CA Chorus™

Product Guide

Version 04.0.00, Third Edition

CA Technologies Product References

This document references the following CA Technologies products:

- CA ACF2™ for z/OS (CA ACF2)
- CA Chorus™
- CA Chorus™ for IMS Database Management
- CA Chorus™ for Security and Compliance Management
- CA Chorus™ for Storage Management
- CA Chorus™ Infrastructure Management for Networks and Systems
- CA Chorus™ Software Manager
- CA Deliver
- CA NetMaster NM for TCP/IP
- CA SYSVIEW
- CA Scheduler
- CA Top Secret® for z/OS (CA Top Secret)
- CA View®
- CA Workload Automation

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 second edition of this documentation:

- Add a Hyperlink Using a CA-Provided CSV (see page 65)—Clarified the example.
- Display an Entity in a Time Series Facility Chart (see page 48)—Added a note to clarify that changing the charting period changes the amount of aggregate data and the Y axis values of the chart.

The following documentation updates have been made since the first edition of this documentation:

- Launch a Custom Link from the Investigator (see page 44)—Added this topic describing how to launch a custom user interface, a Web link, or a wizard to view data related to CA Chorus objects.
- View Alerts (see page 14), Investigation Launcher (see page 21), How to Use the Investigator (see page 25)—Noted that the auto-refresh interval is 30 seconds.

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

- View Default TSF Chart (see page 53)—Added this section to describe how you can view a default TSF chart with predefined metrics.
- Edit Policy Parameters (see page 36)—Added this section to describe how you can edit policy parameters using the Investigator module.
- Launch Investigator (see page 28)—Added this section to avoid repetition and keep the subsequent procedures short and simple.
- Start a Custom Investigation (see page 56)—Added this section to describe how *Custom Investigation* displays details about specific objects with just one-time configuration. A new session restores the saved data for you to resume your task from the same point.
- How to Use the Investigator (see page 25)—Added view access to the Investigator for RACF users.
- Save a Policy (see page 34)—Replaced the Create Policies topic with this new topic. This simplified process migrates all of the Version 3.0 policies to the Version 4.0 format.

- Run a Policy (see page 35)—Added this section to describe how you can run a saved policy.
- Investigation Launcher (see page 21)—Rebranded the Investigator module as the Investigation Launcher and added new search and path functionality details.
- How to Use the Investigation Launcher (see page 23)—Added this section to describe the new features for the Investigation Launcher.
- Search Data (see page 28), Save a Search (see page 30)—Split the View Tabular Data topic into these two new topics.
- Select Existing Policies (see page 71)—Updated with new shopping cart behavior.
- How to Use the Investigator (see page 25)—Added a note to update the functionality of the refresh icon.
- Global—Rebranded the JBoss server to the CA Chorus Application Server.

Contents

Chapter 1: Getting Started

CA Chorus

CA Chorus offers an intelligent management interface enabling greater productivity and simplified knowledge sharing. CA Chorus transforms how administrators:

- Collaborate with colleagues
- Interact with management tools
- Leverage the mainframe

Each CA Chorus discipline represents the domain of management that exists in your data centers. Therefore, within each discipline, you decide which IT roles your staff assumes. Each discipline offers the following usability features:

- Access to state and event data
- In-context domain documentation
- Object-based navigation
- Policy management
- Reporting
- Security data model extension
- Time Series data graphing

This guide details common behaviors. For discipline-specific information, see the applicable *User Guide*.

CA Chorus Architecture

The following diagram details the architecture and data flow for CA Chorus components:

