
CaminoSoft Centera Integration Managed Server HSM

Centera Proven Edition

March 18, 2009

Rev 1.1

Abstract

This document provides an overview of CaminoSoft's Managed Server HSM, Centera Proven Edition integration with EMC Centera

Table of Contents

CaminoSoft Company Overview	3
Introduction and Overview of Integration	3
Revision Information	5
Application Revision and Centera SDK Revision Support	5
Hardware Platform support.....	7
ISV Support Contacts	Error! Bookmark not defined.
CentraStar Support Policy.....	7
Data Model	8
Architectural Overview - Diagram	8
Centera Integration Module Licensing.....	10
Content Type	10
Containerisation.....	10
Single Instancing.....	10
Embedded Blobs.....	11
CDF Structure	11
Storage Device Mode	11
Centera Specific Configuration Details	12
Content Address Management – Generating a list of clips	12
Threading Model.....	13
Streams Usage.....	13
Compliance Support.....	14
Retention Period and/or Retention Classes	14
Advanced Retention Management.....	14
Pool Handling	14
Configuration	14
Access Profiles.....	15
Global Options	15
Archiving Capabilities.....	15
Content Replication.....	17
Failover	17
High Availability (HA).....	17
Deployment Scenarios.....	17
Disaster Recovery.....	17
Miscellaneous API Usage	18
SDK Logging Support	18
Application Registration.....	18
SDK Usage	18
Query API	18
Sample CDF.....	20
SDK Log Sample	20
Error Handling	21
Performance Expectations.....	22
Release Schedule	22

CaminoSoft Company Overview

CaminoSoft Corp. (OTC.BB: CMSF) is a developer and manufacturer of software solutions that address the storage management of vast quantities of data generated in a wide range of businesses and applications. The company's Information Lifecycle Management tiered storage solutions for Microsoft Windows 2000/2003/2008, Novell NetWare, and NetApp Filer environments include comprehensive administrative policies that allow organizations to reclaim storage resources, dramatically reduce backup and recovery time, control file retention, and achieve regulatory compliance. CaminoSoft has established alliances with industry-leading technology partners, and the company markets its Storage Management and High Availability solutions worldwide through commercial distributors, value-added resellers, and systems integrators.

CaminoSoft Corp.
600 Hampshire Road, Suite 105
Westlake Village, CA 91361
Telephone: +1-805-370-3100
Toll-Free: 800-889-8248 (USA)
www.caminosoft.com

Introduction and Overview of Integration

Application introduction and overview

CaminoSoft Managed Server HSM

CaminoSoft's Managed Server HSM™, Centera Proven Edition (MSHSM) provides a simple, cost-effective solution for handling the complexities of data storage management. By applying consistent policies, administrators control the location of files as they move through their lifecycle, beginning with high-performance, costly storage (such as SAN arrays), then to less-expensive disk arrays and near-line devices, and eventually to off-line archives. Files archived in such a way tend to become more "fixed" in nature as they move along the storage hierarchy. Centera's combined features, which include optimized storage for fixed content, real-time data access, high-availability, scalability and lower overall cost of storage, make it

very attractive as an archive solution for controlling the otherwise explosive growth requirements of an organization's primary SAN infrastructure.

Like CaminoSoft's standard Managed Server HSM solutions, the Centera Proven Edition enables the administrator to configure policies and usage of the "virtual" Central Storage Pool and archive files that are infrequently accessed to the pool. Using additional policies, Centera Proven Edition extends this functionality by providing the ability to archive files further along a storage hierarchy to an EMC Centera. The policy-driven engine provides for "lights-out" operation and reduces the time and space to make backups of the files that are actually changing day-to-day, while providing complete, transparent access by users to all of their data.

Markets

- Banking
- Financial Securities
- Accounting
- Federal-State-Local Government
- Education
- Healthcare

ISV Application Centera SDK Integration overview

Managed Server

- *Include details on how content flows from your application to the Centera*

When MSHSM initiates, connections are established with the target storage devices as defined by administrators in the CaminoSoft archiving profiles. In the case of Centera, the provided connection information (IP addresses, PEA file, etc.) is used to establish Centera communication. Based on scheduled scans of the managed server's file system, candidate files are identified subject to rules and policies defined by the administrator. For every candidate file that meets the criteria, a new Centera file object is created (clip and BLOB). Along with the candidate file's data (blob), certain of the file's metadata (source path, size, date stamps, etc.) are stored in the clip to be used for identification purposes and to facilitate disaster recovery (reconstruction of the file stubs that link to the Centera objects).

- *Include a diagram depicting this*

CaminoSoft's Managed Server HSM, Centera Proven Edition, running on a Windows 2000/2003 server, can manage files directly from its own resources to the Centera. It can also act as a gateway to Centera providing access from multiple NetWare, Windows and/or NetApp "managed" servers.

When a file is archived to Centera, Managed Server HSM leaves behind a "stub" which is a fraction of the size of the original file. The "stub" file contains metadata about the original file, as well as the content address returned by the Centera.

Managed Server HSM can also manage MAC files stored by MAC clients on Windows 2000/2003 servers. The application only archives the primary data stream for MAC files, leaving behind the resource fork intact on the primary storage.



