

When Thomas Steed of CasetrackerLaw approached AHPHosting about High Availability for his web applications we took a close look at the application and what we would need to support the applications needs and the end users needs.

CaseTracker Law as setup on a Cluster of Windows File server with a Linux MySQL Cluster backend. The front end or public side of the application is running on Windows 2003 Server with Adobe ColdFusion 8 Enterprise version IDE.

We have several parts to this applications support structure.

- 1} Network & bandwidth both internal and external
- 2} Redundant Network Hardware
- 3} Redundant Power nodes
- 4} Multiple Windows file servers
- 5} Multiple Linux MySQL servers
- 6} Redundant Hardware Load Balancers

The network map is a major part of the cluster operation. To provide complete fail over protection we duplicated the network for the CaseTracker application. Each server is on a separate pair of dedicated layer 2 gigE switches'. What this means is that AHPHosting can suffer a switch failure on the network and the CaseTracker application would not be effected. Each server is powered form separate power nodes in our data facility so we can lose power on any one of the 8 power legs we have and the application still remains online and un-interrupted, even if it is one of the CaseTracker switches or power nodes that fails.

The next part of the network is who the network actually is. AHPHosting retains colocation and bandwidth services from Internap.com. Internap is the leading bandwidth data center in the united states. Colocation and bandwidth are more expensive here than any other provider and for good reason. Internap has one awards for the routing polices they use that put them on the front page for all other data centers to model. Currently Internap's Atlanta Georgia facility is peered with 13 Tier one backbone providers to ensure we have the best routes to end users available.

About the Windows portion of the Cluster first.

AHPHosting deployed a single SuperMicro server with Intel Dual Xeon 3.2 GHz CPUs and 4 GB ECC RAM. We used Western Digital 10k RPM Drives in a RAID 5 Disk Array. The RAID 5 Array gives CaseTracker a total of 210 GB fault tolerant storage with fault tolerance for possible hard drive failures. This is the administrative server. All file updates and application adjustments take place on this server. Updates are pushed out to the front line servers where the public application resides. This server can manage up to 100 instances of the CaseTracker application under load. It is designed however to act as the fail over incase any one of the primary servers fails or has to be taken offline. The High Availability allows for server failure with out interruption to the application from the end users point of view. The administrative server acts as the fail over server for the front line servers that I will discuss in the next paragraph. It will not be publicly accessed unless one of the primary ColdFusion file servers has a failure. At that point the server would only be routing packets for 10 to 15 sites, depending on the front line server that failed.

The next part of the Windows equation is the primary HA servers.

These are the Windows ColdFusion file servers that are on the front line of the public internet. These servers are also enterprise grade SuperMicro servers. We used Intel two Quad Core 5335 2.0 GHz CPUs with 4 GB ECC RAM each. These are also deployed with RAID 5 10K RPM disk arrays for fault tolerance in case of a single disk failure in the array. If a drive fails in a Raid 5 array we can remove that hard drive and replace it with a new drive without even turning the server off. These are known as HOT SWAP drives. Each front line server supports 10 to 15 instances of the CaseTracker application. Each front line server is synchronized with the files and settings on the administrative server. Backups of the front line servers are run on a daily basis and stored up to 14 days in a remote location for security.

File Server Load Balancing...

The Load Balancing is handled by a pair of Kemp technologies LoadMaster Linux servers designed specifically for this application. The Loadmasters' constantly monitor the Windows nodes in the cluster to

confirm availability. If one of the Windows nodes fails for any reason then both the users sessions and packet traffic are diverted to the administrative server until the front line server is recovered. We can even afford to lose a LoadMaster and still continue to route traffic to the CaseTracker application. By using these Loadmasters' AHP can insure that the CaseTracker Law application is always available with a reasonable 99.9999% uptime.

Ahp also had to develop and design the database server architecture to support the CaseTracker application in a high availability mode also.

We deployed 2 Dual Core SuperMicro servers each with 8 GB RAM, 10K RPM drives in RAID 1 Mirror array. We are using all GPL licensed software to include Centos Version 4.5 as the base operating, MySQL community server as the SQL backplane. The next layer in the architecture is Sequoia by Continuent.org which provides load balancing and replication of the databases and queries between the cluster nodes, Myosotis also by Continuent.org acts as the final software layer of the cluster servers and acts as the go between between the Coldfusion Administrator and the cluster backplane. In addition to this we are using a custom written PHP application which is used for configuring the software applications on the fly this PHP application also manages the backup system.

The backup's are run every 2 hours on the hour on each of each individual client databases so in the event of a catastrophic failure the data loss would be contained to a maximum of 2 hours, these backups are stored on both nodes of the cluster and once daily they are encrypted backed up to a terabyte server.

Database traffic flow

Client Request for data -> Coldfusion Administrator -> Myosotis -> Sequoia -> MySQL Server

Traffic Security for the cluster nodes is secured using public private network scheme the cluster nodes which house the data are on the private network and are not available from the outside network. The nodes only respond to request from the designated addresses and ports of the ColdFusion servers.

Greg