

EMC Centera Data Protection Using the Centera Backup and Recovery Module

A Detailed Review

Abstract

This white paper discusses using the EMC[®] Centera[®] Backup and Recovery Module (CBRM) as an additional option for backing up EMC Centera content apart from the traditional Centera Replication Method. This white paper also includes considerations and best practices while configuring CBRM in all supported environments.

September 2007

Copyright © 2007 EMC Corporation. All rights reserved.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS IS.” EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com

All other trademarks used herein are the property of their respective owners.

Part Number H2953

Table of Contents

Executive summary	4
Introduction	4
Audience	4
Terminology	4
Overview.....	5
Supported configurations	5
Data flow diagrams	7
Backup.....	7
Restore	7
CBRM host requirements.....	8
Licensing	8
Backup types	9
Forcing incremental backups using EMC NetWorker or Symantec NetBackup as the DMA.....	11
Backup and restore types supported by IBM's TSM as the DMA.....	12
Backup types.....	12
Restore types	12
Performing parallel backups	13
Factors affecting performance	15
Centera node configuration.....	15
Number of access nodes available and the load on those nodes.....	15
Number of simultaneous queries running on EMC Centera	16
Backup device/media.....	18
Thread count vs. available RAM on the CBRM host	18
Network configuration	20
Factors that affect streaming data	20
Limitations of CBRM.....	20
Conclusion	21
References	21

Executive summary

EMC® Centera® Backup and Recovery Module (CBRM) is a software module that enables EMC Centera storage system content to be backed up to and restored from tape by using industry-standard backup products, also known as Data Management Applications (DMAs). CBRM leverages a customer's existing backup infrastructure. It supports EMC NetWorker® and is qualified to work with Symantec NetBackup and IBM's Tivoli Storage Manager (TSM). CBRM uses the Network Management Data Protocol (NDMP) to provide an open backup infrastructure for protecting EMC Centera data. The DMA initiates the backup and restore operations, maintains the backup indexes, and originates all tape operations.

For added information protection, CBRM provides an important backup and restore capability for EMC Centera customers. While EMC Centera replication is the method of choice for data protection on Centera, many organizations require the ability to back up EMC Centera content to tape for the following:

- Added operational-recovery protection
- Protection of EMC Centera content for single-site organizations that cannot use replication for backup and restore
- Compliance with customer regulatory policies that may require all data in an organization to be saved to tape

Centera Backup and Recovery Module provides an alternative and complementary solution to EMC Centera replication. CBRM is capable of performing full, differential, date range, or individual C-Clip™ backups. CBRM can recover all, a subset of, or individual C-Clip content for additional flexibility and recovery based on the user priorities. The Direct Access Restore feature helps improve efficiencies by reading only the data that's needed for the restore. CBRM supports automated or manual backup capabilities to simplify on-demand backups and backup management by policy.

Introduction

This white paper provides an overview of the EMC Centera Backup and Recovery Module and its supported configurations. It also describes the licensing options and backup types, plus the factors that affect CBRM performance.

Audience

This white paper is intended for customers, EMC partners, and support personnel involved in planning, configuring, and administering EMC Centera backups and restores using Centera Backup and Recovery Module.

Terminology

Centera Software Developer's Kit (SDK): The library that exports the functions to talk to a Centera storage system. These are generic functions that allow applications to integrate with the Centera storage system.

Data Management Application (DMA): In this context, the NDMP-enabled backup client application. The DMA controls the NDMP session. In NDMP there is a master-slave relationship. The DMA is the session master; the NDMP services are the slaves.

Network Data Management Protocol (NDMP): An open protocol for enterprise wide, network-based data management such as backup and recovery. NDMP is a control protocol used to control the NDMP services participating in the session. NDMP specifies the format and means of transmission of messages and payload data between a DMA and an NDMP server, and between two NDMP servers.

NDMP host: The host computer system that executes the NDMP server application. Data is backed up from the NDMP host to either a local tape drive or to a backup device on a remote NDMP host.

NDMP server: An instance of one or more distinct NDMP services controlled by a single NDMP control connection. Thus a data/tape/SCSI server is an NDMP server providing data, tape, and SCSI services.

NDMP session: The configuration of a DMA and NDMP services to perform a data management operation such as a backup or recovery.

NDMP Tape server: The NDMP server with the tape device attached.

Overview

CBRM has two primary components: a Network Data Management Protocol (NDMP) data server and an EMC Centera data module that reads and writes content to the EMC Centera storage system. CBRM requires an NDMP-enabled DMA to back up and restore data stored on the system. The CBRM NDMP server interacts with the DMAs using NDMP and instructs the data module to perform the required tasks of backup and recovery.

In essence, CBRM acts as an NDMP-interfaced proxy to the EMC Centera storage system to enable customers to back up or restore EMC Centera content. The data module uses the Centera Software Developer's Kit (SDK) to interact with the storage system while performing backup and restore operations. Figure 1 shows the basic architecture of the CBRM solution.