- Explain if you use Centera Pool Entry Authorization (PEA) file
Yes
- Give an example of your application connect string when trying to connect to Centera
Specify the path of the .pea file on the managed server, followed by the IP Address while specifying HSM Pool (target).
If IP Address is 192.168.1.1, and the .pea file is located at c:\profiles\strict.pea, the IP Address should be specified as:
192.168.1.1?c:\profiles\strict.pea

Revision Information

Application Revision and Centera SDK Revision Support

Please detail here the current revision of your Application and the Centera SDK supported.

If this is an update to this document please retain the revision history. (Add additional rows if required).

Application Name	Generally Available (GA) Date (if available)	Application Revision	Centera SDK Revision
CaminoSoft Managed Server HSM	December 8, 2008	5.3.3306	3.2.631
CaminoSoft Managed Server HSM		5.2.3254	3.1.477 (EOSL)

CaminoSoft Managed Server HSM	March 3, 2006	5.2.3243	3.1.477 (EOSL)
CaminoSoft Managed Server HSM		5.2.3224	3.1.477 (EOSL)
CaminoSoft Managed Server HSM	January 31, 2006	5.2.3221	3.1.473 (EOSL)
CaminoSoft Managed Server HSM	August 15, 2005	5.2.3205	2.3 SP2 (EOSL)
CaminoSoft Managed Server HSM		5.1	2.3 SP2 (EOSL)

End of Life (EOL): Indicates the date the product was no longer available for sale.

End of Service Life (EOSL): Indicates the date after which standard EMC technical support is no longer available.

Operating Systems Supported

Please list all operating systems supported and revisions (add additional rows if required).

CaminoSoft's Managed Server HSM, Centera Proven Edition runs on Windows 2000/2003, and 2008 servers. It can manage files directly from the Windows server(s) upon which is installed, or act as a gateway to Centera providing access from multiple NetWare, Windows and NetApp "managed" servers.

Operating System <i>(specify if it is 32bit, 64bit)</i>	Revisions Supported
Windows 2000/2003 servers	All revisions including Windows 2003 64 bit
Windows 2008 servers	All revisions including 64 bit

Hardware Platform support

CaminoSoft's Managed Server HSM, Centera Proven Edition runs on Windows 2000/2003 and 2008 servers, and supports Central Storage Pool comprised of one or more off the shelf hardware configurations of network attached storage (NAS) target devices, general purpose file servers, EMC Centera or combinations of the above.

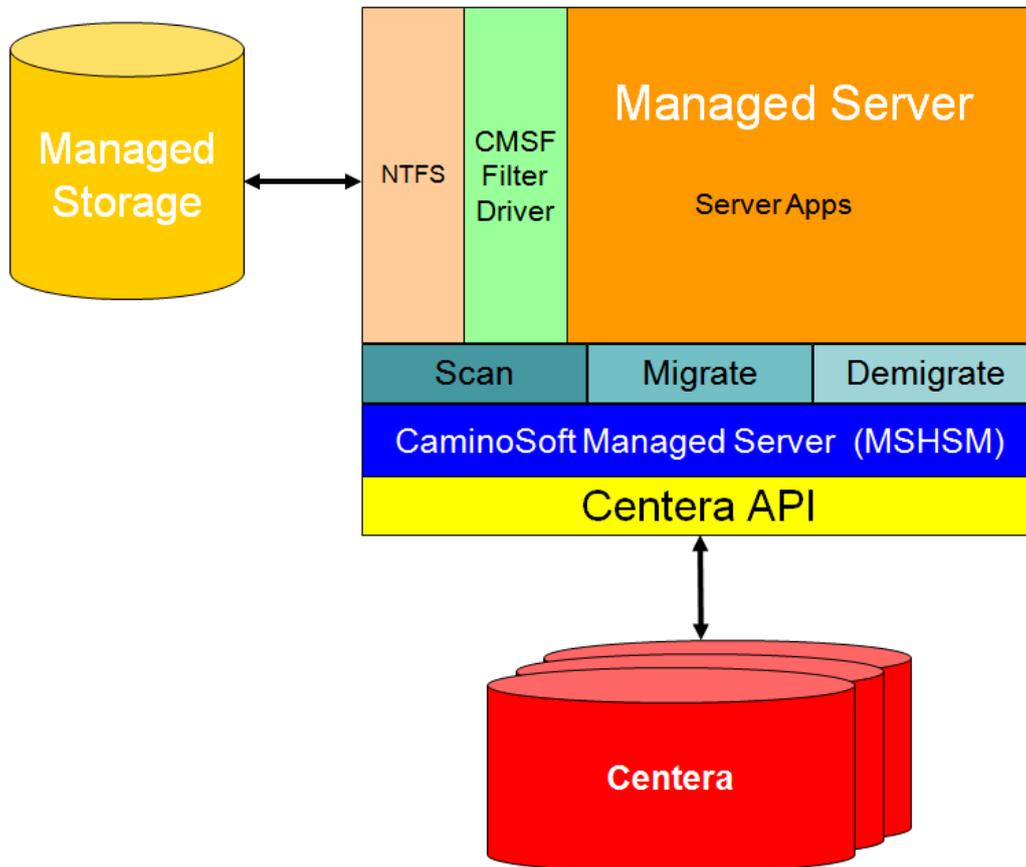
CentraStar Support Policy

CaminoSoft relies on the EMC Centera statement of compatibility between the Centera SDK and CentraStar version.