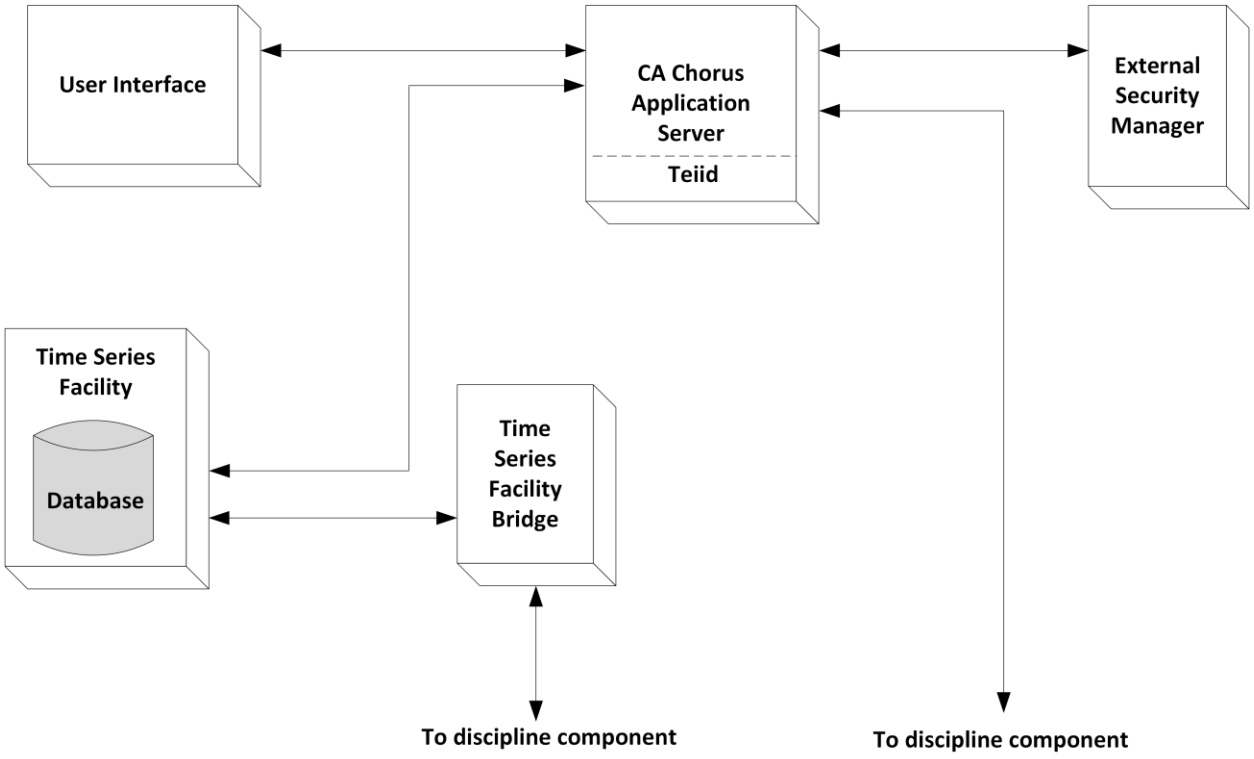

The following list details the components and products that you use with CA Chorus:

User Interface

Provides product access using a browser (Firefox or Internet Explorer). The user interface (UI) uses SQL queries to fetch data using Teiid and different data sources.

CA Chorus Application Server

Hosts the CA Chorus application. This server is a Java-based application server that operates cross-platform.

Teiid

Translates data source content for the CA Chorus Application Server. Teiid is a data virtualization system that lets applications use data from multiple, heterogeneous data stores. Teiid includes tools, components, and services for creating and executing bi-directional data services. Through abstraction and federation, data is accessed and integrated in real time across distributed data sources. This process occurs without copying or otherwise moving data from its system of record.

Time Series Facility (TSF)

Stores data that the CA Chorus products collect and provide. The TSF provides a single point for collection, storage, management, and organization of product performance and other TSF data.

Time Series Facility (TSF) Bridge

Allows TSF client requests to send and store data into the TSF Server.

External Security Manager

Protects resources and access rights to CA Chorus and its disciplines. Each discipline interoperates with CA ACF2, CA Top Secret, and IBM RACF to provide user ID, resource access, and PassTicket security.

How to Configure Your Workspace

All CA Chorus administrators, regardless of discipline, can customize the workspace to monitor performance data. Doing so simplifies the problem diagnosis and resolution. Customization also improves productivity when navigating multiple systems.

This scenario explains how a CA Chorus administrator configures the workspace to monitor performance data. Workspace configuration can include setting up multiple dashboards and modules.

The following illustration shows how a CA Chorus administrator configures the workspace:

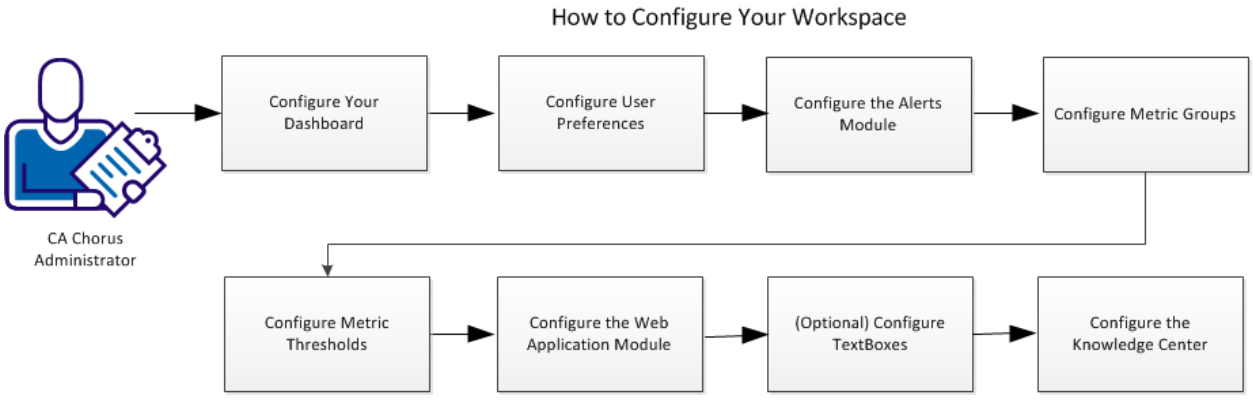

To configure your workspace, complete the following tasks:

1. Configure Your Dashboard (see page 12)
2. Configure User Preferences (see page 13)
3. Configure the Alerts Module (see page 13)
4. Configure Metric Groups (see page 15)

5. (Optional) Configure Metric Thresholds (see page 16)
6. Configure the Web Application Module (see page 17)
7. (Optional) Configure TextBoxes (see page 18)
8. Configure the Knowledge Center (see page 18)

Configure Your Dashboard

The *dashboard* is the primary point of access and view into CA Chorus. A dashboard is a customizable area that contains modules that are necessary for your tasks and projects. You can create multiple dashboards to customize your environments that are based on your roles and assigned tasks and add and remove dashboards as needed.

For example, you can manage all alerts from one dashboard named *Alerts* and can manage all command activities from a dashboard named *Commands*. Additionally, if you support multiple roles in CA Chorus, you can create one dashboard named *Database Administrator* and another named *Security Administrator*.

To log in to CA Chorus and launch a dashboard, enter your mainframe user ID and password at the URL. If you do not know the URL, contact your system administrator.

To create a dashboard, click the plus sign dashboard tab, and follow the prompts. A dashboard name can be up to 255 characters. Numbers (0 through 9), characters (a through z and A through Z), underscores, and blank spaces are allowed. A dashboard name cannot be the same name of an existing dashboard or the name of a shared dashboard.

To delete, edit, or share a dashboard, right-click the dashboard tab and follow the prompts. If you have modified the dashboard after sharing it, share the dashboard again to overwrite the earlier dashboard.

To import a shared dashboard, click the plus sign dashboard tab, select the option for a shared dashboard, select a shared dashboard from the Object Picker, and click Select. The dashboard name changes and a random unique-number gets appended to the name when you do one of the following actions:

- Import a dashboard that you do not own more than once.
- Import your own dashboard.

Configure User Preferences

By default, CA Chorus displays a warning message when you remove items from your view. For example, a module from the dashboard or a note from a module.

From each warning message, you can specify if you want this message to appear in the future. Your selection is saved to the Preferences dialog. You can also set all preferences from this one dialog. The following procedure helps to set your confirmation message preferences.

Follow these steps:

1. Select the Preferences link from the upper-right corner of your dashboard.
2. Select your preferences, and click Save.

   Your selections take effect immediately and are saved across sessions.

Configure the Alerts Module

A module is a container for data or a small program that lets you perform a specific task or set of tasks from your dashboard. The Module Library houses these modules. You can add multiple instances of modules to dashboards to track and manage performance and customize workflows. The Module Library contains default modules ready for use. The library is collapsed by default.

The *Alerts module* lets you monitor and investigate alerts from the dashboard as they are generated. The following procedure helps to configure the Alerts module for the first time. Each CA Chorus discipline requires a separate instance of the Alerts module.

Follow these steps:

1. Add the Alerts module to your dashboard from the Module Library.
2. Click the alerts configuration link.
3. Select the source of alert (for example, your discipline) from the drop-down list.
4. (Optional) Customize the module label. This text appears as the header on your module. This information is useful when you have multiple Alerts modules on the same dashboard.
5. Click Save.

   The Alerts module opens.
6. (Optional) Click the Configuration icon, define which column headings appear in your module and in what order, and click OK.

   The customized Alerts module is available to display alerts.
7. (Optional) Repeat this procedure if you are adding multiple modules to one or more dashboards. Each discipline requires a separate Alerts module.

View Alerts

After you configure the module for your discipline, the Alerts dialog displays the alerts that are already generated. The following screen capture details the Alerts window:

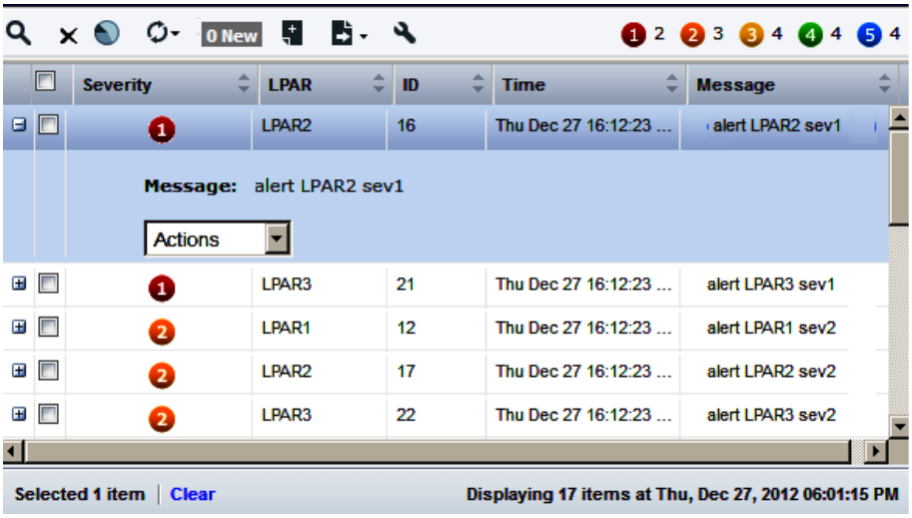

The auto-refresh option is enabled by default for authorized users. When this option is enabled, any changes in back-end data refresh the front-end data once every 30 seconds. For example, when the server generates new alerts, they appear in the Alerts module after 30 seconds. Click Refresh ( ) to disable or enable the auto-refresh option. If you are not an authorized user, click Refresh ( ) to refresh the data manually.

You can view the alert counters based on severity. The alert counters are refreshed when you click the refresh button or the new alerts button or when you delete an alert.

You can customize the tabular view of the dialog data by clicking the wrench icon. You can view details about an alert to investigate an issue on a monitored system.

After you view the alert details, you can analyze the resource or application from which it is occurring to determine the root cause. During the investigation, you can add notes for future reference and can see notes from coworkers related to the resource activity.

To see the available options to manage alerts, hover over the toolbar.

Follow these steps:

1. Select the configured Alerts dialog on your dashboard.
2. (Optional) Click the View Filter icon to search for a specific alert by specifying a filter criteria.
3. Select an alert from the list by clicking the + button. Use the scroll bar to move up and down the alerts list.

4. Select one of the options from the drop-down list named Actions. Other than the following actions, you may also see discipline-specific actions:

   View Detail

   Displays more information (for example, when and where the alert occurred) about the alert.

   Investigate

   Launches the Investigator. The Investigator opens to the system where the alert occurred, and displays information about the relevant performance area.

   Note: The Notification icon in the Investigator toolbar notifies you when messages are available from the back-end products.

5. (Optional) Highlight text from the alert, and click the question mark icon. The Knowledge Center displays content related to the highlighted text, which can be CA Technologies documentation, public notes and paths, your private notes and paths, user-added documentation, and so on.

Clicking the View Filter or View Charts icons maximizes the Alerts dialog. Restoring the dialog hides the filtering or charting panels.

Configure Metric Groups

The *Metrics panel* provides a visual display of key performance metrics for your monitored systems. System metrics are organized in data groups or categories.

You can create metric groups to monitor the performance of a system in the Metrics panel. If a system that you are monitoring is experiencing issues, the performance change is clearly visible in the Metrics panel.

Note: To turn on or turn off the Metrics panel, use the button  or  respectively.

Follow these steps:

1. Click the plus sign (+) control on the Metrics panel.
2. Follow the prompts to populate your Metrics panel.
3. Select the available metrics, and click Finish.

Note: For more information about discipline-specific domains and supported metric entities, see your discipline-specific *User Guide*.

(Optional) Configure Metric Thresholds

Metric thresholds are the visual tools that let you monitor the activity spikes for a metric. You can specify the direction of the spike (above or below the specified value) and the value of the threshold. When an activity reaches the specified value, the metric color changes.

Follow these steps:

1. Click the applicable metric, and click Set Threshold.

   A blue threshold bar appears at the bottom of the metric.

2. Specify the threshold value for the Caution State, Danger State, or both, and click Set Threshold.

   Caution State

   Identifies the level above or below the specified value in the adjacent field that is associated with the Caution state threshold.

   Default: Above

   Limits: Dependent on metric value range

   Danger State

   Identifies the level above or below the specified value in the adjacent field that is assigned to the Danger state threshold.

   Default: Above

   Limits: Dependent on metric value range

   The metric now responds to the specified threshold criteria.

   Note: The value for Caution or Danger can be set to 0 but not when the state is set to below.

View and Investigate Metrics

You can view a metric on the Metrics panel related to resource use on your systems that you want to monitor. You can then investigate the metric directly from the Metrics panel or add it to your dashboard for further monitoring. When you add the metric to the dashboard, the larger view makes it easier to view the activity for the metric. To investigate the metric directly from the Metrics panel, click the applicable metric, and click Investigate.