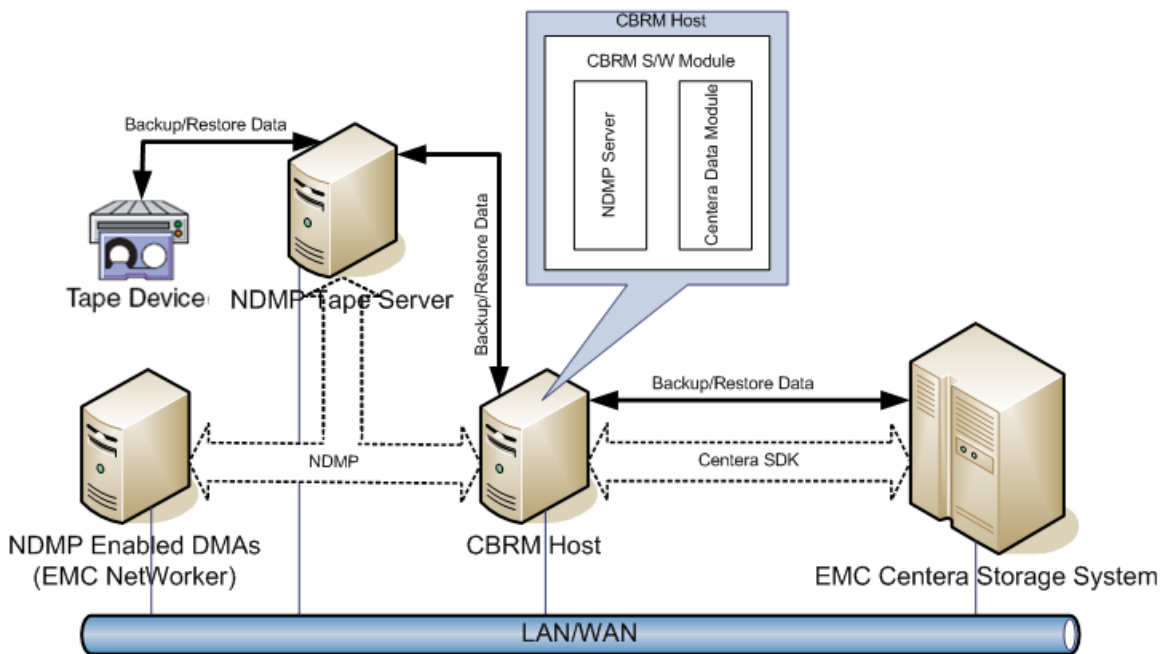


Figure 1. Basic architecture of Centera Backup and Recovery Module solution

Supported configurations

The following configurations are available for using CBRM with EMC NetWorker and Symantec NetBackup:

- **Single host:** CBRM and the DMA software are on the same host. The NDMP Tape server is on a different host (NetWorker only).
- **Dual host:** CBRM and the DMA are on separate hosts. The NDMP Tape server and DMA are on the same host.

- **Three Way:** CBRM, DMA, and the NDMP Tape server are on three separate hosts (required for Symantec NetBackup).

The following applies to the previous configurations:

- NDMP Tape server is a third-party application that provides tape services such as robotic control, tape drive read/write/access control, and others via the NDMP interface.
- CBRM software and NDMP Tape server cannot be on a single host as both of the NDMP servers (CBRM and NDMP Tape server) by default use port number 10000 to communicate with DMAs. CBRM can be configured to use a port number other than 10000 for communicating with DMAs (consult the appropriate configuration requirements that can be supported by the respective DMAs and NDMP Tape servers). However, if a port other than 10000 is used by CBRM, NetWorker will not support this as it expects the NDMP server port to be 10000.

The following configurations are available for using CBRM with IBM Tivoli Storage Manager:

To support IBM Tivoli Storage Manager NDMP requirements, the NDMP Tape server is built in to the CBRM software and is supported only with CBRM deployments using TSM as DMA. TSM is required to control the robotics and the configuration should be as depicted in Figure 2. CBRM's SCSI interface is always opened in TSM environments; however, controlling tape robotics via this interface is not supported or qualified officially.

- **Single host:** CBRM software and DMA are on the same host.
- **Dual host:** CBRM software and DMA are on separate hosts. Configured Jukebox is within the SCSI range of the DMA host.
- **Three Way:** This configuration is not supported.

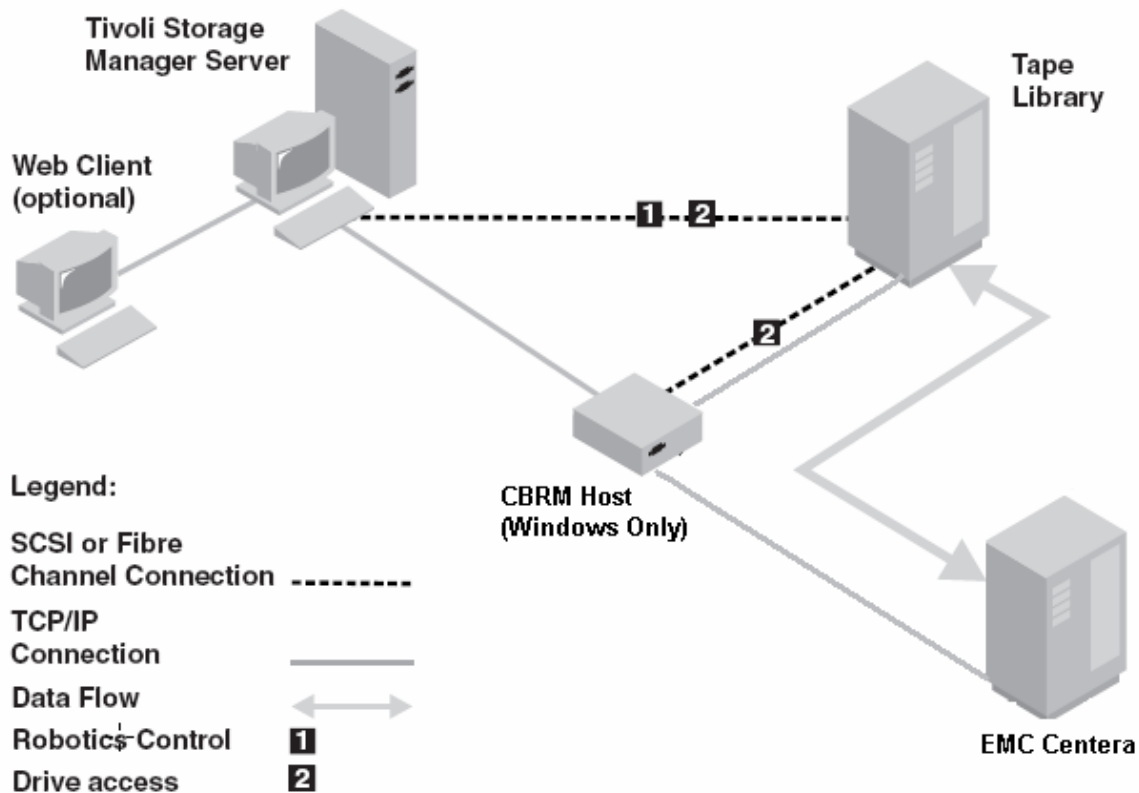


Figure 2. TSM supported configuration with CBRM

Data flow diagrams

The data flows for EMC Centera backup/restores using CBRM resemble traditional NDMP backups/restores except the data from EMC Centera will be read/written over the LAN by the backup module, that is, there must be Ethernet connectivity between the CBRM host and the Centera box. In traditional NDMP backups the backup module is integrated to the NAS OS, which read/writes the data within the NAS box using a LAN free method.