Data Model

Architectural Overview - Diagram

Use the architecture diagram to explain the dataflow.



Detail how content gets from its source to the Centera via your application.

Please include how it is retrieved and your deletion process.

Archiving

- Based on centrally administered rules, policies and schedules, MSHSM scans the managed file systems.
- Candidate files are screened for conformance to rules and policies.
- For each screened candidate
 - If candidate file was previously archived and satisfies rearchive rules, then MSHSM performs re-archive (convert to stub with data truncation – data already resides on Centera, so no data transfer occurs).
 - If candidate was not previously archived and satisfies archiving rules, then file data is copied to the Centera utilizing the Centera API, the file is truncated and a stub is created containing the Centera clip ID.

Recall

- User or application opens archived file.
- MSHSM filter driver hooks open request and sends request to MSHSM to recall file contents from the Centera while user/application waits.

-
- Utilizing the clip ID, MSHSM utilizes the Centera API to copy file content from the Centera into the file's primary stream.
 - Once data is written to the stream, the file is released to user/application.

Deletion

MSHSM provides Deletion Management - a feature whereby the administrator can specify a number of days to hold the target file (of a deleted file stub) before it will attempt to delete it.

- User or application deletes a archived file (stub).
- MSHSM filter driver hooks delete request and sends request to MSHSM to make an entry into the CaminoSoft database table containing target files to be deleted.
- During the deletion cycle, the database table is examined and the dates of each entry are compared to the policy set by the administrator.
- If the number of days since the deletion exceeds the number of days set by the administrator, an attempt is made to delete the target file from the Centera utilizing the Centera API. Retention settings on the physical Centera supersede deletion attempts by MSHSM. In that case, the deletion is ignored and the associated record in the database table is removed.

Centera Integration Module Licensing

Explain if the Centera integration is a separate add-on module to your main application and how you license it.

The Centera integration is embedded into the standard product. Its functionality remains dormant until a specific license key is entered which activates it.

Content Type

Describe the type of content you write to Centera - e.g. office documents, medical images, scanned images, print queues, etc.

In addition, detail the average object size and typical ingest requirements - e.g., You could be archiving medical images that range from 0.5 MB to 20 MB and require being able to ingest 2000 medical images per hour for a typical customer.

Because MSHSM archives content from a file system (system, user and application folders), virtually any type of file (of any size) has the potential of being written to Centera. The archiving process (copy to Centera) typically functions in off-hours. So, performance is generally not critical. Recalls (copy from Centera), however, occur in real time while users and applications wait. Recalls must be as efficient as possible to avoid application/workstation timeouts and remain as transparent as possible to users.

Containerisation

When the content that an application needs to archive to Centera, or any other archive platform, is small e.g., < 1 MB it can make sense to aggregate multiple small files into one large file before archiving to Centera.

When individual objects are required to be read back the application can use the `FPTag_BlobReadPartial()` method specifying an offset and byte size to read so that only the portion of the container you want needs to be retrieved.

Detail how you containerise small objects (e.g. a tar ball, zip file, just append them, etc.) and where you are storing the offset and size of each individual object (e.g. in your database along with the content address, as meta data in the CDF, or as a side file linked to from the CDF.)

MSHSM does not containerize objects. Each archived file constitutes its own object. Policies are set for minimum file size within the application to avoid archiving small files (customers are discouraged from archiving small files).

Single Instancing

Describe whether your application writes content that will be single instanced, i.e., the Centera will ensure that identical content is only stored once but can be referenced many times.

And/Or

Some applications may implement single instancing on their end so that they only send 1 copy of the content to the Centera. Describe whether your application does implement single instancing.

MSHSM writes primary data stream information directly from the source files. Therefore, any duplicate files will have identical primary data streams which Centera's single instancing algorithms will resolve to a single stored copy of the duplicate content.

Embedded Blobs

If an application is writing many small files and does not or cannot aggregate the small files into a container, the Centera SDK embedded blob functionality can be used to alleviate the issues in handling millions of small files.

If this is enabled, the blobs will be Base64 encoded and written to the CDF as a tag attribute. This halves the number of objects being written to the Centera to 1 instead of 2 (CDF and Blob) in the case of 1 blob per CDF. In the case of more than one small blob per CDF the reduction in object use is increased.

Please indicate if you enable embedded blobs and if you do so, what do you set the `FP_OPTION_BUFFER_SIZE` to?

The CaminoSoft's Centera integration does not support blob embedding. However, the application allows for setting a minimum threshold for file archiving. The default threshold for minimum file archiving is 10 KB.

CDF Structure

Please describe the structure of the CDF:

- Is it a simple one tag one blob?

Yes

- Has a hierarchical tag structure?

No. One CDF contains just one tag, which in turn points to one blob.

- Why you create multiple tags? (Explain.)

Multiple tags are not created. Only one tag is created; this tag corresponds one on one to the file being archived to Centera.

- What metadata attributes do you write at the CDF level and at the tag level?

Metadata at Tag Level

- **ServerName:** This contains the name of the server where managed server HSM is running from.
- **FileName:** Fully qualified file name (including UNC part if file Managed Server HSM is archiving a file from remote server (NetApp Filer and Remote Windows Server))

- **Metadata at Clip Level**

Retention Period: As configured by user in archiving profile.

