BusinessObjects Enterprise XI 3.1 Backup and Recovery Best Practices

Applies to:

BusinessObjects Enterprise XI 3.1

Summary

There are multiple ways in which a backup can be performed of the BusinessObjects Enterprise (BOE) XI 3.1 system. This paper outlines the best practices for establishing a backup strategy and provides insight into how to perform cold and hot backups of the system.

Author(s): Michael Stamback

Company: SAP Business Objects

Created on: 30 June 2009

Author Bio

Michael Stamback is a Director of Product Management in SAP Business Objects covering the BusinessObjects Enterprise Platform and associated components.

Table of Contents

Disclaimer: The following white paper describes some backup methodologies that have not been tested in SAP Business Objects labs and are not officially supported, unless otherwise stated. SAP Business Objects is in no way liable for results that may occur from attempting the stated backup methodologies.

Introduction

This document describes procedures for performing a back up of the Web and BusinessObjects Server system and data files, as well as the procedures to be followed to recover from data loss or hardware failure. The plan execution requires an experienced BusinessObjects professional, operating system administrator, and database administrator.

The back up and recovery process is similar for all environments within a staged system; development, test, and production. Therefore, this document does not refer to any specific environment. It is recommended to back up all environments.

Back Up and Recovery Process Concept

A back up and recovery plan consists of precautions to be taken in the event of a full system failure due to a natural disaster or a catastrophic event. The plan aims to minimize the effects of the disaster on the daily operations so that the organization is able to maintain or quickly resume mission-critical functions. It is recommended, as part of a BusinessObjects Enterprise (BOE) disaster recovery plan, to involve an implementation of redundant servers in a back up system, which mirrors the primary system. In the event that the primary system goes down, the back up system is still available and becomes the operational system.

It is a best practice to back up BusinessObjects servers on a daily basis.

The Importance of a Proper Backup Sequence

Without backing up the Enterprise system databases and the file repository, no restoration of the environment is possible.

Good backup integrity requires shutting down certain BusinessObjects services prior to capturing the backup. Under certain conditions, failure to follow a proper service shut down sequence before backing up could result in one or both of the following consequences:

- False sense of security in the quality/integrity of the backup
- Inability to fully recover BusinessObjects if the backup is of poor integrity

Types of Backups

A backup can be conducted in multiple ways:

1. Enterprise Content backup vs. server backup

   A server backup is the deepest type of back up. Typically, this is a full backup that captures every byte on every local hard disk in a server. This type of backup insures the server's operating system as well as any applications and data stored on the server's local drives. Server backups are infrequent, typically only necessary after a stable installation/configuration is achieved and anytime major updates to any of the server components are made.

   Enterprise content backup is a subset of a server backup and focuses on the essential components of BOE only. With an Enterprise content back up, one can recover a BusinessObjects environment from scratch on a different hardware environment if necessary.

   Enterprise content backup does not insure the operating system or any applications, including the Enterprise executables and DLLs; however, it does insure all of the intellectual content that is created and customized by users of the Enterprise system, which is the most important thing. Content backup also will not insure one against corrupted or deleted Enterprise application DLLs/EXEs or other operating system problems, but these can be resolved by performing a re-installation of operating systems or

Enterprise. Enterprise content backups are relatively fast and inexpensive to capture and thus frequent in nature, as content continuously changes.

This document mostly focuses on an Enterprise content back up strategy.

2. Incremental backup vs. full backup

   The choice between incremental backup and full backup is typically a decision made by the disaster recovery team. A full backup captures every targeted byte. Full backups require longer system down time to achieve (in the case of cold backups) but are the safest and least complex type of backup to restore. Because full backups are the slowest and most expensive to capture and maintain, Business Objects advises capturing full backups on a weekly basis, where resources allow.

   Incremental backups can be performed on a more frequent basis. It begins as a full backup, but each subsequent backup thereafter captures only those files changed since the last backup. They are faster and less expensive to capture than full backups, but are less robust. Because of this, incremental backups are usually acceptable for daily backups.

