CA Chorus™ for Storage Management

Site Preparation Guide

Version 03.0.00, Fourth Edition

CA Technologies Product References

This document references the following CA Technologies products:

- CA ACF2™ for z/OS (CA ACF2)
- CA Chorus™ (CA Chorus)
- CA Chorus™ for DB2 Database Management (CA Chorus for DB2 Database Management)
- CA Chorus™ for Security and Compliance Management (CA Chorus for Security and Compliance Management)
- CA Chorus™ for Storage Management (CA Chorus for Storage Management)
- CA Chorus™ Infrastructure Management for Networks and Systems (CA Chorus Infrastructure Management)
- CA Chorus™ Software Manager (CA CSM)
- CA Common Services for z/OS (CA Common Services for z/OS)
- CA Detector® for DB2 for z/OS (CA Detector)
- CA Easytrieve® Report Generator (CA Easytrieve Report Generator)
- CA PDSMAN® PDS Library Management (CA PDSMAN)
- CA Top Secret® for z/OS (CA Top Secret)
- CA Vantage™ Storage Resource Manager (CA Vantage)

Contact CA Technologies

Contact CA Support

For your convenience, CA Technologies provides one site where you can access the information that you need for your Home Office, Small Business, and Enterprise CA Technologies products. At http://ca.com/support, you can access the following resources:

- Online and telephone contact information for technical assistance and customer services
- Information about user communities and forums
- Product and documentation downloads
- CA Support policies and guidelines
- Other helpful resources appropriate for your product

Providing Feedback About Product Documentation

If you have comments or questions about CA Technologies product documentation, you can send a message to techpubs@ca.com.

To provide feedback about CA Technologies product documentation, complete our short customer survey which is available on the CA Support website at http://ca.com/docs.

Documentation Changes

The following list details the changes made since the third edition of this guide:

- Server Requirements (see page 20)—Noted the new feature to automatically configure heap memory.

The following list details the changes made since the first edition of this guide:

- Legal Notices (see page 2)—Updated to reflect public documentation legal disclaimer.
- The installation and configuration process of the storage engine for CA Chorus for Storage Management is simplified. The chapter "Installing and Configuring the Storage Engine (see page 49)" is updated accordingly. (Back-end PTF RO64747 required.)
- You can now launch the storage engine Host Configuration Client (Config Client) from the Storage Management interface (see page 9) available in the Quick Links Module. (PTF RO62885 required.)
- You can now change values in the Customer Site Costs Data object directly in the Investigator. References that you have to edit and resubmit the Cost Analysis E4HI0011 configuration job in the topic Acquire Information for the Cost Analysis Feature (see page 24) is updated accordingly. (PTF RO63077 required.)
- Software Requirements (see page 19)—Removed release-specific FIXCAT references and added a variable instead.

The following list details the changes made since the last release of this documentation:

- System Requirements (see page 21)—Removed the heap memory requirements.
- Server Requirements (see page 20)—Updated the heap memory requirements, noted the parameters to modify the heap size, noted that the values represent real storage, and moved the heap memory requirements to this new topic.
- Software Requirements (see page 19)—Clarified browser support.

Contents

Chapter 1: Architecture and Installation Overview

Terminology

Storage Engine

Refers to the CA Vantage system on the z/OS host. The *storage engine* is the back-end product on the z/OS host that provides information to the CA Chorus for Storage Management user-interface.

Storage Management interface

Refers to the storage engine web client (the CA Vantage web client). In the CA Chorus interface, the Quick Link to the Storage Management interface is "Manage Storage Resources". When selected, the *Storage Management interface* opens in a separate browser window.

Usually, the CA Chorus and the Storage Management interfaces use the same terminology for the same object categories, objects, and options. The Storage Management interface uses storage engine terminology. The following list details a few exception examples:

Object Categories

Some object category names differ in the *CA Chorus Investigator interface* and the *Storage Management interface*. For example, in the *Storage Management interface* object tree, an object category is named CA Vantage Internal Management. In the *CA Chorus Investigator interface*, it is named Storage Platform Administration.

Object Names

Some object names differ in the *CA Chorus Investigator interface* and the *Storage Management interface*. For example, in the *Storage Management interface*, the first object in the Automation and Logging category is named Server (Event Type) Status. In the *CA Chorus Investigator interface*, it is named Automation Event Types.

You can launch the storage engine Host Configuration Client (Config Client) to set storage engine system parameters from the Tools menu in the Storage Management interface.

Storage Engine Windows Client

Refers to the CA Vantage Windows Client. The *storage engine windows client* interface is a stand-alone, Windows-based user interface. CA Chorus for Storage Management users access it to do the following tasks:

- Create storage engine Log Scripts to collect storage metrics data for the Time Series Facility.
- Create storage engine Automation Scripts to generate storage engine messages for Alerts.
- Launch the storage engine Host Configuration Client (Config Client) to set storage engine system parameters. The Config Client can also be launched from the Storage Management interface which can be launched from the Quick Links module in CA Chorus.

How the Installation Process Works

This guide details the tasks that a system programmer and security administrator can complete before starting the installation, deployment, and configuration tasks that are described in the *Installation Guide*. The following diagram provides a high-level overview of the CA Chorus and discipline installation, deployment, and configuration process and the guides that you use.

Important! You must use CA Chorus Software Manager to install CA Chorus and its disciplines.

Important! If you install a discipline, you must deploy and configure it.

Note: For the boxes that indicate work from the discipline *Site Preparation Guide*, repeat this step for each discipline that you are installing.

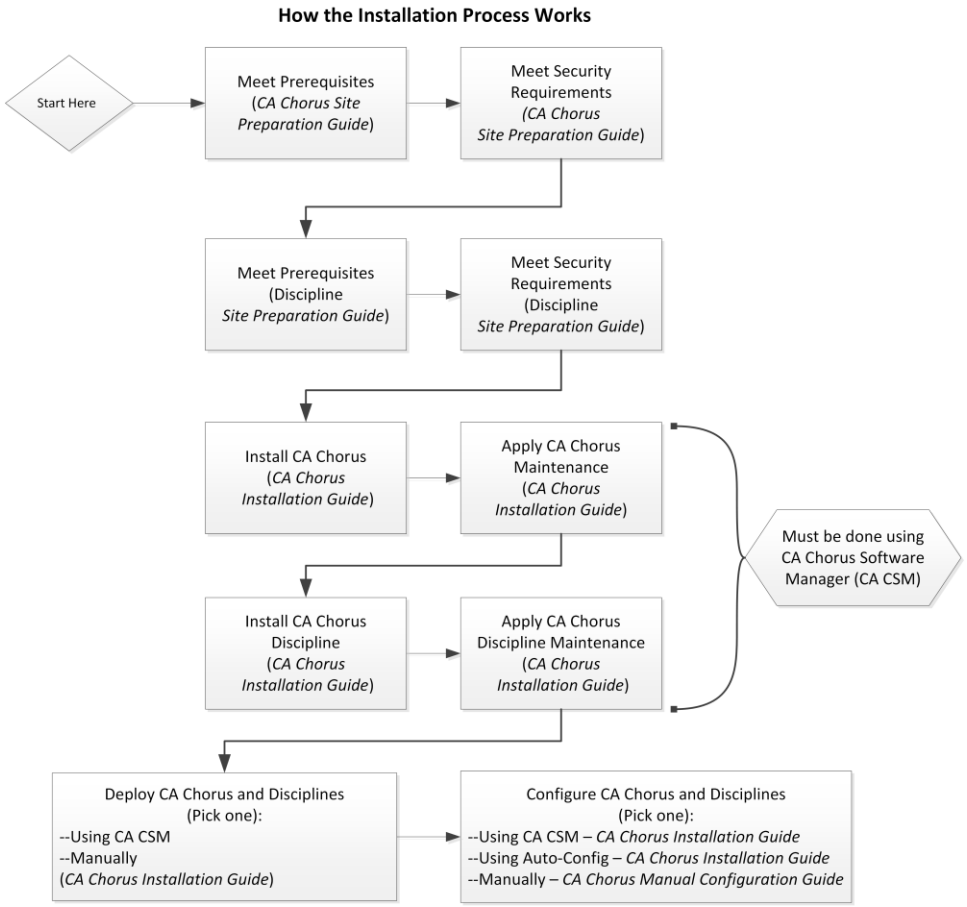

To install, deploy, and configure your CA Chorus and its disciplines, complete the following steps:

1. Meet the software, system, port, and other prerequisites as described in the *CA Chorus Site Preparation Guide*.
2. Meet the security requirements as described in the *CA Chorus Site Preparation Guide*.
3. Use the Prerequisite Validator to confirm that you have set up your system correctly as described in the CA Chorus *Site Preparation Guide*.

4. Meet the software, system, port, and other prerequisites as described in the applicable discipline *Site Preparation Guide*. Repeat this step for each discipline that you are installing.

5. Meet the security requirements as described in the applicable discipline *Site Preparation Guide*. Repeat this step for each discipline that you are installing.