You can manipulate the metrics in a metric group as follows:

- Hover over a metric to view the full metric name, value, and description.
- Move a metric anywhere in its group by dragging and dropping it where you want.
- Click the down arrow in a metric to investigate it or to view a larger graphical representation version of the metric on the dashboard.

Note: Hover over the panel controls to manage the Metrics panel.

Follow these steps:

1. Click the metric you want to investigate, and select Add to Dashboard.

   The metric appears on the dashboard in a larger view for monitoring. The metric group name and metric name appear in the header of the graph. You can hover over the graph to see the value (time, metric value) at any point. Use the slider control to adjust time and to see data during the changed time. The slider control refreshes when the graph is updated.

2. Click Actions, Investigate.

   The Investigator opens at the source of the metric data. You can now review the data source to identify dependencies and relationships.

Configure the Web Application Module

This module lets you quickly add a web application to your dashboard. Doing so lets you further customize your dashboard to view and manage your data in the most efficient manner.

The web application must be IFrame-renderable (Inline Frame). An IFrame is an HTML document that is embedded within another HTML document. The web application address must be in one of the following formats:

- http://www.*address text.URL suffix*
- https://www.*address text.URL suffi*x

  Note: When you specify the web address, it is not necessary to include http:// *and* www. The following examples are valid URLs:

  - http://ca.com
  - www.ca.com
  - http://www.ca.com

To add a web application, add the module to your dashboard and follow the prompts.

Examples

- As a security administrator, you want to stay current with your mainframe peers. So, you add the CA Chorus Security Global User Community as a Web Application module in one of your dashboards. Doing so lets you integrate access to this site with your dashboard.
- As a system administrator, you want to stay current with the latest news from CA Technologies. Therefore, you add http://ca.com/support as a Web Application module in one of your dashboards. Each day when you log in to CA Chorus, you check this site for new vulnerability alerts, announcements, and so on.

Note: For troubleshooting information about the Web Application module, see the topic Errors While Configuring the Web Application Module in the *Troubleshooting Guide*.

Configure Text Boxes

The TextBox module lets you add customizable labels or lists to the modules in your dashboard. Doing so lets you differentiate modules so you can quickly find data.

Examples:

- Add a label to a Policy Light module to differentiate multiple modules.
- Create a To Do list to track in one of your dashboards.

To create a text box, add the module to your dashboard, create your label, and then place it near the applicable module. The TextBox is saved across user sessions.

Configure the Knowledge Center

The Knowledge Center is the repository for all documentation in CA Chorus. Content can include:

- CA product documentation
- User-generated documentation, including:
  - Your paths and public paths that you have saved while working in the Investigator.
  - Your notes and public notes.
- Websites
- Links to third-party documentation
- Chicago-Soft MVS/QuickRef messages

Knowledge Center content can help you perform tasks and become familiar with the product.

We recommend that you configure your repository to show only content that supports your role. Doing so can help improve the relevance of search results.

Example: Security Administrator Options

As a security administrator using CA ACF2 with CA Chorus, you want to configure the Knowledge Center to display content relevant to your role. To do so, select CA Chorus, MVS/QuickRef, User Documentation, and CA ACF2 documentation as your content sources. Do not select CA Top Secret or any other sources that apply to other disciplines.

Follow these steps:

1. Click the question mark icon from any module or dialog.

   The Knowledge Center opens.

2. Click Advanced Search.
3. Select the documentation sources for your Knowledge Center.
4. Click *Always search only these sources*.

   All future searches result in content from the configured sources.

Your workspace is configured with all modules and tools customized for your use.

Chapter 2: Using the Investigator

Investigation Launcher

The Investigation Launcher lets you start new investigations and view saved investigation paths and saved searches.

A path is a series of actions that you perform to accomplish a task. By using the Investigator, you can track an earlier investigation as well as start a new investigation, and save your investigation series as a Public Path or a Private Path. You can save complete and incomplete investigations. You can load, edit, and delete complete paths; and resume incomplete paths and unsaved paths. This module helps you quickly return to a specific step in a task for training, troubleshooting, and daily operations.

The *Investigation Launcher* includes several tools to help you monitor and investigate data and systems:

Investigator

Displays data and systems in tabular and chart views.

Table Viewer

Displays tabular view of data. Each row of the table represents an object or a group of objects.

Topology Viewer

Displays a pictorial view of data. After you select data from an Investigator table, you can launch this tool.

Time Series Facility

Displays line charts of time-based performance data. After you select data from an Investigator table, you can launch the Time Series Facility (TSF).

Notes

Lets you embed notes in rows of data to promote knowledge sharing.

Paths

Houses the saved private and public paths that you have traversed while working in the Investigator. These paths get automatically stored in the Knowledge Center and appear in the Knowledge Center search results.