ClipName: Fixed name of 'CMSFRTDM'.

Storage Device Mode

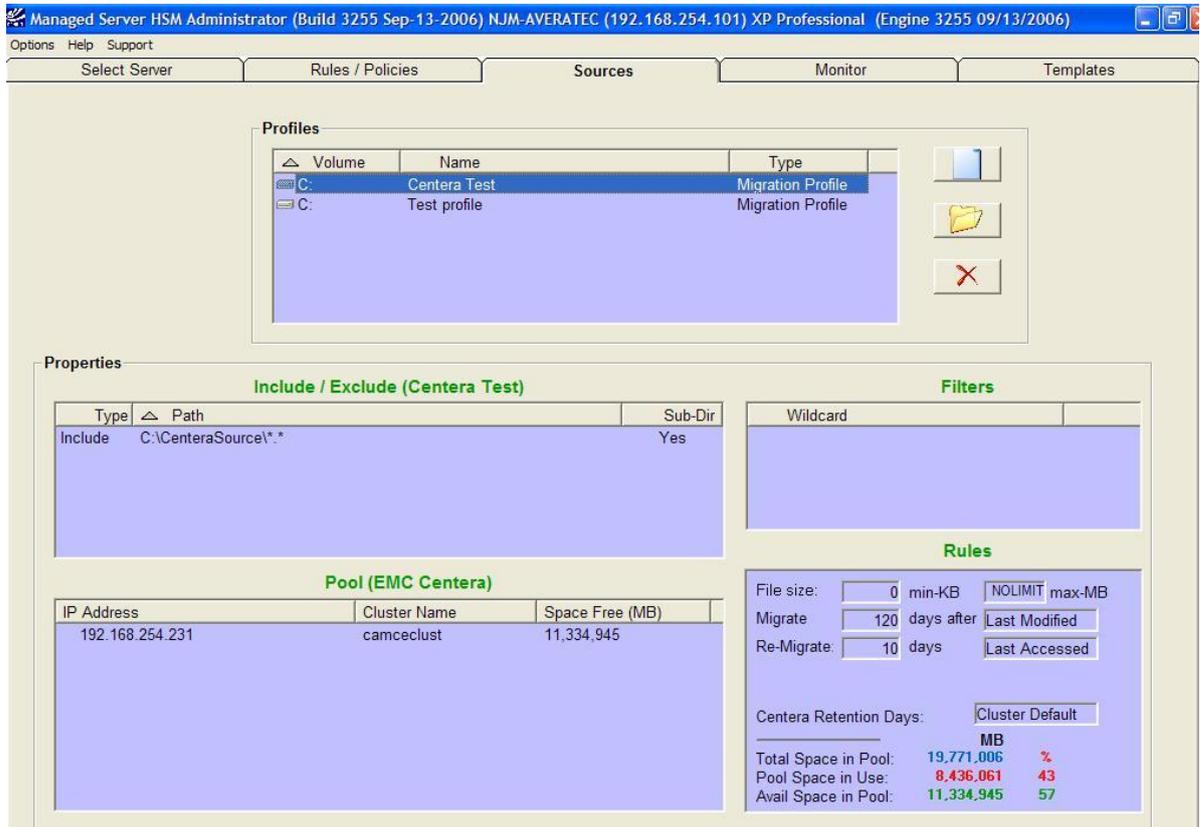
Explain how the Centera appears as a device in your archive module. In particular, can you define multiple Centera as multiple separate storage devices or as one logical storage device mapping to multiple physical Centeras.

The only way to define a Centera in Managed Server HSM product is by its IP Address. More than one IP Address can be defined in one archiving profile. This list of IP Addresses is provided directly to the Pool_Open API to communicate with Centera device(s). The Managed Server HSM product does not know the internal details of how Centera(s) is configured; it simply communicates via the IP Addresses provided to it.

It is up to the Centera Administrators to configure Centera as needed, however, the product merely uses the IP Addresses to communicate, and may not know the specifics about the Centera.

As such, multiple Centera storage devices cannot be mapped as one logical storage device. Each Centera needs to be individually defined in a archiving profile.

CaminoSoft's Managed Server HSM, Centera Proven Edition is configured through the administrator's GUI. The Centera is configured as a target device. The user must specify the IP addresses of the Centera access nodes.



Once the administrator provides the IP Addresses of the Centera target device, the cluster information is automatically gathered and displayed in the **Pool (EMC Centera)** section. Also, when a Centera Pool is specified, the administrator is allowed to specify **Centera Retention Days** or accept the cluster default for the files associated with the selected profile. Since multiple profiles can be specified, multiple retention periods can be specified to be associated with the files included in each profile.

Centera Specific Configuration Details

The Managed Server HSM GUI allows administrator to define multiple IP Addresses for a given profile. It is recommended that the IP Addresses of all the access nodes of both the primary and secondary clusters in the profile be defined. The IP Addresses of the primary access nodes should be defined first, followed by those of the secondary cluster. Such configuration will ensure both network and content failover for Reads and Queries. Furthermore, if the primary cluster is unavailable at the time Centera Pool is created, the secondary cluster will take on the role of primary and will accept Write transactions.