6. Install CA Chorus and the applicable disciplines using CA CSM as described in the *CA Chorus Installation Guide*. This step involves acquiring the CA Chorus software (transporting to your z/OS system) and installing using SMP/E. The installation process creates a CSI environment and runs the RECEIVE, APPLY, and ACCEPT SMP/E steps. The software is untailored.

7. Deploy CA Chorus and the applicable disciplines using CA CSM or a manual process. The *CA Chorus Installation Guide* details both methods.

   This step copies the target libraries to another system or LPAR.

   Important! For deployments from CA CSM, you must deploy CA Chorus and your disciplines at the same time. For example, installing CA Chorus, DBA, and Security, and then deploying only CA Chorus and DBA is not supported.

   Important! To use the CA CSM Software Configuration Service, CA CSM deployment is required.

8. Configure CA Chorus and the disciplines. This step creates customized load modules, bringing the CA Chorus software to an executable state. You configure the product using one of the following methods:

   Note: We recommend one of the first two options as the most efficient method to configure your products.

   CA CSM

   This method lets you use the wizard-based CA CSM tools to configure the product. For this configuration method, a deployment using CA CSM is required.

   The *Installation Guide* includes the CA Chorus and discipline steps for this method.

   Automated Configuration

   This method lets you edit one batch job (ETJICUST) and one configuration file. A Java program then propagates your changes to the applicable members. You then manually submit each job. For this option, we recommend that you configure the platform and disciplines at the same time.

   The *Installation Guide* includes the CA Chorus and discipline steps for this method.

   Manual

   This method lets you manually edit and run each configuration job.

   For this method, configure CA Chorus and its disciplines using the *Manual Configuration Guide*.

   Your CA Chorus system is installed, deployed, and configured.

Architecture and Setup Overview

The following diagrams provide an overview of the architecture and setup of CA Chorus for Storage Management. After installation and setup is complete, you can use CA Chorus to manage storage resources across your z/OS enterprise.

- Single storage engine architecture and setup of CA Chorus for Storage Management:

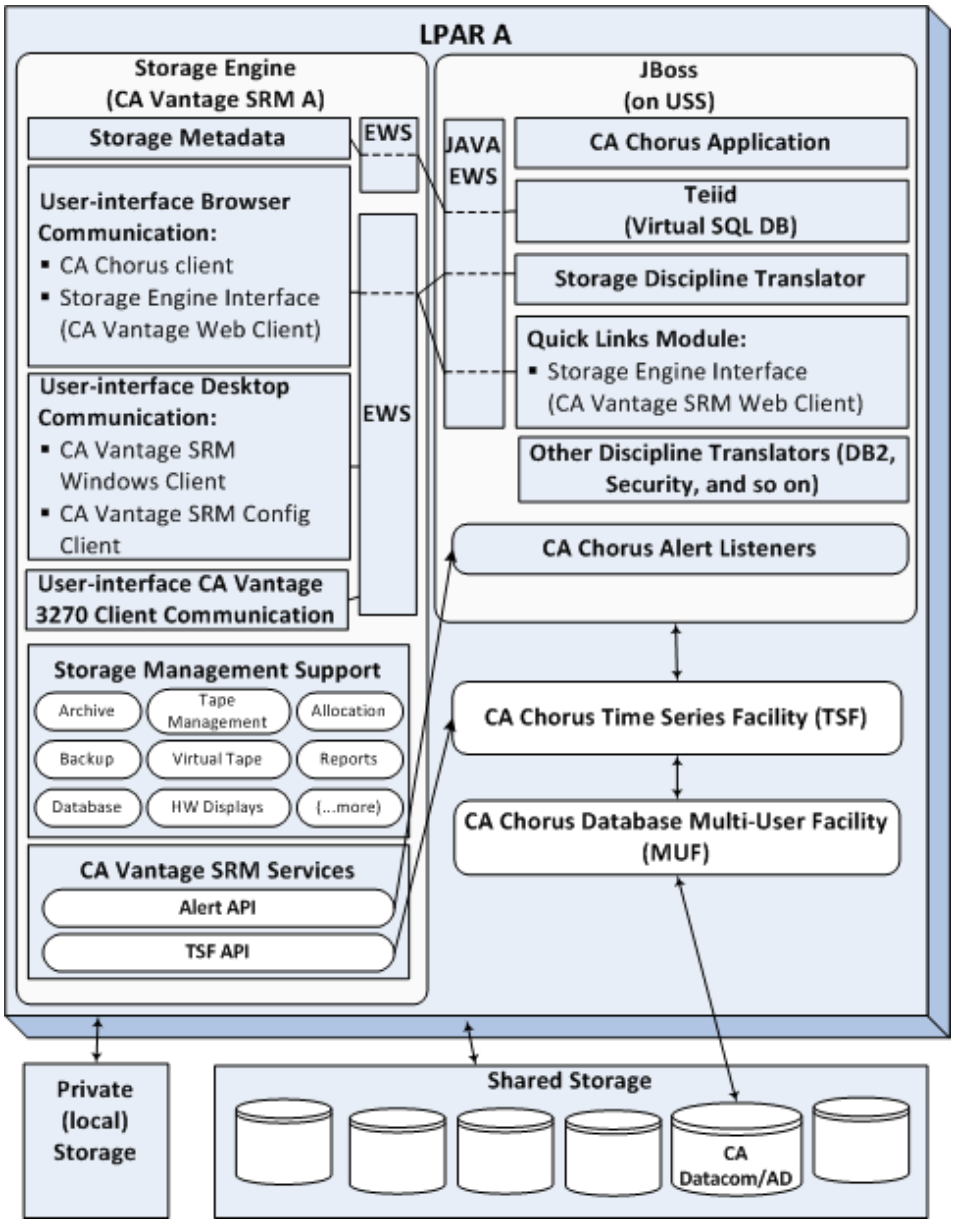

- Multiple storage engine architecture and setup of CA Chorus for Storage Management:

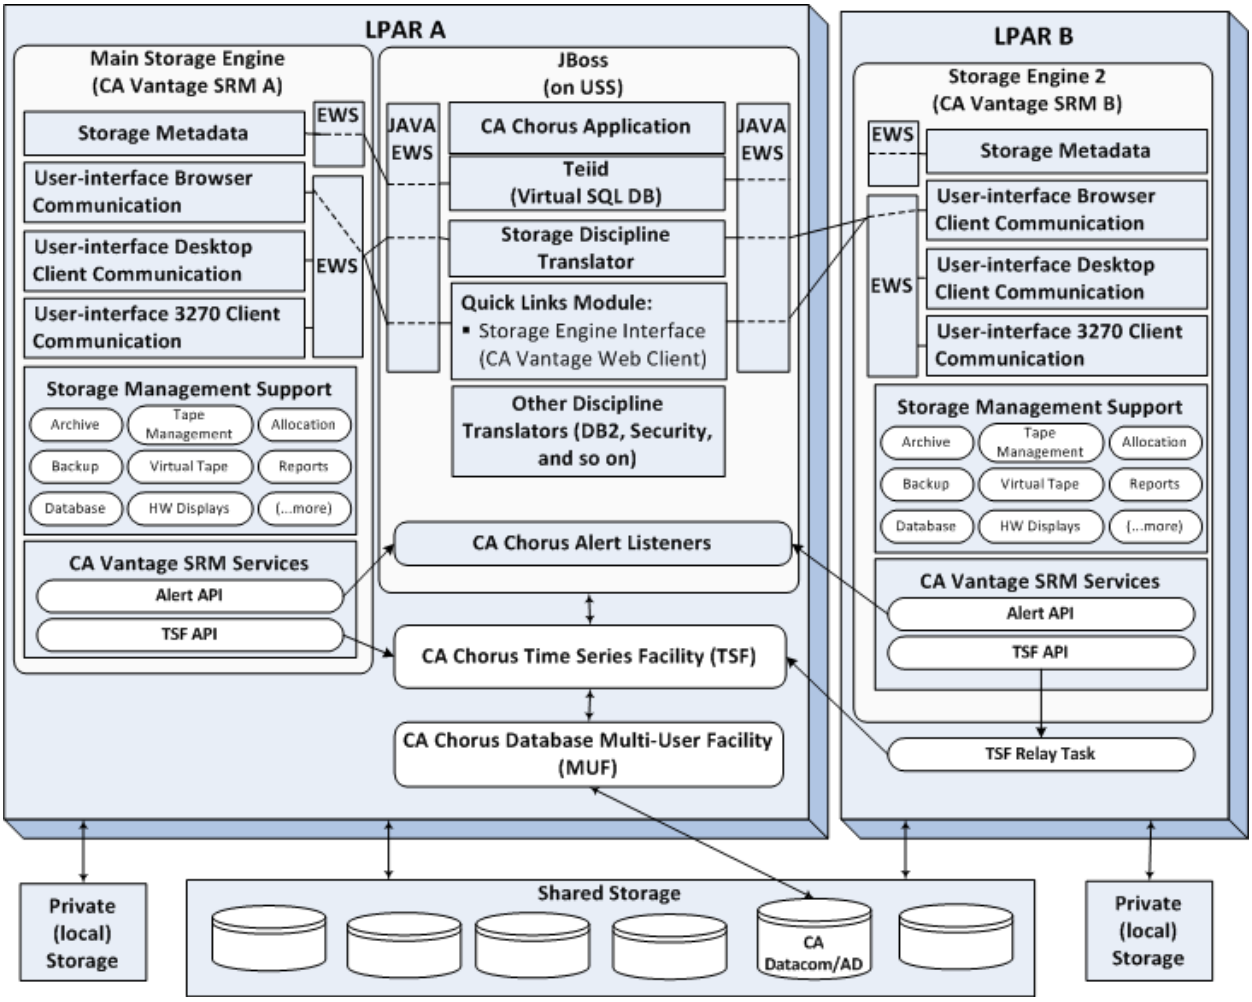

Observe the following points:

- In the single storage engine setup, it is common to have the storage engine and the JBoss server in the same LPAR. However, it is not required to have both in the same LPAR.
- The diagram showing multiple storage engine architecture and setup of CA Chorus for Storage Management shows two LPARs. However, you can have more than two LPARs.

