Maricopa's Blackboard – Disaster Recovery

High Level Abstract/Overview

Background

Maricopa's Blackboard is a mission critical application that servers as a primary teaching and learning platform for the Maricopa County Community College District. Maricopa's Blackboard is available 24x7x365 with the exception of a few scheduled maintenance windows. Currently Maricopa's Blackboard is in use at 8 of the 10 MCCCD colleges.

This document is intended to provide a high level, non-technical overview of the current application environment and some of the plans to restore service in the case of a disaster.

Current Environment

The design of the current Maricopa's Blackboard environment is aimed toward avoiding application downtime. This goal is accomplished by having redundancy of key hardware components and replication of data.

Redundancy

The Blackboard application is made up of a collection of web/app servers, file servers, database servers, storage area networks and the networking infrastructure that connects them all together.

> **Web/App servers** – There are four "identical" servers that run the Blackboard application and provide the user interface to the end users. The user load is spread across the web/app servers by a "load balancer". To bring up the login page of Blackboard the user must go through the load balancer. Each user that logs in is connected to a web/app server. The load balancer remembers which user is sent to which web/app server (while the user is logged in) and keeps that user "stuck" to that web/app server. If one of the web/app servers develops a problem or "goes down" those users will lose their session and will need to login again. When the user brings up the Blackboard login page they will be sent to a web/app server by the load balancer. The load balancer monitors which of the web/app servers are available and only sends users to servers that are available. No user data is stored on any of the web/app servers. There are spare web/app servers in the Blackboard environment that can be activated if needed.
>
> **Collaboration & Integration Servers** – The Blackboard collaboration tools use a collaboration server to provide the virtual classroom and chat services. Blackboard does not have a way to configure multiple collaboration servers. Likewise Maricopa has a dedicated Integration server that is used to push data from the Student Information System (SIS) into Blackboard. These two servers have been configured to be a "warm backup" for each other. If either of these servers were to fail the other could be configured quickly to take on the role of the other. Each server is capable of providing both services with good performance. Interface data stored on the Integration is a copy of data from other servers. Loss of this data would not cause any problems with the production Blackboard system.

File Servers – Blackboard stores content uploaded by the instructors and students into content folders on a file server. For Maricopa's Blackboard the actual files are stored on two independent storage arrays in a "Storage Area Network" or SAN. A storage array is a highly redundant collection of individual disk drives that are accessed by an array controller. The data on the array is "striped" in such a way that if any of the array drives fail the data on the failed drive can be extrapolated from the other still functioning drives. Failed drives in the array can be replaced while the array is operational. When a new drive is installed the array will populate the drive with extrapolated data to fill in what is missing. The array is accessed by a pair of redundant servers that present the drives to Blackboard. The arrays are connected to the servers via two independent high speed "fiber channel" controllers. The "primary" file server tracks all changes to the Blackboard content and all changes are replicated to the other file server. If the primary file server goes down the content on the other file server will be up to date.

Database Server – There are two database servers in the Blackboard environment. One server is the database server that is accessed by the Blackboard application. Using a built in feature of Microsoft SQL server, known as "log shipping", all changes to the primary database server are "shipped" to the backup database server. The logs are "shipped" once per hour. Also, incremental database backups are done every 30 minutes and stored on a different server.

Networking – The Maricopa network that supports Maricopa's Blackboard is highly redundant following Cisco best practice design. Components of Blackboard are distributed across redundant "distribution" and "core" switches in such a way that the failure of one switch would not take down the Blackboard application. For example half of the web/app servers are connected to one network switch while the other half are connected to a different switch. If one of the switches fail, at worst only half of the Blackboard web/app servers are impacted.

Disaster Recovery

The term "Disaster Recovery" can cover a lot of different scenarios from a server hardware failure to a complete loss of the data center. For the purposes of this overview the assumption will be that the data center is intact and that the disaster is limited to one or more individual pieces of hardware.

