CA Chorus™ for Storage Management

User Guide

Version 03.0.00, Second Edition

CA Technologies Product References

This document references the following CA Technologies products:

- CA 1® Tape Management (CA 1)
- CA ACF2™ for z/OS (CA ACF2)
- CA Allocate™ DASD Space and Placement (CA Allocate)
- CA ASTEX™ Performance (CA ASTEX)
- CA Datacom®/AD (CA Datacom/AD)
- CA Datacom®/DB (CA Datacom/DB)
- CA Disk™ Backup and Restore (CA Disk)
- CA Encryption Key Manager
- CA IDMS™/DB (CA IDMS/DB)
- CA LDAP Server for z/OS (CA LDAP Server)
- 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)
- 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 Tape Encryption
- CA TLMS® Tape Management (CA TLMS)
- CA Top Secret® for z/OS (CA Top Secret)
- CA Vantage™ Storage Resource Manager (CA Vantage)
- CA Vtape™ Virtual Tape System (CA Vtape)

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 documentation updates have been made since the last release of this documentation:

November2013

- You can now open the Investigator from Alerts. This option is added to the View Alerts (see page 61) topic. (PTF RO65056 required.)
- Legal Notices (see page 2)—Updated to reflect public documentation legal disclaimer.
- The storage engine installation and configuration for CA Chorus for Storage Management is simplified for new customers. When new customers run the CHRSGMI job, the log scripts that provide the TSF object data are created. The topics How to Collect Storage Metrics for TSF (see page 95) and Create a Storage Engine Log Script to Generate Storage Metrics for Time Series (see page 96) are updated accordingly. (Backend PTF RO64747required.)
- 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 topics How to Analyze and Lower DASD Costs (see page 20) and Cost Analysis (see page 72) are updated accordingly. (PTF RO63077 required.)
- New storage object tree category added, CA SRM Open Systems Objects (see page 71). (Backend PTF RO60081 required.)
- You can now launch the storage engine Host Configuration Client (Config Client) from the Storage Management interface (see page 15) available in the Quick Links Module. Use the Config Client to set or change storage engine system parameters. (PTF RO62885 required.)

June2013

- Describe the two ways to display object performance (TSF) object data (see page 89):
  - Side-by-side charts in the Investigator using Statistics object data (see page 89).
  - In the Metrics panel (see page 93). Configure the Metrics panel to show specific TSF object metrics in the panel, set Metrics thresholds, investigate TSF data, and create your own Metrics dashboard.
- Cost Analysis object names changed from Analysis to Analyses. Additional cosmetic changes to the Cost Analysis objects are reflected in the sections Cost Analysis (see page 72) and the work flow scenario How to Analyze and Lower DASD Costs.
- The word "Storage" is added to the topic title "Perform Storage Actions on Object Data in the Investigator (see page 83)".

- The section Alerts for Storage Conditions (see page 99) is changed to include Policies Module Alerts. A new section Policies Module Alerts (see page 99) is also added.
- New workspace sample section Sample CA Chorus for Storage Management Workspace (see page 17) added.
- New work flow scenario How to Analyze and Lower DASD Costs added.
- The Visualizer in the Investigator is now called the Topology Viewer. All references to the *Visualizer* are changed to refer to the *Topology Viewer*.
- The My Workspace tab in the Investigator is now called My Dashboard.
- Changed product name in Architecture (see page 10) section graphics.
- Workflows changed to Scenarios, see chapter "Introduction".
- New section Cost Analysis (see page 72) added.
- All relevant topics are updated to reflect the new Storage Investigator object tree.

Contents

Chapter 1: Introduction

Architecture

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:

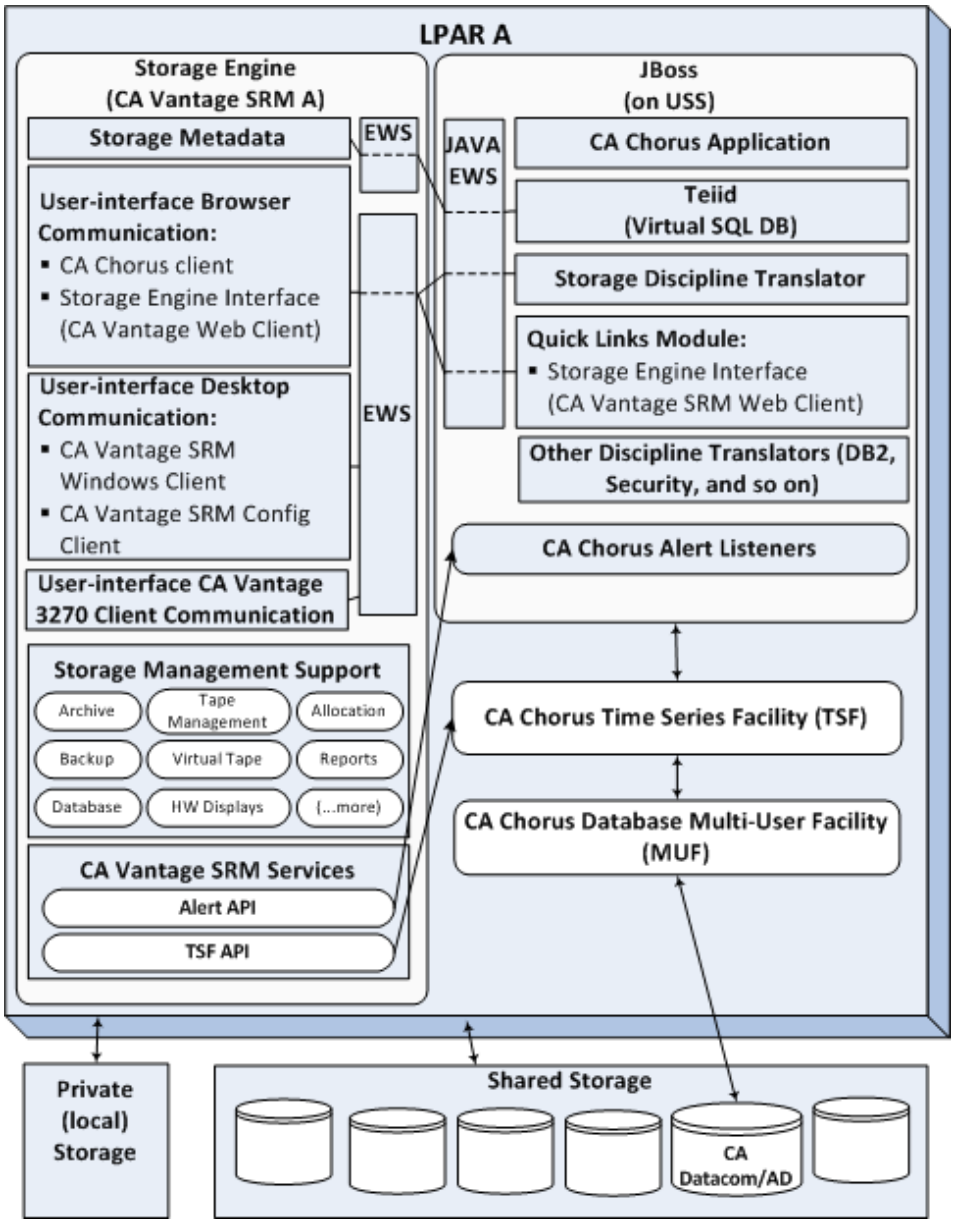

- Multiple storage engine architecture and setup of CA Chorus for Storage Management:

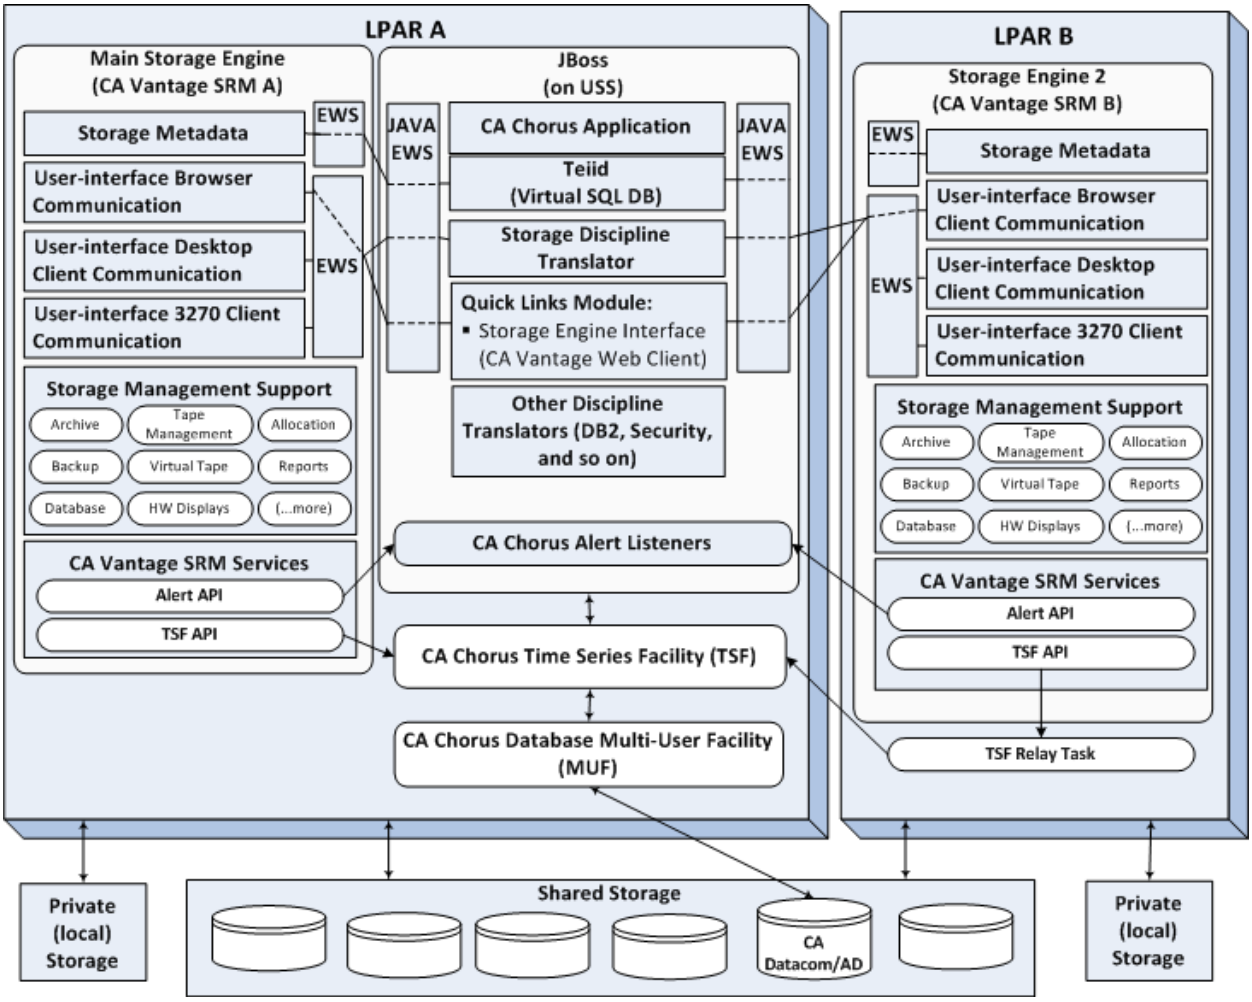

Observe the following points:

- In the single storage engine setup, it is common to have the storage engine and JBoss 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:

CA Chorus Application

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.

  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.

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.

Storage Knowledge Center Best Practices

The Knowledge Center is the repository for all CA Chorus documentation. The Knowledge Center includes online help and guides from CA, user-generated documentation, and links to third-party content. Links to relevant topics appear in the Knowledge Center window when you click the online help icon or by searching. When you request online help, the search engine finds topics that are focused on the task you are performing. The engine also searches based on your location in the interface. This information appears in the Knowledge Center window and is updated whenever you refresh the window or click the online help icon.

We recommend that you add storage-specific documentation to your Knowledge Center. For example, you could add the documentation that is associated with a specific release of IBM z/OS. This best practice helps ensure that your storage administrators have accurate and current storage information.

Note: For the steps to index content, see the *CA Chorus Product Guide*.

We recommend that you add the following content to your Knowledge Center:

- CA Customer Value Program reports. For more information, see http://ca.com/support.
- *IBM Introduction and Release Guide* for your supported z/OS versions.
- IBM z/OS glossaries.
- Storage best practices for your site.

Note: Access to the Knowledge Center configuration is restricted. For details about defining Knowledge Center access permissions, see the *CA Chorus Administration Guide*. To request access, contact your security administrator.

We also recommend that you configure your search settings so that only CA back-end product content specific to your discipline appears in Knowledge Center results. Implementing this recommendation can improve the relevance of search results. For the configuration steps, see the *CA Chorus Product Guide*.

Sample CA Chorus for Storage Management Workspace

With this topic, we demonstrate the possibilities for workspace customization.

The following list details key product touch-point examples that demonstrate why you would use each component for your discipline. CA Chorus offers an intuitive user interface; however, as you review this list and look at the user-interface, if you are unsure how to proceed, see the *CA Chorus Product Guide*. Additionally, you can also click the question mark icon from within a module for procedures and concept help topics.

Note: You can also simply apply a dashboard that a peer has shared; however, for the purposes of this example, we build new ones.

Dashboards

Create three dashboards with the following names:

Performance

You can use this dashboard to contain the Investigator and Alerts modules.

Investigator

Add the Investigator to your *Performance* dashboard.

The Investigator lets you customize details such that you can quickly see site-specific details.

In the Storage object tree, select the object Storage Groups. This object is located in the Volumes & Utilization, Storage Groups, and Space and Other Attributes directory of the Storage object tree. Select the Totals row by clicking anywhere in the row besides the selection box. When you do so, the Details pane opens at the bottom of the field with several default tabs.

Your focus is mainly on the % Alloc, % Free, Volume Count, Capacity Bytes, Allocated Bytes, and Free Bytes fields. Create a new tab and insert these fields into this tab. Doing so lets you customize the view and find your key information quickly.

Storage Group name TOTAL – Lets you see the field data that is shown in this detail tab and is the total for all your online Storage Groups.

% Alloc – Lets you see the total percent amount of allocated space in all your online storage groups.

% Free – Lets you see the total percent amount of free space in all your online storage groups.

Volume Count – Lets you get a quick overview of the total number of volumes in all your online storage groups.

Capacity Bytes – Lets you see the total bytes in all your online storage groups.

Allocated Bytes – Lets you see the total bytes allocated in all your online storage groups.

Free Bytes – Lets you see the total free bytes in all your online storage groups.

Save your new tab using the Save View icon in the Detail pane and give the view a name, for example Online SG Totals. The next time you open the Storage Groups object and select the TOTAL row, you can select this tab and get a quick overview of your total online storage groups.

Alerts Module

Add the Alerts module to your *Performance* dashboard. Follow the wizard to add the Alerts module. Use the label option in the wizard to name it. For example, give it the label Storage Alerts.

Note: Storage alerts can appear in the Alerts module if you have created CA Chorus Policy Light alerts (see page 99) or storage engine alerts (see page 99) and the criteria in the alerts is met.

Metrics

You can use this dashboard to contain metrics graphs from the Metrics panel.

After you configure the Metrics module, you can add storage metrics to the *Metrics* dashboard. Dashboards can have one or many metrics added to them, or you can have more than one *Metrics* dashboard.

Add the following Storage metrics to the Metrics panel and then add them to your *Metrics* dashboard. Doing so lets you see at a glance the status of these three metrics.

- **Storage Groups** – Lets you monitor the ratio between total storage capacity bytes, allocated bytes, and free bytes. This lets you evaluate the space usage compared to the free bytes.
- **Catalogs** - Lets you monitor whether there is enough free space on the volume to allow its catalog to get an extent if needed.

- **Catalog Cache** - Lets you monitor proper catalog cache usage. High Hit% values indicate proper usage of cache.

Use the TextBox module to add labels and more detail above each metric. For example, above the Catalogs metric graph, you could add a TextBox stating your site policy of how much free space on the volumes you should allow.

Tools

As a storage administrator, you want quick access to tools that can help you manage your storage resources. Add the Quick Links module to your *Tools* dashboard. Doing so gives you quick access to the Storage Management interface.

As a storage administrator you can use the Storage Management interface to create storage engine Joined Objects and schedule periodic output reports to servers and send them as attachments to emails. For more information about the Storage Management interface, see its online help.

Knowledge Center

Open the Knowledge Center, and click Advanced Search. Under Show Results From, select MVS/Quick Ref, User Documentation (see page 16), CA Chorus Documentation, CA Vantage Documentation on CA Support, and *Always search only these sources*. Doing so improves search result relevance as you search only the data sources related to your discipline.

Policy Lights

Add an existing policy to this module and add it to your Performance dashboard. The example policy (see page 99) provided in this guide indicates when you have more than 10 data sets online that have been idle for more than 365 days. Consider moving online data sets that are unused for a long period of time to cheaper storage media.

Dashboard Sharing

Now that you have built dashboards by function within your discipline, share these dashboards with your peers to help them start working quickly. To do so, simply right-click the dashboard and follow the prompts.

Summary

This sample shows you one of the many ways you can customize your workspace according to your CA Chorus discipline. This customization can improve productivity and response time for customer issues. Use this sample as a starting point to explore how you can customize the workspace for your user- and site-specific needs.

How to Analyze and Lower DASD Costs

As a storage administrator, you oversee the costs of your storage environment and manage DASD (direct access storage devices) costs. CA Chorus for Storage Management lets you efficiently analyze your storage environment and manage the cost of your DASD.

To demonstrate the analytical tools, consider the Current Analyses object. Each row in the object presents an analytical scenario. Some of the Current Analyses object scenarios are:

- What are the costs of maintaining any excess DASD capacity?
- What are my annual Storage Administration costs?
- What is the overall cost wasted or unnecessary DASD?
- What are the costs of maintaining data sets not referenced in the past year?
- What are my total costs associated with excess and over-allocated DASD?

The questions in this scenario are based upon the current DASD capacity, free space characteristic data, and peak demand targets (DASD free space padding). The DASD free space padding value is expressed as a percentage of the overall capacity of online DASD and represents the desired level of excess capacity needed to handle peak demand, as well as short-term growth in DASD requirements.

The analytical scenarios not only show current DASD costs, but they also show the projected costs for the next two years. The current (first) year cost is calculated as a number of gigabytes of unallocated DASD over the desired free space padding multiplied by the total cost per gigabyte of DASD. The second and third year calculations represent the potential cost savings that might occur if unused idle space is reduced.

The values displayed are calculations based on the values shown in the Customer Site Costs Data object and data provided by the storage engine. The Customer Site Costs Data object is located in the Cost Analysis directory of the Storage object tree.

Note: For more information about changing values shown in the Customer Site Costs Data object, see the *CA Chorus Administration Guid*e.

The following are a few examples of the variables used in the analytical scenarios for this object:

- Floor space cost per square foot
- Cost per kWh for electricity
- Typical cost per GB for DASD
- Average annual erosion of DASD per GB
- Designated Free Space Padding

The following is the storage resource data provided by the storage engine:

- Total DASD capacity (in gigabytes) of all online DASD.
- Total allocated space (in gigabytes) of all online DASD.

The following analyses are performed in this scenario:

- Convert your monthly DASD maintenance cost into an annual cost.
- Determine DASD floor space cost.
- Determine the true cost per gigabyte you are paying for DASD.

Note: When the Current Analyses object is displayed in the Investigator, click the (?) help button in the object table toolbar to display the Knowledge Center. Click the link to open the Current Analyses object help. The object provides a description of the object, and its Actions, Details tabs, and fields (columns).