To install and configure CA Chorus for Storage Management, you must:

- Install and configure the storage back-end engines (CA Vantage) on one or more selected systems (LPARs).
- Install and configure a single CA Chorus to communicate with the storage engines.

The diagrams illustrate the basic architecture and its parts:

LPARs

Identify logical partitions of a mainframe (z/OS system), on which you execute CA Vantage as a back-end engine for CA Chorus for Storage Management. Multiple LPARs are supported.

JBoss

Contains the CA Chorus system. Includes the following items:

JBoss

Provides the browser support and components that communicate with the back-end engines for the various disciplines, such as the storage engine.

Teiid

Provides a data virtualization system that allows applications to use data from multiple, heterogeneous data stores, or back-end storage engines.

Storage Discipline Translator

Provides an abstraction layer between the Teiid Query Engine and the storage data source. The translator converts Teiid issued query commands into storage engine-specific commands and executes them.

Quick Links Module

Launches the Storage Engine Interface (CA Vantage Web Client). If you have multiple CA Chorus disciplines, it also launches their back-end browser interfaces.

Other Discipline Translators

Contains other CA Chorus discipline translators, when you have a multiple CA Chorus discipline setup.

JAVA EWS

Provides JAVA methods for making Enterprise Work Station (EWS) requests to the storage engine. EWS is the proprietary communication protocol over TCP/IP used by the storage engine.

CA Chorus Listeners

Provide the service for receiving Alerts that are sent by the various back-end storage engines.

CA Chorus Time Series Facility (TSF)

Provides the facility for receiving, storing, and querying metrics about objects that are managed by the back-end engines. Because all metrics are stamped with the date and time, a series can be graphed over time to show the past trend, and to project into the future.

CA Chorus Database Multi-User Facility (MUF)

Provides the infrastructure for storing and retrieving TSF data within the physical database.

Storage Engines (CA Vantage Systems)

Identifies your CA Vantage systems that are configured to manage your shared and private storage environments.

Storage Metadata

Provides storage object, field, action, and relationship information displayed in CA Chorus. Storage Metadata is provided by the main storage engine to CA Chorus.

In a multiple storage engine environment, if the main storage engine is not available, CA Chorus selects the next available storage engine as the main storage engine.

Storage Engine EWS

Provides the proprietary Enterprise Work Station communication protocol support over TCP/IP by which clients communicate with the storage engine.

User Interfaces

Consist of:

- Two browser interfaces:
  - CA Chorus client
  - CA Vantage Web Client

    Installed in the CA Chorus for Storage Management system and launched from the CA Chorus Quick Links module.
- Two desktop interfaces:
  - CA Vantage Windows Client

    Installed on PCs separately.
  - CA Vantage Configuration Client

    Included in and launched from the CA Vantage Windows Client and the CA Vantage Web Client.

    Note: For more information about the CA Vantage Windows Client and the CA Vantage Configuration Client, see the CA Vantage documentation.
  - (Optional) CA Vantage 3270 Client

    Included in the CA Vantage installation. For more information, see the CA Vantage documentation.

Storage Management Support

Identifies the storage management functions to be supported, such as: backup, archive, allocation, management of both regular and virtual tape systems, along with displays of hardware devices and much more.

Alert API

Pushes Alerts to the CA Chorus Alert Listeners using TCP/IP connections, which are defined in the %%DSNPFX%%.URLS data set.

TSF API

Pushes TSF data to the CA Chorus TSF started task using TCP/IP connections. From remote Storage Engines, TSF data is pushed to the local TSF Relay Task, which connects through TCP/IP to the CA Chorus TSF.

Shared Storage

Represents all storage devices that are shared between z/OS systems.

Private Storage

Represents storage devices, if any, that are not shared with other z/OS systems.

Chapter 2: Addressing General Prerequisites

Software Requirements

The following software is required for CA Chorus for Storage Management:

- CA Technologies Software—The following software is required:
  - CA Chorus Version 3.0

    For initial site installations, you install CA Chorus and the disciplines at the same time as described in the *CA Chorus Installation Guide*.

    When you are installing a discipline into an existing CA Chorus instance, confirm that the version of your discipline matches the CA Chorus version.

    Note: For more information about installing and configuring CA Chorus, see the *CA Chorus Installation Guide*.
  - CA Vantage (the storage engine) Release 12.6 base system and the CA Vantage Automation Option is installed and configured on your z/OS with the latest CA Recommended Service and PTFs applied from the Fix Category (FIXCAT label):
    - CA.ProductInstall-RequiredService.CA-Mainframe-Chorus.*, where * indicates the version of CA Chorus that you are installing.

    This requirement includes the following tasks:
    - Activating PassTicket support in the storage engine (see page 35).
    - Allowing the CA Chorus started task and users to log in to the storage engine.

Configure the storage engine, the various components, and additional configuration for CA Chorus for Storage Management in accordance with the instructions in the CA Vantage SRM Configuration Guide.

Note: For specific instructions, see the following items:

- For installation instructions, see the *CA Vantage Installation Guide*.
- For security information about logging in to the storage engine, see the *CA Vantage Reference Guide*.
- For configuration instructions, see the *CA Vantage Configuration Guide*.

- IBM Software—The following software must be available on the systems where you install CA Chorus:
  - IBM Ported Tools for z/OS with IBM APAR OA10965. (This APAR provides support for the Xvfb application.) This application is required for the charting utility started task. The charting utility is used to create charts in the Storage Management interface.

  No additional IBM software requirements exist for the CA Chorus for Storage Management discipline.
- PC software that is required for each user:
  - Adobe Flash Player 9.0.124 or above
  - At the release of Version 3.0, CA Chorus supports Microsoft Windows Internet Explorer 9 and Mozilla Firefox 13 through 19. As new browsers are released, we will validate them and post compatibility on the CA Chorus product page under Content Type, Recommended Reading.

    Note: CA Chorus requires a minimum screen resolution of 1024 x 768. If your screen resolution does not meet this requirement, use full screen mode (F11 in most browsers) to include the scroll bar on the display.
  - **Note:** To start the Host Configuration Client, the JAVA Platform SE plug-in must be current and enabled in the user's browser.

Server Requirements

Confirm that your site meets the following requirements:

Real storage

200 MB heap memory for CA Chorus for Storage Management plus 2450 heap memory for CA Chorus

Note: CA Chorus automatically configures the heap memory size based on the disciplines that you install at JBoss startup.

Note: If all disciplines are installed, 3150 MB is required. CA Chorus automatically configures the heap memory size based on the disciplines that you install. This configuration is done during the CA CSM Software Configuration Service (SCS) or in CETJJCL(ETJI0150).

System Requirements

Confirm that your site addresses the following system requirements:

Processor

CA Chorus uses a JavaVM environment on z/OS. So, we *strongly* recommend that you use a zIIP specialty processor for the best performance and better use of resources.

Disk

CA Chorus for Storage Management requires approximately 180 cylinders on the ZFS that holds the pax installation file.

Note: This space can be reclaimed after installation.

Target Libraries

The following table shows the target library Low-Level Qualifier (LLQ) data set space requirements in tracks for CA Chorus for Storage Management:

Data Set Name	Tracks
CE4HJCL	11
CE4HXML	10
CE4HZFS (ZFS directory)	3000

Distribution Libraries

The following table shows the distribution library Low Level Qualifier (LLQ) data set space requirements in tracks for CA Chorus for Storage Management:

Data Set Name	Tracks
AE4HJCL	11
AE4HSHSC	10
AE4HZFS (ZFS directory)	2000
AE4HXML	10

The following table shows the Storage Management interface database (CE4HVDB) and reports (CE4HRPT) LLQ data set space requirements in tracks for the Storage Management interface. The following space allocation values are the values allocated in the installation job:

Note: Over time, users create reports, user-views, and so on, which could fill the space allocated. You may need to increase the space accordingly.

