Hyperion Essbase Backup Strategies, Part 1

Backing up Essbase can be accomplished in a number of ways. Some methods suit some organizational cultures better than others. It is hard to argue that one method is better than another for this reason. Below are two methods, and the pros and cons of each.

There are a number of factors that must be considered. If the environment uses some of the new Hyperion tools, like EPMA, then one must allow consideration for the synchronization of the warehouse that holds the data for EPMA. Where the different Hyperion applications (Shared Services, the web server, etc.) that work together reside is also a factor.

To minimize the complexity of this discussion, only information related to Essbase will be discussed.

Backup the entire server

Pros: An image of the entire server is available in the case of disaster recovery and is normally in sync to that point in time of failure
Cons: Speed, cost, and data availability

Taking an image of the entire server is one option. This will provide the most secure backup strategy. If there is a hardware failure, getting back to the point of failure does not require a server rebuild. This method is probably the quickest solution to restore all Essbase applications. Price, speed, and data availability must be considered with this solution. Taking an image of a server can be very time consuming and quite often, Essbase must be turned off for this to occur without skipping critical files. Because a large amount of data is backed up, a large amount of storage is required. The time Essbase is down can have a significant impact on the people using Essbase. There can be a very expensive price tag for the amount of tape and/or SAN that is required. To effectively image a server without significant downtime, techniques like shadow copy or data mirroring are likely used.

Backup critical Essbase files

Pros: Speed, cost, data availability
Cons: Recovery time is sometimes longer, more effort if a complete system failure occurs, and data from the most recent backup to the point of failure is lost

The files required to be backed up to recover from a catastrophic event are actually very small in size. The bulk of the amount of data related to Essbase is the pag and ind files, or the data and index files. These files, in most environments, consume at least 90% of the total space. If these are ignored during the backup process, the process can be much faster, far less expensive, and Essbase is not required to be off for the backup to occur. Although this method can take longer to restore an entire server, it can be quicker to restore a few applications. In most situations, a faster, cheaper solution, where the availability isn’t negatively impacted, is a far more palatable option. This is only an option if you have either the data that sources the databases or data exports (input or level 0) of the Essbase databases. If these are available, the databases can rebuild the pag and ind files.

Deciding on a backup method

Determining the best option boils down to cost and resources. Taking an image of the server requires at least 2 times more disk space, a more complicated network/hardware infrastructure, and far more resources to build and store sufficient backup versions. What is gained is an up to the minute backup. If the cost associated with this method outweighs the cost of having to rebuild the data that was loaded between the time of failure and the last backup, then this solution is the best option. In my opinion, it is hard to justify the investment in the capital required to support this for what little is gained.

First, disasters rarely happen. With the RAID and SAN solutions today, disk failures that cause data loss are not the main reason a server fails, a hardware component failure is. If the component that fails is replaced, the data doesn’t have to be restored.

Second, if a database becomes corrupt and unusable, a complete reload of the data is required. Many times corruption can exist, unnoticed, in a database for weeks. If the data is not available to reload, it is possible to lose weeks or months of data.

Third, if a disaster does occur, any data sourced from another system can be recreated. Remember, the only data required is the data that has changes prior to the most recent backup, which is normally the previous night. The data loaded by users, either through Hyperion Planning web forms or spreadsheets (Excel Add-In or SmartView), also exists somewhere else. It might be frustrating for users to enter it again, but the data does exist and can be restored, normally with minimal effort. In very large environments, this backup method can save millions of dollars.

Whether the decision is made to mirror the server, backup the critical Essbase files excluding the data consolidations and index files, or some method in the middle, it is wise to test the disaster recovery plan. There is nothing worse than restoring from a backup only to find out that it is useless.

The second installment of this topic will be dedicated to how and what is required to have a secure DR plan if the pag and ind files are ignored in a backup strategy.