Content Address Management – Generating a list of clips

Confirm that you support Content Address up to 65 bytes in length.

Yes. The product uses standard data structure defined in SDK 'FPID'.

Explain where you store the Content Addresses for future use.

Frequently, customers need to be able to retrieve from the application a list of all current Content Addresses that the application knows about. Explain how this can be accomplished by the customer.

- *Are the clip Content Addresses stored in a database, a flat file or a stored somewhere else?*

Content addresses are stored within the source file system in an alternate data stream (file stub) attached to each archived file. MSHSM also stores the addresses in its databases (tables containing archiving and recall activity).

- *To obtain this list of clips is it a SQL query that is run and if so, what is the SQL? Does the customer just access a flat file or is this a Customer Service Request from the customer?*

CaminoSoft has a number of utilities that can provide specific information relating to stub contents (including Centera content addresses) which can be obtained by special request from our Customer Support engineers.

Threading Model

The Centera is designed to be accessed by multiple IO threads either by multiple threads within one process or multiple single threaded processes. Describe how your application writes/reads to the Centera.

- *Do you do multiple threaded writes? What is the default number of threads?*

Yes, writes are multithreaded. Default writes: 40.

- *Multi-threaded reads? What is the default number of threads?*

Reads are multi-threaded and will spawn per request within server limits.

- *Is the number of write/read threads configurable?*

The number of writes are configurable via command line. The number of reads are not configurable, it simply depends on the number of concurrent request to retrieve archived files.

Streams Usage

Detail which of the Centera SDK stream types you have implemented - i.e. file based, memory buffer based or generic streams.

The file based stream is used by utilizing the API `FPStream_CreateFileForInput`.

If you have used generic streams can you please confirm if:

- *You have implemented marker support*
 - Not Applicable / Not Implemented
- *You are capable of rewinding the stream 100MB or just abandon the transfer if any errors occur which would cause the SDK to want to rewind the stream*
 - Not Applicable / Not Implemented

Compliance Support

Retention Period and/or Retention Classes

Do you support assigning retention to content written to the Centera?

If you do support retention, do you support setting retention by period or by retention classes?

If you use retention classes, what are the names and durations of the classes you need to us?

If you support setting retention via retention period, do you support extending the retention period? Also if you do allow the retention period to be extended, how do you manage the deletion of the old CDF that had the original retention period assigned to it?

Archiving Profiles are configured through the CaminoSoft Managed Server HSM administrator's Graphical User Interface (GUI) to control the retention placed on folders of content. The retention period for the folder may be specified in days. The specified retention period is written in metadata contained in the Centera Content Descriptor File (CDF). If no retention period is specified, then the CaminoSoft Managed Server HSM sets a default value of zero in the CDF. Once the retention period is specified for a Archiving Profile, the value cannot be changed.

An administrator can select one the following retention options while creating profile. All clips created as a result of file archiving will have the appropriate retention set:

- No Retention
- Infinite Retention
- Cluster Default
- Custom Retention (in Days)

Advanced Retention Management

Do you support any of the Advanced Retention Management functionality of the SDK?

- *Event Based Retention (EBR)*
- *Litigation Hold (LH)*
- *Min / Max Retention Governor*

CaminoSoft's Managed Server HSM, Centera Proven Edition currently does not implement the Event Based Retention (EBR) and Litigation Hold (LH) functionality available with Centera SDK v3.1.

Pool Handling

Configuration

Do you allow the pool connection string to be configured by the storage administrator?

Yes.

Do you just pass the connection string un-interpreted to the FPPool_Open() method?

Yes. Multiple connection strings are concatenated per SDK guidelines.

Do you allow the setting of the Lazy Pool open strategy?

No.

Do you check that the pool opened has the correct capabilities for you to work at FPPool_Open() ahead of any IO to the cluster- e.g. checking that you have write capability before erroring on every write to the cluster if the storage admin has refused to grant this.

No. All the writes will fail if write capability is not configured.

How many and how frequently do you open pool connections?

The following processes open pool with Centera:

- **Archiving**
One pool is opened at the beginning of the scan, and closes at the end. The archiving frequency is configurable by Administrator, however typically this happens once a day.
- **StubCheck**
One pool is opened at the beginning of scan and closes at the end of scan. This process is started after the archiving process completes.
- **TargetDelete**
One pool is opened at the beginning of process and closes at the end. This process is started after the stubcheck process completes.
- **Recall**
One pool is opened when the very first request to retrieve data from Centera is received. The connection is open for the full duration of the program. This program typically runs 24x7 on the server, thereby causing the connection to be opened virtually all the time.

Access Profiles

Do you support Access Profiles? If you do not, please explain why not?

CaminoSoft's Managed Server HSM, Centera Proven Edition supports the use of Access Profiles, and the use of Virtual Pools on Centera clusters running CentraStar v3.0 and above.

Global Options

Which SDK Global Options do you set before opening a pool?

None.

Archiving Capabilities

Frequently, customers will have legacy archive systems and they will wish to move content from these legacy systems to their Centera.

Detail the utilities or processes available to customers to perform this type of Archiving activity.