Data Set Name	Tracks
CE4HRPT	150
CE4HVDB	150

TSF Space Requirements

The Times Series Facility (TSF) requires allocated space. The installation process sets a default of the allocated space. Over time you might find that more space is needed depending on the amount of data logging you do for this feature. For instructions how to determine the amount of space you need and change the amount of space that is allocated for TSF, see the *CA Chorus Administration Guide*.

Port Requirements

CA Chorus has port requirements for the JBoss server and Time Series Facility (TSF) components. For more information, see the *CA Chorus Installation Guide*.

The CA Chorus for Storage Management discipline requires one additional TCP/IP port. This port is used for the IBM Xvfb application. The application is used to create charts in the Storage Management interface that is launched from the CA Chorus Quick Links module. The default port that is used by IBM is 11.

Note: For more setup information about the Xvfb application for the Storage Management interface, see the *CA Chorus Configuration Guide*.

Common Services Requirements

The following CA Common Services for z/OS (CCS) are used with CA Chorus for Storage Management:

- CAICCI
- CAIRIM
- CA LMP
- CAMASTER

  Important! The CAMASTER address space must be running. If it is running, the following message is in the z/OS syslog as part of the IPL portion of messages:
  CAMS101I CAMASTER INITIALIZATION COMPLETE.
  CAMASTER is a noncancelable started task that provides system services and storage resources for various CA products and CCS. CAMASTER uses minimal CPU, and cannot be stopped or restarted.

- CA Easytrieve service to generate batch reports

Note: These services are delivered and installed with CA Common Services for z/OS. If other CA Technologies products are installed at your site, these services could be installed. If these services have not been installed, do so now. For more information about the installation and configuration of these components, see the CCS documentation. To use the CA Easytrieve service with CA Chorus, no additional configuration is required in CCS after the SMP/E installation steps.

Additional common services can be used with the products that must be installed to support CA Chorus for Storage Management.

Note: For more information about those requirements, see the related product installation and implementation documentation.

LMP Key Requirements

The CA License Management Program (CA LMP) tracks licensed software in a standardized and automated way. CA LMP reports on activities that are related to the license, usage, and financials of CA Technologies products.

CA Chorus and its disciplines are licensed with an LMP key. You acquire the LMP key with one of the following methods:

- From your product media
- From CA Support

During startup, CA Chorus for Storage Management license validation is performed by calling the CA LMP service of the CAIRIM component. CAIRIM is a component of CA Common Services for z/OS. The CAIRIM component verifies the product LMP keys. When a CA Chorus user logs in, the CA Chorus server makes LMP checks against the appropriate product keys. For example, if CA Chorus for Storage Management and CA Chorus for Security and Compliance Management are installed, the LMP check is performed for both disciplines, and the seat count for these disciplines is increased.

To report on and to verify compliance with CA Technologies product seat license concurrent usage, SMF type 89 records are collected over the required reporting period. A registered software usage report can be produced from the SMF type 89 records using the IBM z/OS product usage report utility program, IFAURP. The IFAURP report utility supports specification of a software vendor such that only CA products are included in the generated reports. Use member CHSMFRPT in *your_chorus_hlq.*CETJJCL to extract SMF type 89 records from the installation SMF collection data sets. After submission, this job sorts the extracted records and produces the software usage reports.

Note: For more information about CA LMP, CAIRIM, and generating software product registration reports, see the CA Common Services for z/OS documentation.

Acquire Information for the Cost Analysis Feature

The CA Chorus Cost Analysis feature comes with a set of standard system analysis scripts that provide various cost analysis of your environment to help you manage your data center costs.

One of the values used in the analysis calculations is the Cost/Value field value shown in the Customer Site Costs Data object. The Cost Analysis feature E4HI0011 configuration job is provided with default Cost/Value field values. The default Cost/Value field values attempt to represent typical site values. During the installation process, you can run the job with the default values or customize Cost/Value field values in the job to better reflect your site and then run the job. The instruction to run the job the first time, is documented in the *CA Chorus Installation Guide*.

Note: You can also change configured Cost/Value field values at any time after installation by editing the Cost/Value field values in the Details tab of the Customer Site Costs Data object, or by editing and resubmitting the E4HI0011 configuration job. For more information, see the *CA Chorus Administration Guide*.

In preparation for the configuration, obtain values for the following variables in the Cost Analysis feature configuration job.

Observe the following points:

- USD is the only currency supported. For monetary values, you should provide values that are expressed in dollars. The monetary values in the Cost Analysis feature appear with the dollar sign. For example, $10.00. However, if you want to use another currency, use that currency consistently for all monetary values. If you enter monetary values in a different currency consistently, calculations are then consistent using the common currency basis, but appear with a $ sign in front of them in the Cost Analysis feature.
- The variables which you can acquire information for in the job already have default values. Review each of these values to ensure that they are representative of your environment and your data center.
- Do not use punctuation characters in any values other than a decimal point which should be specified using the dot or period (.) character. That is, do not use dollar sign ($) or comma (,) characters when specifying any of the values for the variables.

The following variables are available in the job which you can acquire information for:

CURRENCY_BASIS_ISO_4217

Specifies the currency for any of the other variables that require a *monetary* value. Currently, USD is the only currency supported.

Default: USD

DASD_COST_PER_GB

Specifies the cost that you pay per gigabyte of Direct Access Storage Device (DASD) for your typical DASD purchase. That is, actual or averaged over a range of typical DASD devices which can be found in your data center.

Default: 500.00

DASD_CAPACITY_IN_GB

Specifies the total capacity in gigabytes of your typical DASD purchase. It should include actual or averaged over a range of typical DASD devices that may be found in your data center. When set properly, DASD_CAPACITY_IN_GB multiplied by DASD_COST_PER_GB should result in the average price that you pay for a typical DASD device.

Default: 2100

DASD_MONTHLY_MAINT_COST

Specifies the average monthly maintenance license fee or cost for the typical DASD device that may be found in your data center. This value, combined with other information about your typical DASD, is used to determine the total cost of ownership of your typical DASD devices in your data center.

Default: 126.00

DASD_FLOORSPACE_AREA_PER_UNIT

Specifies the average area of floor space, in square feet, for a typical DASD device. It should include the total floor space requirements, including any dead floor space that may be required for things such as cooling, access to maintenance panels, and cabling.

Default: 15.00

DASD_KVA_RATING

Specifies the average power consumption rating for your typical DASD devices. The power consumption rating should be listed in the hardware specifications documentation for your devices, or may be obtained by contacting your hardware vendor. The value must represent the power consumption in units of kilovolt-amperes.

Default: 5.0

DASD_KBTU_RATING

Specifies the average heat radiation rating for your typical DASD devices. The heat radiation rating should be listed in the hardware specifications documentation for your devices, or may be obtained by contacting your hardware vendor. The value must represent the heat radiation rating in units of kilo-BTU's (British Thermal Unit).

Default: 16.00

DASD_ANNUAL_COST_EROSION_PERCENTAGE

Specifies the average reduction in the cost of purchasing DASD devices that you have experienced year-over-year, on average, as a percentage.

It is normal for the cost-per-unit for hardware devices to go down each year due to various factors such as newer technologies and vendor competition. While this erosion in the cost of these devices may not be realized every year, on average prices generally do go down over time. This value should reflect the average reduction in the cost of purchasing DASD devices that you have experienced year-over-year on average as a percentage. For example, if the cost of DASD is dropping by 5 percent on average each year, set this to a value of 5.

Default: 20

FLOORSPACE_COST_PER_AREA_UNIT

Specifies the total overhead cost per square foot of maintaining the raised floor area of your data center. Primarily, this cost includes the cost of leasing the data center building space as well as any recurring cost such as building maintenance or the cost of leasing any aspect of the raised floor space itself. This value should represent the cost per square foot of raised floor space.

Default: 200.00

FLOORSPACE_AREA_UNITS

Specifies the unit of measure that is used in defining the floor space cost per unit value. This should be set to SQFT and the floor space values specified should all be based upon using square feet as the unit of measure.

Default: SQFT

ENERGY_COST_PER_KWH

Specifies the average energy cost per kilowatt-hour that you pay for the energy consumed by your data center.

Default: 0.05

CPU_COST_PER_MIPS

Specifies the average cost of purchasing and maintaining your mainframe processors calculated as a cost per million instructions per second (MIPS).

Default: 10500.00

CPU_COST_PER_1000_IO

Specifies the CPU cost per 1000 of input or output operations (I/Os). At some sites, such as those that might involve on-demand capacity or use-based pricing, part of your total cost of ownership or leasing may involve a fee that you are charged that is based upon the number of I/Os that were serviced by your processors. If this is the case for your data center, provide this cost using this variable.

Default: 0.0055

OVERHEAD_COST_PER_TAPE_MOUNT

Specifies the data center operations cost for supporting physical tape mount activities averaged down to a rough cost per physical tape mount. That is, your data center operator overhead for salaries and wages that can be tied to the activities around handling physical tapes, reduced down to an average cost on a per tape mount basis.

Default: 0.45

PERCENT_MOUNTS_FROM_SAME_TAPE

Specifies the percentage of your tape mounts that are for the same tape, where no operator intervention is required.

