not only in the sheer size and volume of the data, but also the fact that the data must be accessed by 10,000 physicists who are dispersed around the globe," said Julian Bunn, who is responsible for the development and testing of the prototype Tier 2 centers at the Center for Advanced Computer Research (CACR) at Caltech.

Data Storage and Access a Major Problem

Reliable storage and rapid access are major concerns for Caltech as it prepares to deal with several petabytes of data. After considering other data storage systems, Bunn settled on RAID systems because they provide state-of-the-art performance in high-speed storage and access of data. Initially, his group began building a couple of RAID systems themselves by buying SCSI disks and controllers from Adaptec. They also tried some fiber channel devices, as well as some devices from Sun.

In order to know exactly what performance to expect when the project goes online in 2006, the speed of the storage systems was tested using benchmarks similar to the applica-



tions and operating systems that would be attached to the RAID arrays. These included a simple "C" application and a Java application. Bunn found the results were "pretty mediocre." Using RAID 0, the best his team saw with an Adaptec 3000S PCI RAID controller and four Seagate Cheetah 15,000 rpm Ultra3 disks was read at 23 MB/sec and write at 14 MB/sec. In another test, using Windows 2000 software striping, and an array of Seagate disks, Bunn and his team read at 55 MB/sec and wrote at 35 MB/sec using default settings without further experimentation. Using RAID 5 for a more realistic test, they only observed a sustained write of 2.8 MB/sec and 6.0 MB/sec with four and eight drives respectively.

The Caltech team then tried a Fiber Channel from nStor since Fibre

Channel was reputed to transfer data at 100 MB/s. Bunn described his experience, "The devices from nStor performed poorly and were very unreliable. The best we had ever seen out of the Fibre Channel was 30 MB/s while we were expecting 100 MB/s."

Winchester Systems FlashDisk - the "Fastest in the World"

It was after these experiences that the Caltech group came across an advertisement for Winchester Systems FlashDisk. What attracted them initially was the claim that FlashDisk was the fastest system in the world. Bunn said, "We jokingly, and somewhat facetiously, said to one another at the time, 'Well, our quest is over - these guys actually say they are the fastest RAID in the world. We need look no further.'"

The Numbers Tell the Story

Bunn continued, "Sure enough, when we tested a FlashDisk from Winchester Systems, it turned in the best performance figures that we had seen by a wide margin. It wasn't just a little better, but a whopping 50 per-

"FlashDisk... a whopping 50 percent faster reading, and 100 percent faster writing than anything else that we measured..."

Julian Bunn

cent faster reading, and 100 percent faster writing than anything else that we had measured at that time. With the Wide Ultra2 SCSI, we found sustained 67 MB/sec. read and write. It didn't make any difference whether it was reading or writing. It was 67 MB/sec. That was a very impressive number at the time because up until then we had been seeing very poor performance from all the other systems we tried." With the Ultra160 SCSI, these speeds were increased to sustained read at 97 MB/sec. and sustained write at 89 MB/sec.

Flash Disk Provides Phenomenal Reliability While Being Hammered

It was not just the speed of FlashDisk that impressed the Caltech team, but its impressive reliability as well. "The reliability of these disks is essential," said Bunn. "The data from this very complicated experiment is

like the crown jewels. The total cost in implementing one of these experiments is billions of dollars, so it is absolutely essential that we don't lose any of the raw data. It would be disastrous. The reliability of the Winchester Systems disks is phenomenal. We have had no failures since installing FlashDisk. And, it is not as if they are sitting spinning idle, they are really being hammered very, very hard, within two or three weeks, a terabyte is almost full." This stands in contrast to Bunn's colleagues at CERN who spend a lot of time investigating various types of RAID solutions and who have built many of their own arrays. In their experience with other systems, they have had terrible management and reliability problems keeping all of their disks going all of the time.

Other Centers Migrating to Winchester Systems

Bunn concluded, "After learning of our success with the FlashDisk array, San Diego has installed a Winchester Systems RAID array. Moreover, our success has generated a lot of interest within the high-energy physics community. It is well known that we are very pleased with Winchester Systems FlashDisk."

For more information on FlashDisk
OpenRAID and
Winchester Systems go to:

www.winsys.com

WINCHESTERSYSTEMS®
Enterprise Storage Solutions

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