The following illustration shows how a storage administrator analyzes and manages the cost of DASD:

How to Analyze and Lower DASD Costs

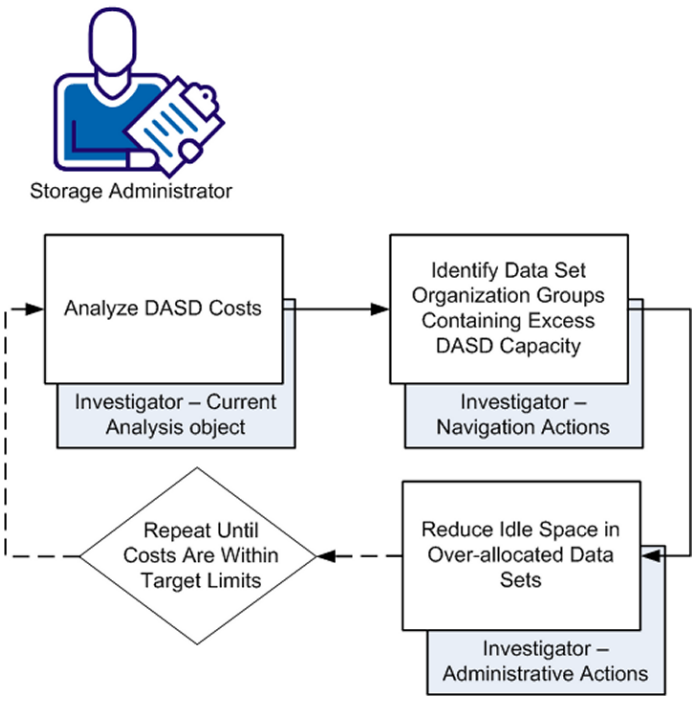

To analyze and manage DASD costs, do the following:

1. Analyze DASD costs (see page 22).
2. Identify data set organization groups containing excess DASD capacity (see page 26).
3. Reduce idle space in over-allocated data sets (see page 27).
4. Repeat steps 1 through 3 until DASD costs are within your target limits.

Analyze DASD Costs

This topic shows how a storage administrator can analyze DASD to determine if current costs are within target limits. DASD that sits above target limits can mean increased costs for your site.

Follow these steps:

1. Log in to CA Chorus.
2. Add the Investigator module to your dashboard from the Module Library, and click Start New Investigation.

   The Investigator opens.
3. Select Storage from the discipline drop-down list.
4. Expand the Cost Analysis directory in the Storage object tree, and open the Current Analyses object.

   Each row poses a question about your storage environment. We recommend that you review those with a severity code. In the following example, the first row shows a severity one, and the second row shows a severity two. The severity column is blank in the other rows because the scenario results are within accepted (target) limits. The following illustration is an example of the Current Analyses object table for a single host:

Notes	Severity	Final Result	Description	Host Name
	1	$44,451.53	What are the costs of DASD space wasted by over-allocated data sets?	S0000 on P111
	2	$8,032,478.00	What are the costs of maintaining any excess DASD capacity?	S0000 on P111
		$523,224.00	What are my annual Storage Administration costs?	********
		$226,364.60	What are the costs of maintaining data sets not referenced in the past year?	S0000 on P111
		$343,426.44	What is the overall cost wasted or unnecessary DASD?	S0000 on P111

5. Click anywhere in the row besides the check-box in the severity one row.

   The Actions pane on the right-side of the table and the Details pane at the bottom of the Investigator window expand, as shown in the following example.

   Note: The Details tab in the Details pane is selected (by default).

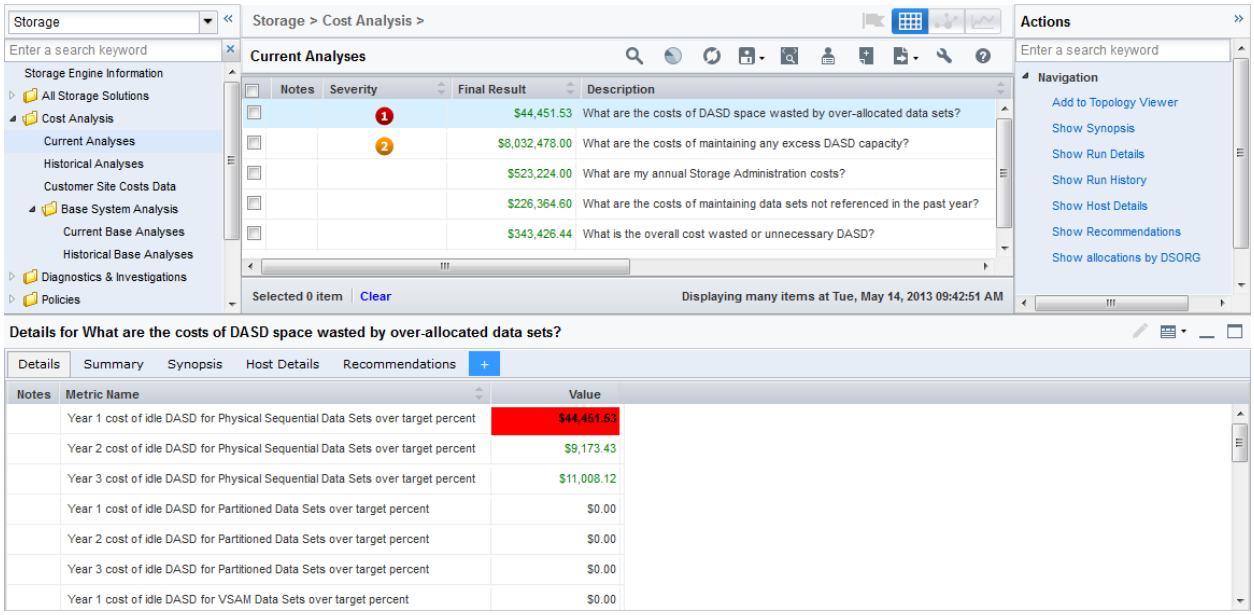

6. Click the action *Show allocations by DSORG*, to open the Alloc/Idle Bytes Dsorg (DTOCDSOR) summary object.

   The Alloc/Idle Bytes Dsorg (DTOCDSOR) summary object opens. This object contains the storage resource data used in the analytical calculations. This object summarizes data sets according to their DSORG (data set organization group) field. The number of data sets for each data set organization is shown, with other summary statistics about the allocated space, idle space, and number of extents for each data set organization group.

7. Customize the *Show allocations by DSORG* table view by completing the following steps:

   a. Click the Customize your Tabular Data View (Wrench) icon in the table toolbar, to display the Select and Reorder Columns dialog.

   b. Click the column title Tot Idle Bytes in the All Selected Columns pane and click the up-arrow continuously until the column title is situated under the column title Nbr Of DSNs.

      This step changes the column sequence in the object display so that you can see the number of data sets right after the total amount of idle bytes.

c. Verify that the check-box *Always use this column configuration* is selected, and click Save.

   The Show allocations by DSORG table is updated and the Tot Idle Bytes column is now displayed after the Nbr Of DSNs column in the table.

d. Click the Tot Idle Bytes column heading cell two times, to sort the table in descending order by total number of idle bytes.

   The table now shows the data set organization types with the highest number of total idle bytes at the top of the list.

This customization helps you identify which data set organization groups have over-allocated data sets (the groups with highest amount of idle space), which you can reduce.

8. Click the Current Analyses object in the object tree to open it.

9. Click anywhere in the row besides the check-box in the row *What are the costs of DASD space wasted by over-allocated data sets?*.

   The Actions pane expands and the Details pane at the bottom of the Investigator window is populated with the Details tab selected.

   The information displayed in the Details tab is the same information that is displayed when you select the Navigation Action *Show Run Details*. Both options show the total cost by year, for the different data set organization groups.

   - Year 1 represents the costs of your current over target percent.
   - Year 2 is calculated using the costs of your current over target percent and an erosion factor.
   - Year 3 is calculated using the costs of your second year and the same erosion factor.

   Note: The erosion factor used is the percent value of the DASD_ANNUAL_COST_EROSION_PERCENTAGE variable. You can see this value in the Customer Site Costs Data object.

   The cells containing values that are above your accepted target value have colored backgrounds.

The following illustration shows the Details pane with the Details tab selected.

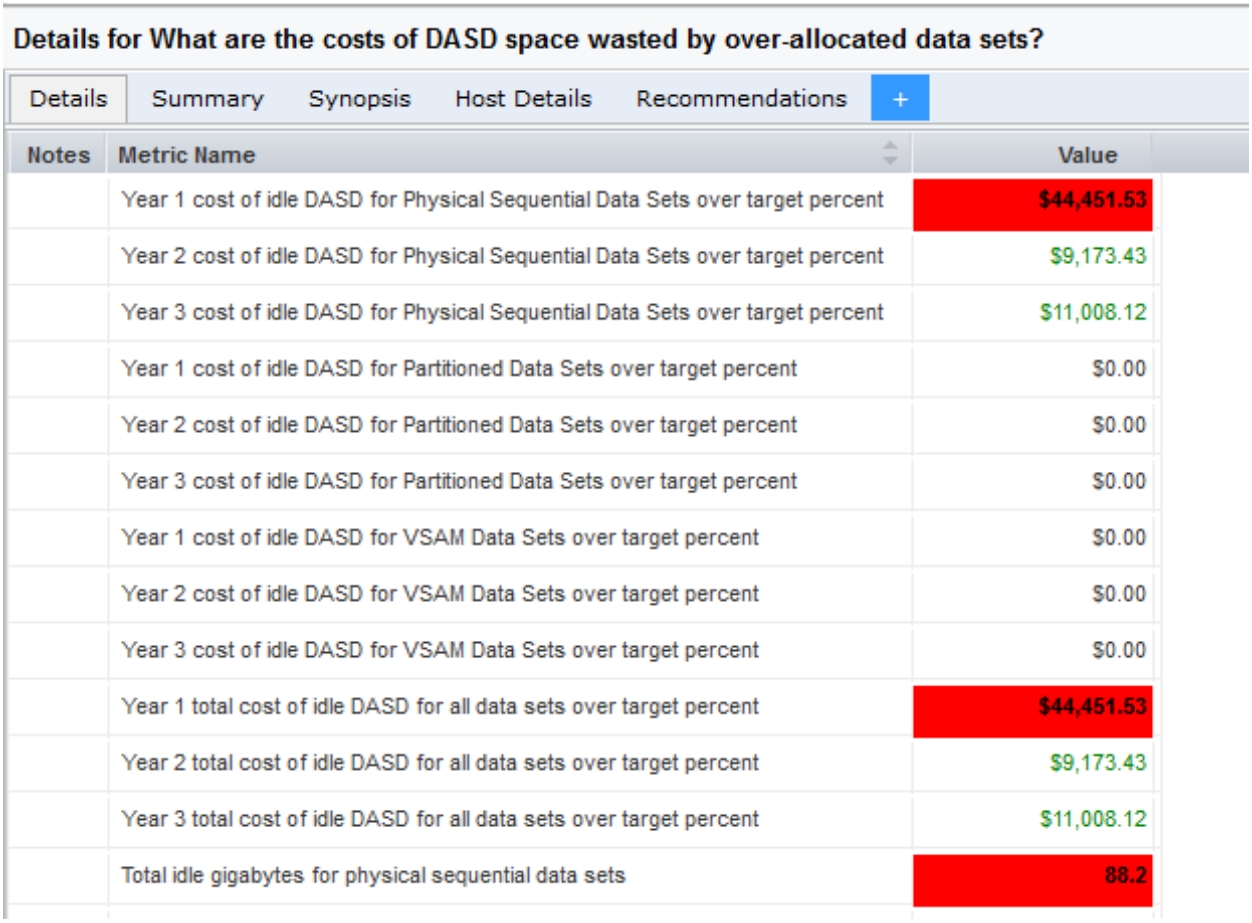

Details for What are the costs of DASD space wasted by over-allocated data sets?

Details | Summary | Synopsis | Host Details | Recommendations | +

Notes	Metric Name	Value
	Year 1 cost of idle DASD for Physical Sequential Data Sets over target percent	$44,451.53
	Year 2 cost of idle DASD for Physical Sequential Data Sets over target percent	$9,173.43
	Year 3 cost of idle DASD for Physical Sequential Data Sets over target percent	$11,008.12
	Year 1 cost of idle DASD for Partitioned Data Sets over target percent	$0.00
	Year 2 cost of idle DASD for Partitioned Data Sets over target percent	$0.00
	Year 3 cost of idle DASD for Partitioned Data Sets over target percent	$0.00
	Year 1 cost of idle DASD for VSAM Data Sets over target percent	$0.00
	Year 2 cost of idle DASD for VSAM Data Sets over target percent	$0.00
	Year 3 cost of idle DASD for VSAM Data Sets over target percent	$0.00
	Year 1 total cost of idle DASD for all data sets over target percent	$44,451.53
	Year 2 total cost of idle DASD for all data sets over target percent	$9,173.43
	Year 3 total cost of idle DASD for all data sets over target percent	$11,008.12
	Total idle gigabytes for physical sequential data sets	88.2

10. Complete the analysis by selecting the following Details tabs:

    Summary

    Displays the details of the Costs of DASD space wasted by over-allocated data sets scenario details in a *form* format.

    Synopsis

    Shows how the value of the scenario is calculated.

11. Click anywhere in the row besides the check-box in the *What are the costs of maintaining any excess DASD capacity?*.

    Note: The contents of the Actions pane and the Details pane pertain to the scenario selected by clicking anywhere in the row besides the check-box of the scenario row. Selecting the check box on a row does not expand the Details pane.

    The Actions pane expands and the Details pane at the bottom of the Investigator window is populated with the Details tab being selected (by default).

12. Review the costs associated with maintaining the excess DASD capacity in the Details pane by clicking the various tabs available. In our example where the scenario is a severity two, the tabs provide the following information:

    Details

    Details of the values used to perform the cost analysis. It also provides costs for the current year and the next two years (projected costs).

    Summary

    Provides an analysis of the cost of excess DASD capacity over and above what is required for peak demand.

    Synopsis

    Provides an explanation of the analysis.

    Recommendations

    Provides a recommendation of how to reduce costs.

    In this example, the information in the analysis indicates you are above your target limits and you need to reduce your DASD related costs. You now have an overview of the actual costs of your excess DASD capacity. You also have an overview of the costs of maintaining the excess DASD capacity and a recommendation of how to reduce it.

Identify Data Set Organization Groups Containing Excess DASD Capacity

As a storage administrator, you want to identify the data set organization groups containing excess DASD capacity so you can reduce excess DASD capacity and your costs of wasted unused DASD space. Doing so can save your company money.

Follow these steps:

1. Expand the Cost Analysis, Base System Analysis directory in the Storage object tree of the Investigator, and open the Current Base Analyses object.

   This object shows quantitative values for the same analytical scenarios displayed in the Current Analyses object which primarily shows financial values (see page 22).

2. Click anywhere in the row besides the check-box in the row *How much excess unused DASD capacity is being maintained?* and click the Details tab.

   The details of the Final Result value calculation are displayed. The Final Result value is the total amount of excess DASD capacity in Gigabytes.

3. Click the Synopsis tab to see how the value is calculated.

   The synopsis tells you that this scenario analyzes the current DASD free-space relative to the overall DASD and compares that to the DASD free-space padding goal for your local site (DASD free-space padding).

4. Click anywhere in the row besides the check-box in the severity two row, to display actions and details for it.
5. Review the values in the Details tab to see the details of the Final Result value calculation. Items that need your attention have colored backgrounds, where a red background highlights the items that have the highest values outside your target limits.
6. Click the Recommendations tab to see the recommendation of what you should do to improve your storage resource allocations.

   If the scenario *How much excess unused DASD capacity is being maintained?* has a severity one grading, the recommendation suggests you free-up over allocated space in one or more of your data set organization groups.
7. (Optional) Investigate historical analysis by reviewing the Historical Analyses and Historical Base Analyses objects. These objects show the history of the same scenarios of the Current Analyses and Current Base Analyses. These objects can help you identify your historical costs and amounts of wasted DASD space wasted by over-allocated data sets.
8. Select the Navigational action *Show Synopsis* in the Actions pane or the *Synopsis* tab at the bottom of the table pane, to see a summary of the major points of the selected analytical scenarios.

   You have identified the data set organization groups that contain excess DASD capacity so you know the groups from which you can reduce unused space.

Reduce Idle Space in Over-allocated Data Sets

As a storage administrator, you want to reduce idle space in data sets to reduce your costs of excess wasted DASD space.

In this example we show you how to reduce idle space in physical sequential (PS) and partitioned (PO) data set organization groups.

Follow these steps:

1. Expand the Diagnostics & Investigations, Data Set Management, All directory in the Storage object tree in the Investigator, and open the Data Sets for System (All) object.

2. Create a filter to display PS data sets with a large amount of idle space and click Search.

   For example, create the following filter statement:

   ```
   % Idl >= 80
   AND Dso = PS
   ```

   Note: The greater-than or equal-to 80 percent idle space is a suggested value. Use a filter value according to your policy. Sites can release idle space that is based on a lower or higher idle space percentage value.

   The table is updated and lists only the PS data sets containing idle space greater-than or equal-to 80 percent. If data sets do not appear, none meet the search criteria. Consider changing the percent idle space filter to a lower value. If data sets are listed and you want to release idle space, go to the next step.

3. Click the % Idle column header until the table rows are sorted by the % Idle column in descending order with the highest % idle listed at the top of the table.

4. Select the check box of the rows of the over-allocated data sets from which you want to release idle space.

5. Select Idle Space Release in the Administrative Actions section of the Actions pane.

6. Select the following options at the bottom of the Actions dialog to simulate the action:

   Note: We suggest that you first simulate the line action. When you are satisfied with the results, perform the real line action by selecting the Live Mode (Y,N) field.

   Live Mode (Y,N)

   Clear this check box to indicate N for simulate mode.

   Convert Tracks (Y,N)

   Select this check box to indicate Y, which indicates cylinder allocation attributes are converted to track allocations. Clear this check box to indicate N, which indicates cylinder allocation attributes are not converted to track allocations.

   Note: Consider if the data set receives many updates and a large amount of allocated space is needed. If that is the case, then you should probably not convert the cylinder allocation attributes to track allocations.

   Pct Idle After Space Release

   Specify the amount of idle space to remain in the data set after the release occurs, according to your policy.

7. Click Submit to simulate the storage engine line action.

   A simulated result of the action is displayed.

8. Choose one of the following options:
   - A Release message that includes the data set names is displayed when there is no error. If you do not have errors, you are ready to officially release the idle space. Go to step 12.
   - If an error occurs, a message is displayed indicating the error condition. If you have errors, go to step a.
     a. Do one of the following tasks to resolve the error:

        - Unselect the error items in the Investigator so they are not included in the line action.

        - Change the Pct Idle After Space Release value to a lower value.

     b. Click Submit to simulate the line action again.

     c. If error messages persist, repeat steps a and b until no error messages appear.
9. Select the Live Mode (Y,N) check box to indicate that you want to perform the action.
10. Click Submit to release space.

    Idle space is released according to your specifications. The idle space values that you specified for Pct Idle After Space Release remains in the data sets.
11. Release space from PO data sets by repeating steps 2 through 10. Change the filter in step 2 to display PO data sets. The following filter statement is how the new filter looks:

    ```
    % Idl >= 80
    AND Dso = PO
    ```

    You have identified the PO data sets with idle space above your policy and released idle space from them.
12. (Optional) Click the notes icon in the Investigator tool bar, and add a note to the items for which you reduced idle space to indicate what you did and when.
13. (Optional) Repeat the cost analysis of your DASD (see page 22) to see if you need to release additional idle space.

    Your DASD costs are within your target limits. You have reduced unused over-allocated space in your DASD so it is available for use. These changes let you reduce or defer the cost of adding new DASD to your site.

How to Migrate Old Data to Free Online Space

As a storage administrator, you review and manage storage space usage in Storage Groups to avoid space abends. We recommend that you archive or migrate older data to lower-cost media instead of occupying online DASD space in critical storage groups. This scenario shows how a storage administrator identifies potential space problems within Storage Groups and then migrates old data to free online space.