3. Cold backup vs. hot backup

   Using a cold or hot backup strategy is typically dependent on whether a system can tolerate the downtime required to perform a cold backup. Prior to running a cold backup, the Central Management Server (CMS) service must be stopped so that a back up of the CMS system database and File Repository Server (FRS) can occur. Stopping the CMS will prevent users from accessing the Enterprise system, a factor which must be taken into consideration when scheduling cold back ups. This assures an accurate snapshot of the system occurs, since no transactions can occur during the backup procedure. Cold backups must be done in off peak hours, thereby limiting the frequency at which the backups can occur and sometimes requiring choreography of schedules with other dependencies.

   Hot backups tend to be the preferred method in global or highly available environments where there is not a feasible window of downtime, as backups can occur while a system is still running, thereby keeping usage of the system unaffected. This enables the narrowing of the gap between the period of time when the last backup was run and when the system experiences failure, as backups can occur more frequently, reducing the amount of work lost in the event of a failure. Because the backup occurs on a live system, a hot back up exposes the possibility that the CMS and Performance Management (PM) tables and FRS will be captured in a state of change; furthermore, the FRS, CMS and PM tables could be captured in an out-of-sync state, a situation which could compromise the integrity and usefulness of the back up.

It is recommended to back up the CMS system database and PM repository once daily with incremental backups and full backup only once per week. The frequency of backups may be revised based on an acceptable amount of time for recovery to satisfy your organization's requirements. This rest of this paper will focus on the procedures necessary for conducting a backup of enterprise content.

What to Back Up

There are several content-oriented elements that must be backed up daily in order to recover from a disaster. Routinely backing up all of the following content elements will enable you to recover from virtually any type of disaster (virus, hardware failure, natural disaster, and so on).

Database

Central Management Server (CMS) System Tables

The CMS contains all the user rights and metadata information about reports and universes, as well as data connection information. The CMS is the heart of the BusinessObjects environment, so it's critical to frequently back up this database. The environment cannot be restored without a properly backed up CMS database.

Performance Management (PM) System Tables

The PM database stores metrics, key performance indexes (KPI), metadata, and key relationships that drive dashboards and scorecards. In many Enterprise environments, theses tables will be stored in the same physical database as the CMS system tables, making it easy to capture during a routine backup of the CMS tables. This database should be captured at the same frequency as the CMS system database.

File Repository Server (FRS)

This is a standard OS file share that contains all the report templates and instances for the environment. The file repository typically ranges in size from 1 GB to 100 GB depending on the size and complexity of the deployment. Since the FRS is designed to exist synchronously with the CMS tables, it should be backed up at exactly the same time as the CMS system tables.

Failure to capture the file repository and CMS system tables simultaneously, when the CMS and FRS services are stopped, could result in poor back up integrity. This is because of the increased risk for orphaned report objects or report pointers in the CMS system database if a database restore becomes necessary. Orphaned report pointers are those records in the CMS system database that do not point to a valid input or output file in the Enterprise Input or Output FRS. If a user were to select a report object or report instance that was orphaned, an error message would occur and they would not be able to access that object or instance. Orphaned input or output files in the file repository are a benign problem compared to orphaned pointers in the CMS tables. Only the Input and Output subfolders need to be backed up (the Temp folder can be ignored).

Local audit log files

When auditing servers is enabled, each BusinessObjects server writes audit records to a log file local to the server. At regular intervals, the CMS communicates with the audited servers to request copies of records from the local log files.

Auditing Tables

This database contains usage statistics and auditing information for the environment. A back up of this database is not necessary for recovery from a disaster. However, recovery of a CMS database without recovery of the auditing tables from the same backup might result in a corrupt auditing database due to the presence of duplicate event IDs. It's highly recommended that the auditing tables be included as part of the backup procedure.