Backup

With backups, the data from EMC Centera will be read by CBRM over the LAN and sent to the NDMP Tape server over the LAN, which in turn writes the backup data to the backup device, preferably to tape.

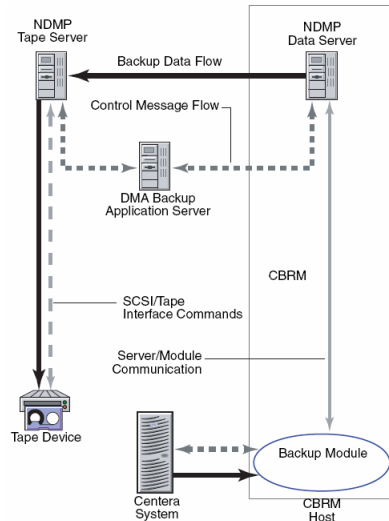


Figure 3. Data flow for backup workflow

Restore

With restores, the data from the backup device (tape) will be read by the NDMP Tape server and sent over the LAN to CBRM, which will then write the data to EMC Centera over the LAN.

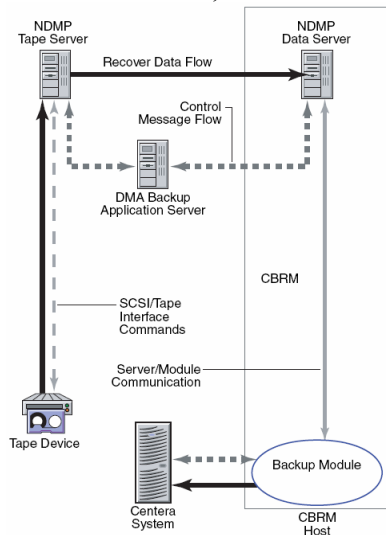


Figure 4. Data flow for recovery workflow

CBRM host requirements

The minimum Microsoft Windows operating system and hardware requirements for CBRM are:

- Windows 2000, or Windows 2003 with the most current service pack installed (Server, Advanced Server, and Datacenter Editions)
- 2.5 GHz Pentium P4 or faster Pentium-compatible processor
- 512 MB RAM
- 20 MB of disk space for installation and additional space for logs and diagnostic files (approximately 1 GB of available hard-disk space). Additional space is needed for backup and recovery logs following installation.

The minimum Solaris operating system and hardware requirements for CBRM are:

- Solaris 8 or Solaris 9
- 512 MB RAM
- 4 GB disk drive

Solaris CBRM installations are not supported with IBM Tivoli Storage Manager.

CBRM is capable of using multiple threads to improve the throughput of backup/restore operations. Each thread uses approximately 200 MB of physical RAM that is available at the time of the backup/restore operation. It is advised to choose a CBRM host machine with sufficient RAM for running multithreaded backups/restores.

Licensing

CBRM licensing is based on the EMC Centera configuration. The basic CBRM license entitles backup/recovery of a four-node Centera configuration. The customer will be provided with a 20-character (including two hyphens) license key for each four-node configuration. If a customer has a configuration with more than four nodes, they must purchase multiple four-node license keys to match their Centera node configuration.

Prior to CBRM version 2.1, licensing was based on a basic eight-node configuration. Additional licenses keys were provided in increments of eight nodes. A separate, unique license key was issued for each of the Centera configurations shown in Table 1.

Table 1. Licensing in CBRM versions 2.1 and earlier

LIC-001013	Backup and Recovery Module (8 nodes or less)
LIC-001014	Centera Backup and Recovery Module (9 to 16 nodes)
LIC-001015	Centera Backup and Recovery Module (17 to 24 nodes)
LIC-001016	Centera Backup and Recovery Module (25 to 32 nodes)

From CBRM 2.1 and onward, the license models are as shown in Table 2. Each of these licenses is for a basic four-node Centera configuration. Multiple four-node licenses will be issued to match a customer configuration.

Table 2. Licensing in CBRM version 2.1 and later

456-100-136	CBRM license for NetWorker - 4 Centera nodes
456-100-137	CBRM license for Symantec NBU - 4 Centera nodes
456-100-138	CBRM license for IBM TSM - 4 Centera nodes

The license keys will be put in to a license file “cbm.lic” under the “etc” directory of CBRM and will be read and evaluated for all backup operations.

Backup types

CBRM supports the following backup types:

- Data range backups. These backups are initiated using the BEGIN_DATE and END_DATE variables in the CBRM resource file.
- Level backups, which are split into the following three categories:
 - Level 0 (full) to Level 9
 - Incremental*
 - Differential* (only supported with IBM's TSM as DMA)

* Incremental and differential backups must be preceded by a full/level 0 backup.

Table 3 distinguishes among the backup types. Table 4 shows the BEGIN_DATE and END_DATE selection criteria.

Table 3. CBRM backup types

Backup type	Function
Date Range	Selective backups by using BEGIN_DATE and END_DATE variables in the CBRM resource file.
Full (Level 0)	Backs up all C-Clips stored on EMC Centera from beginning (installation) until either 1) the end date (if specified) or 2) to the present Centera cluster time.
Level (1 – 9)	<p>Backs up files that have been added to EMC Centera since the last backup with a lower-numbered backup level. For example:</p> <ul style="list-style-type: none"> • A level 1 backup backs up all C-Clips that have added to EMC Centera since the most recent full backup (considered a level zero). • A level 3 backup backs up all C-Clips that have added to Centera since the most recent backup of level 2, level 1, or full backup. For example, if the most recent backup was at level full, then a level 3 backup backs up all the C-Clips that have added to EMC Centera since the full backup. However, if the most recent backup was at level 2, then a level 3 backup backs up only those C-Clips that have added to Centera since the level 2 backup. • A level 9 backs up all C-Clips that have added to EMC Centera since the most recent backup of any level except level 9.
Incremental	Backs up C-Clips that have added to EMC Centera since the last backup, regardless of level.
Differential (IBM TSM only)	Backs up C-Clips that have added to EMC Centera since the last full backup.