The Storage Groups object describes the space usage in each defined Storage Group. Each Storage Group is a collection of one or more volumes. The following is a list of the types of Storage Groups:

MVS (Multiple Virtual Storage)

Storage Groups defined within your MVS System Generator (sysgen) or IO Generator (IO Gen).

SMS (System Managed Storage)

Storage Groups defined within your IBM Data Facility Storage Management Subsystem (DFSMS).

VAM (Volume Allocation Manager)

Storage Groups defined within your CA Allocate system.

USR (User)

Storage Groups you defined within your storage engine.

The following illustration shows how a storage administrator identifies potential space problems within Storage Groups and then migrates old data to free online space.

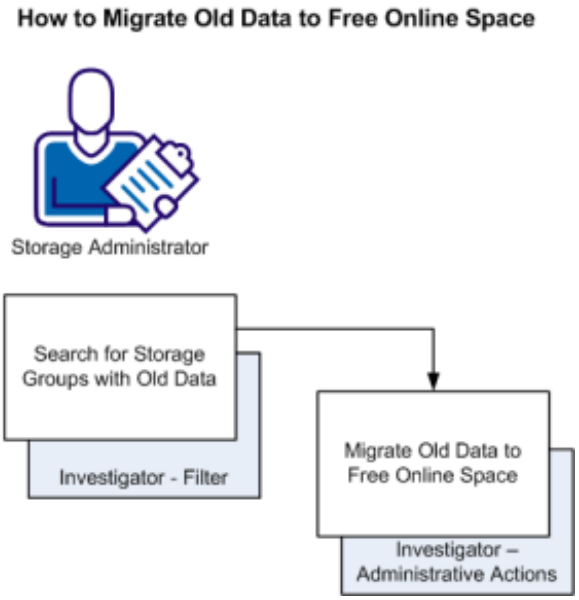

This procedure shows how to free online space in SMS Storage Groups.

Follow these steps:

1. Log in to CA Chorus.
2. Add the Investigator module to your dashboard from the Module Library, and click Start New Investigation.

   The Investigator opens.
3. Select Storage from the discipline drop-down list.
4. Expand the Volumes & Utilization, Storage Groups, Space and Other Attributes directory in the Storage object tree, and open the Storage Group object.
5. Create the following filter to display only SMS Storage Groups and click Search:

   `Type = SMS`

   The object table updates and displays only the SMS Storage Groups.
6. Click the column heading for % Allo All Volumes until it is sorted in descending order and then review the highest percent candidates.

   If the Storage Groups with the highest percent allocated are candidates for migration, record them for further investigation in step 8.
7. Expand the Diagnostics & Investigations, Data Set Management, All directory in the Storage object tree, and open the Data Sets for System (All) object.
8. Create the following filter, to display one of the Storage Groups that you noted in step 6 and click Search:

   `SG-Name = storage_group_name`

   The object table updates and displays the data sets in the Storage Group specified.
9. Click + to add the following filter statement to the existing filter, and click Search.

   `AND DaysUn > number_of_days`

   For example, enter 365 to identify data sets that have not been used for more than 365 days.

   The complete filter statement in the filter pane is:

   ```
   SG-Name = storage_group_name
   AND DaysUn > number_of_days
   ```

   The object table updates and displays the data sets that have not been used for the specified number of days, which qualifies them for migration according to your policy.
10. Sort the object display on the number of days the data set has not been used, by clicking the DaysUn column heading.

    You have identified old data sets that you can migrate to alternative storage media.

11. Select the rows that you want to migrate.
12. Select the appropriate Archive action in the Administrative Actions section of the Actions pane.

    The archive action that you select depends on your policy for the selected data sets and the systems you have. For example, your Archiving actions could be one of the following:

    CA Disk Archive

    Instructs CA Disk to create a backup copy of an online data set, to index the copy in the CA Disk files data set, and to dispose of the online copy.

    FDR/ABR Archive

    Generates a deferred archive request and inserts it into the ABR queue for processing using Innovation Data Processing's FDRABR system.

    Hsm Migrate Lvl1

    Migrates a data set to level 1 storage using IBM's DFSMShsm system.

    Hsm Migrate Lvl2

    Migrates a data set to level 2 storage using IBM's DFSMShsm system.

    The Actions dialog opens.

    Note: If you select CA Disk Archive, you must enter the appropriate parameters in the Actions dialog. You can simulate the CA Disk Archive action by selecting the check box next to SET MODE=SIMULATE - Simulated execution in the Actions dialog.
13. Click Submit to archive the data.

    Note: The CA Disk Archive action is a batch job that is put in the job queue for execution. When the job runs depends on how many jobs are in the queue.

    The data is migrated offline upon completion of the action job.
14. Repeat steps 8 through 13 to investigate and migrate each of the Storage Groups that you noted in step 6.

    You have migrated data offline to free online space after the archive action jobs have run. Additional online DASD space in critical storage groups is available for use.

    The storage engine pool-scan process updates the data displayed in the Storage Group. To verify the data is migrated offline, open the Storage Group object in the Investigator after the next pool-scan is completed and verify that the data sets are archived.

How to Scratch Uncataloged Duplicate Data Sets

Storage administrators must identify and remove uncataloged duplicate data sets, which can lead to incorrect input to production jobs. Uncataloged duplicate data sets are defined as data sets with the same name on different volumes that are not portions of a multiple volume data set.

This scenario shows how a storage administrator identifies uncataloged duplicate data sets and removes (scratches) uncataloged duplicate data sets.

Duplicate data sets tend to appear on all systems. The cataloged versions are usually the correct ones. The uncataloged versions are usually the invalid ones created when a problem arose. The leftover uncataloged versions can cause more problems. Uncataloged duplicate data sets build up over time and can occupy (waste) a significant amount of disk space. They can also cause allocation errors even though the requests are valid, but they select volumes where the old uncataloged data sets with the same names are found. If such allocation failures are within important applications, the impact and recovery can be severe.

Important! If you have multiple z/OS systems available, most of your z/OS system data sets on the various system packs appear as duplicates; at least one set for each system. The data sets connected to the running system appear cataloged to that system, while all others do not. Exclude intended duplicates from analysis and cleanup.

The following illustration shows how a storage administrator identifies uncataloged duplicate data sets and scratches the duplicate data sets.

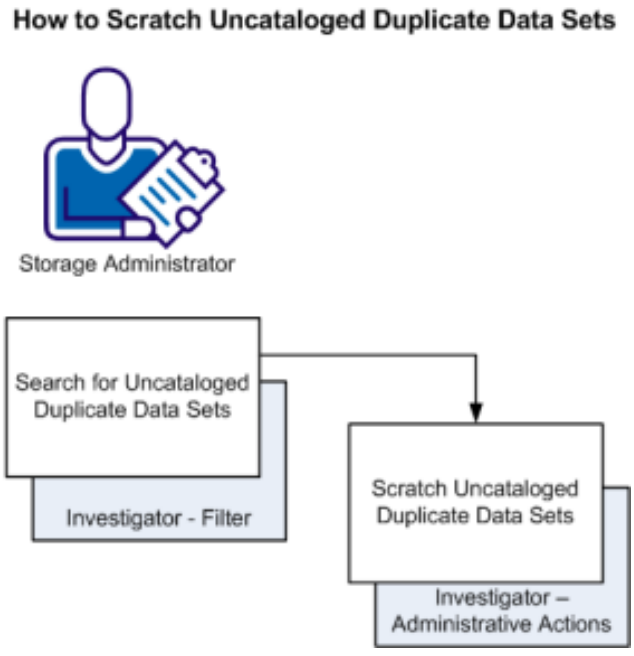

Follow these steps:

1. Log in to CA Chorus.
2. Add the Investigator module to your dashboard from the Module Library, and click Start New Investigation.

   The Investigator opens.
3. Select Storage from the discipline drop-down list.
4. Expand the Diagnostics & Investigations, Data Set Management, Duplicate Data Sets directory in the Storage object tree, and open the object Duplicate Data Sets for System.

   In the Duplicate Data Sets object, you expect to see at least two entries, if not more, for each Data Set Name (DSN). When only one entry is shown, a filter has been used which excludes the entry or entries on the other volumes. This filter can lead to a single DSN entry for a duplicate data set, because its duplicates reside on the excluded volumes.
5. Create the following filter so that only data sets that are not cataloged to the volume are displayed, and click Search:

   `Cat = N`

   The object displays with only the duplicate data sets in your system that are not cataloged. You have identified uncataloged duplicate data sets that can be potentially scratched.
6. Review and identify if the uncataloged duplicate data sets listed are candidates for scratching. If you identify scratch candidates, continue to the next step.
7. Select the uncataloged duplicate data sets that you want to scratch.
8. Click Scratch Data Set in the Administrative Actions section of the Actions pane.

   The Actions dialog opens.

   Important! When you submit this action, the storage engine scratches the uncataloged duplicate data sets.
9. Click Submit.

   A message appears in the action dialog saying the uncataloged duplicate data sets are scratched. The scratched items remain listed in the Duplicate Data Sets object until the data set table of contents (DTOC) refresh occurs, which is dependent on your site setup.
10. (Optional) Confirm that the uncataloged duplicate data sets are scratched by attempting to rescratch them (by performing steps 7 to 9 again). If they are scratched, the action dialog advises that they are previously scratched.

You have scratched uncataloged duplicate data sets, freed wasted space, and minimized the likelihood of allocation errors. Because uncataloged duplicate data sets can build up again over time, we recommend that you perform this scenario periodically.

How to Convert Data Sets from Multiple to Single Volume

As a storage administrator, you want to convert multiple volume data sets to single volume candidates to reduce confusion and simplify management of these data sets. This scenario shows how a storage administrator identifies multiple volume data set candidates and then converts them to single volume data sets.

The following illustration shows how a storage administrator identifies multiple volume data sets and then converts them to single volume data sets.

How to Convert Data Sets from Multiple to Single Volume

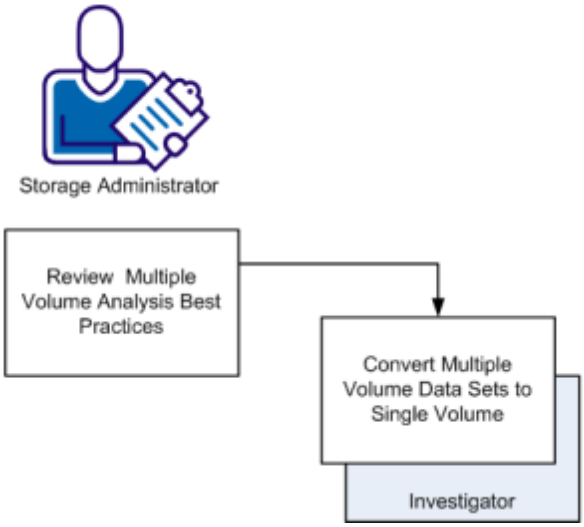

To convert data sets from multiple to single volume, do the following:

1. Review Multiple Volume Analysis Best Practices (see page 35).
2. Convert multiple volume data sets to single volume (see page 37).

Multiple Volume Analysis Best Practices

Multiple volume detail segments are defined as data sets with the same volume table of contents (VTOC) name which are found on different volumes and are diagnosed as parts of the same data set.

For Virtual Storage Access Method (VSAM), this diagnosis means that the data component itself or an index component itself is multiple volume. In the case where a VSAM key-sequenced data set (KSDS) has its data component on one single volume and its index on another volume, it is not considered multiple volume by the storage engine. Each component is single volume.

Observe the following points when analyzing multiple volume data sets:

- In the Multivol Data Sets For System object, the Volume sequence numbers greater than 1 indicate it is part of a multiple volume data set. For example, a value of 3 means it is part 3 of a multiple volume data set.

  The Multivol Data Sets For System object is located in the Diagnostics & Investigations, Data Set Management, Multi Volume Data Sets directory in the Storage object tree in the CA Chorus Investigator. The Seq Nbr field displays the Volume sequence number.

- The Summary of Multvol Data Sets object displays the total number of parts spread across all Volumes in the Vols field. For example, if a data set has parts that are spread across eight volumes, the Vols field shows a value of 8.

  The Summary of Multvol Data Sets object is located in the Diagnostics & Investigations, Data Set Management, Multi Volume Data Sets, Summaries directory in the Storage object tree in the CA Chorus Investigator.

- Multiple volume data set segments form the detail input to the Summary of Multvol Data Sets object.

  This object combines the relevant fields data, such as space allocated, from all the multiple volume segments (parts) that are found on different volumes. For example, if a data set has parts that are spread across eight volumes, this object shows a line which is the *total* for all eight parts of the data set.

- All parts of a multiple volume data set must exist.

  This point means that all sequence numbers should be found. If a sequence number is missing, something is probably wrong. However, if volumes have been excluded from processing, the excluded volumes can contain the missing portion of a multiple volume data set.

  You can list all the parts, (to see that all the sequence numbers exist), from the Summary of Multvol Data Sets object by selecting the item in the table, then selecting the Navigation action *Show Multivol Data Sets for System*.

- The VTOC entry for parts 2, 3, and so on, should all identify volume 1.

  When a data set extent is going to another volume (making it a multiple volume data set), DADSM functions record the VOLSER of the first volume within the VTOC entry for all the multiple volume parts. That is, the VTOC field DS1DSSN in parts 2, 3, ... N of the data set all identify the volume where part 1 should be found. By zooming to a dump of the VTOC entry, you can see the DS1DSSN field identifying volume 1 where part 1 should be found.

- In most environments, the catalog should reflect the correct volumes.

  The volumes in an IDCAMS listcat should match the volumes in the sequence number order indicated in the object display. If they do not, something is probably wrong. However, non-SMS data sets do not have to be cataloged, and the non-SMS environment allows VOLSERs to reference an uncataloged data set in the JCL.

- Clean bad multiple volume data sets, but do so cautiously.

  As evident from the preceding points, identifying good from bad multiple volume data sets is not always easy. However, bad entries can easily waste disk space.

- Parts of multiple volume data sets can disappear.

  The parts can disappear for many reasons. The following are some of the common reasons:

  - A full-volume restore can back level a volume to a state before part N of the multiple volume data set being created on it. All parts on other volumes are still found, but the part on the restored volume is now missing.
  - Reinitializing a volume can cause a part to disappear. Historically, multiple volume data sets were not often used or understood. As a result, some vendor and homegrown utilities do not handle them correctly, which can lead to VTOC and catalog errors for these data sets.

Convert Data Sets from Multiple to Single Volume

Use this procedure to identify multiple volume data sets and then convert them to single volume data sets.

Follow these steps:

1. Log in to CA Chorus.
2. Add the Investigator module to your dashboard from the Module Library, and click Start New Investigation.

   The Investigator opens.
3. Select Storage from the discipline drop-down list.
4. Expand the Diagnostics & Investigations, Data Set Management, Multi Volume Data Sets, Summaries directory in the Storage object tree, and open the Summary of Multvol Data Sets object.
5. Review the content of the object display for potential candidates to convert from a multiple volume data set to a single volume data set.

   You have identified multiple volume data set candidates for conversion to single volume data sets.

   Note: Before you convert the data set, confirm that sufficient space exists on the Target Volume. The A-Trks field indicates the total number of tracks that are needed on the target volume. Record the value for the next step.

6. Identify a volume candidate with enough space. To find a candidate volume object with the needed target space available, follow these steps:

   a. Expand the Volumes & Utilization, Volumes(DASD), Online, Space and Other Attributes directory in the Storage object tree, and open the Volume Space Usage object.

      The volumes available are displayed in the object table.

   b. Click the FrCyls and FrTrks fields column headings to find volumes with the largest space available.

   c. Record the volume name that contains the needed target space available which you want to move the multiple volume data sets to in step 10.

7. Expand the Diagnostics & Investigations, Data Set Management, Multi Volume Data Sets, Summaries directory in the Storage object tree, and open the object Summary of Multvol Data Sets object.

   The Summary of Multvol Data Sets object opens.

8. Select the row of the data set that you want to convert to single volume data set.

   The row of the data set that you want to convert is selected.

9. Select Move in the Administrative Actions section of the Actions pane.

   Note: The Move *action* moves a data set from one volume to another, catalogs it on the new volume, and deletes the original data set. For VSAM, the cluster is copied, not just the component.

   The Move *action* displays the Action dialog which provides the following options:

   New data set name (optional)

   Specify a new name to give to the data set.

   For VSAM, specify the new cluster name, not a component name.

   To Volume

   Specify the new target volume for the data set. (The volume name that you recorded in step 7.)

   Catalog (Y/N)

   Select this check box for Y. This indicates that you want the moved data set to be cataloged.

   Replace (Y/N)

   Clear this check box for N. This indicates that you do *not* want to replace the data set if it exists on the target volume.

10. Click Submit.

    The specified *action* is performed.

11. Repeat steps 6 through 10 for each data set that you want to convert to single volume.

    Multiple volume data sets are converted to a single volume. You have simplified management of these data sets and reduced the likelihood of confusion.

How to Release Idle Space in Data Sets

As a storage administrator, one of your responsibilities is to release idle space in over-allocated data sets. By releasing idle space, other applications can use the space, and you can avoid space abends. In both instances, you improve the efficiency of your site. This scenario includes the procedures to achieve this goal.

Idle space can be released from physical sequential (PS), partitioned (PO), and virtual storage access method (VSAM) type data sets. This procedure includes an example to release idle space from PO and PS data sets.

The following illustration shows how a storage administrator identifies idle space in over-allocated data sets and then releases the idle space.

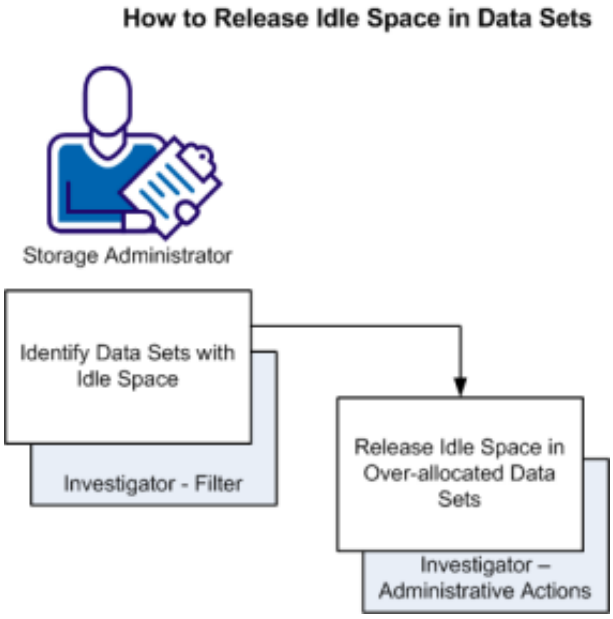

Follow these steps:

1. Log in to CA Chorus.
2. Add the Investigator module to your dashboard from the Module Library, and click Start New Investigation.

   The Investigator opens.

3. Select Storage from the discipline drop-down list.
4. Expand the Diagnostics & Investigations, Data Set Management, All directory in the Storage object tree, and open the Data Sets for System (All) object.
5. Create a filter to display PS data sets with a large amount of idle space and click Search.

   For example, create the following filter statement:

   ```
   % Idl >= 80
   AND Dso = PS
   ```

   Note: The greater-than or equal-to 80 percent idle space is a suggested value. Use a filter value according to your policy. Sites can release idle space that is based on a lower or higher idle space percentage value.

   The table is updated and lists only the PS data sets containing idle space greater-than or equal-to 80 percent. If data sets do not appear, none meet the search criteria. Consider changing the percent idle space filter to a lower value. If data sets are listed and you want to release idle space, go to the next step.
6. Click the % Idle column header until the table rows are sorted by the % Idle column in descending order with the highest % idle listed at the top of the table.
7. Select the rows of the over-allocated data sets from which you want to release idle space.
8. Select Idle Space Release in the Administrative Actions section of the Actions pane.
9. Select the following options at the bottom of the Actions dialog to simulate the action:

   Note: We suggest that you first simulate the *line action*. When you are satisfied with the results, perform the real *line action* by selecting the Live Mode (Y,N) field.