It is very common for an application to unmount and then immediately remount the same tape. When this occurs, normally there is no need for data center operator interventions which means that the overhead cost per tape mount does not apply. This value should provide a rough estimate of what percentage of your tape mounts are for the same tape and therefore do not require operator intervention.

Default: 10

TAPE_COST_PER_CARTRIDGE

Specifies the average purchase cost of one of your tape cartridges.

Default: 50

TAPE_CARTRIDGE_CAPACITY_IN_MB

Specifies the average capacity of your physical tape media, expressed in megabytes, for your typical tape devices.

Default: 2400

DASD_FREESPACE_PADDING

Specifies how much excess DASD capacity (padding) is optimal in your data center, as a percentage of your overall capacity, to service peak demands.

The amount of DASD space that is required for your day-to-day operations probably fluctuates day-to-day, week-to-week, and month-to-month. At certain times of day, week, month, or year, you probably experience peaks in demand for DASD space. This value should represent how much DASD is needed in order to handle these peak periods and is expressed as a percentage of your overall DASD capacity. For example, if normal demand is for 2000 GB of DASD and during peak periods that demand grows to 2200 GB, then you would need a minimum of 200 GB of excess DASD capacity in order to handle these peak times (or 10 percent padding in your capacity to service these peak demands). This value should be used to identify how much excess DASD capacity is optimal in your data center, as a percentage of your overall capacity, to service the peak demands while not also maintaining far more DASD capacity than is necessary for servicing these peak periods.

Default: 20

DASD_PS_DSORG_FREESPACE

Specifies the percentage of the average over-allocated DASD space that is acceptable as dedicated to over-allocated Physical Sequential (PS) type data sets.

It is normal to over-allocate data sets to avoid space-related abends. Some times, it might not be known exactly how much space a data set ultimately needs. To avoid these space-related abends, users often err on the side of over-allocating data sets far more than is needed. When data sets are allowed to remain greatly over-allocated for long periods of time, this over-allocated DASD space is effectively wasted space that your data center has paid for and is paying for when this DASD could be returned for use, or DASD purchases could be deferred if this wasted space were reclaimed. This value should be set to a percentage that represents the average DASD space that is dedicated to over-allocated PS data sets at any given time as a percentage of the total space allocated to all physical-sequential data sets. For example, if it is acceptable to have as much as 10 percent of the DASD space that is allocated to PS data sets to be unused space (over-allocated), then this value should be set to 10.

Default: 0

DASD_PO_DSORG_FREESPACE

Specifies the percentage of the average over-allocated DASD space that is acceptable as dedicated to over- allocated partitioned (PO) type data sets.

It is normal to over-allocate data sets to avoid space-related abends. Some times, it may not be known exactly how much space a data set ultimately needs. To avoid these space-related abends, users often err on the side of over-allocating data sets far more than is needed. When data sets are allowed to remain greatly over-allocated for long periods of time, this over-allocated DASD space is effectively wasted space that your data center has paid for and is paying for when this DASD could be returned for use, or DASD purchases could be deferred if this wasted space were reclaimed. This value should be set to a percentage that represents the average DASD space that is dedicated to over-allocated PO data sets at any given time as a percentage of the total space allocated to all PO data sets. For example, if it is acceptable to have as much as 20 percent of the DASD space that is allocated to PO data to be unused space (over-allocated), then this value should be set to 20.

Default: 20

DASD_VSAM_DSORG_FREESPACE

Specifies the percentage of the average over-allocated DASD space that is acceptable as dedicated to over- allocated virtual storage access method (VSAM) type data sets.

It is normal to over-allocate data sets to avoid space-related abends. Some times, it may not be known exactly how much space a data set will ultimately need. To avoid these space-related abends, users often err on the side of over-allocating data sets far more than is needed. When data sets are allowed to remain greatly over-allocated for long periods of time, this over-allocated DASD space is effectively wasted space that your data center has paid for and is paying for when this DASD could be returned for use, or DASD purchases could be deferred if this wasted space were reclaimed. This value should be set to a percentage that represents the average DASD space that is dedicated to over-allocated VSAM data sets at any given time as a percentage of the total space allocated to all VSAM data sets. For example, if it is acceptable to have as much as 20 percent of the DASD space that is allocated to VSAM data to be unused space (over-allocated), then this value should be set to 20.

Default: 25

STORAGE_ADMIN_ANNUAL_SALARY

Specifies the average annual salary for a storage administrator for your data center. This value is used to factor in the cost of storage administration activities, where applicable, in analyzing your environment.

Default: 105000

HOURS_PER_DAY_STANDARD_WORK_DAY

Specifies the number of hours that employees work in a normal work day in your data center.

Default: 8

HOURS_PER_WEEK_STANDARD_WORK_WEEK

Specifies the number of hours that employees work per week in a normal work week in your data center.

Default: 40

HOLIDAY_DAYS_PER_YEAR

Specifies the total number of business days, per year, that are not worked due to local holidays.

Holidays can mean different things in different countries and different cultures. For the purposes of this job, a holiday is a day put aside by local customs or laws to commemorate or celebrate a particular event. This value should be set to the total number of business days per year that are not worked due to local holidays.

Default: 12

VACATION_DAYS_PER_YEAR

Specifies the average number of business days, per year, that are not worked due to vacation.

In this case, a vacation day is defined as a day off from work for pleasure, rest, or relaxation other than a holiday. This value should be set to the average number of business days per year that are not worked as an observance of a vacation day or personal day off, if personal days off are not counted as holidays.

Default: 14

HOURS_PER_WEEK_ON_REPORTING

Specifies the average number of hours per week that your storage administrators spend on writing or generating storage management reports.

Default: 50

HOURS_PER_WEEK_MONITORING_DASD_SPACE

Specifies the average number of hours per week that your storage administrators spend on monitoring your DASD volumes and pools to ensure that there is adequate DASD space for your various systems and applications.

Default: 3

HOURS_PER_WEEK_MONITORING_VSAM

Specifies the average number of hours per week that your storage administrators spend on monitoring your production VSAM data sets to determine when a VSAM reorganization might be necessary.

Default: 8

HOURS_PER_WEEK_MONITORING_FRAGMENTATION

Specifies the average number of hours per week that your storage administrators spend on monitoring your DASD volumes for free space fragmentation.

Default: 0

HOURS_PER_WEEK_ON_DASD_SPACE_ABENDS

Specifies the average number of hours per week that your storage administrators spend on addressing any DASD space-related production or test abends.

Default: 2

HOURS_PER_WEEK_OFF_HOURS_ON_STORAGE_ISSUES

Specifies the average number of hours per week outside of the normal work day addressing any storage-related requests or issues.

Default: 3

HOURS_PER_WEEK_ON_STORAGE_HARDWARE

Specifies the average number of hours per week that your storage administrators spend on addressing any hardware-related errors.

Default: 4

HOURS_PER_WEEK_ADHOC_STORAGE_PROCEDURES

Specifies the average number of hours per week that your storage administrators spend addressing spur of the moment (ad-hoc) storage management requests and procedures.

Default: 0

HOURS_PER_WEEK_ADHOC_STORAGE_QUESTIONS

Specifies the average number of hours per week that your storage administrators spend researching and answering ad-hoc storage-related questions.

Default: 14

HOURS_PER_WEEK_HSM_REPORTING

Specifies the average number of hours per week that your storage administrators spend on writing or generating Hierarchical Storage Management (HSM) related reports.

Default: 2

HOURS_PER_WEEK_HSM_ISSUES

Specifies the average number of hours per week that your storage administrators spend on addressing HSM-related problems/issues.

Default: 0

HOURS_PER_WEEK_HSM_AUDITS

Specifies the average number of hours per week that your storage administration staff spends on auditing the HSM archives.

Default: 2

HOURS_PER_WEEK_TAPE_MEDIA_ISSUES

Specifies the average number of hours per week that your storage administration staff spends dealing with physical tape media issues.

Default: 140

HOURS_PER_WEEK_TAPE_REPORTING

Specifies the average number of hours per week that your storage administration staff spends on writing or generating tape media or tape catalog reports.

Default: 21

HOURS_PER_WEEK_TAPE_UTILIZATION

Specifies the average number of hours per week that your storage administration staff spends on issues dealing with physical tape media utilization.

Default: 5

HOURS_PER_WEEK_TAPE_SILO_AUDITS

Specifies the average number of hours per week that your storage administration staff spends on auditing tape silos.

Default: 0

HOURS_PER_WEEK_LOST_TAPES

Specifies the average number of hours per week that your storage administration staff spends on addressing missing or lost tapes.