Table 4. BEGIN_DATE, END_DATE selection criteria

Level	BEGIN_DATE Specified	END_DATE Specified	Resulting Backup		Comments
			Level	Date Range	
Level 0 (full backup)	Yes	Yes		✓	Ignore level (Default setting for time range backups)
	No	Yes	✓ *		BEGIN_DATE will set to epoch (start of time).
	Yes	No		✓	END_DATE will set to current time.
	No	No	✓		BEGIN_DATE will set to epoch; END_DATE will set to current time.
Level (1 – 9)	Yes	Yes		✓	Ignore level, proceed with time range backup.
	No	Yes	✓ *		BEGIN_DATE will set to most recent lower level backup date as explained above.
	Yes	No		✓	Set END_DATE to present time.
	No	No	✓		Set END_DATE to present time. BEGIN_DATE will set to most recent lower level backup date as explained above.
Incremental	Yes	Yes		✓	Ignore level, proceed with time range backup.
	No	Yes	✓ *		BEGIN_DATE will set to most recent backup date (irrespective of level) as explained above.
	Yes	No		✓	END_DATE will set to current time.
	No	No	✓		Set END_DATE to present time. BEGIN_DATE will set to most recent backup date (irrespective of level) as explained above.
Differential	NA	Yes	✓		BEGIN_DATE if specified will be ignored; END_DATE will set to the specified value.
	NA	No	✓		END_DATE will set to the present time.

* This is a deviation from the actual level backup's definition. The END_DATE is user-defined so current time will be overridden by the given value.




Figure 5. Interrelationship between backup levels

In Figure 5, for example, if the user performs a full backup for the first time, all data on the EMC Centera will get backed as explained in Table 3. For the subsequent level 1, 3, 5, and 2 backups the data that has been added to EMC Centera from the previous lower level will get backed up. For the fourth backup (level 2) in the above sequence (level 1 > 3 > 5 > 2), note that the backup data constitutes data that has been added since the previous level 1 backup. For the subsequent incremental backups, the data that has been added from the previous backup, regardless of the backup level, will get backed up. For the last full backup in Figure 5 all data on EMC Centera will get backed up, including data of the previous full, level, and incremental backups.

Forcing incremental backups using EMC NetWorker or Symantec NetBackup as the DMA

By specifying FORCE_INCR=y in the CBRM resource file, all the level backups (1 through 9) behave as incremental backups and ignore the actual level definition. If FORCE_INCR=y (default value “n”) is specified, irrespective of the level, all the C-Clips that have been added to EMC Centera will be backed up.

Note: Performing NDMP incremental backups are not supported by IBM Tivoli Storage Manager.

Backup and restore types supported by IBM's TSM as the DMA

Backup types

- Full backup image, which includes all C-Clips within the EMC Centera storage system
- Differential backup image, which includes all C-Clips that have changed since the most recent full backup
 - If differential backup is specified and full backup not found, a full backup will be performed.
 - During differential backup, the server links new differential image with correct full image in the TSM database.

Restore types

- Full file-system image
- Full file-system image plus one differential image
 - Restore of differential image is automatically preceded by restore of corresponding full image.
- Selected files (selected directories *not* supported)
- Restore to most recent backup or point in time

Note: Date Range backups are not supported with TSM as the user-specified BEGIN_DATE will be ignored and the server (TSM) specified value will be considered as the beginning point for the current backup. However, END_DATE can be specified to limit the backup data range.

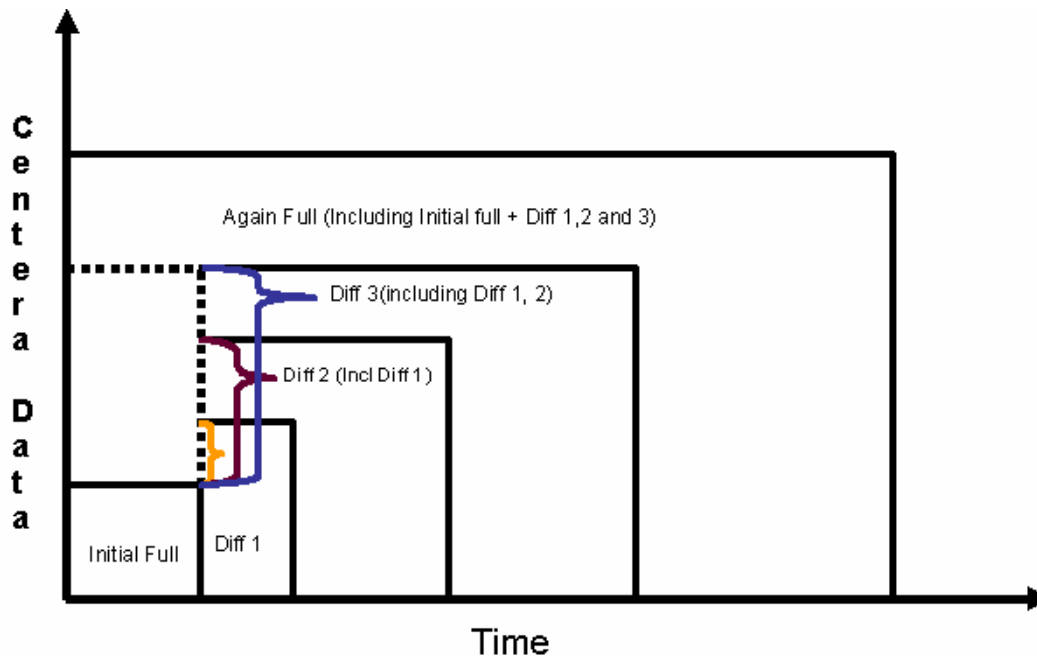


Figure 6. Interrelationship between full and differential backups in IBM TSM environments

In Figure 6, for example, if the user performs a full backup for the first time, all data on the EMC Centera will get backed up. For the subsequent differential backups (Diff1, Diff2, and Diff3), the data that has been added to EMC Centera from the previous full backup will get backed up. Note that for differential backups Diff2 and Diff3, the data that is part of the previous differential backups will also get backed up.