   Live Mode (Y,N)

   Clear this check box to indicate N for *simulate mode*.

   Convert Tracks (Y,N)

   Select this check box to indicate Y, which indicates cylinder allocation attributes are converted to track allocations. Clear this check box to indicate N, which indicates cylinder allocation attributes are not converted to track allocations.

   Note: Consider if the data set receives many updates and a large amount of allocated space is needed. If that is the case, then you should probably not convert the cylinder allocation attributes to track allocations.

   Pct Idle After Space Release

   Specify the amount of idle space to remain in the data set after the release occurs, according to your policy.
10. Click Submit to simulate the storage engine *line action*.

    A simulated result of the action is displayed.

11. Choose *one* of the following options:

    - A *Release* message that includes the data set names is displayed when there is no error. If you do not have errors, you are ready to officially release the idle space. Go to step 12.
    - If an error occurs, a message is displayed indicating the error condition. If you have errors, go to step a.

        a. Do *one* of the following tasks to resolve the error:

        - Unselect the error items in the Investigator so they are not included in the *line action*.

        - Change the Pct Idle After Space Release value to a lower value.

        b. Click Submit to simulate the *line action* again.

        c. If error messages persist, repeat steps a and b until no error messages appear.

12. Select the Live Mode (Y,N) check box to indicate that you want to perform the action.

13. Click Submit to release space.

    Idle space is released according to your specifications. The idle space values that you specified for Pct Idle After Space Release remains in the data sets.

14. Release space from PO data sets by repeating steps 3 through 13. Change the filter in step 5 to display PO data sets. The following filter statement is how the new filter looks:

    ```
    % Idl >= 80

    AND Dso = PO
    ```

    You have identified the PO data sets with idle space above your policy and released idle space from them.

15. (Optional) Click the notes icon in the Investigator tool bar, and add a note to the items for which you reduced idle space to indicate what you did and when.

    Other applications can now use the released space, which can help to avoid space abends. You have improved the efficiency of your site.

How to Manage Scratch Tape Availability and Status

As a storage administrator, you know how many tapes the daily cycle uses. You regularly verify that enough scratch tapes are available in your scratch tape pool to meet your site requirements. In this context, the term *scratch tape* refers to a tape that has scratch status, and it is available for applications to write data on. *Scratch tape pool* refers to a library of scratch tapes.

Some applications are set up to send data to tape. Tape management systems typically do not allow applications to write data at the beginning of the tape unless the tape has scratch status. If a tape has scratch status, any data it contains can be over-written.

An insufficient number of scratch tapes for the daily processing cycle can prevent your applications from completing on time. For sites with a heavy dependence on tape, insufficient tape can be a critical problem.

This scenario shows how a storage administrator identifies existing scratch tapes and then adds additional scratch tapes if not enough scratch tapes are available.

The following illustration shows how a storage administrator identifies existing scratch tapes and then adds additional scratch tapes.

How to Manage Scratch Tape Availability and Status

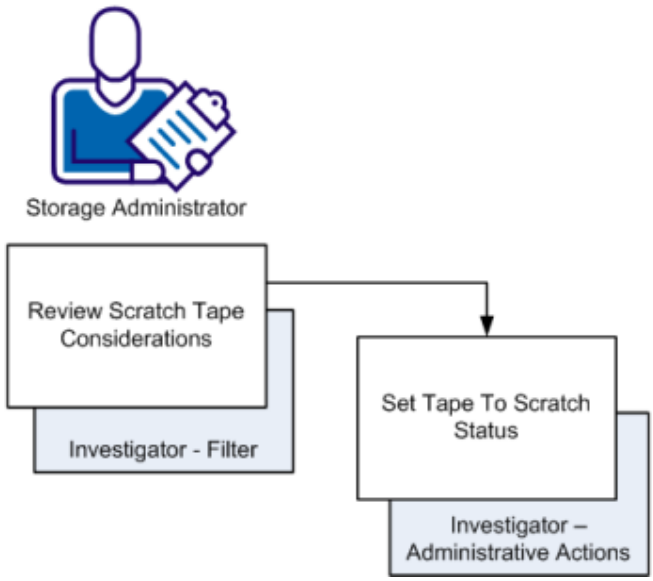

To manage scratch tape availability and status, do the following:

1. Review Scratch Tape Considerations (see page 43).
2. Set tape to scratch status (see page 45).

Review Scratch Tape Considerations

Storage administrators can proactively use CA Chorus for Storage Management to review the tape situation if you use CA 1, CA TLMS, or IBM's DFSMSrmm tape management systems. This review includes the amount of scratch tapes needed for a daily cycle and the amount of scratch tapes available for it. A scratch tape is a tape that has scratch status, and it is available for applications to write data on. Consider special tape usage patterns for weekly, monthly, or yearly processing. You can monitor the number of scratch tapes used in each of these periods.

To analyze usage patterns, search the CA 1, CA TLMS, or DFSMSrmm volume records. In the storage engine, *tapes* are referred to as *volumes* in the tape management system objects. Use filters to identify the number of volumes you need for your daily cycle and the number available.

Review expired tapes available for setting to scratch status. Tapes can be expired by the tape management system, and are considered available for reuse. If a tape has scratch status, any data it contains can be overwritten. A best practice is to erase data on tapes when they are expired or before you set them to scratch status.

The following example is for customers that have CA TLMS. If you use CA 1 or IBM's DFSMSrmm tape management systems, the process is similar. Instead of using the CA TLMS Volumes object, for CA 1 use the CA 1 Volumes object. For DFSMSrmm, use the DFSMSrmm Volumes object.

Use this procedure to identify the number of scratch tapes needed for your daily cycle and the number of tapes available for scratching.

Follow these steps:

1. Log in to CA Chorus.
2. Add the Investigator module to your dashboard from the Module Library, and click Start New Investigation.

   The Investigator opens.
3. Select Storage from the discipline drop-down list.
4. Expand the Diagnostics & Investigations, Tape Resource Management, Tape Management System, CA TLMS, and Volumes directory in the Storage object tree, and open the CA TLMS Volumes object.

   A message is displayed advising you must create a filter.
5. Create the following filter, and click Search:

   `Create Date = date`
   `Create Time > start_time_of_daily_cycle`

   Volumes that are created during the specified daily cycle are displayed.

6. (Optional) Create the following filter to identify tapes scratched on a specific date, and click Search:

   ```
   Scratch = true
   Scratch Date = date
   ```

7. (Optional) Create the following filter to identify scratched tapes that are scratched during a specific date range, and click Search:

   ```
   Scratch Date > date
   And Scratch Date < date
   ```

8. (Optional) Create the following filter to identify tapes that are created during a specific date range, and click Search:

   ```
   Create Date > date
   And Create Date < date
   ```

9. (Optional) Create the following filter to isolate the tapes that are created on a specific day, by a specific application, and click Search:

   ```
   Create Date = date
   AND Jobname = job_name
   ```

10. (Optional) Create the following filter to identify the use of different types of tapes, and click Search:

    ```
    Tape Type = tape_type
    ```

    Note: Tape types are defined and maintained by your tape administrator.

    If your site identifies physical or virtual by the volume name prefix, you can create a similar filter but use the starts with operand. For example, create the following filter:

    ```
    Volume starts with virtual_tape_prefix
    ```

11. (Optional) Create the following filter to list only scratched tapes within a volume range, and click Search:

    ```
    Scratch = true
    And Volume starts with start_of_volume_ID
    ```

    You have filtered the CA TLMS Volumes object to list tapes created during your daily cycle, tapes scratched on a specific date or during a date range, tapes created by an application, and by tape types, and ranges. These results indicate the number of different types of scratch tapes that you need during a time period.

Set Tape to Scratch Status

The following example is for customers that have CA TLMS. If you use CA 1 or IBM's DFSMSrmm tape management systems, the process is similar. Instead of using the CA TLMS Volumes object, for CA 1 use the CA 1 Volumes object. For DFSMSrmm, use the DFSMSrmm Volumes object.

Use the following procedure to manage scratch tape availability and set tape status to scratch using CA TLMS.

Follow these steps:

1. Add the Investigator module to your dashboard from the Module Library, and click Start New Investigation.

   The Investigator opens.

2. Select Storage from the discipline drop-down list.

3. Expand the Diagnostics & Investigations, Tape Resource Management, Tape Management System, CA TLMS, and Volumes directory in the Storage object tree, and open the CA TLMS Volumes object.

   A message is displayed advising that you must create a filter.

4. Click the Customize Your Tabular Data View (Wrench) icon on the Investigator toolbar, to open the Select and Reorder Columns dialog.

5. Select Scratch from the drop-down list in the *Summarize by* field, (optional) clear the check-box next to *Always use this column configuration*, and click Save.

   The object updates and provides the total number of tapes with scratch status (true) and the total number of tapes with no-scratch status (false).

6. Do one of the following:

   - If the number of scratch tapes available is within your available scratch tape policy limits, you have completed the investigation.
   - If the number of rows (records) is not within your available scratch tape policy limits, perform the next step.

7. Click the expansion-arrow in front of "false (*number*)" to list the tapes that do not have scratch status.

8. Identify candidates that can be set to scratch status. For example, look for the object field that indicates expiration date/type.

9. Select the rows of the items you want to set to scratch status.

   Important! If a tape has scratch status, any data it contains can be overwritten. Do not set a tape to scratch status unless you are sure the data it contains can be overwritten.

   The lines representing the tapes that you want to scratch are selected.

10. Select Scratch Status in the Administrative Actions section of the Actions pane.

    The Actions dialog opens.

11. Click Submit.

    The tapes are set to scratch status.

12. (Optional) Click the notes icon in the Investigator toolbar to add a note to the items that you set to scratch status, to indicate what you did and when.

You have identified and set tape candidates to scratch status. You have enough scratch tapes available to meet your needs.

How to Free Up Storage Space by Executing An Automation Script

As a storage administrator, you want enough storage space available in crucial storage groups.

Consider the following example: During working hours, you receive an Alert that tells you a storage group is getting exceptionally low on space. You know that you have already set up a storage engine General Object Automation (GOA) script that is scheduled to execute at midnight. The GOA script frees up additional space on the item referenced in the Alert. You want to free up additional space immediately so you know that enough space is available before you leave for the day.

This scenario shows how you can execute the GOA script online from the Investigator that frees up additional space immediately.

The following illustration shows how a storage administrator locates the General Object Automation Scripts object and then fires the script that frees up space in a storage group immediately.

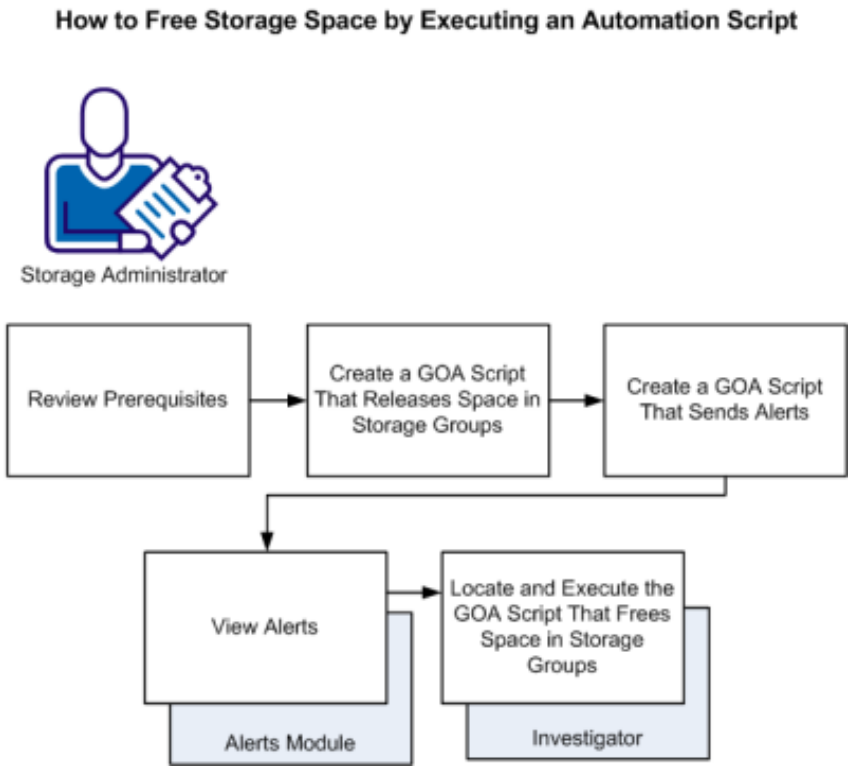

To free up storage space immediately by executing an automation script, do the following:

1. Review Prerequisites (see page 47).
2. Create a GOA script that frees storage space (see page 48).
3. Create a GOA script that sends Alerts (see page 50).
4. View Alerts in the Alerts Module (see page 52).
5. Execute the GOA Script that frees space in Storage Groups (see page 52).

Review Prerequisites

The storage engine windows client must be installed on a PC to create storage engine GOA scripts. .

Note: For more information about installing the windows client, see the *CA Vantage Installation Guide*.

Create a GOA Script That Frees Storage Space

You can create a storage engine GOA script to monitor available space in Storage Groups and free space automatically.

Use this procedure to create a GOA script to monitor available storage space and free storage space automatically.

Follow these steps:

1. Open the back-end storage engine windows client by double-clicking its icon on your PC.
2. Connect to the back-end storage engine system host that is setup for CA Chorus for Storage Management.
3. Click z/OS, Object Scripting, Automation Scripts, General Object Automation in the windows client main menu.

   The Automation Scripts Builder opens.

   Note: Click the Help button on any page of the wizard to get a description of the options on the page.
4. Click Next repeatedly until the Source Object Selection page opens.
5. Select the Space and Other Attributes object in the Storage Groups folder and click Next to open the Script Name page.
6. Select *Create a script with this name:*, enter MIGSTGRP in the *Script Name:* field, and click Next to open the Title and System Execution page.
7. Enter Migrate From Storage Group in the *Title* field, enter Migrates unused data sets from Storage Groups in the *Description* field, and click Next to open the Timing for Events page.
8. Select Weekday Schedules, click Add to display the Schedule the Action dialog.
9. Select All (Every day), enter 2400 in the *Time* field, and click Set to specify you want the script to run on weekdays at midnight.
10. Click Next repeatedly until the Filter/Sort page opens, and click Filter to open the Filter dialog.
11. Create the following filter and click Apply and Execute:
    ```
    "% Allo" > 80
    ```
    This filters storage groups that have an allocated percentage of greater-than 80 of the space available. Adjust the value according to your policy.
12. Click Next repeatedly until the Different Object Selection page opens.
13. Select *Base action on a different object*, *Zoom object*, and click Next to open the Zoom/New Object page.
14. Select DTOC4POL and click Next to open the Filter/Sort for New/Zoom Object page.

15. Select the check-box Filter (optional), click Filter, create the following filter, and click Apply and Execute:

    ```
    DaysUn > 60
    ```

    This specifies you want to identify the data sets that have not been used for the last 60 days. Adjust the number of days in the filter according to your policy.

16. Click Sort and do the following:

    a. Click DaysUn in the Available Fields column.

    b. Click the Move arrow so it moves to the Sort Order column.

    c. Click the sort icon so that it sorts in descending order and click Apply and Execute.

    This specifies you want to migrate the data sets that have not been used for the longest period first.

17. Click Next to open the Record Inclusion page.

18. Select *Until Free Space % is Percentage* and enter 25.

    This specifies you want to free space until 25 percent of the Storage Group is free. That is, the Storage Group is 75 percent allocated instead of 80 percent allocated. Adjust the value according to your policy.

19. Click Next repeatedly until the Action Selection page opens, select Perform Object Action, and click Next to open the Perform Object Action page.

20. Scroll down to Hsm Migrate Lvl1 and select it.

21. (Optional) Enter 100 in the *Block selection of lines (maximum number of rows available for Object Action)* field.

    Adjust the value according to your policy.

    This specifies you want to migrate unused data sets in Storage Groups until either 25 percent of the space in the Storage Group is freed or a maximum of 100 data sets are migrated from the Storage Group. If you do not specify a maximum number of data sets to be migrated, then the system could migrate many more smaller data sets than desired to meet the maximum threshold of 75 percent allocated.

22. Click Next, select the *Save, Activate and Execute the script* option, and click OK.

    The script is saved, activated, and executes at midnight each workday.

23. Close the storage engine windows client.

    You have created a GOA script that automatically migrates data sets to free space in Storage Groups that are at least 80 percent allocated.

Create a GOA Script That Sends Alerts

To receive alerts in the Alerts module for various storage conditions, storage engine GOA scripts are required to send the alerts when conditions are met.

Use this procedure to create a GOA script that sends CA Chorus for Storage Management an alert when critical storage group space usage reaches the maximum capacity level.

Follow these steps:

1. Open the storage engine windows client by double-clicking its icon on your PC.
2. Connect to the storage engine system host that is set up for CA Chorus for Storage Management.
3. Click z/OS, Object Scripting, Automation Scripts, General Object Automation in the windows client main menu.

   The Automation Scripts Builder opens.

   Note: Click the Help button on any page of the wizard to get a description of the options on the page.
4. Click Next until the Source Object Selection wizard page opens.
5. Select the Space and Other Attributes object in the Storage Groups folder and click Next to open the Script Name page.
6. Select *Create a script with this name:*, enter SGALERT in the *Script Name:* field, and click Next to open the Title and System Execution page.
7. Enter Monitor SG in the *Title* field, enter Monitor Storage Groups for Alert Threshold of 80 percent in the *Description* field, and click Next to open the Timing for Events page.
8. Select Daily Cycles, specify 120 minutes, and click Next to open the Options page.
9. Select the following options:
   - Load Script
   - Enable Script when Loaded (for 24 hour scripts only).
   - Use the value of the STMONSCR system parameter to control if the script should be monitored with the Internal Status Monitor.
10. Click Next repeatedly until the Filter/Sort page opens.
11. Click the Filter button and create a filter with the following filter statement in the Filter dialog:

    ```
    "% Allo" >= 80
    ```

    This filters storage groups that have an allocated percentage of 80 or higher of the space available.

12. Click Apply and Exit to save the filter and close Filter dialog.
13. Click Next to open the Record Selection page.
14. Select the option *Select filtered records* and click Next repeatedly until the Record Inclusion page opens.
15. Select the option *Select All Records* and click Next to open the Conditions for Actions page.
16. Select Equal to, enter 1 in the *number of records* field, and click Next to open the Action Selection page.
17. Select the option *Perform Object Action* and click Next to open the Perform Object Action page.
18. Select *Send Chorus Alert* and enter the following:
    - Enter 1 in the *Severity Value of 1 to 5* field.

      Note: A severity value 1 is the highest severity.
    - Click in the *Alert Text* field and do the following:
      a. Type SG and then space.
      b. Click the Variables button and double-click Storage Group Name from the Substitution Variables window.
      c. Click after %%POOLNAME%% in the text field, and add a space.
      d. Type "is now " (with a space after the word now).
      e. Double-click the Pct Allocated variable in the Substitution Variables window.
      f. Click after %%POOLUPCT%% in the text field, add a space, and type "pct filled. Action required!".

         The Alert Text field contains the following text:

      ```
      SG %%POOLNAME%% is now %%POOLUPCT%% pct filled. Action required!
      ```

         This alert message displays as follows in the Alerts Module:

         SG *storage_group_name* is now *percent_value* pct filled. Action required!
19. Click Next to open the Saving and Activation page, select the *Save, Activate and Execute the Script* option, and click OK.

