

Community College Learns From SAN Disaster

February 2007

Cuesta College, a large community college in California, stored all of its critical data on a redundant SAN. It was satisfied that the SAN's redundancy would protect the data in the event of a component failure. What it did not anticipate was a failover fault. The SAN controller failed, and the failover was faulty. The result – two days of hard downtime and several weeks cleaning up stray problems. Now a dually redundant SAN provides the data protection the College originally sought.

Cuesta College

Cuesta College (www.cuesta.edu) is a public community college offering dozens of associate degree and certificate programs. The subject areas covered by these degree and certificate programs range from a variety of technologies to marketing, business, social sciences, paralegal, nursing, music, fashion, and many others.



Cuesta College is located in San Luis Obispo in central California on a 150-acre campus near the Pacific Ocean. It has an additional campus in Paso Robles, California, and centers in the Arroyo Grande and Nipomo High Schools, also in California.


The college has a ninety-year history, having started out as a division of San Luis Obispo High School in 1916. In 1964, it became part of the San Luis Obispo County Junior College District and offered evening courses at a California National Guard facility. In 1965, the college was officially named Cuesta College; and it broke ground for its current campus in 1970.¹

The College now serves over 10,000 students at its two campuses in San Luis Obispo and Paso Robles. It has over 500 full-time and part-time faculty members and is part of the California Community Colleges system. It is accredited by the Western Association of Schools and Colleges.

The Growth of IT Services

As with most enterprises, Cuesta's dependence on data processing technology has evolved over the years to the point that it is now totally dependent upon its systems. The data which supports these operations are measured in terabytes. Various services are provided by twenty-odd heterogeneous servers representing several different platforms and operating systems.

¹ Wikipedia, "Cuesta College."



A few years ago, the College consolidated its data processing operations by acquiring a two-terabyte storage area network (SAN). The SAN is linked to the various servers via fiber channel and provides storage for all of the mission-critical files and databases that had previously been directly attached to the servers. These mission-critical applications include administrative systems, infrastructure services, and faculty data.²

The intent of the SAN was to consolidate the College's data in order to provide better protection and availability than had been achieved when the data was distributed among the multiple servers. The SAN is fully redundant to protect the data, which is so necessary to the operation of the College. It uses RAID arrays for the storage media and uses dual SAN controllers in an active/backup configuration. If a disk fails, the RAID array continues in operation. The failed disk can be replaced while the system is operational and can be brought online by restoring its data from the other disks

in the array. If the primary controller fails, the standby controller immediately takes over the control of the disk arrays.

The Cuesta IT staff had done its homework. The College's data was protected by a redundant central storage facility that could be easily managed. Data loss was not an issue. Or so it thought.

The Failover Fault

The IT staff's confidence was rooted in the knowledge that should a component fail, there was another one available to instantly take its place. This is called *failover*. What it hadn't counted on was the possibility, although remote, that the failover wouldn't work. This is known as a *failover fault*.³

In the early 2000s, the primary SAN controller failed. But it didn't fail over to the backup controller. Rather, it appeared to the servers on the network that it was still working. Instead, it was sporadically spewing data all over the disk array.

When the problem was finally discovered, the system was shut down. It took two full days to restore it to service. But that wasn't the end of it. The IT staff spent another two weeks to stabilize the network and several more weeks cleaning up stray problems related to the failure.

Losing access to the system's data for two full days and then having only sporadic access seriously impacted important school operations. Many applications such as email were totally down. Servers running Active Directory, DNS (Domain Name Service), clustering services, and SQL 2000 were affected. Faculty and administrators could not get to their data. Fortunately, students were not affected since it was the school's practice not to provide students with accounts on the system.