Performing parallel backups

Performing parallel backups might provide some performance gain. To reduce the backup window especially when performing a complete backup that requires backing up all the data stored on EMC Centera, split the data set across the backup sessions that are running in parallel.

CBRM supports parallel backups with the following prerequisites:

- Operate each backup session on a separate resource file.
- Configure the thread count in each session so they are balanced across the number of parallel backup sessions (versus the total RAM available on the CBRM host).
- Configure the backup data set using the `BEGIN_DATE` and `END_DATE` variables.

Parallel backups are tested with EMC NetWorker internally but they are not officially qualified and parallel backups are not supported with IBM's TSM and Symantec NetBackup.

Access profiles are used by CBRM to authenticate to an EMC Centera cluster, and by clusters to authenticate to another cluster for replication or restore connections. EMC Centera system administrators can create access profiles that consist of a profile name, a secret (password), and a set of capabilities and roles. The access profile for CBRM can be exported as a Pool Entry Authorization (PEA) file and can be provided to CBRM. CBRM resource files are configuration files that must be created and associated with each backup/restore operation.

More information on how to configure CBRM resource files can be found in the *EMC Centera Backup and Recovery Module Version 2.1 Installation and Administration Guide*. More information on creating PEA files can be found in the EMC Centera online help.

Parallel backups on the same CBRM host

Parallel backup sessions can be run on the same CBRM host by configuring separate CBRM resource files (with different names to distinguish) for each session. The same PEA file can be shared among the backup sessions. Care has to be taken while configuring the `THREAD_COUNT` variable in each CBRM resource file. Configure the `THREAD_COUNT` variable in each CBRM resource file to balance the total number of backup/restore threads that can be spawned with the total physical memory available on the CBRM host.

Refer to the section “Thread count vs. available RAM on the CBRM host” for more information on setting the `THREAD_COUNT` variable in the CBRM resource file.

Figure 7 shows two parallel sessions running on the same CBRM host backing up a Centera pool named Pool For Mail using the same PEA file to access the Centera cluster. Note that there are two separate CBRM profiles (resource files) used for each backup session. Individual values for `BEGIN_DATE` and `END_DATE` can be set in each resource file to control the backup data range in each session.

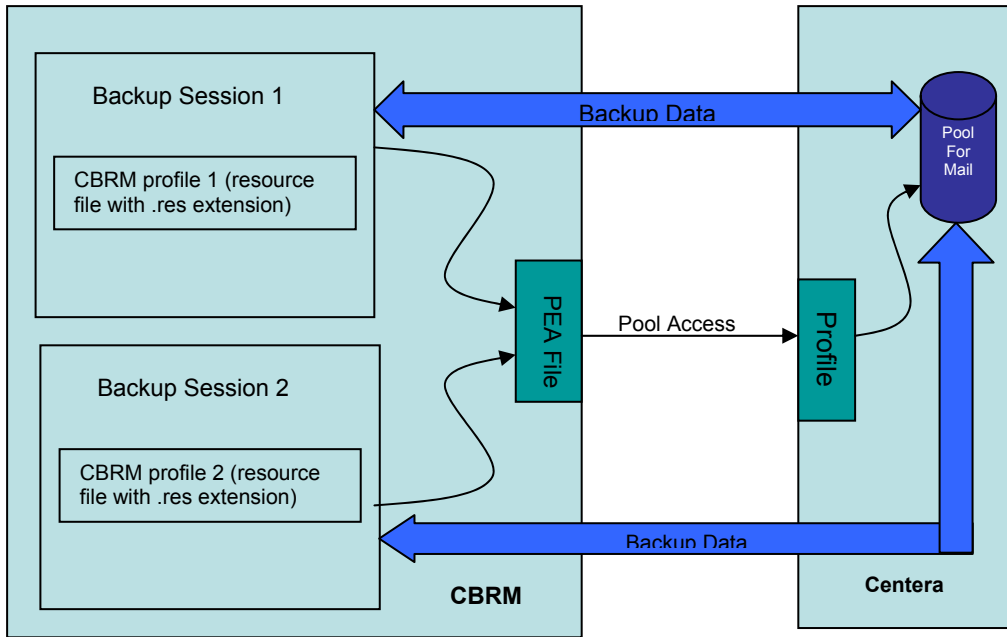


Figure 7. Parallel backups on the same CBRM host

Parallel backups using different CBRM hosts

Parallel backup sessions can be run on different CBRM hosts by configuring separate CBRM resource files (but these resource files can have the same name) for each session. A separate PEA file has to be provided for each backup session that can access the same pool.

Figure 8 shows two parallel sessions running on different CBRM hosts backing up a Centera pool named Pool For Mail using separate PEA files to access the Centera cluster. Individual values for BEGIN_DATE and END_DATE can be set in each resource file to control the backup data range in each session. All constraints that impose on the THREAD_COUNT variable can also be applied based on the physical memory available on each CBRM host.

Refer to the section “Thread count vs. available RAM on the CBRM host” for more information on setting the THREAD_COUNT variable in the CBRM resource file.