Searches

Houses the saved searches that you have traversed while investigating data. You can run, edit, and delete the searches.

The auto-refresh option is enabled by default for authorized users. When this option is enabled, any changes in the back-end data refresh the front-end data once every 30 seconds. For example, whenever the server generates new alerts, they appear on the Alerts module after 30 seconds. Click Refresh (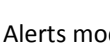) to disable or enable the auto-refresh option. If you are not authorized for auto-refresh, click Refresh () to refresh the data manually.

How to Use the Investigation Launcher

The following screenshot details the Investigation Launcher:

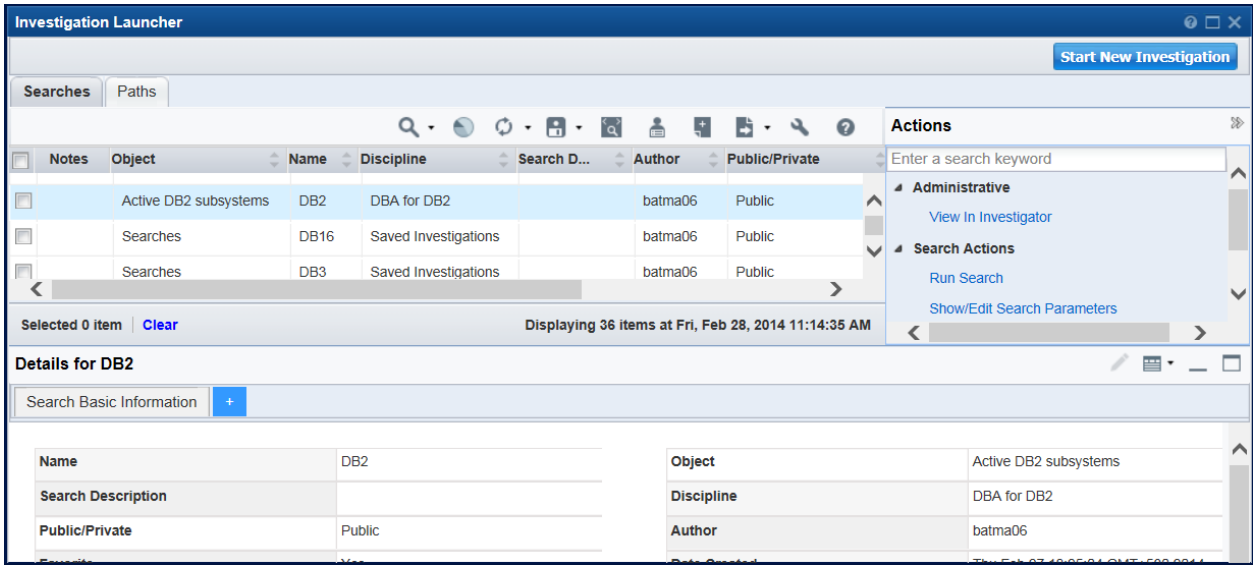

The Investigation Launcher provides you with access to searches and paths across all disciplines. The search actions let you view search result and prepopulated search criteria. The Path actions let you load complete paths and resume incomplete paths.

The Data, Actions, and Details behavior are the same in the Investigation Launcher and the Investigator. To learn more about this functionality, see the applicable entries in How to Use the Investigator (see page 25).

Investigator

The Investigator presents a cohesive view of data. The Investigator helps you view and analyze information in discipline-specific data repositories. The Investigator provides multiple panes for data management.

After you launch the Investigator (see page 28), view tabular data or display desired data in the Topology Viewer or Time Series Facility (TSF). The Topology Viewer provides a pictorial overview of data, which lets you quickly identify relationships. TSF presents performance metrics as line charts for periods of time that you can specify. To switch to these views, use the Table View ( ), Topology View ( ), and the Time Series ( ) buttons in the toolbar. You can also save Investigator search queries as a JCL batch job that you can use to generate reports.

The Investigator lets you operate data in the following modes:

Live

Displays current data and enables you to work with the Investigator. When you launch the investigator, it opens in Live mode.

Note: The Investigator opens in Historical mode when you launch it from a saved path (complete or incomplete).

Historical

Displays an investigation path as a series of breadcrumbs, with each breadcrumb representing a step. The Investigator generally displays "live" data, and enters Historical mode only when you perform one of the following tasks:

- Click a breadcrumb.
- Launch the Investigator from a saved path. You can select a saved path from one of the following locations:
  - The Investigation Launcher > Public Path or Private Path.
  - The Investigator > Saved Investigations pane.

In Historical mode, the background color of all panes except the History pane turns gray. Also, the last breadcrumb in an incomplete path is in the selected state, and the Investigator toolbar is frozen. You are limited to the following tasks:

- Click the Refresh icon to go to the Live mode.
- Click the Help icon.
- Click the Clear Path button.