    Note: To receive an earlier *warning* alert that a storage group is reaching 70 percent of its capacity level, consider making another GOA script. This script would be similar, but with a filter with the following expression: "% Allo" > 70 AND "% Allo" < 80. Give this a Severity 2 condition with the message: "Warning! SG %%POOLNAME%% is now %%POOLUPCT%% pct filled."

20. Close the storage engine windows client.

    You have created a GOA script that sends CA Chorus for Storage Management an alert when critical storage group space usage reaches 80 percent or higher of the maximum capacity level.

View Alerts in the Alerts Module

Use this procedure to view Alerts for storage groups.

Follow these steps:

1. Log in to CA Chorus, and add the Alerts module to your dashboard.
2. Select the *configure link*, Storage Alerts, and click Save.

   If any of your storage groups have equal-to or greater-than 80 percent usage, the Alerts module window displays them as severity level 1 alerts.
3. Click the + symbol next to the Alert icon.

   The Alert expands.
4. Click the View Detail link.

   The full Alert message is displayed.
5. (Optional) Select the item in the Alert window after you have read the message and click Delete.

The alert indicates that you might need to reduce the percent used. Consider releasing idle space, archiving data, or allocating more volumes to the storage group.

Execute the GOA Script That Frees Space In Storage Groups

Use this procedure to fire the GOA script that frees up space in storage groups immediately.

Follow these steps:

1. Add the Investigator module to your dashboard from the Module Library, and click Start New Investigation.

   The Investigator opens.
2. Select Storage from the discipline drop-down list.
3. Expand the Diagnostics & Investigations, Automation and Logging, General Object Automation (GOA) directory in the Storage object tree, and open the Object Automation Scripts object.

   The Object Automation Scripts object opens in the Investigator.

4. Locate the GOA script that frees up space in the storage group, and select it.
5. Click the Administrative Action *Fire*.

   The Action dialog opens.
6. Click Submit.

   A confirmation dialog that the script is submitted opens.
7. Click OK.

   Note: The Storage Group object is a Last Interval object, which means it is refreshed based on a schedule in the storage engine. To force an immediate refresh, do the following:

   a. Expand the Diagnostics & Investigations, Storage Platform Administration directory in the Storage object tree and open the Commands object.

   b. Create the following filter and click Search:

      `Data Record starts with refresh`

   c. Select the line with the *Refresh,Pools* command, and select the appropriate Administrative Action:

      - Exec On Cur Host
      - Exec On Sel Hosts

      The Actions Dialog opens.

   d. Click Submit.

      The storage engine refreshes the object data.
8. Expand the Volumes & Utilization, Storage Groups, Space and Other Attributes directory in the Storage object tree, and open the Storage Group object to verify that enough space is now available.

   Adequate space is available in the storage group for applications to store data.

How to Configure the Metrics Panel to Monitor Storage Group Space Usage

This scenario shows how a storage administrator configures the Metrics Panel to monitor critical storage group space usage and create a metrics dashboard. Monitoring this data is important because when you reach the threshold, you must allocate more space to avoid abends.

The Metrics panel displays storage group statistics metrics. From the Metrics panel you can view the data in a chart on your dashboard and launch the Investigator for deeper analysis.

In this scenario, the storage administrator monitors critical storage group space usage in a metrics chart that indicates when usage nears or exceeds the maximum defined threshold value. When usage nears or exceeds the maximum defined threshold value, the allocated space in the critical storage group can be adjusted using the appropriate Administrative Action in the Investigator.

The following illustration shows how a storage administrator configures the Metrics panel, sets thresholds, and creates a metrics dashboard in the CA Chorus interface.

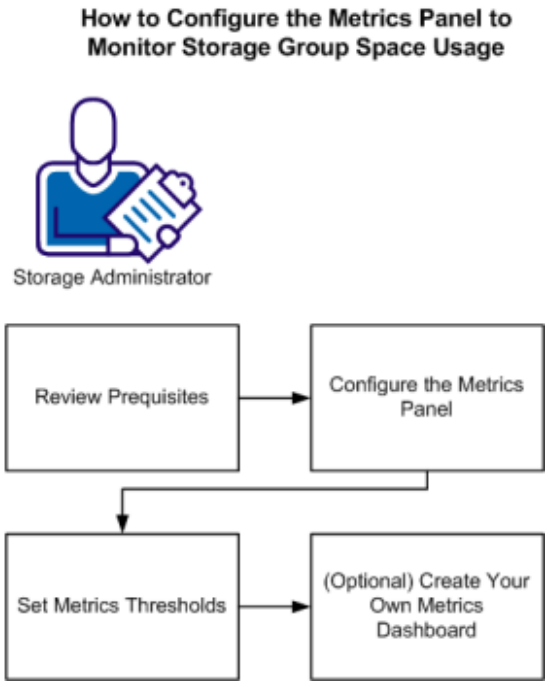

To configure the Metrics panel to monitor storage group space usage, do the following:

1. Review Prerequisites (see page 54).
2. Configure the Metrics panel (see page 55).
3. Set Metrics thresholds (see page 55).
4. (Optional) Create your own Metrics dashboard (see page 56).

Review Prerequisites

To view storage metrics in CA Chorus you must first set up the storage engine to collect storage metric data for the Time Series Facility (TSF). For more information, see How to Collect Storage Metrics for TSF (see page 95) in the *CA Chorus for Storage Management User Guide*.

Configure the Metrics Panel

Use this procedure to configure the Metrics panel. Repeat this procedure any time that you want to add a metric. If you have multiple disciplines installed, the Metrics panel can display data across all disciplines.

Follow these steps:

1. Log in to CA Chorus.
2. Click the plus sign control (+) on the Metrics panel.

   The Select Metrics Domain view of the Add New Metrics panel opens.
3. Click the down arrow in the *Metric Source* field, and select Storage.
4. Click the Shows tree icon next to the *Select Entity* option.

   The Statistics category of the storage navigation tree opens. This category lists the storage objects available for the TSF.
5. Click the pointer next to the Storage Groups category, drill down until you can select the entity of the critical storage group, and then select it.
6. Click Next.

   The Select Metrics view of the Add New Metrics panel opens.

   The left pane shows the Available Metrics (fields of the Storage Groups object that metrics data is available for).
7. Click % Allo to display the percentage of space that is allocated of the total space available in the Storage Group. You can also select more fields by holding down the Ctrl button and clicking them. For example, the Allo Bytes, Cap Bytes, and Free Bytes fields.
8. Click the single arrow pointing to the Selected Metrics pane and click Finish.

   The selected items are listed in the Selected Metrics pane. The new items are added to the Metric panel. They appear as small graph icons in the row towards the top of the interface.

(Optional) Set Metrics Thresholds

Use this procedure to set the caution and danger thresholds according to your site policy.

Follow these steps:

1. Stop the scrolling feature, if necessary.
2. Click the down-arrow in the bottom right-corner of the % Allo metric in the main interface. Hover your cursor over the small % Allo metric to display the down-arrow.

3. Select Set Threshold from the drop-down list.
4. Click the down arrow in the *Caution State: Data Spikes* field and select *Above*. Enter 70 in the field next to Above.

   This creates a 70 percent allocation (% Alloc) *caution* threshold line in the metrics graph.
5. Click the down arrow in the *Danger State: Data Spikes* field and select *Above*. Enter 80 in the field next to Above.

   This creates an 80 percent allocation (% Alloc) *danger* threshold line in the metrics graph.
6. Click Set Threshold.

   The metric now responds to the specified threshold criteria. The threshold bar and the metric color changes to yellow when the caution criteria is met and the color changes to red when the danger criteria is met.

   Note: The value for Caution or Danger can be set to zero, but not when the state is set to below.

(Optional) Create Your Own Metrics Dashboard

Use this procedure to create your own metrics dashboard to display the critical storage group percent allocated Metrics graph (or multiple Metrics graphs).

Follow these steps:

1. Click the + (My Dashboard) tab in the CA Chorus interface and click Create a Dashboard.
2. Enter a tab name (maximum characters that are allowed are 21). For example, the name of the critical storage group or "Critical SG" (or "Metrics" if you plan to add multiple Metrics graphs to it).
3. Enter a description and click Add Dashboard.

   A new and empty dashboard opens.
4. Click the down-arrow in the bottom right-corner of the % Allo metric and select Add to Dashboard. To add multiple Metrics graphs to this dashboard, repeat this step for each Metric graph you want to add.

   The critical storage group percent allocated Metrics graph displays in the new dashboard.

   Every time that you open the interface and select this new dashboard tab, the critical storage group percent allocated Metrics graph is displayed with the actual used percentage of the storage group, and your threshold lines. If the percentage used is between or above the threshold lines consider releasing idle space, archiving data, or allocating more volumes to the storage group.

5. (Optional) To investigate the Storage Group Statistics object item, click the down arrow in the Metrics graph title bar and click Investigate.

   The Investigator opens and displays the item in the Storage Group Statistics object.

   Note: You cannot perform Administrative Actions in the Storage Group Statistics object. If you want to perform corrective action open the item in the Storage Groups object and perform the action.

   You have configured the Metrics Panel to monitor critical storage group space usage and created a metrics dashboard. You can view the metrics chart that indicates when usage nears or exceeds the maximum defined threshold value. Monitoring this data is important because when you reach the threshold, you must allocate more space to avoid abends.

How to Create an Alert to Monitor Storage Group Space Usage

As a storage administrator, you want to be alerted when a critical Storage Group space usage is reaching the maximum capacity level. This indication is important because when the usage level reaches the maximum, you can allocate more space to avoid space abends.

This scenario shows how a storage administrator sets up an alert to advise when critical Storage Group space usage reaches the maximum capacity level. You can set the maximum capacity level according to your policy.

The following illustration shows how a storage administrator creates an alert to monitor Storage Group space usage using the storage engine windows client and then view and manage it in the Alerts module.

How to Configure Alerts for Storage Group Space Usage

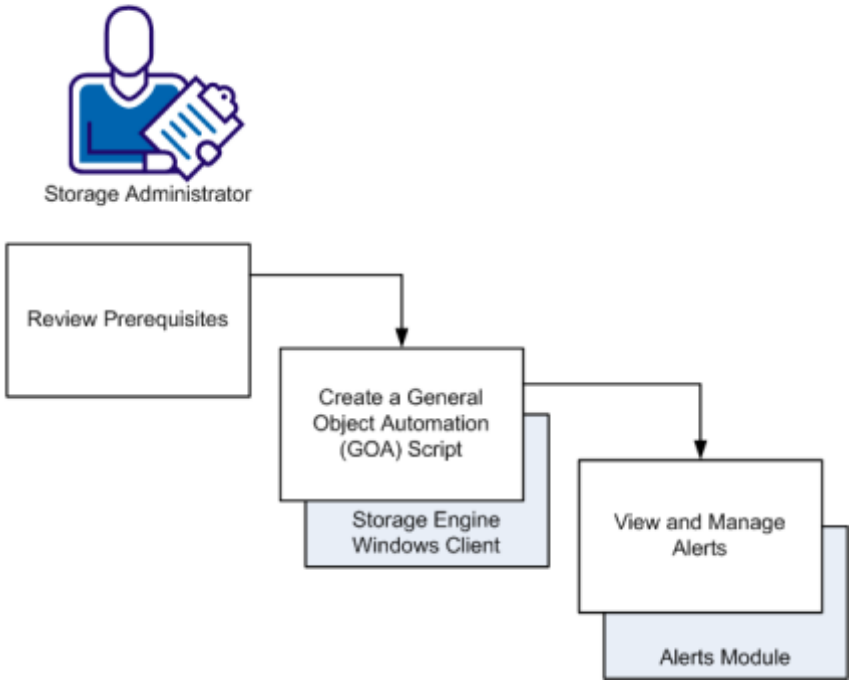

To create an alert to monitor storage group space usage using the storage engine windows client and view it in the Alerts module, do the following:

1. Review Prerequisites (see page 58).
2. Create a General Object Automation (GOA) script (see page 59).
3. View and manage alerts in the CA Chorus Alerts module (see page 61).

Review Prerequisites

The storage engine windows client must be installed on a PC to create the storage engine GOA script.

Note: For more information, see the *CA Vantage Installation Guide*.

Create a General Object Automation (GOA) Script

General Object Automation (GOA) scripts can be created to monitor storage objects and initiate corrective actions.

To receive alerts in the Alerts module for various storage conditions, storage engine GOA scripts are required to send the alerts when conditions are met.

Use this procedure to create a GOA script to send CA Chorus an alert when critical Storage Group space usage reaches the maximum capacity level.

Follow these steps:

1. Open the storage engine windows client by double-clicking its icon on your PC.
2. Connect to the storage engine system host that is setup for CA Chorus for Storage Management.

   Note: For instructions about how to create a host connection definition and how to connect to a host, see the *CA Vantage Windows Client Guide*. This guide is available in the windows client and in the CA Chorus Knowledge Center.
3. Click z/OS, Object Scripting, Automation Scripts, General Object Automation in the windows client main menu.

   The Automation Scripts Builder opens.

   Note: Click the Help button on any page of the wizard to get a description of the options on the page.
4. Click Next until the Source Object Selection wizard page opens.
5. Select the Space and Other Attributes object in the Storage Groups folder, and click Next to open the Script Name page.
6. Select *Create a script with this name:*, enter SGALERT in the *Script Name:* field, and click Next to open the Title and System Execution page.
7. Enter Monitor SG in the *Title* field, enter Monitor Storage Groups for Alert Threshold of 80 percent in the *Description* field, and click Next to open the Timing for Events page.
8. Select Daily Cycles, specify 120 minutes, and click Next to open the Options page.
9. Select the following options:
   - Load Script
   - Enable Script when Loaded (for 24 hour scripts only).
   - Use the value of the STMONSCR system parameter to control if the script should be monitored with the Internal Status Monitor.
10. Click Next repeatedly until the Filter/Sort page opens.

11. Click the Filter button and create a filter with the following filter statement in the Filter dialog:

    ```
    "% Allo" >= 80
    ```

    This filters storage groups that have an allocated percentage of 80 or higher of the space available.

12. Click Apply and Exit to save the filter and close Filter dialog.
13. Click Next to open the Record Selection page.
14. Select the option *Select filtered records*, and click Next repeatedly until the Record Inclusion page opens.
15. Select the option *Select All Records*, and click Next to open the Conditions for Actions page.
16. Select Equal to, enter 1 in the *number of records* field, and click Next to open the Action Selection page.
17. Select the option *Perform Object Action* and click Next to open the Perform Object Action page.
18. Select *Send Chorus Alert* and enter the following:
    - Enter 1 in the *Severity Value of 1 to 5* field.

      Note: A severity value 1 is the highest severity.

    - Click in the *Alert Text* field and do the following:
      a. Type SG and then space.
      b. Click the Variables button and double-click Storage Group Name from the Substitution Variables window.
      c. Click after %%POOLNAME%% in the text field, and add a space.
      d. Type "is now " (with a space after the word now).
      e. Double-click the Pct Allocated variable in the Substitution Variables window.
      f. Click after %%POOLUPCT%% in the text field, add a space, and type "pct filled. Action required!".

         The Alert Text field contains the following text:

         ```
         SG %%POOLNAME%% is now %%POOLUPCT%% pct filled. Action required!
         ```

         This alert message displays as follows in the Alerts Module:

         SG *storage_group_name* is now *percent_value* pct filled. Action required!

19. Click Next to open the Saving and Activation page, select the *Save, Activate and Execute the Script* option, and click OK.

    Note: To receive an earlier *warning* alert that a storage group is reaching 70 percent of its capacity level, consider making another GOA script. This script would be similar, but with a filter with the following expression: "% Allo" > 70 AND "% Allo" < 80. Give this a Severity 2 condition with the message: "Warning! SG %%POOLNAME%% is now %%POOLUPCT%% pct filled."

20. Close the storage engine windows client.

    You have created a GOA script that sends CA Chorus for Storage Management an alert when critical Storage Group space usage reaches 80 percent or higher of the maximum capacity level.

View and Manage Alerts

You can view alerts for storage groups and then take actions from within the Investigator module to resolve issues.

Follow these steps:

1. Log in to CA Chorus and add the Alerts module to your dashboard.
2. Follow the prompts to configure the module.

   After you configure the module, the alert window opens.

   In this example, if any of your storage groups have equal-to or greater-than 80 percent usage, the Alerts module window displays them as severity level 1 alerts.
3. Click + at the beginning of the row of the alert you want to view and manage.

   The alert row expands and the Actions field box is displayed.
4. Click the down arrow in the Actions field box and do one of the following options, from the Actions drop-down menu:

   - Click View Detail to view the alert details.
   - Click Investigate to resolve issues in the Investigator module.

     Note: Possible actions to resolve this example issue include releasing idle space, archiving data, or allocating more volumes to the storage group.
5. (Optional) Select the alert check-box and then click the Delete icon, to remove an alert from the Alerts module.

   You viewed and managed your alerts.

Chapter 2: Viewing Storage Object Data in the Investigator

View Storage Data in the Investigator

The Investigator gives you access to a storage repository. This repository is a replication of the information from the storage database, which is accessible using the storage engine. However, the information in the storage repository is designed to service storage resource activities. The Investigator simplifies your access to this information. The data in the Investigator is from the external storage engine.

The system groups Storage objects by categories in the expandable and collapsible storage object tree. The categories appear as folders in the storage object tree much like a directory. Categories contain common types of storage objects. Storage objects represent physical and conceptual components that are associated with one or more z/OS hosts.

When you select a storage object from the object tree, the storage engine collects instances of that object and the Investigator displays the data. After an object appears, select a row of data in the object to display more granular information in the Details pane. The Details pane opens at the bottom of the Investigator.

Follow these steps:

1. Log in to CA Chorus.
2. Add the Investigator module to your dashboard from the Module Library, and click Start New Investigation.

   The Investigator opens.
3. Select Storage from the discipline drop-down list.
4. Expand a storage category (directory) to list the storage objects included in that category.

   For example, expand the Expand the Volumes & Utilization, Storage Groups directory to see the objects in the Storage Groups directory and its sub directories.

5. Select a storage object.

   For example, expand the Expand the Volumes & Utilization, Storage Groups, Space and Other Attributes directory, and select the Storage Groups object. The Investigator displays the Storage Groups object table.

   Note: When opening very large storage objects, a message can be displayed stating that you must create a filter first. If this occurs, create a filter and click Search. The object is displayed with more manageable data.

6. Scroll through the data table or filter data by clicking the View Filter icon (Magnifying Glass icon at the top of the table).

   Note: To use a wildcard to filter on a string, enter the string and then the % sign.

7. (Optional) Select a row or multiple rows.

   The Actions panel opens to the right of the object table and the Details pane expands under the object table:

   - The Actions pane lists navigational (see page 83) and administrative (see page 85) type actions available for the line item selected. If you select multiple rows, the Actions pane displays only the navigational and administrative type actions available for all the selected rows.
   - The Details pane displays a number of tabs dependent on the navigational (see page 83) type actions available for the selected line item. The first tab displays the selected line item details. The remaining tabs display information from other objects pertaining to the selected line item. If you select multiple rows, the Details pane displays the line item details and navigational type actions available for the last row selected.

As you drill down into storage object data, the Investigator table header includes information to indicate how you arrived at a piece of data. If you filter data, those values appear as header information in your results. As you drill into data, use the Actions pane on the right to identify dependencies and relationships among the records.

The object data toolbar provides additional options. Hover your cursor over the icons on the toolbar to see the available options.

For most objects, the number of records that are found is shown in the "Displaying *number_of* items at *date_time*" field at the bottom of the object. However for some storage objects, the quantity of records is shown as "Displaying *many* items at *date_time*" because the total number of records is not known.

If you have a multiple storage engine environment in CA Chorus for Storage Management, storage object rows display records from each storage engine available. A column is displayed in the storage object listing the storage engine the record data comes from. If a storage engine is not running or if the connection to a defined storage engine is not available, the data from the storage engine is not displayed in the object. You receive a warning message when you open the storage object if a defined storage engine is not available. Open the Storage Engine Information (see page 67) object to see storage engine information, status, and response statistics.