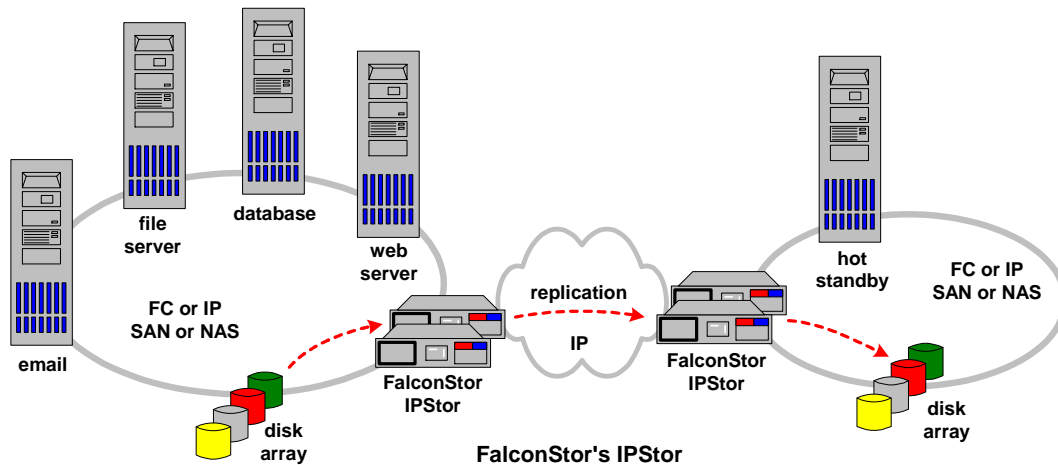
² The bulk of the information herein has been taken from the article by Alan Radding, [Through disaster, school learns valuable backup lessons](#), *Computerworld Storage Networking World Online*; August 30, 2004.

³ See Chapter 5, [The Facts of Life, Breaking the Availability Barrier: Survivable Systems for Enterprise Computing](#), available at amazon.com, for an analysis of the potentially significant impact of failover faults on system availability.

It was realized that future failures such as this would only be worse since the school's data was growing at a rate of approximately 25% per year. A better solution for database protection had to be found.

FalconStor to the Rescue

The IT staff found its solution in IPStor, a product from FalconStor (www.falconstor.com). IPStor is a fully redundant virtualization appliance for implementing SANs and NASs across multivendor and multiplatform environments. It provides a centralized management interface for these environments.



Not only does IPStor provide redundant controllers, it provides dual connections (one from each controller) to the disk arrays. It is hardware-agnostic, running on a variety of platforms under Linux or Solaris. IPStor supports mirroring, replication, snapshot copies and several other functions.

Cuesta initially used IPStor's mirroring capabilities. It purchased two IPStor appliances to provide additional redundancy in case one of the redundant appliances failed. This configuration, in fact, provided four independent connections to the disk arrays. Data was now mirrored on separate disk arrays for added protection.


As Murphy would have it, the central SAN array, which had been repaired and certified for return to service, failed again before the full installation of the FalconStor arrays had been completed. However, this time the result was far different. The mirror copies were promoted to primary storage, and service was restored in about four hours. Some data was lost, but the impact was negligible. The school since has added new SAN disks to augment the initial troublesome arrays.

Since then, the SAN's primary array failed for a third time. This time, IPStor was fully operational; and failover occurred with no user impact. Best of all, no data was lost. The only way the IT staff knew that there had been a failure was from the logs.

Improved Backup as a Plus

With the ever increasing amount of data, Cuesta found that it was struggling with its backup procedures to tape. As data volume grew and as access requirements began to extend into the evening hours, the IT staff found it increasingly difficult to complete its tape backups in the ever-shrinking overnight window.

The solution was found in FalconStor's snapshot capabilities. Cuesta purchased FalconStor's TimeMark, which takes efficient incremental snapshots and stores them on a SAN storage array. Typically, a transaction-consistent dump of the database is taken on the weekend when traffic is light. Thereafter, TimeMark can be invoked to quickly take a snapshot of only the data that has changed. A TimeMark snapshot is fast enough that these incremental snapshots can be taken quite frequently, such as hourly.



Cuesta uses TimeMark to take snapshots during the ever shortening overnight window. As a consequence, it can now leisurely back up its database to tape from the snapshot even during heavy usage times without affecting production services.

Summary

Cuesta College learned the hard way about failover faults in redundant systems that aren't supposed to fail (at least, hardly ever). In just a handful of years, it has experienced three faults in its redundant SAN hardware.

The first one was disastrous. Since then, the College has added redundant backup to its redundant systems. The subsequent failures were handled without any substantial impact. The second failure occurred before the implementation of its dually redundant solution was complete. Even in that circumstance, the College experienced a delay of only a few hours and little data loss.

The third failure occurred after the implementation was complete. That failure was hardly noticed and caused no data loss.