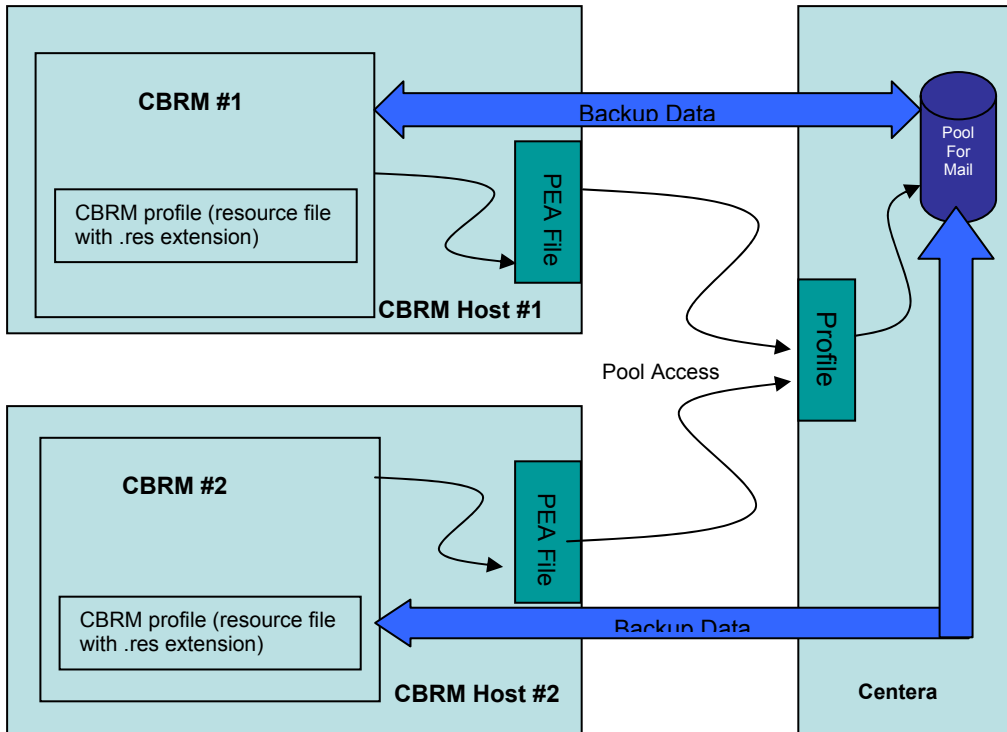


Figure 8. Parallel backups using two different CBRM hosts

Factors affecting performance

CBRM backup/restore performance is affected by combinations of the following general conditions:

- EMC Centera node configuration
- Backup device and media
- Thread count
- C-Clip size
- Network configuration
- EMC Centera API streaming options

Every EMC Centera environment is unique and therefore “best case” performance will be different. This section is intended to educate so that optimum CBRM performance and operation can be achieved through a reasonable amount of planning. It is not intended to set expectations for performance because each CBRM experience will be unique.

Centera node configuration

Number of access nodes available and the load on those nodes

The maximum number of threads per node with the access role (“access node”) on an EMC Centera cluster is 20. A cluster with four access nodes thus supports 80 threads. In this scenario, a single backup session can use up to 80 threads, and two or more backup sessions can use 80 threads combined.

The number of application servers (including CBRM) writing/reading to a single EMC Centera cluster influences the overall number of files stored/retrieved and bandwidth utilization of the cluster. In a test using one, two, four, and six application servers, with each application using 10 threads to write, the impact

of multiple application servers was measured on an EMC Centera configured with four access nodes and 28 storage nodes. Performance scales significantly as additional pairs of application servers are added.

A good network stack is also important. Dedicated application servers that operate with two, four, or eight processors also deliver a better performance.

Number of simultaneous queries running on EMC Centera

Centera queries can take significant time to complete. For example, a cluster with 3 million C-Clips with a typical throughput of 100 results will take 3 hours to complete. It is therefore not uncommon for a query to be interrupted before it completes. For example, if a full backup is run on EMC Centera using CBRM, based on the clip sizes that are being backed up, CBRM itself might interrupt the query. Or, the network connection between the CBRM server and the access node might be lost, in which case the cluster recognizes the problem and aborts any ongoing queries. CBRM can resume an interrupted query by launching a new query specifying the start time as the value of the timestamp returned by the last valid query result.

In the case of resuming queries, there will be a significant performance impact as the older queries on EMC Centera need to be purged, which releases all of the allocated resources. A fresh query needs to be started.

It is recommended that while performing full backups from an EMC Centera make use of the query end date to limit the number of query results. This helps the EMC Centera system to avoid query retries. Full queries on EMC Centera are resource-consuming and take considerable time to complete.

For example, if a full backup were to be performed on EMC Centera that has 1 million C-Clips to back up, and the average C-Clip size is in the range of 1 GB - 5 GB, backing up C-Clips with this average size would take a considerable amount of time based on the thread count and the network connectivity between the CBRM server, Centera, and backup device server. CentraStar[®] will clear off the resources allocated for a query process if there are no query-related activities (reading query results) happening on the client (CBRM). To deal with this situation, make use of the end date to limit the number of query results that fall in to the full backup criteria and divide up the query in chunks that are manageable by EMC Centera. In this case divide the queries in increments of, say, 10,000 each and perform a full backup followed by multiple incremental backups.

The following steps will give insight on how to perform a complete backup of an EMC Centera by splitting the backup data based on the number of C-Clips and their average size.

Note: Don't confuse the term "complete" backup with the traditional "full" backup. "Complete" backup here implies backing up the whole data on EMC Centera by performing multiple backups operations, but the traditional "full" backup implies backing up the whole data on EMC Centera in one backup operation.

1. Before performing the backup:
 - a. Estimate the total number of C-Clips on the source (to be backed up) Centera.
 - b. Estimate the average size of the C-Clip that will be stored on to EMC Centera. This can be estimated based on what kind of data that is being stored on EMC Centera and can be specific to an application that is being used to write data. Examples include email (range from 10 KB to 500 KB) and X-ray images (10 MB to 1 GB).
 - c. Estimate the size of the data that will be archived to EMC Centera over a specified time interval (weekly, monthly, yearly, and so on).
2. Split the backup data set based on the number of C-Clips/approximate size of the backup data.

Example: If the C-Clips written to EMC Centera in a week is around 10,000 and the approximated data size (assuming an average C-Clip size as 500 KB) would be 10,000 * 500 KB=4.7 GB, the approximate data of a month's C-Clips will be 4.7 * 4=18.8 GB. The value 18.8 GB for a backup data size is more reasonable to fit into the backup window of, say, 1 hour backing up at an approximate rate of 5.5 MB/s and the query range has manageable results (4*10,000), which can be handled easily by CBRM as well as EMC Centera when compared to that of a full query of all C-Clips. In this example, backing up EMC Centera data in increments of monthly data is more feasible.