Note: In a multiple storage engine environment:

- The Main Engine is the storage engine that provides CA Chorus with storage object attributes, which include storage object table header information, the available Administrative Actions for the object, relationships to other objects, and so on.
- User-defined Summary Objects that exist in the Main Engine at the time of the JBoss server startup are included in the storage object tree. Use the Summary Objects Distribution Wizard in the storage engine windows client to distribute user-defined Summary Objects to your storage engines. For more information about the Summary Designer and the Summary Objects Distribution Wizard in the storage engine windows client, see the *CA Vantage User Guide*. This guide is available in the Knowledge Center.

More Information

Viewing Storage Object Performance Data in the Investigator (see page 89)
Storage Engine Information (see page 67)
How to Display Object Help (see page 65)
Filter Required for Very Large Storage Objects (see page 66)
View Storage Object Relationships in the Topology Viewer (see page 80)
View Storage Object Performance Data in the Time Series Facility (see page 91)
Storage Resource Objects (see page 67)

How to Display Object Help

You can select the Help (?) icon in the Investigator toolbar to find out more about the Storage object, including descriptions of the object, fields, and actions available.

The object online help displays in a separate window. This object information comes from the storage engine. The object help topic can contain information about the object that is not available in all of the storage engine user-interfaces. For example, the help can describe features, actions, and action parameters that are not available in CA Chorus, but they are available in the storage engine windows client or web client user-interfaces.

The following fields are available in most storage objects but are not described in their object help:

SysplexVan

Displays the name of the sysplex the storage engine is running on.

LparVan

Displays the name of the LPAR the storage engine is running on.

SubsystemVan

Displays the subsystem name of the storage engine.

Maximum Number of Rows Displayed in Storage Objects

The default value for the maximum number of rows displayed in storage objects in the Investigator is 5000 rows. You can set the maximum number of rows displayed in storage objects. The setting for the maximum number of rows is specified in the Dchorus.storage.numberOfRowsLimit variable in the ENVETJ member in the *chorus_runtime_hlq*.CETJOPTN job. For more information, see the *CA Chorus Installation Guide*.

Filter Required for Very Large Storage Objects

When opening very large storage objects, a message can be displayed stating that you must first create a filter (regardless of the maximum rows setting (see page 66)). If this message appears, create a filter, and click Search. The object is displayed with more manageable data.

The following is a list of some of the storage objects that require a filter:

- Data Sets for System (All)
- VSAM Data Set Attributes (VVDS)
- CA Vtape objects: CA Vtape Virtual Volumes and CA Vtape Virtual Volume Data Sets.
- CA Disk objects: Dsnindex Records, Archvols Records, Archcmds Records, Dasdspcb Records, Dmspools Records, Restcmds Records, Retexcld Records, and Migrecat Records.
- DFSMShsm Control Data Sets objects
- STK ACS Volumes
- BCS Data Set Entries
- Tape Management Systems: Volumes and Files objects.

Storage Resource Objects

The Storage Investigator tree lists storage objects by category (folders) and sub folders much like a file directory. Click the arrow next to the category folder and sub folders to expand the tree to display the list of objects in them. Click the object to open the object table in the interface.

Note: The Storage Engine Information (see page 67) object is displayed above the main category folders in the Storage Investigator tree.

Storage Engine Information

The Storage Engine Information object provides information about the storage engine (hosts) you have set up for CA Chorus for Storage Management. It provides storage engine identification, status, and statistical information.

When you select a row in this object the Details pane displays the row details in the following column group tabs:

Storage Engine Information

Provides identification information and status.

Query Statistics

Provides the high-level summary for all queries.

Storage Engine Query Statistics

Provides response summary for only the queries that are directed to the storage engine.

Hourly Query Distributions (Storage Engine)

Provides the number of queries per hour that is directed to the storage engine.

Hourly Query Distributions (All)

Provides the number of queries per hour for all queries.

When a user clicks an object in the user-interface, one or more queries are made to the data source to obtain the requested information. The queries are divided into the following types:

Constant Queries

These types of queries return the specified constant (the number 1 or a literal string, for example) in all rows for the appropriate column. These queries are sometimes used to find the number of rows in a requested object without having to transmit any data from the object.

Data Queries

These types of queries return rows of data from a requested object.

Metadata Queries

These types of queries return the attributes for an object, its fields, its actions (if any) and its related objects (if any).

Internal Queries

These types of queries return the status and statistics regarding each storage engine connection.

Note: To find out more about the columns and information displayed in the Storage Engine Information object, see the Storage Engine Information object help (see page 65).

Examples of when to use this object.

- You open a storage object in the Investigator, and you receive a message that a storage engine connection is not available. The storage object does not display rows of data from the storage engine that is not available. Open the Storage Engine Information object and verify the status of the host which is not displaying data. If it has status *not available*, contact your system administrator to verify communication with and status of the storage engine. There could be different reasons for the *not available* status. A couple examples are as follows:
  - The storage engine task needs a restart.
  - The communication between the storage engine and CA Chorus is lost.
- You open a storage object, and you notice that it is taking an exceptionally long time to display the data. Open the Storage Engine Information object and review the response times. If the response time is much higher on one particular storage engine (host), ask your system administrator to verify that the storage engine has enough resources and that the communication link between the storage engine and CA Chorus is working properly.

All Storage Solutions

The All Storage Solutions category folder contains all the Summary Objects available in the storage object tree. They are divided into two sub folder categories:

Factory-defined

Contains Summary Objects supplied with the storage engine from the factory.

User-defined

Contains Summary Objects that you have created. Use the Summary Designer in the storage engine windows client to create your own Summary Objects.

In a multiple storage engine environment, user-defined Summary Objects that exist in the Main Engine at the time of the JBoss server startup are included in the storage object tree. Use the Summary Objects Distribution Wizard in the storage engine windows client to distribute user-defined Summary Objects to your storage engines. If you create new Summary Objects after the JBoss server startup, you must restart the JBoss server to display them in the Investigator.

Note: For more information about starting and stopping the JBoss server, see the *CA Chorus Administration Guide*. For more information about the Summary Designer and the Summary Objects Distribution Wizard in the storage engine windows client, see the *CA Vantage User Guide*. Both guides are available in the Knowledge Center.

Summary Objects summarize the data present in any Source Object. For example, you want to know how many items match certain filter criteria, the average value of certain fields, or the maximum and minimum values applicable to a field. Each row in a Summary Object provides statistics on a Summary Group. A Summary Group is simply a subset of the Source Object data that you define using a *method*. The statistics available include counts, maximum values, minimum values, average values, and totals. These statistics are provided for each Summary Group and for all the records in the object table.

Summary Objects have an object name, an entry on the Object tree and behave the same way as all other objects. Summary Objects can take advantage of all storage engine services such as report distribution, historical logging, trending, and automation. These services let you distribute reports, generate alerts, take automated actions, and identify trends, which are all based on the summarized statistical information present in Summary objects.

The following methods are available to define the Summary Groups within an object:

By Value

This method defines Summary Groups that are based on the value of up to five fields that combine to produce a key. All records that have the same values in this key make up a Summary Group and are summarized in a single Summary Object row. A separate row is produced for every change or break in the value of the key. A single Summary Object row is included that provides overall statistics for the entire Summary Object.

An example of the Value Method is to summarize DASD volumes by the vendor field. In this example, there would be one Summary Group for each vendor (for example, IBM, EMC, HDS, and StorageTek) with selected statistics calculated for each vendor's share of the DASD volumes. The Value Method also provides overall statistics for the entire list of DASD volumes.

By Step

This method defines Summary Groups that are based on a numeric field and a step size. A single Summary Object row is produced for all the rows that fall into each Summary Group that is established by the field name and step size. A single Summary Object row is included, that provides overall statistics for the entire Summary Object.

An example of the Step Method is to summarize data sets by the block size field, with a step size of 4096. In this example, there would be one Summary Group for all data sets having block sizes from 0 to 4095, from 4096 to 8191, from 8192 to 12287, and so on. The selected statistics are calculated for the data sets having block sizes that fall into these ranges. Overall statistics for the entire list of data sets are provided.

By Filter

In this method, the content of each Summary Group is determined by a complete filter expression. Up to 1000 filters are allowed. All records that match each filter make up a Summary Group and are summarized in a single Summary Object row. If any records do not match any of your filters, the summary of those records is given a Summary Group name of UNASSIGNED and is displayed on a separate row. A grand total for all of the records (including UNASSIGNED) is displayed in the ALL RECORDS row. The Unassigned row, All Records row, or both can be turned off so that they do not appear in the table.

An example of the Filter Method is to summarize tape volumes by megabytes used. Where one filter is set for all tape volumes using less than 50 MB and another filter for tape volumes using more than 50 MB.

By Node

This method is similar to the Value method. But instead of using the entire value of a field to define the contents of a Summary Group, only a portion of a field value is used. The "." and "/" characters are used as delimiters. Each node is separated by a delimiter. The common example of a Node field is the data set name or directory name. A Summary Object Row is produced for all the rows that have the same values in the specified nodes.

For example, a data set name summarized by node 1 identifies all the unique high-level qualifiers being used, and produces a Summary Object row for each of them. A subtotal or summary line is produced for all the rows that have the same values in the specified nodes.

Note: For more information about Summary Objects, see the *CA Vantage User Guide*. This guide is available in the Knowledge Center.

CA SRM Open Systems Objects

The CA SRM category displays the CA SRM open systems objects. These objects provide an overview of your open systems storage environment. These objects are available if you have licensed CA SRM and configured the storage engine to display them.

Note: For more information about CA SRM object configuration, see the *CA Vantage Configuration Guide*.

The objects represent your open systems storage:

- RAID disks
- Disk arrays
- Logical disks of your disk array
- RAID groups
- All open systems storage system groups
- Total open systems storage by file system
- Total open systems storage by operating system
- Application Groups
- All Applications
- DB2 Databases and tables

Cost Analysis

The Cost Analysis provides automatic analysis of systems, highlights anomalies that might be occurring within them, and presents information from the analysis in both quantitative metrics about the technologies, and also applies cost information to the analysis.

Cost Analysis comes with a set of standard system analysis scripts that provide various cost and quantitative analysis of your environment to help you manage your data center costs. The calculations of the analysis are based on the values shown in the Customer Site Costs Data object and data provided by the storage engine.

The following is a list of the Cost Analysis factory supplied objects in the Cost Analysis folder:

Current Analyses

This object displays the current cost analyses results of your site. The cost metrics displayed are based on the values from the Customer Site Costs Data object and information obtained from the storage engine. Each row in the object is a current cost analytical scenario. Some of the Current Analysis object scenarios are:

- What are the costs of maintaining any excess DASD capacity?
- What are my annual Storage Administration costs?
- What is the overall cost wasted or unnecessary DASD?
- What are the costs of maintaining data sets not referenced in the past year?
- What are my total costs associated with excess and over-allocated DASD?

Historical Analyses

This object displays the historic cost analyses results of your site. The cost metrics displayed are based on the values from the Customer Site Costs Data object and information obtained from the storage engine. Each row in the object is a historical cost analytical scenario. That is, the table shows the same scenarios as the Current Analyses object but it has repeated scenario entries, one for each time period that the information is calculated.

Customer Site Costs Data

This object displays the values that the Cost Analysis uses along with object data obtained from the storage engine to calculate quantitative and cost metrics for the Cost Analysis objects. You configured the Cost/Value field values during the discipline installation in the Cost Analysis E4HI0011 configuration job. You can change the Cost/Value field values by editing them in the Details pane of this object, or by editing and resubmitting the E4HI0011 configuration job. For information about how to change these values, see the *CA Chorus Administration Guide*.

The following is a list of the factory supplied objects in the Base System Analysis subfolder under the Cost Analysis folder:

Current Base Analyses

This object displays the current quantitative analyses results of your site. The quantitative metrics displayed are based on the values from the Customer Site Costs Data object and information obtained from the storage engine. Each row in the object is a current quantitative analytical scenario. Some of the scenarios available in the Current Base Analysis object are:

- How much excess unused DASD capacity is being maintained?
- How much DASD space is being wasted by over-allocated data sets?
- How much DASD is being held by data sets not referenced in the past year?
- How much of my overall DASD usage is wasted or unnecessary?

Historical Base Analyses

This object displays the historic quantitative analyses results of your site. The quantitative metrics displayed are based on the values from the Customer Site Costs Data object and information obtained from the storage engine. Each row in the object is a historical quantitative analytical scenario. That is, the table shows the same scenarios as the Current Base Analyses object, but it has repeated scenario entries, one for each time period that the information is calculated.

Navigating Between Cost Analysis Object Scenarios and Storage Engine Objects

Cost Analysis comes with a set of standard system analysis scripts that provide various cost and quantitative analysis of your environment to help you manage your data center costs. The calculations of the analysis are based on the values shown in the Customer Site Costs Data object and data provided by the storage engine.

Each row in the Cost Analysis objects is an analytical scenario. You can select an analytical scenario row and then open the storage engine object that provided the data used with the information in the Customer Site Costs Data object to calculate the information displayed in the scenario. You can hop from the scenario to the storage engine object by using the Navigation Actions provided for the scenario. When you select a Navigation Action to open the corresponding storage engine object, you can return to the Cost Analysis object scenario by selecting the line item in the storage engine object and the appropriate Navigation Action.

Example of navigating between a Cost Analysis object scenario and storage engine objects

The following is an example of hopping between the Current Analysis object scenario: *What are the costs of maintaining any excess DASD capacity?* and the storage engine object Storage Groups. And then returning to the Current Analysis object scenario *What are the costs of maintaining any excess DASD capacity?*.

Follow these steps:

1. Log in to CA Chorus.
2. Add the Investigator module to your dashboard from the Module Library, and click Start New Investigation.

   The Investigator opens.
3. Select Storage from the discipline drop-down list.
4. Expand the Cost Analysis folder.
5. Click the Current Analyses object to open it.
6. Click the selection box at the beginning of the scenario (row) *What are the costs of maintaining any excess DASD capacity?*, to open the Actions pane on the right side of the Investigator.
7. Click the Navigation action *Show total online DASD Capacity*, to open the Storage Groups object.

   The Storage Groups object opens with the following filter applied:

   ```
   'SG Name' = TOTAL
   ```

   The Storage Group with the fixed name of TOTAL and type ONL is generated by the storage engine and contains the totals for all online volumes. You can scroll to the right and sort the table to find the information you are interested in.

   The displayed row is selected which causes the Action pane to open on the right side of the Investigator.
8. Click one of the following Navigation actions:

   Show current cost summaries

   Opens the Current Analyses object with the current scenario *What are the costs of maintaining any excess DASD capacity?* being displayed.

   Show history cost summaries

   Opens the Historical Analyses object with all the historical scenarios *What are the costs of maintaining any excess DASD capacity?* being displayed.

9. (Optional) Click the Current Analyses object in the Storage tree, to display all the scenarios in the object.

   You have navigated between the Current Analyses object and the Storage Groups object. Viewing the Storage Groups object allows you to see space availability and usage details in your storage groups.

Diagnostics and Investigations

The Diagnostics & Investigations category folder contains sub category folders which contain diagnostic and investigation type objects. These objects provide information about your storage resources.

Automation and Logging

The Automation and Logging sub category folder has the following sub folder categories:

- General Object Automation (GOA)
- Message Automation (MA)

The objects in the Automation and Logging category folder let you view the automatic actions that are taken when user-defined schedules, events, and conditions in the z/OS environment are met. In addition, you can view storage engine system activity that is related to automation and logging events.

You can execute General Object Automation (GOA) scripts immediately by selecting it in the Object Automation Scripts object and selecting the *Fire* Administrative Action.

Note: For an example how to execute GOA scripts immediately, see How to Free Up Storage Space Add Storage Space by Executing An Automation Script (see page 46).

Backup and Archive Management

The Backup and Archive Management sub category folder contains objects that display information related to back up and archive for the following major archival and backup systems:

- The CA Disk system
- The IBM DFSMShsm system

Catalog Management

The Catalog Management sub category folder contains objects that let you view information about the catalog defined in the system to which you are connected.

Data Set Groups (Applications)

The Data Set Groups (Applications) sub category folder contains objects that display:

- Group Space on Primary Storage
- Group Space on Secondary Storage
- Group Space on Tape Storage

These objects give you a complete view and analysis of the space usage by each Data Set Group or application defined in the storage engine.

Data Set Management

The Data Set Management sub category folder contains objects that display data set information.

The storage engine gathers attributes of data sets from different sources, such as the VTOC and VVDS. The storage engine combines them into a memory-resident Data Set Table Of Contents (DTOC). This combination provides almost instant access to information about all of your online data sets. The DTOC data is indexed and analyzed, such that your data sets can be viewed in multiple ways, including:

- All Data Sets
- Duplicate Data Sets
- HFS and PDSE (POE) Data Sets
- Multiple Volume Data Sets (Summaries)
- Multiple Volume Data Sets (Details)
- Uncataloged Data Sets
- Cataloged Data Sets Not On Disk

Database Management

The Database Management sub category folder contains objects that display space-related information for the following systems, including:

- Adabas
- CA Datacom/DB
- CA IDMS/DB
- DB2

Encryption Key Management

The Encryption Key Management sub category folder contains CA Encryption Key Manager objects. If you use CA Encryption Key Manager, you can work with CA Encryption Key Manager objects that display the full lifecycle of your cryptographic keys.

Software-based cryptographic applications and hardware-based cryptographic devices are supported. The following information is provided:

- The creation, use, and retirement of cryptographic keys
- The creation and export of digital certificates

Storage administrators typically use this information to manage the keys used in cryptographic processing. Proper management is critical to the usefulness and reliability of a cryptographic solution.

Hardware Errors - Disk/Tape

The Hardware Errors - Disk/Tape sub category folder contains objects that display temporary and permanent I/O errors on DASD and tape.

Performance Management (DASD)

If CA ASTEX is installed and configured on the LPAR to which you are connected, the *Performance Management* sub category folder is available. The *Performance Management* sub category folder contains objects that display direct access storage device (DASD) and cache levels of the storage hierarchy.

Storage Platform Administration

The Storage Platform Administration sub category folder contains objects that display storage engine internal information such as memory usage, CPU usage, debug information, and so on.

Tape Resource Management

The Tape Resource Management sub category folder contains objects that display tape processing activity on your local and virtual tape systems.

If activated on the storage engine system to which you are connected, audit of tape usage can also be found in this category.

z/OS System Resources

The z/OS System Resources sub category folder contains objects that display different system aspects in the LPAR where you are connected. These objects include APF, Linklist, and enqueue information. In addition, it contains objects that display z/OS-defined system symbols and parameters.

Reports and Views Utilization

The Reports & Views Utilization category folder contains the following *modeling type* objects:

Archive Data Set Modeling

This modeling object provides the number of data sets and the space that would be archived if you adopted an archive policy based upon the value in the Days Unused field.

Alloc/Idle Bytes Blksize

This modeling object summarizes data sets according to the block size (blksize) field. Data sets with the same block size are grouped together and analyzed. The number of data sets of each block size is shown, along with other summary statistics about the allocated space, idle space and number of extents for each block size group.

Alloc/Idle Bytes Dsorg

This modeling object summarizes data sets according to their data set organization type (Dso or Dsorg) field. The Dsorg can be physical sequential (PS), partitioned (PO), virtual storage access method (VS or VSAM), and so on type data sets. The number of data sets of each Dsorg is shown, along with other summary statistics about the allocated space, idle space, and number of extents for each Dsorg group.

Statistics

The Statistics category folder displays the object performance data for the Time Series Facility (TSF). TSF stores data that is collected and provided by the back-end storage engine.

Note: For more information, see View Storage Resource Statistics in the Investigator (see page 89) and View Storage Object Performance Data in the Time Series Facility (see page 91).