Default: 4

Chapter 3: Addressing Security Requirements

This section contains the following topics:

Security Requirements

Before you begin the installation process, verify that the CA Chorus for Storage Management installer user ID has the following security privileges defined:

- For UNIX System Services:
  - (Optional) Ability to manipulate zFS data sets. This ability requires UPDATE authority to the appropriate entities within the FSACCESS class. Commented out by default.
    - FSACCESS lets you secure access to a ZFS file system container (that is, a data set). The resource name is the ZFS file system name.
    - For example, if you defined a ZFS file system named OMVS.ZFS.WEBSRV.TOOLS and then created directories U1 and U2 with files in the directories, a resource check for class FSACCESS resource OMVS.ZFS.WEBSRV.TOOLS would occur when a user tries to access a file in directory U1 or U2 in the ZFS file system. For more details, see the applicable security product documentation.
  - A valid OMVS definition and the installer user ID has a valid UID that is *not* UID(0).
  - Superuser authority.
  - READ access to the following resources in the FACILITY class:
    - (Optional) BPX.SUPERUSER
    - (Optional) BPX.FILEATTR.APF
    - (Optional) BPX.FILEATTR.PROGCTL
    - (Optional) BPX.FILEATTR.SHARELIB

- (Optional) BPX.SERVER
- SUPERUSER.FILESYS.PFSCTL profile in UNIXPRIV resource class

Important! All CA Chorus users, including the installer, must have access to the UNIX group ID associated with the CA Chorus started tasks. The default group is CHORGRP.

- For z/OS:
  - Authority to create, update, and execute from the installation data sets and libraries.
  - Authority to execute the commands to manipulate the external security manager (CA ACF2, CA Top Secret, or IBM RACF) database.

In addition, set up *each* storage engine security to let the CA Chorus started task and users log in to the storage engine. Any new or existing storage engine (CA Vantage) user requires CHORTHD, CHORADM IDs in their security profile. All new CA Chorus for Storage Management users need the CA Vantage FACILITY SAMSFAC.

Note: For more information about setting up the storage engine security, see the *CA Vantage Reference Guide*.

APF authorization and other security requirements that must be performed through an external security product are defined during the CA Chorus configuration process, as described in the *CA Chorus Installation Guide*. You complete those tasks during the installation because you need to access various jobs and members from the installation package.

User Requirements

To authorize CA Chorus for Storage Management users to work in CA Chorus, complete the following tasks:

Note: If you have completed a task as part of configuring a different CA Chorus discipline, you do not have to redefine it for this discipline.

1. Authorize CA Chorus for Storage Management users to access USS. Verify that each user has an OMVS segment defined, which includes a home directory, default shell program, UID, and GID or group.

   Note: Confirm that the UID has READ/WRITE access to the HOME directory. For more information about access to USS resources, see the *CA Chorus Site Preparation Guide*. Sample commands for CA ACF2, CA Top Secret, and IBM RACF are provided.

2. Define READ access to the following resources to authorize users to access CA Chorus resources:

   - CHORUS.ROLE.STORAGE—Controls access to the CA Chorus base application for CA Chorus for Storage Management users.
   - (Optional) CHORUS.SETTINGS.KNOWLEDGECENTER—Controls access to the documentation indexing feature of the Knowledge Center component of CA Chorus. Commented out by default. This resource includes access to the Settings option (the wrench icon in the Knowledge Center window) to index documents, URLs, and to upload documents.
   - (Optional) CHORUS.SETTINGS.AUTOREFRESH—Controls access to the auto-refresh option of CA Chorus. Commented out by default.

   Note: CA Chorus defines security resources in class CAMFC. You can remove the CAMFC resource class and user access permissions to restrict access and prevent specific users from logging in to CA Chorus. For more information about managing user access and sample commands to enable this access, see the *CA Chorus Site Preparation Guide*.

3. Configure the following PassTickets (see page 35) using CA ACF2, CA Top Secret, or IBM RACF to authenticate (log in) and authorize users:

   Note: If you have already completed the pass ticket support during storage engine installation, then you do not need to do it again as described here.

   Important! We recommend that an experienced security administrator performs these tasks.

   - Activate PassTicket sign-on in the security system for *each* storage engine started task to accept PassTickets from the following applications:
     - CA Chorus for Storage Management
     - Storage Management interface
   - Authorize the CA Chorus for Storage Management and Storage Management interface applications to generate PassTickets. By default, passtickets are generated automatically for all users. You can change this setting and generate passtickets to give access to individual users.

PassTicket User ID Authentication

PassTickets are required for users to access the z/OS components and products that CA Chorus and its supported disciplines use. A *PassTicket* is a temporary encoded and encrypted substitute for the user password that can be used to access a specific application. The PassTicket must be used within a few minutes of the time it is generated.

Using PassTickets enables the z/OS components and products to authenticate a user ID without sending z/OS passwords through the network. Instead, the user is authenticated after they first log in with a valid z/OS user ID and password. The following process occurs when the user selects a function that accesses a z/OS component:

- The CA Chorus web service calls the z/OS security product to generate a PassTicket for access authorization.
- The PassTicket is sent with the user request to the component, possibly on a different z/OS system.
- The component calls the z/OS security product to authenticate the user using the PassTicket as a password substitute before processing the request.

Configuration information for local and remote systems is provided in PassTicket Configuration for CA Chorus Systems (see page 36).

PassTickets are needed for the following functions:

Note: The CA Chorus for Storage Management and Storage Management interface applications are on the same JBoss server.

- To let the CA Chorus for Storage Management application connect to the storage engine. This requirement is necessary, for example, for the Investigator to get information from the storage engine when storage objects are displayed. For more information, see How to Activate PassTicket Support for CA Chorus for Storage Management (see page 37).
- To let the Storage Management interface application connect to the storage engine. The Storage Management interface is launched by clicking the Manage Storage Resources link in the Quick Links module in CA Chorus. For more information, see How to Activate PassTicket Support for the Storage Management Interface (see page 43).

Review PassTicket Requirements

The CA Chorus server generates PassTickets that permit users to access the various back-end products that the CA Chorus disciplines use. As users access components, PassTickets are generated to validate the requests.

The CA Chorus PassTicket configuration includes the following systems:

- One z/OS system running the JBoss server and the back-end products (like CA Detector, CA Compliance Manager, CA Vantage, and CA NetMaster NM for TCP/IP) that are required for the CA Chorus disciplines on the same system. This type of system is a CA Chorus server system.

- Additional z/OS systems running only the products and components that the CA Chorus disciplines require. This type of system is known as a CA Chorus remote system.

The CA Chorus server system provides the entry point for CA Chorus users. Users can then access all of the CA Chorus remote systems that they have been authorized to use in your network of z/OS systems.

The PassTicket configuration for the security product must be done on each z/OS system that is hosting a component that CA Chorus uses. Configure PassTickets in your z/OS security products to enable the generation and validation of connections that are required for CA Chorus disciplines. If your site meets the following criteria, no additional security setup is required on the remote systems:

- The security products in your z/OS configuration are using a shared security database.
- You want to add one or more remote systems, only the CA Chorus server system setup is required.
- If the requisite products and components exist on a remote system that does not share the security database, additional security setup is required on the remote systems.

How to Activate PassTicket Support for CA Chorus for Storage Management

Use the following jobs to generate PassTickets for the CA Chorus Quick Links storage management interface.

- E4HI095A for CA ACF2 for z/OS
- E4HI095R for IBM RACF for z/OS
- E4HI095T for CA Top Secret for z/OS

PassTickets are required to enable CA Chorus for Storage Management to obtain information from the storage engine. For example, the CA Chorus for Storage Management Investigator obtains storage object information from the storage engine.

Note: If you have installed the storage engine and already completed the PassTicket support, then skip this section.

To activate PassTickets, perform the following steps:

1. Activate PassTickets in the security system (CA ACF2, CA Top Secret, or IBM RACF) for *each* storage engine started task.
2. Authorize the CA Chorus for Storage Management application to generate PassTickets with the security system.

To activate PassTickets, use the following example procedures. The following syntax applies to all the example procedures:

0123456789ABCDEF

Specifies the variable representing the secret session key that is used in the examples.

CA Chorus for Storage Management Application

Specifies the CA Chorus for Storage Management application started task running on the JBoss web server on USS.

Storage Management Interface Application

Specifies the Storage Management interface application started task running on the JBoss web server on USS.

storage engine host

Specifies the backend storage engine application (CA Vantage) started tasks running on the z/OS host.

JBoss_userID* or *uid-of-stc-userid

Specifies the user ID given to the CA Chorus application and Storage Management interface application when JBoss is installed on USS. The *JBoss_userID* and *uid-of-stc-userid* are different terms that are used in the samples and are typically the same ID. This ID is the started task user ID created in ETJI095*x* in *your_chorus_hlq*.CETJJCL. This ID must be able to generate PassTickets for any user. The default is **CHORADM**. One of the functions in the started task is to generate PassTickets for users accessing CA Vantage. CHORADM is the user ID generating the PassTicket. CHORTHD is the user ID used to access CA Vantage during startup. CHORADM generates a PassTicket on behalf of CHORTHD so that CHORTHD can log in to CA Vantage. The ETJI095*x* job is part of the CA Chorus installation, where *x* equals A for CA ACF2, T for CA Top Secret, and R for IBM RACF. For more information, see the *CA Chorus Installation Guide*.

Target_userID

Specifies the user ID of the end user using the CA Chorus for Storage Management or Storage Management interface applications on an internet browser.