Microsoft Windows

The CaminoSoft Managed Server HSM for Windows 2000/2003 engine periodically scans for candidates on the NTFS file systems and archives them to the specified Centera. A connection is established at the beginning of archiving (scan) cycle, and this connection handle is used through the rest of the scan to perform archiving in a multi-threaded fashion. The connection is closed at the end of the archiving cycle.

A thread is launched for every qualifying file to be archived to Centera. A maximum of 40 (configurable) threads can be created at a given time. The scan engine waits for thread availability prior to proceeding with further scan.

For every file to be archived, a clip is created first, its retention is then set, top tag is fetched, new tag is created, image stream is created, and at the end blob is written. The application also stores a custom string attribute called as server name and filename on the clip as meta-data. This information is used to restore stubs on the server during disaster recovery. After successful creation of clip/blob, data is truncated from the source file, and stub is created with information such as Centera's IP address and Clip ID. The file is eventually marked offline on the Windows Server.

Novell NetWare

The implementation requires a gateway Windows Server to act as middle-tier between the NetWare file server(s) and Centera. The NetWare files are archived by using CaminoSoft Managed Server HSM for NetWare to the intermediate Windows Server via CIFS protocol. The files are later picked up by Managed Server HSM for Windows, Centera Edition to be archived to Centera. Windows archiving to Centera is discussed in the prior section.

The target file created on the Windows Server, as a result of archiving of source file archived from NetWare Server, consists of a header and source data of the original file. When file is archived from NetWare to Windows, a unique header is added at the top of the target file, making it unique in the system. In order to take advantage of the single-instancing feature of Centera, the Windows HSM separates the header from the source-data and stores them into two different blobs pointing to the same clip. A pure Windows file is archived to just one clip and one blob pair.

NetApp Filer / IBM N Series Filer

Utilizing NetApp Fpolicy protocols, files residing on NetApp filers, Vfilers, and N Series filers can be managed and similarly archived to Centera(s) and/or other target devices. Policies for all supported platforms (Windows, NetWare and NetApp/N Series) are administered via the same graphical management console.

Legacy data to Centera

This feature would require stub conversion. CaminoSoft Professional Services can provide assistance. A "stub" file archiving utility is built-in.

Centera to Centera

EMC Centera Professional Services can assist with Centera to Centera archiving.

Centera to XAM

Centera to XAM archiving is still to be determined.

Content Replication

In terms of replication, CaminoSoft generally suggests Centera bidirectional replication. It further recommends that replicated deletes should not be configured. This recommendation is made with the intent of preserving a full history of data on the target replica cluster. CaminoSoft encourages all customers to maintain this configuration at least during initial pilot deployment of the joint solution. CaminoSoft solution also provides a mechanism for removing the orphan clips from the target replica cluster once it is determined that they are no longer required.

Failover

Do you make any changes to the default failover strategies for reads?

No. There is no explicit support for failover (see below).

If the primary Centera is down do you reopen the pool connection, specifying the replica IP address so that you can continue writing to a Centera?

No. The recommended practise is that the Administrator provides IP Addresses of both primary and secondary Centeras while setting up connection in the product. The product passes all the IP Addresses in the order as specified by Administrator to the pool open call.

SDK attempts to connect with Centeras in the order as specified in the connection string. SDK does not try to connect any further once it finds the Centera to which it can successfully connect. However, in due course of time, if the connected Centera becomes unavailable, the SDK tries to connect with the subsequent Centeras as defined in the connection string.

The above feature is utilized by the product to provide failover support. Once the primary Centera(s) become unavailable, SDK automatically connects with the secondary (fail over) Centera(s) and performs both reads and writes using the newly connected Centera(s).

High Availability (HA)

Can your application be configured for HA? If so, briefly describe how this is done.

The product can be configured under standard Microsoft Clustering Services. Both (all) nodes in the cluster needs to have Managed Server HSM installed. The node with the shared (managed) volume mounted is the one that performs all the interaction with Centera. Once this node becomes unavailable, the managed volume is mounted on the other node via Microsoft Clustering services. After the failover, the Managed Server HSM running on the failed-over node automatically recognizes the failed over volume and starts managing. There are no additional steps required during fail over or fail back once the initial configuration is completed.

Deployment Scenarios

The CaminoSoft and Centera integration supports clustered servers. One Target may be configured per Archiving Profile. Centera replication may be used in Disaster Recovery configuration. CaminoSoft also provides a recovery utility that can query the Centera cluster and rebuild the "stub" files that may be lost in the event of a disaster.

Disaster Recovery

CaminoSoft's DR Tool for Centera can be used to recreate stubs on the production Windows Server. The tool allows specific clips or specific files to be recreated. Date parameters can be optionally included to restrict the DR query range for improved performance.

Miscellaneous API Usage

SDK Logging Support

Does your application support the dynamic SDK logging available from SDK 3.2?

No.

Application Registration

Detail the two strings (AppName, AppVer) you pass to the FPPool_RegisterApplication() method which you call before the FPPool_Open().

This API is not used.

SDK Usage

Detail all Centera API methods used.