The Investigator switches from Live to Historical mode when you perform one of the following tasks:

- Click the Refresh icon in the Investigator toolbar.
- Select a node in the tree.
- Select an action.

How to Use the Investigator

The following screenshot details the Investigator:

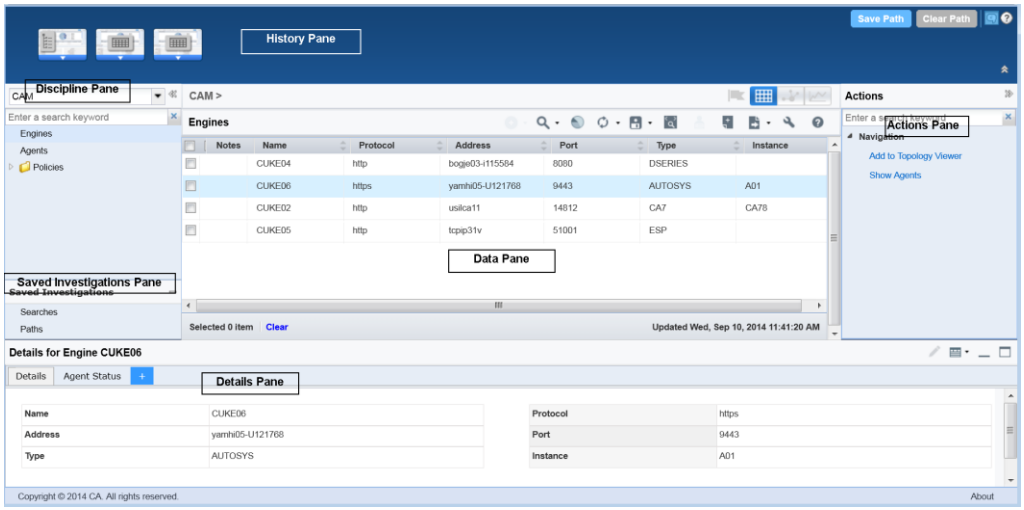

The Investigator provides a wider view and access to searches and paths across all disciplines. The Investigator window has the following panes:

Discipline

Displays discipline-specific objects and helps you retrieve data for querying, analyzing, maintaining, and monitoring systems.

The Discipline drop-down list also includes the available disciplines and an option to add custom objects to the Investigator. To access shared dashboards, maps that are saved from the Topology Viewer, and policies, select Chorus from the Discipline drop-down list.

The Investigator lets you save new policies that are specific to your needs.

Note: For more information about saving new policies, see Save a Policy (see page 34).

Saved Investigations

Displays links to saved Searches and Paths, which let you view search history and investigation paths in the Data pane. For a list of actions for both options, see the Actions pane.

Data

Displays all data for the selected object in the tree in a table view by default. The table contains all properties of a path displayed in columns. When you select a path in the table, detailed path information appears in the Details pane. You can view the SQL query behind the data in the table view by clicking the View SQL icon (). The auto-refresh option is enabled by default for authorized users. When this option is enabled, any changes in back-end data refresh the front-end data once every 30 seconds. For example, when the server generates new alerts, they appear in the Alerts module after 30 seconds. Click Refresh () to disable or enable the auto-refresh option. If you are not an authorized user, click Refresh () to refresh the data manually.

You can select rows or can highlight a row in the Data Pane. The following screenshot illustrates a selected row:

The following screenshot illustrates a highlighted row:

The following screenshot illustrates a highlighted and selected row:

The following table differentiates selecting and highlighting:

Selecting	Highlighting
Select the check box on each row to select those rows. Select the Select All check box to select all rows. After using this option, you cannot cancel the selection of individual rows.	Click anywhere on the row except the check box to highlight that row. You can highlight only one row at a time.
Cancel the selection of a row by clearing the check box on that row or by clicking the Clear button on the status bar. When you select the Select All check box to select all rows, you cannot cancel the selection of individual rows. Cancel the selection of all rows by clearing the Select All check box or by clicking the Clear button on the status bar.	Cancel the selection of a row by highlighting the same row again, by selecting the same row using the check box, or by clicking the Clear button on the status bar.
For both single-row and multi-row selections, the Details pane shows no data.	See the details of the highlighted row in the Details pane. When you highlight a row that is already selected, the Details pane shows the details.

Selecting	Highlighting
When a single row is selected and no row is highlighted, the Action pane shows actions corresponding to the selected row. When multiple rows are selected and no row is highlighted, the Action pane shows actions that are common to the selected rows. When a single row or multiple rows are selected and a row is highlighted, the Actions pane shows the actions corresponding to the highlighted row.	See the actions corresponding to the highlighted row in the Actions pane.
The Notes icon () is disabled.	The Notes icon is enabled.
The Status bar displays the count of selected rows.	The Status bar does not display the count of highlighted rows.