Custom Java applications/code

Any programmatic customizations to InfoView or other custom user interfaces should be backed up as frequently as they change. During active development periods, these items should be backed up daily. Thereafter, they can be backed up weekly or monthly, or as often as the code is modified. Currently, Business Objects Integrated System consists only of Port Director Dashboards and does not contain any custom code.

The back up and recovery process should be revised once more projects are integrated into it.

Database Connections (ODBC DSNs)

Special database connections (such as ODBC DSNs) can not be easily captured and restored via normal back up processes. Because of this limitation, a back up and recovery plan should cover connectivity parameters for all known data sources. At a minimum, it would cover:

- The name of the ODBC DSN
- The name of the target database
- The type of target database in Oracle
- The user ID and password used to connect to the database
- A listing of the reports that rely on the data source

- Any other pertinent information that an administrator can use to recreate the data source manually if necessary

Query Data Source – Data Warehouse

Although technically not Enterprise components, all reporting databases and cubes should be backed up on a daily basis or at least as frequently as the data changes.

Overview of BOE XI 3.1 Backup

BOE XI 3.1 backups may be undertaken in "cold" or "hot" mode. There are best practices associated with performing either that ensure a fully stable and consistent system following a restore.

To determine what method of backup, hot or cold, is most appropriate for your organization, consider the following questions:

- How often do you need to backup your content?
- How frequently will you require a restore?
- How much time exists to complete the backup?
- Do you require high availability of the BOE system?

The answer to these questions should determine your choice of cold vs. hot backup. If the frequency of backup is daily or longer and there is a reasonable window in which the system can be unavailable, then cold backup will be the preferred choice. If backups are required more frequently than daily or a reasonable window for downtime does not exist, then hot backup will be the preferred choice.

Regardless of which method you choose, a few best practices should be followed.

1. Communicate the planned backup windows so users can plan activities outside of that timeframe accordingly.
2. Do not schedule reports to run during the backup window in order to ensure a complete backup.
3. Do not schedule federation jobs to run during the backup window in order to ensure federated content gets included in the backup.
4. Do not run a promotion process for a report template during the backup window. This will ensure the template gets captured as part of the backup procedure, and any reports running against the template, in the event of a restore, will succeed.

Cold Backup

A cold backup is the recommended approach to backups by BusinessObjects, as it is the most straightforward way to ensure a complete backup occurs. A cold backup requires a shutdown of the BOE services, thereby making the system unavailable. As such, cold backups should be scheduled during off peak hours. All parties affected should receive communication of the timeframe in which the system will be unavailable, thereby avoiding potential loss of productivity with the system.

A cold backup strategy requires the shutting down of BusinessObjects services prior to capturing the back up. The high-level sequence for backup is

1. Stop the Central Management Server (CMS) and all job-processing and Performance Management (PM) servers, preferably using an automated script. This will:
    i. Release the connection of the CMS to the system database
    ii. Release PM's connection to its tables in the CMS system database
    iii. Prevent users from accessing BOE.

iv. Prevent report jobs/analytics from executing and terminate any report job/analytics already in execution

v. Assure the continued integrity of the FRS.

If the sequence is not followed the result can be a false sense of security in the quality/integrity of the backup or an inability to fully recover Business Objects if the back up is of poor integrity.

To aid in stopping and restarting Enterprise XI services before and after a cold back up occurs, it is recommended that two scripts be used, one to stop the relevant services and the other to restart them. These scripts can be executed before and after the database back up process begins. The UNIX distribution of BOE includes the ccm.sh script to assist in the start, stop and restart process.

Hot Backup

As stated earlier, the advantage to a hot backup is that the system does not have to be brought down to perform the backup operation. This enables the narrowing of the gap between the periods of time when the last backup was run and when the system experiences failure, as backups can occur more frequently, reducing the amount of work lost in the event of a failure. Many organizations prefer the use of a hot backup strategy in general due to the global nature of their business. Systems may be accessed globally, which leaves little to no window for downtime, and thus a hot backup strategy is required.