- **Open**

FPPool_Open

FPPool_GetClusterTime

- **Write**

FPStream_CreateFileForInput

FPTag_BlobWrite

FPTag_SetStringAttribute

- **Read**

FPClip_Open

FPClip_GetName

FPClip_GetTopTag

FPTag_GetFirstChild

FPTag_Close

FPStream_CreateFileForOutput

FPTag_BlobRead

FPClip_GetTotalSize

- **Delete**

FPClip_Delete

- **Retention**

FPClip_SetRetentionPeriod

- **Error**

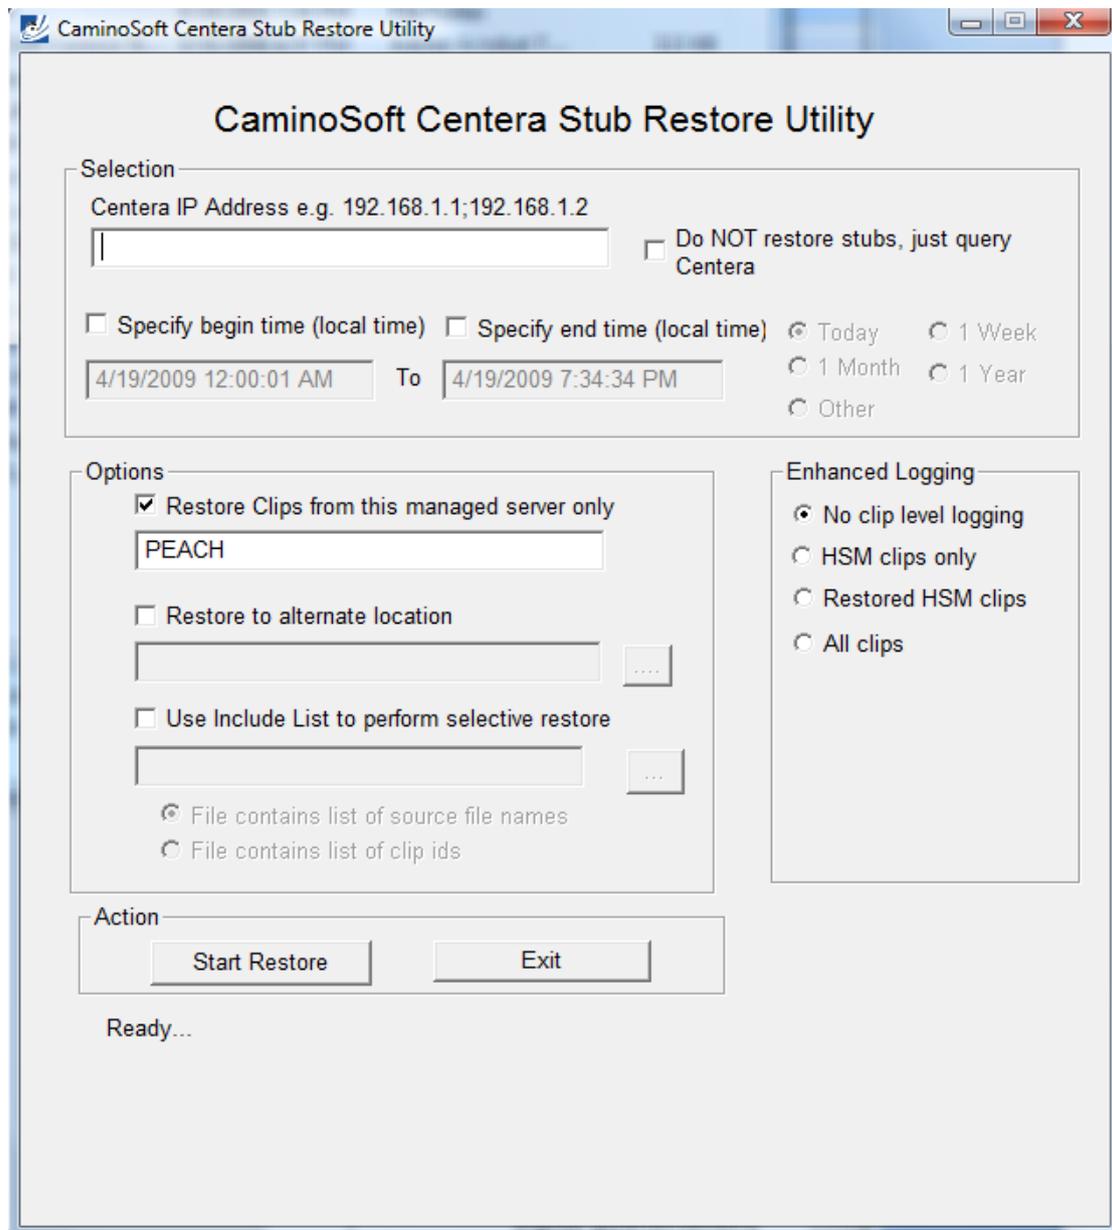
FPPool_GetLastError

Query API

Do you use the Query API? If so, why? Detail how you use the Query API?

Yes, Query API is used. These APIs are used for Disaster Recovery purposes. These set of APIs are used to enumerate all the clips (created by the product) stored on Centera. The application obtains the meta data of these clips to re-create stubs.