Volumes and Utilization

The Volumes & Utilization category folder contains sub category folders which contain storage resources and utilization type objects. The following sections are the sub category folders included in this category folder.

Allocation Control

The Allocation Control sub category folder provides allocation information if you have CA Allocate. If you have CA Allocate then the following objects are available in the CA Allocate sub folder category:

CA Allocate Quota

The Quota component of CA Allocate allows you to establish, monitor, and enforce disk space quota limits for individuals or groups of users. The Quota Table is dynamically maintained in virtual storage by CA Allocate as data set allocation activity occurs.

CA Allocate Spare Volumes

The relation between CA Allocate and CA Vantage is characterized by the following interactions:

- CA Vantage can read the storage groups (pools) defined by CA Allocate and present them through its presentation managers.
- CA Vantage can monitor CA Allocate storage groups for threshold violations and take corrective actions by submitting batch jobs.
- CA Vantage can add spare volumes to a CA Allocate storage group that has been detected to be in violation.

DFSMS Constructs

The DFSMS Constructs sub category folder contains objects that display information related to constructs defined and used in the IBM Data Facility Storage Management Subsystem (DFSMS) environment.

RAID Devices

The RAID Devices sub category folder contains objects that display in-depth information about the following Redundant Array of Independent Disks (RAID) devices:

- IBM
- HDS
- EMC SMI-S APIs

Storage administrators typically use this information to determine an individual RAID subsystem performance and capacity data.

Storage Groups

The Storage Groups sub category folder contains objects that display storage capacity on storage groups on the LPAR to which you are connected.

Supported storage groups are the ones that are defined in MVS, DFSMSMS, CA Allocate, and user-defined (USR) storage groups.

Storage administrators typically use this information to analyze storage group space usage.

UNIX for zSeries

The UNIX for the zSeries sub category folder contains objects that display directory and file usage statistics from zSeries Linux File Servers.

Volumes (DASD)

The Volumes (DASD) sub category folder contains objects that display storage capacity on online volumes on the LPAR you are connected to. In addition, it lets you view information about offline volumes.

Storage administrators typically use this information to analyze volume space usage.

View Storage Object Relationships in the Topology Viewer

CA Chorus for Storage Management lets you view storage object relationships using the Topology Viewer. The Topology Viewer provides a graphical view of storage resources in your system and their relationships. This view can simplify your ability to identify relationships as you manage your storage resources.

A graphical view is useful when you want to see the relationships of storage resources and their dependencies in your storage environment. As you drill down in the Topology Viewer, your storage environment is displayed in hierarchal, graphic format.

You can launch the Topology View from any object, by selecting a line item in a storage object and clicking the *Replace Topology Viewer* or *Add to Topology Viewer* navigational actions from the Action pane in the Investigator. However, not all Storage objects include topology relationships. That is, while any object line item can be added to the Topology Viewer, some objects do not have *child* relationships to other storage resources which can be drilled down to.

For a top down topology view of your storage resources and their relationships in your system complex (sysplex) or the LPARs you are connected to, start the Topology Viewer from the Sysplex Information Display or LPAR Information Display objects. These objects are located in the Diagnostics & Investigations, z/OS System Resources directory in the Storage object tree.

The Topology Viewer remembers the graphical view you last displayed in the Topology Viewer (with the expansions, filters, zooms, and so on, you applied the last time you opened the Topology Viewer). If you want to start with a fresh display, select the Replace Topology Viewer Navigation action instead of the Add to Topology Viewer action from the Action pane in the Investigator.

Note: The Topology Viewer comes with default configuration. You can change this configuration to better display your storage resource environment. For example, the number of DASD volumes that are installed in most environments is very large so the Topology Viewer cannot realistically provide a visual representation of each of these volumes on a single mapping. You can change the Topology Viewer configuration to sub-divide and group online DASD volumes. You can also customize online DASD maps where DASD is logically grouped according to various characteristics of the volumes such as by application, line-of-business, and so on. For configuration information, see the *CA Chorus Administration Guide*.

Example:

You want to get a graphical overview of your storage resources in your sysplex so you can easily see your storage resource architecture.

Follow these steps:

1. Log in to CA Chorus.
2. Add the Investigator module to your dashboard from the Module Library, and click Start New Investigation.

   The Investigator opens.
3. Select Storage from the discipline drop-down list.
4. Expand the Diagnostics & Investigations, z/OS System Resources directory in the Storage object tree, and open the Sysplex Information Display object.

   The Sysplex Information Display object opens and displays the sysplex in the object table.
5. Click the sysplex row if it is not already selected.

6. Click one of the following actions in the Actions panel:
   - Replace Topology Viewer

     Opens the Topology Viewer and replaces existing hierarchal, graphics with the selected item.
   - Add to Topology Viewer

     Opens the Topology Viewer and adds the selected item to existing graphics.

   The selected item is displayed in the Topology Viewer.
7. Right-click the Sysplex Information Display box in the Topology Viewer, and select the LPAR Information Display box.

   The LPAR Information Display diagram is displayed in hierarchal, graphic format under the System Information Display diagram. The text box informs you of the number (in parenthesis) of LPARs available in the sysplex .
8. Right-click the LPAR Information Display diagram, and click Expand from the pop-up menu, to display the LPARs in the sysplex in hierarchal, graphic format.
9. Right-click one of the LPAR diagrams you are connected to (if you have a multiple host environment), to display a list of the storage resources available in the LPAR.
10. Select a storage resource from the list to add it to your topology diagram.
11. (Optional) Continue to drill down to the storage resources in this fashion, to add items to the topology diagram and view their relationships in the sysplex.

    The Topology Viewer continues to display the upper part of the hierarchal graph. You can, for example, repeat steps 9 through 11 and select a different storage resource available in the LPAR, to start a new branch in the hierarchal graph. You can remove a branch, by right-clicking a storage resource in the branch and selecting Remove.
12. (Optional) Use the Topology Viewer toolbar icons to customize the diagram, zoom in and out, print, and so on. Hover your cursor over the toolbar icon to see a description of the tool.

    A graphical overview of your storage resources in your sysplex is shown, which lets you easily see your storage resource architecture.

Perform Storage Actions on Object Data in the Investigator

The two types of Actions available when a storage object row is selected in the Investigator are:

Navigation

The navigational types of actions available include:

Add to Topology Viewer

Adds the selected item to the Topology Viewer. For more information, see View Storage Object Relationships in the Topology Viewer (see page 80).

Add Entity to Time Series

Adds the selected item to the Time Series. This action is only available for object line items of objects in the Statistics folder. For more information, see View Object Performance Data in the Time Series Facility (see page 91).

Show other (related) objects to the row that is selected

Displays associated data in another object. To obtain more information about a selected record (row), use this to navigate to another object, if the selected record has associated data available in another object. For example, the Storage Groups object lists the storage group name and their space attributes. you can select a record and navigate to the Volumes object to see all the Volumes associated with the selected Storage Group.

Administrative Actions

An Administrative Action is a storage engine function or service that can be performed on an object or object attributes (object row). Depending on the object and the line item that is selected, the Administrative Actions listed are storage engine actions that can be performed on that line item, and in some cases the whole object. The selectable actions that are listed relate to specific functionality of the object being accessed.

How to Perform Cross-discipline Data Set Investigation with Security Discipline

You can perform the following cross-discipline Navigation actions with CA Chorus for Security and Compliance when you select a storage discipline object that contains a data set name in any format:

- Show Users of selected data sets.

  Opens the security discipline object that lists users and their access privileges that have access to the data set in the object row selected.

- Show Violations on selected data set.

  Opens the security discipline object that lists violations for the data set in the object row selected. Violations can be users that tried to access the data set but did not have authority.

The cross-discipline actions are available at sites where the security discipline and storage discipline are implemented and are up and running.

Example of how to use the navigational action *Show Violations on selected data set* in the Investigator.

1. Select Storage from the discipline drop-down list.
2. Expand the Diagnostics & Investigations, Data Set Management, All directory in the Storage object tree, and open the Data Sets for System (All) object.

   The Data Sets for System (All) object opens in the Investigator.
3. Create a filter to find the data set you want to investigate violations on.
4. Select the data set.

   The Action pane opens.
5. Select the navigational action *Show Violations on selected data set*.

   The security discipline object opens that lists violations for the data set in the selected object row. If the violations are a concern, contact your security administrator for further investigation and remediation.

   Note: You can return to the storage discipline object by clicking the appropriate storage discipline object bread crumb that is displayed at the top of the Investigator window.

Show Related Object Navigational Actions

The Storage Object Navigation Actions category includes the Show in *object_name* navigational type action. This action displays the selected item in another object. In the storage engine, these options are referred to as *zooming* to another object. When you select a row in an object, the Navigation section of the Actions pane lists *show actions* for the other storage objects the item attributes are available in. You can then select a Navigation *show action* to drill down to the related object based on the information in the selected row.

For example, if you start with an object display of online DASD volumes, and you select the row with one of the volumes of interest. The Actions pane opens with a Navigation *show actions* list of *related objects*. Such as the data sets (that reside on that volume) or the VTOC attributes (for that volume).

You can also select the *show action* tab in the Detail pane that expands at the bottom of the Investigator window when you select an object line item. When you select a *show action* tab in the Detail pane, the details of the item in another object is displayed.

The difference between selecting *show actions* in the Details pane and the Actions pane is:

- When you select the *show action* in the Actions pane, the information the selected line item information from the object is displayed in the Investigator table pane.
- When you select the *show action* in the Details pane, the related line item information from the selected *show action* object is displayed in the Details pane and the table in the Investigator table pane does not change.

You can also perform a cross-discipline data set investigation with security dole data set object information, for more information see How to Perform Cross-discipline Data Set Investigation with Security Discipline (see page 84).

Storage Object Administrative Actions

An Administrative Action is a storage engine function or service that can be performed on an object or object attributes (object row). Depending on the object and the line item that is selected, the Administrative Actions listed are storage engine actions that can be performed on that line item, and in some cases the whole object. The selectable actions that are listed relate to specific functionality of the object being accessed.

Observe the following notes:

- CA Chorus users require authority in the storage engine to perform Storage Administrative Actions in CA Chorus (see page 87).
- Not all objects have Administrative Actions.

- The available Administrative Actions differ from one object to the next, the procedure for performing an action is similar for all objects.
- Some Administrative Actions are simple, requiring only selecting an option button or require text entries. For example, when renaming a data set in the Action dialog.
- Some Administrative Actions have parameters, and you are required to enter the parameter data in the Action dialog. For example, the Release Idle Space action asks for the percent of idle space you want to keep.
- If an Administrative Actions requires parameters and you select multiple rows, the action is not available. You can only select one row at a time to perform these types of actions. For example, if you want to *Rename* data sets, you use the *Rename* action. When using the *Rename* action, you can only select one data set row at a time, and you can rename it. However, you can select multiple rows and can use the *Delete* action to delete multiple data sets.

More Information

User Authority to Perform Storage Object Administrative Actions (see page 87)

How to Perform Storage Administrative Actions

An Administrative Action is a storage engine function or service that can be performed on an object or object attributes (object row). Depending on the object and the line item that is selected, the Administrative Actions listed are storage engine actions that can be performed on that line item, and in some cases the whole object. The selectable actions that are listed relate to specific functionality of the object being accessed.

To perform an Administrative Action

1. Select a line item in the object table displayed in the Investigator.

   The Action pane is displayed on the right side of the Investigator window with a list of Navigation and Administrative Actions.

   Note: In some cases, if an action requires parameters and you select multiple rows, the action is not available. That is, it does not appear in the Actions pane. If you want to perform an action on multiple rows, select one line item and the Action pane opens where you can see the action that you want to perform. Select another row and see if the action disappears from the Action pane. If it does disappear, then you can only perform that particular action one line at a time.

2. Select an Administrative Action from the list.

   An Action dialog pertaining to the selected action is displayed.

3. Enter the parameter values appropriate for the action you want performed or skip this step if no entries or selections are required.

4. Click Submit.

   The storage engine performs the action and a confirmation dialog is displayed.

Examples: Administrative Actions available with Data Set objects

When working with Data Sets objects, some of the actions available are:

- Migrate
- Release Idle Space
- Rename
- Move
- Copy
- Delete

Note: To see examples that use these actions, see the chapter "Introduction".

User Authority to Perform Storage Object Administrative Actions

When security support is active in the storage engine, the storage engine lets users request one or more Administrative Actions against any selected object or object row. For example, compress a data set or scratch a tape. The storage engine uses the SAF interface to verify each user authority to perform the requested Administrative Actions.

Note: For more information, see the *CA Vantage Installation Guide* and the *CA Vantage Reference Guide*. Both guides are available in the Knowledge Center.

Chapter 3: Viewing Storage Object Performance Data in the Investigator

Storage Resources Performance Data

CA Chorus for Storage Management lets you view, monitor, and compare storage resource statistics in the Investigator.

There are two ways to show object performance data in the Time Series Facility (TSF):

- You can view storage resource Statistics (TSF object data) in side-by-side charts in the Investigator (see page 89). The Statistics category in the object tree displays the object performance data in the TSF. TSF stores data that is collected and provided by the back-end storage engines.
- You can display object performance data in the Metrics panel (see page 93). You can configure the Metrics panel to show specific TSF object metrics in the panel, set Metrics thresholds, investigate TSF data, and create your own Metrics dashboard. The data shown in the Metrics panel comes from the objects in the Statistics category (folder) in the object tree.

Note: For a list of storage objects available for collecting data for TSF, see the section View and Investigate Storage Metrics (see page 93).

View Storage Resource Statistics in the Investigator

Use the following example procedure to review performance statistics in the Investigator for your storage data in table format or by creating side-by-side charts. Reviewing the information in side-by-side charts helps you identify trends.

Observe the following points about the Filter for TSF objects:

- The Filter in the Investigator for TSF objects is case-sensitive when you use some of the operators. For example, when you use the = (equal sign), less than - greater than (<> sign), and *Like* operators. Any test value that you enter in a filter where you use these operators for a TSF object must be in the same case as it is displayed. For storage TSF objects, field values are usually all upper case. For example, if a volume in TSF is shown as TSO123, it must be entered in the = (equal sign) operator filter query as TSO123 (all upper case) to find the match. Entering tso123 results in "no matches", as does any other entry that contains some lower case characters. In all non TSF storage objects (those objects not in the Statistics category), the Filter is not case-sensitive.

- Some TSF Storage Group objects may have spaces between the name and before its type (MVS, SMS, VAM, or USR) in the Storage Group Name field. When you use the Filter to search specifically for these names with spaces, use the *Starts With* or *Contains* operators with just the first or last part of the Storage Group Name (without the blanks). For example, to search for the DISK MVS storage group, use the following filter statement:

  ```
  Storage Group Name Starts With DISK
  ```

  You can also use the % wildcard character which means *any characters* when placed between two words, if you are not sure if there are no spaces, one space, multiple spaces, or other characters between two words. For example, you could create the following filter statement:

  ```
  Storage Group Name Like DISK%MVS
  ```

  The *Like* operator means *includes*, so this filter would find all storage groups that have DISK, then followed with no characters or any characters, and then followed with MVS, anywhere in the Storage Group Name string.

Note: You can also display TSF object performance data in the Metrics panel (see page 93).

Example: Monitor Storage Groups statistics charts in the Investigator

Your team monitors available space on storage groups so applications have enough space and to avoid space abends. Having sufficient space on storage groups improves the efficiency of your site. Retrieve a baseline of acceptable levels so you can compare current conditions against the baseline and measure performance.

Follow these steps:

1. Log in to CA Chorus.
2. Add the Investigator module to your dashboard from the Module Library, and click Start New Investigation.

   The Investigator opens.
3. Select Storage from the discipline drop-down list.
4. Expand the Statistics, Storage Groups subcategory, system, LPAR directory in the Storage object tree.

   Storage Group objects that have been set up to collect storage object metrics for TSF (see page 95) are displayed.
5. Select the Storage Group metrics object for which you want to view the metrics.

   The metrics object table opens.

6. Perform *one* of the following sequences:
   - To review the data using the Time Series Facility (TSF):
     a. Select the record in the table.
     b. Click Add Entity to Time Series in the Actions pane on the right-hand side.
     c. Select the Entity and Metrics for comparison, and click Perform Charting.
     d. Click the icon next to the Contributors button to duplicate the chart.
     e. Modify the End Date/Time or Period for each chart to compare a *baseline* view with a more current view.
   - To review all of the data in a tabular view while in the table view:
     a. Select the wrench icon.
     b. Add all of the Storage Group fields to the grid, and click Save.
     c. The Storage Group statistics appear in the table in the center pane.
     d. Click the export icon in the upper-right corner of the center pane.
     e. Save the table data to a comma-separated value (CSV) file.
7. Repeat these steps periodically and compare the data in TSF graphs or the CSV file. Keep the Storage Group statistics at acceptable levels.

More Information

How Storage Metrics Are Aggregated (see page 94)

View Storage Object Performance Data in the Time Series Facility

The *Time Series Facility* (TSF) stores data that is collected and provided by CA products. TSF provides a single point for collection, storage, management, and organization of the product data. When you request a Time Series chart from the Investigator, TSF provides the data content for the chart.

The TSF displays storage statistics. That is, data from the objects in the Statistics folder in the object tree. When you add a storage object to TSF, it becomes an entity, and you can use it to create a chart. Doing so lets you determine trends and identify valid anomalies. The TSF charts allow up to four lines.

An *entity* is an application performance object that has been added to TSF. You can add the system entity to TSF with the CA Chorus for Storage Management discipline.

Note: For detailed common TSF concepts and procedures, see the *CA Chorus Product Guide*.

Note: You can also display TSF object performance data in the Metrics panel (see page 93).

Example: Analyze storage group usage in the Investigator

In this example, the Storage Groups object is added to the TSF and the total space per group is analyzed. By monitoring the ratio between total storage capacity bytes, allocated bytes, and free bytes; you can evaluate the space usage compared to the free bytes.

The Storage Groups object provides a numeric description of space usage in each Storage Group. Each Storage Group is a collection of one or more volumes. They can be MVS esoteric names, DFSMS, CA Allocate, or storage engine system user-defined Storage Groups.

The Storage Group with the fixed name of TOTAL ONL displays totals for all volumes.

Follow these steps:

1. Log in to CA Chorus.
2. Add the Investigator module to your dashboard from the Module Library, and click Start New Investigation.

   The Investigator opens.

   Select Storage from the discipline drop-down list.
3. Expand the Statistics, Storage Groups subcategory, system, LPAR directory in the Storage object tree.

   Storage Group metric objects that have been set up to collect storage object metrics for TSF (see page 95) are displayed.
4. Select the Storage Group metrics object to which you want to view the metrics.

   The metrics object table opens.
5. Select the record in the table.
6. Click Add Entity to Time Series in the Actions pane on the right-hand side.

   The system is added to the TSF and the TSF panel appears.
7. Select the Entity and Metrics, and click Perform Charting.

   TSF produces a chart for the selected metric. The selected entity becomes a line on the chart.

8. Select an entity from the Contributors drop-down list and click Contributors.

   The Entities panel becomes the Base Entities panel and shows the original criteria that is passed by the Investigator for the selected entity.

9. Click the Contributors drop-down list.

   The drop-down list shows the valid contributors.

   Note: To exit the contributors function, click Back to Entities.

10. Select a contributor type.

    A list of all the contributors for the selected entity and metric combination appears. If no contributors appear, then clear one or more base entities.

    Note: All Storage metrics are defined as “averages”. When multiple values are rolled up for a time period, or across multiple entities, the average value is represented on the Time Series charts.

11. Click Perform Charting.

    The new chart is generated from the selected entities.

    The chart provides an overview of storage group attributes and allows you to easily see historical and current attribute values to evaluate the space usage compared to the free bytes.