Central to the Disaster Recovery ability is the tape backup and archival system. This system is an enterprise class data storage and retrieval system that not only schedules the backups but also catalogs what data has been backed up from what server, when and what tapes it is stored on.

Another integral component of the Disaster Recovery plan is off-site storage of data tapes. At Maricopa, the weekly "full" backup tapes are copied or "cloned" and sent off-site to a data storage vendor. With the advent of the New Student System tape cloning of critical systems and their storage off-site will be will be done daily. Having the tapes stored off-site ensures that the tapes are safe from both accidental erasure and from any disaster that may impact the data center.

Disasters can come in many shapes and complexities, here are a few scenarios and their possible solutions.

Database Server Failure – In the case that the primary database server fails the failure will be investigated. If the server cannot be brought back on-line in a few hours the replicated data will be validated and the backup database server will be put into production. The backup database server is not identical to the primary database so there may be an impact on performance while the primary database server is being repaired. If the data on the backup server does not verify as being error free the data from the previous night backup will be restored and the incremental backups and SQL change logs applied up to the point just before the data corruption starts.

File Server Failure – If the primary file server hardware fails the replicated data will be checked to see that it is valid and the secondary file server will be put into production. This process takes only a few minutes but does require a "reboot" of the Bb application. If the replicated data was corrupted by the failure of the primary file server the Bb content files will need to be restored from tape. Restoration from tape will take several hours.

Storage Area Network Failure – The SAN is designed to be highly available and uses built in redundancy to prevent failure, but like any piece of technology it can (and has) failed. In the event of an array failure, the secondary copy of data on the second array can be brought online in a few hours with only minutes of data loss. If data is corrupt on both arrays then tape and storage log files (from direct attached disk drives) can be used to reconstruct data. This can take several hours with data loss less than an hour. If both arrays fail and the data cannot be reconstructed the previous night backup tape can be loaded. Restoration from take will take several hours.

Individual Web/App Server Failure – This does not pose a threat to data. If one of the web/app servers goes down the other three will carry the load until a spare web/app server can be verified and then put into production. If it is during peak usage there may be a performance impact, but at all other times three of the four web/app servers can support the users.

Networking – There is enough redundancy built into the Maricopa network that it would take several pieces of equipment failing at the same time to bring down Maricopa's Blackboard.

Data Center Power Failure – The Maricopa District Office building that houses the data center is equipped with a backup generator. The computer room has its power supplied via an uninterruptible power supply (UPS). In the case of a loss of building power the UPS will continue to supply power to the servers. With the loss of building power the backup generator will automatically start and power will be supplied to the building before the batteries in the UPS are exhausted. If this event occurs several people in both ITS and building facilities will be paged by the paging system. If the generator does not

start ITS personnel will need to shut down non-essential applications to save battery power. Eventually if the generator does not start Maricopa's Blackboard will need to be shut down manually to prevent data loss. Recently the ITS Operations department moved to a 24x5 schedule. The Operations staff will be trained to perform an orderly shut down of systems to prepare for when that may becomes necessary.

To help ensure that Maricopa is ready in the case of a disaster the backup tapes of Maricopa's Blackboard are tested about every two weeks. For the test a copy of the database is restored from tape to a server in the Blackboard Test environment. Files from the file server backup tapes are also restored to the Test file server. The restored Bb environment can then be accessed via the Test environment web/app server to ensure that the restored system is intact.

It is also important to note that instructors have the ability to backup each of their courses using the Blackboard "Archive" function. An archive of a Blackboard course contains all of the content uploaded by the instructor, the enrollments of the students, all work submitted by students and all grades in the grade book. The archive function was designed to contain all data relevant to the course. If the course were to be deleted or somehow corrupt the course could be restored to the previous archive. Instructors are encouraged to archive their course during the course of the semester. On a smaller scale instructors can also export the grade book from each of their Blackboard courses. The grade book export captures only high level data, but still is a valuable tool for saving Blackboard data outside of the Blackboard environment. Instructors are encourage to export the grade books from their Blackboard course often.