Observe the following steps:

- You can set up PassTickets so that a generic user ID can generate the PassTicket. If you set up a generic user ID, you do not need to enable (register) each individual end-user ID to generate the PassTicket.
- When using PassTickets, the storage engine identifies itself to the security systems, CA ACF2, CA Top Secret, or IBM RACF, with the application ID of **VANTAGE**.
- The security definitions for the storage engine started task and the CA Chorus for Storage Management application are on the JBossserver.
- If the security database is not shared between systems, issue the same commands on any other system.

Important! The examples are provided as a guideline. A Security Administrator familiar with PassTickets configuration must manage this configuration.

Sample: Use CA ACF2 to Activate PassTicket Support for CA Chorus for Storage Management

This example shows how to use CA ACF2 to configure PassTickets for connecting to the storage engine to obtain storage information for the CA Chorus for Storage Management application. An experienced security administrator must perform this procedure.

Note: The commands in this section are samples. For detailed information about using these commands, see the CA ACF2 product documentation.

Follow these steps:

1. Define the session key by entering the following commands:

   ```
   SET PROFILE(PTKTDATA) DIVISION(SSIGNON)
   INSERT VANTAGE SSKEY(0123456789ABCDEF) MULT-USE
   F ACF2,REBUILD(PTK),CLASS(P)
   ```

   Note: This example demonstrates a session key value of 16 hexadecimal digits (creating an 8-byte or 64-bit key). Your key must consist of 16 random hexadecimal digits that are different from the values shown in this example. Each application key must be the same on all systems in the configuration and the values must be kept "secret."

   MULT-USE

   Permits reuse of the same PassTicket multiple times.

   The storage engine host is set up to accept PassTickets from CA Chorus for Storage Management.

2. Enable the started task user ID (CHORADM by default) to generate PassTickets for the VANTAGE application:

   ```
   SET RESOURCE(PTK)
   RECKEY IRRPTAUTH ADD(VANTAGE.- UID(uid-of-stc-userid)
   SERVICE(UPDATE,READ) ALLOW)
   ```

   If you do not want to generate PassTickets automatically for all users, give access for individual users with the following command syntax:

   ```
   SET RESOURCE(PTK)
   RECKEY IRRPTAUTH ADD(VANTAGE.userid UID(uid-of-stc-userid)
   SERVICE(UPDATE,READ) ALLOW)
   ```

   The CHORADM started task ID can generate PassTickets for the specified user ID.

   The CA Chorus for Storage Management application (on the JBoss server) is authorized to generate PassTickets.

   PassTickets for connecting to the storage engine to obtain storage information for CA Chorus for Storage Management data are configured.

Sample: Use CA Top Secret to Activate PassTicket Support for CA Chorus for Storage Management

This example shows how to use CA Top Secret to configure PassTickets for connecting to the storage engine to obtain storage information for the CA Chorus for Storage Management application. An experienced security administrator must perform this procedure.

Note: The commands in this section are samples. For detailed information about using these commands, see the CA Top Secret product documentation.

Follow these steps:

1. Update the Resource Descriptor Table (RDT) to define the PTKTDATA class by entering the following command:

   ```
   TSS ADD(RDT) RESCLASS(PTKTDATA) ACLST(ALL,READ,UPDATE) MAXLEN(26)
   ```

2. Update the Node Descriptor Table (NDT) to set the session key:

   ```
   TSS ADD(NDT) PSTKAPPL(VANTAGE) SESSKEY(0123456789ABCDEF)
   SIGNMULTI
   ```

   Note: This example demonstrates a session key value of 16 hexadecimal digits (creating an 8-byte or 64-bit key). Your key must consist of 16 random hexadecimal digits that are different from the values shown in this example. Each application key must be the same on all systems in the configuration and the values must be kept "secret."

   SIGNMULTI

   Permits reuse of the same PassTicket multiple times.

   The storage engine host is set up to accept PassTickets from CA Chorus for Storage Management.

3. Enable UPDATE access to the following resource for the CA Chorus for Storage Management application (on the JBoss server) to generate PassTickets:

   ```
   TSS PER(JBoss-userID) PTKTDATA(IRRPTAUTH.VANTAGE.)
   ACCESS(UPDATE)
   ```

   If you do not want to implement a generic user ID, set permissions to allow users access to the IRRPTAUT resource:

   ```
   TSS PER(JBoss_userID) PTKTDATA(IRRPTAUTH.VANTAGE.Target_userID)
   ACCESS(UPDATE)
   ```

   The CA Chorus for Storage Management application (on the JBoss server) is authorized to generate PassTickets.

   PassTickets for connecting to the storage engine to obtain storage information for the CA Chorus for Storage Management application are configured.

Sample: Use IBM RACF to Activate PassTicket Support for CA Chorus for Storage Management

This example shows how to use IBM RACF to configure PassTickets for connecting to the storage engine to obtain storage information for the CA Chorus for Storage Management application. An experienced security administrator must perform this procedure.

Note: The commands in this section are samples. For detailed information about using these commands, see the IBM RACF product documentation.

Follow these steps:

1. Define the VANTAGE application by entering the following commands:

   ```
   RDEFINE APPL VANTAGE UACC(READ)
   SETROPTS CLASSACT(APPL)
   ```

   Note: If you want to implement a generic user ID, enter the following additional command:

   ```
   SETROPTS GENERIC(PTKTDATA)
   ```

2. Activate the PTKTDATA class if it is not currently active:

   ```
   SETROPTS CLASSACT(PTKTDATA) RACLIST(PTKTDATA)
   ```

3. Define a profile for the application and specify the session keys:
   ```
   RDEFINE PTKTDATA VANTAGE SSIGNON(KEYMASKED(0123456789ABCDEF))
   APPLDATA('NO REPLAY PROTECTION')
   ```

   Note: This example demonstrates a session key value of 16 hexadecimal digits (creating an 8-byte or 64-bit key). Your key must consist of 16 random hexadecimal digits that are different from the values shown in this example. Each application key must be the same on all systems in the configuration and the values must be kept "secret."

   APPLDATA('NO REPLAY PROTECTION')

   Permits reuse of the same PassTicket multiple times.

   The storage engine host is set up to accept PassTickets from CA Chorus for Storage Management.

4. Enable UPDATE access to the following resource for the CA Chorus for Storage Management application on the JBoss server to generate PassTickets:

   ```
   RDEFINE PTKTDATA IRRPTAUTH.VANTAGE.* UACC(NONE)
   PERMIT IRRPTAUTH.VANTAGE.* CLASS(PTKTDATA) ID(JBoss_userID)
   ACCESS(UPDATE)
   ```

   Note: If you do not want to implement a generic user ID, enter the following commands:

   ```
   RDEFINE PTKTDATA IRRPTAUTH.VANTAGE.Target_userID UACC(NONE)
   PERMIT IRRPTAUTH.VANTAGE.Target_userID CLASS(PTKTDATA)
   ID(JBoss_userID) ACCESS(UPDATE)
   ```

5. Refresh the APPL and PTKTDATA classes:

   ```
   SETROPTS RACLIST(APPL) REFRESH
   SETROPTS RACLIST(PTKTDATA) REFRESH
   SETROPTS GENERIC(PTKTDATA) REFRESH
   ```

   The CA Chorus for Storage Management application (on the JBoss server) is authorized to generate PassTickets.

   PassTickets for connecting to the storage engine to obtain storage information for the CA Chorus for Storage Management application are configured.

How to Activate PassTicket Support for the Storage Management Interface

PassTickets are required to enable the Storage Management interface application to use standard JAVA calls to obtain information from *each* storage engine.

The Storage Management interface application is the storage engine web-based user interface installed in the same JBoss as CA Chorus for Storage Management. The Storage Management interface is launched by clicking the Manage Storage Resources link in the Quick Links module. The Storage Management interface is displayed in a separate browser window.

You must authorize the Storage Management interface application to generate PassTickets in your security system.

To activate PassTickets, use the following example procedures. The following syntax applies to all the example procedures:

0123456789ABCDEF

Specifies the variable representing the secret session key that is used in the examples.

CA Chorus for Storage Management Application

Specifies the CA Chorus for Storage Management application started task running on the JBoss web server on USS.

Storage Management Interface Application

Specifies the Storage Management interface application started task running on the JBoss web server on USS.

storage engine host

Specifies the backend storage engine application (CA Vantage) started tasks running on the z/OS host.

JBoss_userID* or *uid-of-stc-userid

Specifies the user ID given to the CA Chorus application and Storage Management interface application when JBoss is installed on USS. The *JBoss_userID* and *uid-of-stc-userid* are different terms that are used in the samples and are typically the same ID. This ID is the started task user ID created in ETJI095*x* in *your_chorus_hlq*.CETJJCL. This ID must be able to generate PassTickets for any user. The default is **CHORADM**. One of the functions in the started task is to generate PassTickets for users accessing CA Vantage. CHORADM is the user ID generating the PassTicket. CHORTHD is the user ID used to access CA Vantage during startup. CHORADM generates a PassTicket on behalf of CHORTHD so that CHORTHD can log in to CA Vantage. The ETJI095*x* job is part of the CA Chorus installation, where *x* equals A for CA ACF2, T for CA Top Secret, and R for IBM RACF. For more information, see the *CA Chorus Installation Guide*.