Hot backups do, however, surface the possibility of inconsistencies between the FRS and CMS. This is due to the fact that the FRS is typically large (on the scale of gigabytes), while the CMS is relatively small (on the scale of 10s of megabytes), resulting in the backup of the FRS taking longer than the CMS. Because the system remains accessible during a hot backup, transactions can continue to occur during the backup process. Any transactions occurring during this time period will not be part of the backup, resulting in the possibility that the system becomes out of sync on a restore. The following are potential events that can cause the FRS and CMS to be out of sync in a non-atomic hot backup scenario:

- User (or a job) deleted a document or document instance (Webi, CR, etc)
- User (or a job) added a document or document instance (Webi, CR, etc)
- User (or a job) modified a document

The effects of these events in the worst case are as follows:

- Some documents will appear in InfoView but cannot be opened because files will be missing (in case of “delete”)
- In case of “modify” the version of file will be exactly as it was when the file was backed up.
- In case of new document addition after database backup is completed and before FRS backup is completed the new document will not appear in InfoView. There are no functional side effects here except orphan files may exist in the FRS.

The above listed effects will in no way cause the system as such to be unstable. The only victim is the specific user of the document. In order to avoid that the user is presented with above described potential inconsistencies, it is possible to correct them using the repository diagnostic tool (RDT), in which case the users will not experience any abnormal behavior.

There is an alternative approach to hot backup that uses a combination of database hot backup and hardware backup solutions. Deploying the FRS to a storage area network (SAN) that has hardware shadowing to a backup SAN will minimize the potential synchronization issues between the FRS and CMS. With this approach, the CMS database backup can be started at the same time as the FRS shadow copy is peeled away and backed up. After the shadow backup has completed, rejoin it to the primary SAN to resynchronize the two. While it is still possible for an object to be inserted or updated during the time between severing the shadow for backup and the CMS backup initiation, it is highly unlikely. If using this technique, remember that the CMS backup and FRS backup must be started as close to simultaneously as possible, but the FRS backup should never be started before the CMS.

Recovering BOE Content

There are two types of scenarios where recovery of BOE content might be necessary:

- Recovery of BOE content in the event of a system failure
- Selective recovery of BOE content in the event of an accidental deletion, breakage from deployment of a new version of content, etc

While recovery of BOE content in the event of a system failure is far more important, a selective recovery of content is far more common. Both recovery procedures are covered below.

Disaster Recovery

Recovery of BOE in the event of a failure is relatively straightforward, regardless of whether a cold or hot backup procedure was chosen. If a hot backup procedure was used, however, it will be critical to run the Repository Diagnostic Tool (RDT) to ensure the CMS and FRS are synchronized.

High Level Recovery Sequence

In the event a recovery is necessary, follow these steps to restore BOE:

1. If necessary, rebuild and restore the system, ensuring that the number and size of disk volumes are the same or larger than the previous system. If you must rebuild a system by starting with an empty hard drive, install the OS on the same disk as before, then recreate the partitions and volumes as they were on the damaged system. If recovering only BOE content, make sure to install BOE on the same drive as on the original system.
2. Restore the backup of the CMS system database.
3. Restore the backup of the PM repository.
4. Install/configure the data source client software to point to the restored database and report source.
5. Restore the Input and Output File Repositories (FRS)
6. Run the Repository Diagnostic Tool

For replicating on a second server, conduct the installations as above and use the Business Objects Import Wizard to import the required information from one Enterprise system to another.

Individual files from a specific date can be restored from the file system back up tapes. If it is necessary to restore entire file systems, the most recent full system back up tape should be restored first, followed by the first incremental back up, then the second incremental back up, and so on until the file system is fully restored.