3. After estimating the backup data size and number of C-Clips that fall in to a backup (in the previous example data stored on EMC Centera over a month that has around 4*10,000 C-Clips with an approximate total size of 18.8 GB), perform a full backup by specifying END_DATE to limit the backup data set to fall in to the range of all C-Clips that are stored over a month. Once the initial full backup with specified END_DATE is done, perform incremental backups for every month's data by specifying the appropriate END_DATE. Repeat this until you have backed up all data on Centera.

Example: Backing up the complete data that is stored on EMC Centera in increments of monthly data (assuming the Centera in use from 10/01/2007).

Table 5. Sample values for BEGIN_DATE and END-DATE variables for the previous example

Backup type	BEGIN_DATE	END_DATE	Comments
Full/Level 0	01/01/2007,00:00:00	01/30/2007:23:59:59	In case of full/level 0 backups, BEGIN_DATE value can be skipped as CBRM will query EMC Centera from the first C-Clip stored on Centera till the END_DATE
1 st Incremental	NA	02/30/2007:23:59:59	For incremental backups BEGIN_DATE will be decided by CBRM based on the timestamp of the last C-Clip that's backed up as part of the previous backup.
2 nd Incremental	NA	03/30/2007:23:59:59	
3 rd Incremental	NA	04/30/2007:23:59:59	
n th Incremental	NA	NA	For the last incremental backup, END_DATE can be skipped. This will back up all the clips that were stored till the date (that data on which the last incremental backup is performed) from the previous incremental backup. At this point all the data stored on EMC Centera is backed up.

Note: The format for BEGIN_DATE and END_DATE is "mm/dd/yyyy, hh:mm:ss"

Queries with the same criteria always return the same result set, even if content was added to the cluster between successive queries. The one exception is if the end time is not specified in the CBRM resource file, the value for the end date will be set to the "now" value (-1). In this case, content added to the cluster is included in the query results. If some content is not returned due to unavailable nodes, the query result is reported as incomplete and returns a result code FP_QUERY_RESULT_CODE_INCOMPLETE. In case of query failures, CBRM retries five times before bailing out from a backup operation.

The default number of simultaneous queries that can run on the cluster is 10. However, your EMC Centera may be configured differently. The number of simultaneous queries is configurable but is not a customer setting.

In case of any size mismatch occurring for any C-Clip a warning message will be displayed and the C-Clip will be re-read and will be backed up. This retry will happen for a maximum of five times. If the retry limit is reached an error message will be displayed for that C-Clip.

Note: Query operations examine all C-Clips on a cluster, including C-Clips written by other applications.

Backup device/media

The throughput of EMC Centera backup/restores using CBRM also depends on the type of storage device configured. Depending on the media type and the storage device read/write bandwidth, the backup/restore performance is affected. CBRM supports Direct Access Restores (DAR) of individual C-Clips or multiple C-Clips, which reduces the restore times by reading only data required for restores. This requires a faster seek capability of the storage device to locate the required data within the backup image.

Table 6. Sample performance rates of Centera full backups using LTO-3 as backup media

Backup setup	Backup device	Thread count	Total C-Clips	Total backup size	Avg C-Clip size	RAM on CBRM host	Observed throughput (MB/s)
Single host (Solaris)	LTO – 3	6	385998	389.75 GB	1.03 MB	2 GB	10.17
Single host (Windows)	LTO – 3	6	385998	410 GB	1.03 MB	2 GB	10.00
Dual host	LTO -3	NA	NA	NA	NA	NA	NA
Dual host	B2D	NA	NA	NA	NA	NA	NA

Thread count vs. available RAM on the CBRM host

To optimize EMC Centera backup/restore performance, EMC recommends the use of multithreaded backups/restores to increase the maximum transfer rate. The EMC Centera architecture is highly parallel and supports multiple parallel activities. A single thread backup/restore cannot take advantage of the multiple nodes with combined roles (access and storage). However, it is important to optimize concurrent activity on the EMC Centera system because too much activity can lead to tasks competing for resources. For example, tasks requiring disk resources on nodes with the storage role can create bottlenecks that cause performance degradation.

Because threads are distributed evenly over available nodes, the number of nodes influences the number of threads that can be supported. Typically, a 32-node cluster performs better than a 16-node cluster. Determining how many threads to use depends on the number of access nodes. There is a capacity of 20 threads per node with the access role.

Generally, if you are running your CBRM backups against like Centera configurations, you can increase the throughput of small files by increasing the number of threads to optimize read and write performance. For example, a 50 KB file can increase the number of files(objects)/second if written with 20 concurrent threads instead of with a single thread. Likewise, if you have 1,000 objects that need to be transferred to the EMC Centera, it is more efficient to create several threads and perform write transactions simultaneously than it is to perform all 1,000 writes one at a time in the same thread.

In evaluating performance EMC Centera’s architecture is a factor as well as the CBRM server configuration. The memory on the CBRM server is important, especially if backup/restores are running in

multithreaded mode. Each backup/restore thread requires approximately 200 MB of physical RAM at the time of backup/restore operation.

If there is inadequate RAM available to spawn the threads specified via thread count, CBRM reduces the thread count to a value that guarantees successful creation of threads.