Target_userID

Specifies the user ID of the end user using the CA Chorus for Storage Management or Storage Management interface applications on an internet browser.

Observe the following steps:

- You can set up PassTickets so that a generic user ID can generate the PassTicket. If you set up a generic user ID, you do not need to enable (register) each individual end-user ID to generate the PassTicket.
- When using PassTickets, the storage engine identifies itself to the security systems, CA ACF2, CA Top Secret, or IBM RACF, with the application ID of **VANTAGE**.
- The security definitions for the storage engine started task and the Storage Management interface application are on the JBoss server.
- If the security database is not shared between systems, issue the same commands on any other system.

Important! The examples are provided as a guideline. A Security Administrator familiar with PassTickets configuration must manage this configuration.

Sample: Use CA ACF2 to Activate PassTicket Support for the Storage Management Interface

This example shows how to use CA ACF2 to configure PassTickets for connecting to the storage engine to obtain storage information for the Storage Management interface. An experienced security administrator must perform this procedure.

Note: The commands in this section are samples. For detailed information about using these commands, see the CA ACF2 product documentation.

Follow these steps:

1. Define the session key by entering the following commands:

```
SET PROFILE(PTKTDATA) DIVISION(SSIGNON)
INSERT VANTAGE SSKEY(0123456789ABCDEF) MULT-USE
F ACF2,REBUILD(PTK),CLASS(P)
```

Note: This example demonstrates a session key value of 16 hexadecimal digits (creating an 8-byte or 64-bit key). Your key must consist of 16 random hexadecimal digits that are different from the values shown in this example. Each application key must be the same on all systems in the configuration and the values must be kept "secret."

MULT-USE

Permits reuse of the same PassTicket multiple times.

Perform this step for each storage engine host if you have a multiple host environment. All your storage engine hosts are set up to accept PassTickets from the Storage Management interface.

2. Enable the started task user ID (CHORADM by default) to generate PassTickets for the VANTAGE application:

   ```
   SET RESOURCE(PTK)
   RECKEY IRRPTAUTH ADD(VANTAGE.- UID(uid-of-stc-userid)
   SERVICE(UPDATE,READ) ALLOW)
   ```

   If you do not want to generate PassTickets automatically for all users, give access for individual users with the following command syntax:

   ```
   SET RESOURCE(PTK)
   RECKEY IRRPTAUTH ADD(VANTAGE.userid UID(uid-of-stc-userid)
   SERVICE(UPDATE,READ) ALLOW)
   ```

   The CHORADM started task ID can generate PassTickets for the specified user ID.

   The Storage Management interface application (on the JBoss server) is authorized to generate PassTickets.

   PassTickets for connecting to the storage engines to obtain storage information for Storage Management interface are configured.

Sample: Use CA Top Secret to Activate PassTicket Support for the Storage Management Interface

This example shows how to use CA Top Secret to configure PassTickets for connecting to the storage engine to obtain storage information for the Storage Management interface. An experienced security administrator must perform this procedure.

Note: The commands in this section are samples. For detailed information about using these commands, see the CA Top Secret product documentation.

Follow these steps:

1. Update the RDT to define the PTKTDATA class by entering the following command:

   ```
   TSS ADD(RDT) RESCLASS(PTKTDATA) ACLST(ALL,READ,UPDATE) MAXLEN(26)
   ```

2. Update the NDT to set the session key:

   ```
   TSS ADD(NDT) PSTKAPPL(VANTAGE) SESSKEY(0123456789ABCDEF)
   SIGNMULTI
   ```

   Note: This example demonstrates a session key value of 16 hexadecimal digits (creating an 8-byte or 64-bit key). Your key must consist of 16 random hexadecimal digits that are different from the values shown in this example. Each application key must be the same on all systems in the configuration and the values must be kept "secret."

   SIGNMULTI

   Permits reuse of the same PassTicket multiple times.

   Perform the steps for each storage engine host if you have a multiple host environment. All your storage engine hosts are set up to accept PassTickets from the Storage Management interface.

3. Enable UPDATE access to the following resource for the Storage Management interface application (on the JBoss server) to generate PassTickets:

   ```
   TSS PER(JBoss-userID) PTKTDATA(IRRPTAUTH.VANTAGE.)
   ACCESS(UPDATE)
   ```

   If you do not want to implement a generic user ID, enter the following commands:

   a. Update a Division/Department Accessor ID (ACID):

      ```
      TSS ADD(tssdept) PTKTDATA(IRRPTAUTH)
      ```

      The IRRPTAUTH resource in the PTKTDATA class is owned.

   b. Set permissions to allow users access to the IRRPTAUTH resource:

      ```
      TSS PER(JBoss_userID)
      PTKTDATA(IRRPTAUTH.VANTAGE.Target_userID) ACCESS(UPDATE)
      ```

   The Storage Management interface application (on the JBoss server) is authorized to generate PassTickets.

   PassTickets for connecting to the storage engines to obtain storage information for the Storage Management interface are configured.

Sample: Use IBM RACF to Activate PassTicket Support for the Storage Management Interface

This example shows how to use IBM RACF to configure PassTickets for connecting to the storage engine to obtain storage information for the Storage Management interface. An experienced security administrator must perform this procedure.

Note: The commands in this section are samples. For detailed information about using these commands, see the IBM RACF product documentation.

Follow these steps:

1. Define the VANTAGE application by entering the following commands:

   ```
   RDEFINE APPL VANTAGE UACC(READ)
   SETROPTS CLASSACT(APPL)
   ```

   Note: If you want to implement a generic user ID, enter the following additional command:

   ```
   SETROPTS GENERIC(PTKTDATA)
   ```

2. Activate the PTKTDATA class if it is not currently active:

   ```
   SETROPTS CLASSACT(PTKTDATA) RACLIST(PTKTDATA)
   ```

3. Define a profile for the application and specify the session keys:

   ```
   RDEFINE PTKTDATA VANTAGE SSIGNON(KEYMASKED(0123456789ABCDEF))
   APPLDATA('NO REPLAY PROTECTION')
   ```

   Note: This example demonstrates a session key value of 16 hexadecimal digits (creating an 8-byte or 64-bit key). Your key must consist of 16 random hexadecimal digits that are different from the values shown in this example. Each application key must be the same on all systems in the configuration and the values must be kept "secret."

   APPLDATA('NO REPLAY PROTECTION')

   Permits reuse of the same PassTicket multiple times.

   Perform the steps for each storage engine host if you have a multiple host environment. All your storage engine hosts are set up to accept PassTickets from the Storage Management interface.

4. Enable UPDATE access to the following resource for the Storage Management interface to generate PassTickets:

   IRRPTAUTH.VANTAGE.*Target_userID*

   Target_userID

   Specifies the target user ID resource.

   a. Enter the following commands:

      ```
      RDEFINE PTKTDATA IRRPTAUTH.VANTAGE.Target_userID UACC(NONE)
      PERMIT IRRPTAUTH.VANTAGE.Target_userID CLASS(PTKTDATA)
      ID(JBoss_userID) ACCESS(UPDATE)
      ```

      Note: To implement a generic user ID, enter the following commands:

      ```
      RDEFINE PTKTDATA IRRPTAUTH.VANTAGE.* UACC(NONE)
      PERMIT IRRPTAUTH.VANTAGE.* CLASS(PTKTDATA) ID(JBoss_userID)
      ACCESS(UPDATE)
      ```

   b. Refresh the APPL and PTKTDATA classes:

      ```
      SETROPTS RACLIST(APPL) REFRESH
      SETROPTS RACLIST(PTKTDATA) REFRESH
      SETROPTS GENERIC(PTKTDATA) REFRESH
      ```

   The Storage Management interface application (on the JBoss server) is authorized to generate PassTickets.

   PassTickets for connecting to the storage engine to obtain storage information for the Storage Management interface are configured.

Chapter 4: Installing and Configuring the Storage Engine

Follow the instructions in this section to install and configure the backend storage engine for CA Chorus for Storage Management.

Follow these steps:

1. Install the storage engine by following the instructions in the *CA Vantage Installation Guide*.

   Note: If you have installed the storage engine for CA Vantage because you are already a CA Vantage customer or if you have installed CA GMI, skip this step.

2. Configure the storage engine, the various components, and additional configuration for CA Chorus for Storage Management in accordance with the instructions in the chapter "Configuring for CA Chorus for Storage Management" in the *CA Vantage Configuration Guide*.

   Note: If you are an existing CA Vantage customer and have already configured the base system, CA GMI enabled products, and the *CA Vantage* Automation Option, there are additional items to configure so that your existing system is configured for CA Chorus for Storage Management. For more information, see the chapter "Configuring for CA Chorus for Storage Management" in the *CA Vantage Configuration Guide.*

   The storage engine is Installed and configured for CA Chorus for Storage Management.