Below is the screenshot of the application that uses Query API to restore stubs:



Administrator can provide the following optional fields to narrow the query operation:

- Begin and End Dates
- List of clips
- Additional meta data as stored in the clips/tags
- Following APIs are used to perform query:
 - FPQueryExpression_Create,
 - FPQueryExpression_SetType,
 - FPQueryExpression_SetStartTime,
 - FPQueryExpression_SetEndTime,
 - FPQueryExpression_SelectField,
 - FPPoolQuery_Open,
 - FPQueryExpression_Close,
 - FPPoolQuery_FetchResult,
 - FPQueryResult_GetResultCode,

-
- FPQueryResult_GetClipID,
 - FPQueryResult_Close

Sample CDF

Provide a copy of a typical CDF that your application will write to the Centera

```
<?xml version='1.0' encoding='UTF-8' standalone='no'?>
<ecml version="3.0">
<eclipdescription>
<meta name="retention.period" value="0"/>
<meta name="type" value="Standard" offset="0"/>
<meta name="name" value="CMSFRTDM"/>
<meta name="creation.date" value="2009.04.17 18:10:33 GMT"/>
<meta name="modification.date" value="2009.04.17 18:10:43 GMT"/>
<meta name="creation.profile" value="anonymous"/>
<meta name="modification.profile" value="anonymous"/>
<meta name="numfiles" value="1"/>
<meta name="totalsize" value="48712"/>
<meta name="refid" value="ER5439CKEEI7EBJM5TR6H9CHIO"/>
<meta name="clusterid" value="109b2ba4-1dd2-11b2-83e9-df862e5050d3"/>
<meta name="prev.clip" value=""/>
<meta name="clip.naming.scheme" value="MD5"/>
<meta name="numtags" value="1"/>
<meta name="sdk.version" value="3.1.477"/>
</eclipdescription>
<eclipcontents>
<CMSFRTDM_DATA
ServerName="XPTEST"FileName="E:\_src\3286\emc\subfolder\nwelcome.bmp">
<eclipblob md5="9INEF7GIJ30O4xDJ50DN91UK5UJG4142QKBU0A0EBJM5TR6H9CHIO"
size="487
12" offset="0"/>
</CMSFRTDM_DATA>
</eclipcontents>
</ecml>
```

SDK Log Sample

Provide an SDK log of a sample run of your application writing, reading and deleting content from a Centera.

SDK Logging can be enabled by setting the environment variable FP_LOGPATH to the pathname of the logfile you wish to create. More details on SDK logging can be found in the Centera API Programmers Reference Guide.

Attached log files:

Write: centera_sdk_log_write.txt

Read: centera_sdk_log_write.txt

Delete: centera_sdk_log_write.txt

Error Handling

The 10 most common API errors are listed below. How are they handled in your code? Please describe and/or supply a code fragment.

The product always checks the return code from all the API calls into Centera SDK. When an error is detected, the error is logged in the product's error database. Usually additional identifying data such as Centera's IP Address or file name (if applicable) is also recorded in the application. The date/time of the error and the application that caused the error is also recorded.

Error	How does the application handle this error
FP_NO_POOL_ERR (-10020) - No connection with any pool	An error is logged. No explicit attempt is made to reconnect. Usually, SDK is capable of reconnecting in the next call.
FP_CLIP_NOT_FOUND_ERR (-10021) - ClipFile (CDF) is not found in the pool	An error is logged. This usually occurs when the application is trying to retrieve data from Centera. The application returns "access denied" to the application or user trying to open a archived file.
FP_FILE_NOT_STORED_ERR (-10014) - Blob could not be stored on import or could not be found	An error is logged. This may happen during "write" of file's data to Centera. The archiving of this file is failed. The file is attempted to archive again in the next archiving scan.
FP_OPERATION_NOT_SUPPORTED (-10012) - Operation not supported	An error is logged. Appropriate action is taken depending on the origin of error (archiving or recall).
FP_BLOBIDMISMATCH_ERR (-10036) - BlobID mismatch between client & server, blob is corrupt	An error is logged. Appropriate action is taken depending on the origin of error (archiving or recall).
FP_OPERATION_REQUIRES_MARK (-10201) - The application requires marker support	An error is logged. Appropriate action is taken depending on the origin of error (archiving or recall).
FP_SERVER_ERR (-10005) - The server reported an error from the operation	See FP_NO_POOL_ERR.
FP_SOCKET_ERR (-10101) - Error on network socket	See FP_NO_POOL_ERR.
FP_NO_SOCKET_AVAIL_ERR (-10034) - No new socket is available for the transaction	See FP_NO_POOL_ERR.
FP_BLOBBUSY_ERR (-10040) - The blob on the server is busy and cannot be read or written	An error is logged. Appropriate action is taken depending on the origin of error (archiving or recall).

Performance Expectations

If you have conducted any performance tests please give details of the results here.

Following is the summary of performance test performed by Veritest during Centera Proven Certification:

10 - 20K Write Clips:

- 15.92k/sec

10 - 20k Read Clips:

- 143.36k/sec

40 MB Write Clips:

- 4Mb/sec

40 MB Read Clips:

- 10Mb/sec

Release Schedule

Detail how frequently major and minor releases are scheduled for your application. In addition, confirm if you will take advantage of any Centera SDK release as a matter of course as you issue a major/minor release.

Minor releases are typically scheduled quarterly while major releases take place once a year or less. In addition, interim builds are generated as needed to address bug fixes and minor system enhancements. Generally, an upgrade to Centera SDK release is considered warranted if it addresses specific issues our customers may be experiencing, offers increased stability or provides improved performance.