The algorithm is as follows:

```
Input: USER_SPECIFIED_THREAD_COUNT;
Output: EFFECTIVE_THREAD_COUNT

Calculate: TOTAL_FREE_RAM_REQUIRED = USER_SPECIFIED_THREAD_COUNT*200;
Calculate: TOTAL_FREE_RAM_AVAILABLE;

If (TOTAL_FREE_RAM_REQUIRED > TOTAL_FREE_RAM_AVAILABLE)
    /* Leave 200 MB of RAM for other processes to consume, and then calculate effective
       thread count based on the available free RAM. */
    Calculate: EFFECTIVE_THREAD_COUNT =
        (TOTAL_FREE_RAM_AVAILABLE-200)/200
    If (EFFECTIVE_THREAD_COUNT <= 0)
        Warning: Too low memory to spawn specified
                 USER_SPECIFIED_THREAD_COUNT
                 threads, setting thread count to 1 (default value);
        /* In this case backup/restores might fail with errors as there is not enough RAM
           available to start even a single thread. Upgrade the CBRM host RAM as
           required or free the RAM by closing some of the opened applications that are
           running on the CBRM host. */
        Return (1);

    If (EFFECTIVE_THREAD_COUNT < USER_SPECIFIED_THREAD_COUNT)
        Warning: Insufficient/Low memory to spawn specified
                 USER_SPECIFIED_THREAD_COUNT threads, setting thread count
                 to EFFECTIVE_THREAD_COUNT;
        /* In this case user specified thread count reduced to thread that can be spawned
           based on the available free RAM. */
        Return (EFFECTIVE_THREAD_COUNT);

/* If enough free RAM is available to accommodate USER_SPECIFIED_THREAD_COUNT
threads */

Return (USER_SPECIFIED_THREAD_COUNT);
```

Example:

Total Physical RAM available on CBRM host: 2048 MB (2 GB)

Total Free Physical RAM available at the time of CBRM backup/restore operation: 1536 MB (1.5 GB)

Thread Count specified by user: 10 (requires approximately $10*200 = 2000$ MB of free RAM)

The effective Threads spawned: $(1536 - 200)/200 = 6$

To run a successful backup or restore operation with a thread count set to 4 requires a minimum 800 MB of free RAM available on the CBRM server.

Note: EMC recommends that you keep the total number of threads < 20 times the number of nodes with the access role when writing small files (< 500 KB), or < 10 times the number of nodes with the access role when writing large files (> 500 KB).

Network configuration

A good network stack is also important to improve the backup/restore throughput. The number of application servers including CBRM writing to a single EMC Centera cluster influences the overall number of files stored and bandwidth utilization of the cluster. For example, using one, two, four, or six application servers, with each application using 10 threads to write, the impact of multiple application servers writing/reading to the same EMC Centera cluster degrades performance significantly as additional pairs of application servers are added.

Network connectivity between the EMC Centera system, CBRM server, and the server to which the storage device is attached will also affect the backup/restore throughput. High speed Ethernet connectivity with fewer network boundaries will deliver better performance.

Factors that affect streaming data

CBRM makes use of the generic streaming methods provided as part of the EMC Centera Application Programming Interface. The following factors affect the handling and performance of streaming data, depending on your system environment and EMC Centera configurations.

- ID calculation— CBRM uses `FP_OPTION_CLIENT_CALCID_STREAMING` while writing data to Centera and this value cannot be changed by the user. CBRM using Centera SDK calculates the content address while sending the data to the Centera. The client sends the data even if the cluster already contains the data. This option is equivalent to `FP_OPTION_DEFAULT_OPTIONS` and is optimal when writing large files (10 MB or larger), when using many threads, or when identical data is unlikely to exist on the cluster.
- To use the proper storage protection scheme — CPM or CPP — the EMC Centera server needs to know the size of the data before storing it. The SDK will prefetch the data and store it temporarily to memory. The `FP_OPTION_PREFETCH_SIZE` option is used if the total size of a stream is unknown; the Centera SDK reads the blob bytes into memory up to the prefetch size to determine what the stream length might be. The SDK uses this option value to know how best to proceed with other decision processes. The default size is 32 KB. The maximum size is 1 MB. This value can be tuned by configuring the CBRM resource file variable `PREFETCH_SIZE` with the appropriate value.

More information on how to configure `FP_OPTION_PREFETCH_SIZE` can be found in the *EMC Centera Version 3.1 API Reference Guide* and *EMC Centera Version 3.1 Programmer's Guide*.

Limitations of CBRM

- CBRM does backup of EMC Centera by performing raw read/writes of C-Clips and all associated blobs, but CBRM does not back up mutable metadata (MMD), which will get created by Litigation Hold and Event Based Retention.
- CBRM does not back up EMC Centera configuration information. This means that CBRM and CBRM-created backups cannot be used to initialize and recover an EMC Centera system.
- CBRM associates backups with a specific EMC Centera profile. This means that CBRM cannot restore C-Clips to any other EMC Centera system, as in a disaster recovery scenario.
- Generally, the performance profile of a CAS-to-CBRM-to-tape backup and recovery system is much different from traditional backup systems performing from network file devices or SAN block devices. Customers are cautioned to consider the impact on performance created by the volume of data stored to tape and the bandwidth of the CBRM solution before relying solely on CBRM for complete restoration of an EMC Centera system.
- CBRM is not a compliance solution because tape is not WORM-compliant media.

Conclusion

EMC Centera Backup and Recovery Module (CBRM) leverages a customer's existing backup infrastructure to back up EMC Centera content to tape for added operational recovery data protection and to comply with customer regulatory policies. Centera Backup and Recovery Module provide an alternative and/or complementary solution to EMC Centera replication.

CBRM supports EMC NetWorker and is qualified to work with Symantec NetBackup and IBM's Tivoli Storage Manager (TSM). CBRM uses the Network Management Data Protocol (NDMP) to provide an open backup infrastructure for protecting EMC Centera data.

CBRM offers a range of backup options including full, differential, date range, or individual C-Clip backups. CBRM can recover all, a subset of, or individual C-Clip content for additional flexibility. The Direct Access Restore feature identifies only the data that's needed for the restore. CBRM supports automated or manual backup capabilities to simplify on-demand backups and backup management by policy.

CBRM's functionality and supported Data Management Applications make it an ideal choice for added data protection of EMC Centera content.

References

- EMC NetWorker Version 7.x Administrator's and Installation Guides
- *EMC Centera Backup and Recovery Module Version 2.1 Installation and Administration Guide*
- *EMC Centera Version 3.1 API Reference Guide*
- *EMC Centera Version 3.1 Programmer's Guide*
- IBM Tivoli Storage Manager Administrator's and Installation Guides
- NDMP website (<http://www.ndmp.org>)