More Information

How Storage Metrics Are Aggregated (see page 94)

View and Investigate Storage Metrics

You can send various metrics about storage usage from the storage engine to CA Chorus for Storage Management for display in TSF (trend) graphs. That is, you can see how one or more storage metrics vary over a selected time period – the historical usage trend for each metric.

Storage metrics are available for certain fields in the following Storage objects, including:

- Volumes (DASD) – Space and Other Attributes
- Storage Groups – Space and Other Attributes
- Catalog Object – On Volume (Lspace Information)
- Catalog Object – CAS Catalog I/O Statistics
- Catalog Object – CAS Catalog Cache Report
- Data Set Management Object – Application Space Usage on Primary Storage

To add TSF Metrics to your workspace, select “Click here to configure the Metrics panel” at the top of the interface window and follow the prompts.

Note: For an example of how to configure the Metrics panel, set Metrics thresholds, and create your own Metrics dashboard, see How to Configure the Metrics Panel to Monitor Storage Group Space Usage (see page 53).

You can also view storage resource Statistics (TSF object data) in table format or side-by-side charts in the Investigator (see page 89).

More Information

How Storage Metrics Are Aggregated (see page 94)

How Storage Metrics Are Aggregated

Time Series uses aggregation techniques to display metrics across large time intervals when long period values are chosen for the charting action. The charting tool uses this aggregated data to represent one or more actual data values that are sampled across the time interval that is dictated by the selected period.

When you select a long time period value to chart, the interval is adjusted to display information in a more readable way. The bigger the interval, the more potential data that is sampled and used for the aggregated value. This value displays as a single plot point at the end of the time interval. When multiple sampled data points fall into an interval window, those points are aggregated using a calculation method of averaging or summing.

All storage metrics are defined to use an averaging method to calculate the aggregated data points. When multiple values are rolled up for a time interval or across multiple entities, the average value is represented on the Time Series charts. So, for long period time ranges, the displayed plot values could be aggregating one or more samples plotted at the end of the time interval.

Example 1 Storage Interval Aggregation

The selected chart period dictates that hourly intervals are plotted. For the time range 0100-0200, Volume ABC001 has two data samples for the selected metric (100, 150). For the metric chosen, these two values are averaged and represented as 125 on the chart.

Example 2 Storage Contributor Base Entity Aggregation

The user has accessed the TSF Contributors. The user unchecked the Volume Name in the Base Entity criteria to chart a metric across all volumes for a system.

For the time range 0100-0200, each volume has only one value. For ABC001=100, ABC002=200, ABC003=150, ABC004=50, these values are averaged to 125. If the user plots contributors, an individual contributor, such as ABC003, could plot above the chart for the combined result.

How to Collect Storage Metrics for TSF

To create the storage metrics for TSF, the following items must be in place:

- The storage engine system parameter CHTSFSUF must be set. Your system administrator typically sets this parameter when the storage engine was configured for this discipline.

  Note: For more information, see the *CA Vantage Configuration Guide*. This guide is available in the Knowledge Center.

- Create storage engine log scripts (see page 96) for the objects listed in the section View and Investigate Storage Metrics (see page 93). Verify that they run as often as desired. When a log script runs, the new metrics for that object are automatically sent to the TSF in the Investigator. The storage administrator typically creates log scripts after the storage engine is configured.

  Note: If your system installer performed the startup configuration steps for new CA Chorus and new CA Vantage customers by running the CHRSGMI job, then the storage engine log scripts are already created. For more information, see the *CA Vantage Configuration Guide*.

- The target CA Chorus for Storage Management TSF systems must be up (ready to receive the metrics).

Create a Storage Engine Log Script to Generate Storage Metrics for Time Series

To generate storage metrics for TSF, a storage administrator must create storage engine Log Script for the object. When the log script runs, the new metrics for that object are automatically sent to TSF. The storage administrator typically creates Log Scripts after the storage engine is configured.

Observe the following notes:

- The storage engine systems must be up and running to generate storage metrics for TSF. If they are not up and running, the metrics are not created and sent to the target CA Chorus systems.
- Use the storage engine windows client to create the scripts.

  Note: For storage engine windows client acquisition, requirements, and installation instructions, see the *CA Vantage Installation Guide*. This guides is available in the Knowledge Center.
- If your system installer performed the startup configuration steps for new CA Chorus and new CA Vantage customers by running the CHRSGMI job, then the storage engine log scripts are already created. For more information, see the *CA Vantage Configuration Guide*.

Follow these steps:

1. Open the storage engine windows client.

   Note: For installation instructions, see the *CA Vantage Installation Guide*.
2. Connect to the back-end storage engine host that is set up for CA Chorus for Storage Management.

   Note: For instructions about how to create a host connection definition and how to connect to a host, see the *CA Vantage Windows Client Guide*.
3. Click z/OS, Object Scripting, Object Logging.

   The Log Scripts Builder opens.

   Note: Click the Help button on any page of the wizard to get a description of the options on the page.
4. Click Next until the Source Object Selection page opens.

   The Source Object Selection page displays the object tree.

5. Expand the object tree and locate the object that you want to create the log script for.

   Note: Only certain storage objects are set up for TSF. For CA Chorus for Storage Management, you can make log scripts for *only* the following storage objects:

Windows Client Object Tree Folder - Subfolder	Object Name
Volumes (DASD) - Online	Space and Other Attributes
Storage Groups	Space and Other Attributes
Catalog Management – Defined Catalogs	On Volume (Lspace Information)
Catalog Management – Catalog Address Space (CAS)	Catalog I/O Statistics
Catalog Management – Catalog Address Space (CAS)	Catalog Cache Report
Data Set Groups (Applications)	On Primary Storage

6. Select the object for which you want to create the log script.

   The object is highlighted.

7. Click Next.

   The Title and System Execution page opens.

8. Enter a title and optionally a description. If you want the script to run on multiple systems, click the Systems button and specify them.

9. Click Next.

   The Timing for Events page opens. Use this page to specify the schedule when the script is run.

   Note: For more information about log scripts and optimal configuration strategies, see the section Avoid CPU and memory Constraints in the *CA Vantage Best Practices Guid*e. This guide is available in the Knowledge Center.

10. Click Next.

    The Options page opens.

11. Select the following options:

    - Load Script
    - Enable Script when Loaded

12. Click Next.

    The Event Procedure Section page opens.

13. Click Next.

    The Filter page opens.

14. (Optional) Click Filter to create a filter. Then click Define Filter to display the Filter dialog. Define a filter to list only the items you want to create the log for. Click Save to save the filter.

    The filter is created and the Filter dialog closes.

15. Click Next.

    The Log file Creation page opens.

    Use this page to specify:

    - How often a new log file is created (every time, every day, every week, every month, or only when the log is full).
    - Specify the maximum number of logs per file.
    - Specify if you want to log data, object count, or both.
    - Specify whether to use the data in one of the following scenarios:
        - Acquired during the last scanning interval. Select Base snapshot on last Interval.
        - Collect the data on demand before taking the snapshot of the object. Select Base snapshot on real-time data.

    Do not select *Write to SMF*.

16. Click Next.

    The Local Database properties page opens.

    Do not select *Save in Local Database*.

17. Click Next.

    The Saving and Activation page opens with the following options selected:

    - *Execute the script every time it is activated.*
    - *Save and Activate the script.*

18. Click OK.

    The log script for the object is created, saved, and activated.

    Note: Repeat this process to create a storage engine Log Script for each object you want to generate storage metrics for the TSF.

19. Close the windows client after you have made all the log scripts needed.

    The log scripts are created and collect data in the schedules specified.

    Note: You can reopen the windows client to create new scripts, change existing script settings, and delete scripts at any time.

Alerts for Storage Conditions

An alert is a text- and visual-based indicator of abnormal behavior on your logical partitions (LPARs). The Alerts module lets you monitor and investigate alerts from the dashboard as they are generated. When a defined processing limit is reached or exceeded on an LPAR, an alert is issued.

To view alerts in the Alert module, add it to your CA Chorus dashboard and follow the prompts to configure it.

Alerts can be generated two ways:

- Alerts generated by the CA Chorus Policies module (see page 99).
- Alerts generated by the General Object Automation (GOA) component of the storage engine system (see page 101).

Policies Module Alerts

Use the Policy Status Light module to create simple policy alerts. These types of Alerts are better suited for when one field in an object row meets a certain condition of which you want to be notified of.

Example set up of idle data sets Policy Light and Alert

As a storage administrator you want to be alerted of any online data sets that have not been used for a lengthy period of time so you can move them to cheaper storage media. Data sets that have not been used for a long period of time waste online space. This example explains how to set up a Policy Light and an alert when more than 10 online data sets have not been used for more than 365 days. Change values in the example to meet your site specific storage management policy

Follow these steps:

1. Log in to CA Chorus.
2. Add the Policy Status Light module to your dashboard.
3. Configure the Policy Status Light module.
4. Select Storage from the discipline drop-down list and then select Chorus Personal Policies.
5. Click Create New Policy in the Actions pane.

   The Details pane expands with the General Details tab selected.

6. Enter *Idle Data Sets GT 365 days* in the Name field and in the Description field enter *Number of data sets that have been unused for at least 365 days*, select Public, and click Next.

   The Objects tab is selected in the Details pane.

7. Click Next and then Modify to open the Object Picker window and the Storage tree.

8. Expand the Storage tree to the directory Diagnostics & Investigations, Data Set Management, and Summaries.

9. Select the object Archive Data Set Modeling to open it.

   This object is an archive modeling table. It lists the number of data sets and the space that would be archived if you adopted an archive policy based upon the value in the Days Unused field.

   Note: This object provides a list of unused data sets in increments of days. In this example we use 365 days, but you can choose the unused days row that best meets your site's storage policy.

10. Scroll down the table and select the check-box in the row with 365 in the Days Unused column.

11. Click the Select button in the menu bar.

    The line item is added to the Objects tab in the Details pane of the Object Picker.

12. Click Next to open the Rules tab in the Details pane.

13. Click the down-arrow in the Parameter field and select Number of Dsn (number of data sets).

14. Click the down-arrow in the Threshold field and the > (greater-than) operator, enter 10 in the value field, and click Next.

    Note: Enter a value that best meets your site's storage policy.

    The Actions tab opens in the Details pane.

15. Do the following:

    - Select Change Traffic Light Color and Red.
    - Select Send Alerts to Alerts Module and click Configure (it displays when you select Send Alerts to Alerts Module). In the Alert Details dialog select Severity 1 from the Severity drop-down list and enter the message *Number of data sets that are idle for more than 365 days exceeds 10* and click Save.
    - In the Action Management section enter 1 in the Count of Events field and 24:00:00 in the Duration field.
    - In the Action Suppression section enter 24:00:00 in the Suppression Duration field.

16. Click Finish to add the Policy to the Chorus Personal Policies object.

    Note: You can select the policy in the object and select the Edit icon in the Details pane to make changes to the policy.

17. Select the check-box next to the policy and click Save to close the Object Picker window.

    The Edit Title/Add Description dialog opens where you can change the name and enter the description *Investigate data sets that are idle for more than 365 days*.

18. Click Save.

    The Policy Light is added to the dashboard. If the light is green, the number of data sets that have been idle for 365 days or more is equal-to or less-than 10. If the light is red, the number of data sets that have been idle for 365 days or more is greater-than 10.

19. Add the Alerts module to your dashboard, and follow the prompts to configure it.

20. Select Chorus Alerts in the Alert Source field and Idle Data Sets Alerts in the Module Label field.

    When the Policy Light is red, the related Alert appears in the Alerts module labeled Idle Data Sets.

    You have created a light policy and CA Chorus Alerts appear when the conditions defined are met.

Storage Engine Alerts

Storage engine Alerts can send alert notices about various storage conditions to the Alerts module. To send alert notices, implement the General Object Automation (GOA) component of the storage engine system.

The alerts appear in the Alerts module, which lets you quickly see system issues. For example, GOA scripts can be created to send alerts about the following conditions:

- Monitor the allocated space on one, several, or all disk volumes. And send an alert to the Alerts module for each volume when its allocated space exceeds your customizable high threshold.
- Monitor the same volumes to raise an alert when allocated space is below a minimum usage (low) threshold.
- Monitor any storage object for any simple or complex condition, and send an alert to the Alerts module when that condition is found.

To set up storage engine Alerts for the CA Chorus for Storage Management discipline, create a storage engine GOA Script for objects you want to monitor for alert conditions (see page 102). Storage administrators can create GOA scripts after CA Chorus for Storage Management and the storage engine are installed, deployed, and configured. Typically, creating GOA scripts is ongoing. That is, storage administrators typically create or maintain them as and when needed.

Observe the following notes:

- The following storage engine configuration must be in place:
  - Provide the URL for the CA Chorus for Storage Management Alert Listener service to the storage engine.
  - Update the storage engine //STEPLIB DD statement.

  Note: For more information, see the *CA Vantage Configuration Guide*. This guide is available in the Knowledge Center.
- The storage engine systems must be up and running. If a storage engine is not running, it cannot create alerts and cannot send them to the target CA Chorus for Storage Management systems.
- The target CA Chorus for Storage Management systems (associated with the URLs) must be running and ready to receive the alerts. The storage engine does not store the alerts. If the target CA Chorus for Storage Management systems are not running and set to receive the alerts, then the Alerts are lost.

To view storage engine alerts in the Alert module, add them to your CA Chorus dashboard and follow the prompts to configure it.

Create a Storage Engine GOA Script to Generate CA Chorus Alerts

To receive alerts in the Alerts module for various storage conditions, storage engine GOA Scripts are required to send the alerts when conditions are met. Create the scripts for all objects that you want to monitor for *alert conditions*. The scripts must specify the storage engine line action *Send Alert to Chorus*, and the scripts must be scheduled to run as often as needed. Whenever the conditions specified in your scripts are met, the alert is sent to CA Chorus for Storage Management. These alerts appear in the Alerts module.

Use the storage engine windows client to create the scripts.

Note: For acquisition, requirements, and installation instructions, see the *CA Vantage Installation Guide*.

Follow these steps:

1. Open the storage engine windows client.

   Note: For storage engine windows client installation instructions, see the *CA Vantage Installation Guide*.

2. Connect to the back-end storage engine system host setup for CA Chorus for Storage Management.

   Note: For instructions about how to create a host connection definition and how to connect to a host, see the *CA Vantage Windows Client Guide*. This guide is available in the Knowledge Center.

3. Click z/OS, Object Scripting, Automation Scripts, General Object Automation.

   The Automation Scripts Builder opens.

   Note: You can click the Help button on any page of the wizard to get a description of the options on the page.

4. Click Next until the Source Object page is displayed. This page of the wizard displays the object tree.

5. Expand the object tree and locate the object that you want to create the alert script for.

6. Select the object and the data collection mode. Click the Help button in the builder to get a description of the modes available.

7. Click Next.

   The Script Name page opens.

8. Select *Create a new script with this name*, if it is not already selected and type in a script name up to eight characters long.

9. Click Next.

   The Title and System Execution page opens.

10. Fill in the page as needed. If you want the script to run on multiple systems, click the Systems button and specify them.

11. Click Next.

    The Timing for Events page opens. Use this page to specify the schedule for the alert script.

12. Click Next.

    The Options page opens.

13. Select the following options:

    - Load Script
    - Enable Script when Loaded

14. Click Next.

    The Event Procedure Section page opens.

15. Click Next.

    The Filter/Sort page opens.

16. Click Filter and create a filter to list only the items you want to create the alert for.

17. Click Next continuously until the Action Selection page opens.

    The Action Selection page opens.

18. Select the action Perform Object Action.

    The Perform Object Action page is available for selection in the navigation pane.

19. Click Next.

    The Perform Object Action page opens.

20. Select *Send Chorus Alert* in the Action pane.

21. Specify the Severity Type and Alert Text as follows:

    Severity Type

    The Severity Type is displayed in the alert which is displayed in CA Chorus. The Severity Types range from 1 to 5. Where 1 indicates the highest severity level and 5 indicates the lowest severity level.

    Click in the Severity Value of 1 to 5 field and type in a value of 1 to 5.

    Alert Text

    The text entered in the Alert Text field is displayed in the alert which is displayed in CA Chorus. Insert variables in the text to give the alert more clarity. To insert text and variables, perform the following steps:

    a. Click in the Alert Text field and type in the start of your alert message.

    b. Click in the field where you want to insert the variable, and click the Variables button.

       The variables available for the object are displayed in the Variables Substitution window.

c. Double-click the variable that you want inserted.

The variable is inserted accordingly.

You can insert text and multiple variables by repeating these steps.

Example: To create an alert indicating that critical space usage is reaching a maximum level of 80 percent, enter the following text in the Alert Text field:

```
SG %%POOLNAME%% is now %%POOLUPCT%% percent filled. Action required.
```

Where:

- %%POOLNAME%% is the variable for the Storage Group Name field.

- %%POOLUPCT%% is the variable for the %Allo field.

- Everything else is text that you typed in the field.

The resulting alert message would be shown as "SG *storage_group_name* is now *percent_allocated* percent filled. Action required." in the Alerts module.

d. Close the Variables Substitution window.

22. Click Next.

The Saving and Activation page opens with the Save and Activate the script option selected.

23. Click OK.

The alert script for the object is created, saved, and activated.

Note: Repeat this process to create a storage engine GOA Script for each object that you want to monitor for *alert conditions*.

24. Close the windows client after you have made all the alert scripts needed.

The automation scripts are created and alerts are sent to the Alerts module according to the conditions specified.

Note: You can reopen the windows client to create new scripts, change existing script settings, and delete scripts at any time.

Chapter 4: Administering Storage Resource Definitions

Storage Management Interface Overview

The Storage Management interface lets you administer storage resources using the back-end storage engine product. This interface includes many logically grouped storage objects, much like the Investigator. Launch the Storage Management interface from the Quick Links module.

The Storage Management interface is the CA Vantage web client.

You can perform Administrative Actions (see page 85) on storage resources using the Investigator. You can also perform these types of actions using the Storage Management interface. In this interface, these types of actions are referred to as storage engine *line actions*.

Some objects appear in the interface object tree but not in the Investigator Storage tree or some are renamed. A few examples are as follows:

- The category CA SYSVIEW and all child objects have been removed from the Investigator Storage tree.
- The Statistics category appears in the CA Chorus Investigator Storage tree, but not in the Storage Management interface. The Statistics category contains statistics object metrics for the Time Series Facility.
- In the Investigator, a category is named Storage Platform Administration. In this interface object tree it is named CA Vantage Internal Management.

The interface has two online help systems. The Navigation Help guides you with Storage Management interface features and options. The Object Help provides a description of the object, fields (columns), and the storage engine *line actions* that you can perform.

More Information

Access the Storage Management Interface (see page 108)
Set User Global Options in the Storage Management Interface (see page 108)

Access the Storage Management Interface

The Quick Links module contains links to CA Chorus discipline-specific interfaces. This module expedites the time that it takes for you to access your data and to respond to requests or troubleshoot issues.

Follow these steps:

1. Log in to CA Chorus.
2. Add the Quick Links Module to your dashboard.
3. Click Manage Storage Resources in the Quick Links window.

   A Storage Management interface instance opens in a separate browser window, and you are logged in based on your CA Chorus credentials.

Set User Global Options in the Storage Management Interface

You can personalize the Storage Management interface to meet your needs, monitor size, or the way you want the displays to appear. The Global Options feature is provided so you can personalize the interface.

The host definition settings for Maximum Lines Selected, Automatic View Execution, and Display Unavailable Objects are defined in the Host Definition. These settings are the default settings for all end users. If these fields are set to *Use Global Options* in the Host Definition for a host, your *Global Options* settings for these fields apply when you display an object from that host. The *Number of rows displayed* field does not exist in the Host Definition; therefore, the setting that you enter for this field in the Global Options dialog applies to all of your object displays. To find out how to set your Global options, launch the interfaced from the Quick Links module and see the Navigation online help system.