The Repository Diagnostic Tool

BOE XI 3.1 provides a Repository Diagnostic Tool (RDT) for BI administrators to keep the CMS and FRS in a stable/consistent state, thus helping to reduce down time. It is used to scan, diagnose, and repair inconsistencies that may occur between the CMS and the FRS. RDT scans the CMS system database and identifies inconsistencies, repairs the logged inconsistencies, and reports the repair status and completed actions. It can be run against a live system, but in a recovery situation, it should be used after a restoration and prior to starting the BOE services to make sure no inconsistencies exist, especially in situations where a recovery is being performed from a hot backup. This allows for the restoration of the backed up CMS and FRS with no downtime and with consistent behavior.

The table below describes the list of inconsistencies that could potentially occur as a result of a backup and what action the RDT will take to resolve them. Note, the RDT may be setup to merely warn about inconsistencies and allow the administrator to take manual action.

Inconsistency	Description	Repair Actions
InfoObject exists, but no file	It is possible that an InfoObject exists in	Delete the InfoObject, unless otherwise

	the CMS, but there is no file FRS	told
File exists but no InfoObject	It is possible the file exists but there is no corresponding InfoObject	User is notified to republish the object
Invalid Parent ID	An InfoObject can potentially have an invalid parent reference	The object and its children will be moved into a folder call 'Repair'.
Last Successful Instance	The reference to the last successful scheduled instance could be invalid	Remove the ID and let the CMS automatically recalculate it
Invalid Target ID	A shortcut could be pointing to an invalid object	The shortcut object will be deleted, unless told otherwise.
File size is wrong	There can be information discrepancies between the InfoObject and actual file	Update the InfoObject
Empty folders in the system	There may be empty folders due to old objects	Remove the empty directories, unless otherwise told.

The RDT is a command line tool that is run by the administrator. Prior to running the RDT, make sure the user account the RDT will be running on has full access to inputfrsdir, outputfrsdir, and outputdir. The RDT can then be run using the following command:

- For UNIX users, the RDT can be run within a */usr/bin/sh* compatible shell calling *./bobje/enterprise120/linux_x86/boe_reposcan*
- For Windows users, the RDT can be run from <*BOE installation dir*>*\BusinessObjects Enterprise12\win32_x86\reposcan.exe*

Log files are created in XML format for both the scan and repair. A default XSL file is provided to administrators so that they can use it with the XML file to produce an HMTL page for viewing the log files. The XSL file is stored in the following directory <*BOE installation dir*>*\BusinessObjects Enterprise12.0\reposcan*.

Following the running of the RDT as described above, the customer will experience a system without errors that reflects the state of the system just prior to the backup.

Selective Content Recovery

BOE XI 3.1 provides a new BusinessObjects Enterprise Lifecycle Manager (LCM) tool that provides a far easier method for restoring selective content in the event the entire content store does not need to be restored. The LCM tool's primary usage is for promotion of content from one environment to another, for example from development to test. Promotion jobs are created and promote content held within the repository (such as reports/documents, universes, connections, users, groups, etc). LCM also enables

version control of content with the ability to roll back content either fully or partially to previous versions in the event of an error.

There may be some cases where a restore from a backup is not necessary, nor desired, as only a specific content item needs to be restored. In cases such as these, content could be recovered using LCM by simply reapplying a promotion job to the target environment. Additionally, content held in version management can be retrieved by rolling back to a previous version should the current content no longer be valid or is removed. It is not recommended, however, to rollback more than one version.

For more information on how to use LCM, see the LCM user guide on http://help.sap.com.

Summary

When selecting a backup strategy for BOE, keep the following best practices in mind:

- Cold backups are the recommended approach by BusinessObjects and should always be used when a suitable window of time exists that the BOE system can be unavailable
- If such a window does not exist, there are multiple ways in which a hot backup can be performed
- The RDT should be used to fix inconsistencies that may occur during a hot backup
- FRS and CMS backups should be started as near to simultaneously as possible, but the FRS backup should never start before the CMS backup
- CMS and FRS backups should be atomic to minimize inconsistencies. 3rd party backup tools should be used for this.
- In the event a selective restore of content is needed, use the Lifecycle Manager tool to either reapply a promotion job or rollback to previous versions of a content item.

Copyright