Actions

Displays actions based on your selection about discipline and object. You can manipulate data or can drill further into data for analysis. Actions appear as active links under different categories. You can resume an incomplete path, load a complete path, and delete a saved path.

Details

Displays detailed information of a highlighted path in the table.

Note: If you customize the Data pane to include columns from a related table, the new columns do not appear in the Details pane.

As a CA ACF2 or CA Top Secret user, you can edit the details available in the Details pane by clicking the edit icon (). When you click the edit icon, all the tabs in the Details pane enter into edit mode. Make the required changes and click Save. The edit option is not available for an IBM RACF user.

Note: You can edit only the fields of the objects that are set as editable.

Add a tab to the Details pane by clicking the **+** icon on the tab header. Right-click inside a tab and select Insert Field to add fields to the tab. To delete a field or a tab, right-click and select the Delete option. The Details pane must have at least one tab on it, so you cannot delete all the tabs. To change the name of a tab, double-click the name. The drag-and-drop option helps you rearrange the fields. To select multiple fields, use the Ctrl or Shift keys. After customizing the Details pane, you can save the latest view. Saving a view helps you or others reuse the view. To save a view, click the View icon, click Save View, provide a name and description for the view, select the permissions, and click Save. To reuse a view, click the View icon, click Manage Views, click the + sign on the view, and select Apply. The owners can edit and delete their views.

Note: The view that you have saved or applied becomes your default view. To restore the initial view of the Details pane, use the System Default button.

History

Displays the steps that you take in an Investigator session. A breadcrumb in this pane represents one step, and the breadcrumb for the current step is highlighted. This pane lets you follow the steps you or other users made to reach a certain result. These steps form a path. You can save paths as Private to review later or set as Public to let others view.

Launch Investigator

Follow these steps:

1. Log in to CA Chorus.
2. Add the Investigation Launcher from the Module Library to your dashboard.
3. Click Start New Investigation.

   The Investigator opens in a new window.

   Note: The View in Investigator link on the Actions pane of the Investigation Launcher also lets you launch the Investigator with searches and paths selected in the Saved Investigations pane.

You can bookmark the Investigator in your browser. On relogin, you can launch the Investigator from the browser bookmark. When you do so, the Investigator prompts you to log in. If you have already logged in, you can launch the Investigator from the browser bookmark.

Note: You can also launch the Investigator from the Alerts module, Metrics panel, or from the Policy Status Light module. When launching the Investigator from other modules, the Investigator displays contextual data.

Search Data

To find data in the Investigator quickly, use the search option. The Investigator displays filtered content that exactly matches your search criteria, which can include alphabetic or numeric characters.

Note: This procedure applies only to the new investigations.

Follow these steps:

1. Launch the Investigator (see page 28).
2. Select an option from the drop-down list in the discipline pane.

3. (Optional) Filter the objects by typing in the text box below the drop-down list.

   The tree displays filtered information.

4. Select an object from the tree to see its data in the Data pane.

5. Click the Start a new search icon (), if the search panel is not loaded.

6. Specify your search criteria, and click Search.

   Note: Narrow your search by adding criteria. To adjust the criteria, use the plus (**+**) and minus (**–**) buttons. To have the system convert your text-inputs in upper case without using the keyboard Caps Lock, select the Uppercase check box.

7. (Optional) View the SQL query behind the data in the table view by clicking the View SQL icon () in the toolbar.

   Note: To hide or show the search panel, click the Start a new search icon in the Investigator toolbar.

   For the columns with a date as the value, select the search value from the calendar icon. For the columns with time as the value, select the search value from the drop-down list. For the columns having date and time as the value, select the search value using the calendar icon and the drop-down list.

   When you select a time-based object in the tree, the search panel displays date and time ranges. For time-based objects, by default, the table view presents data based on the default values displayed in the Select Time Range panel. You can change the default date and time range using the calendar icon and the drop-down list. If you change the date and time range during a search, the changed values are available for further searches during that Investigator session.

Investigator Search Enhancement

Required filters exist for some of the objects for which you must enter the values and search.

While performing any search on any of the columns in the object, if you navigate to any other object, the search criteria remain the same if there are any common columns.

Save a Search

To make your search available for future reference, use the Save Search option. The Search link in the Saved Investigations pane lets you view the saved searches.

Follow these steps:

1. Identify data based on an existing search (see page 28).
2. Click the Save icon () in the toolbar, and click Save Search.

   The Save Search dialog opens.
3. Enter a name and (optional) description for the search.

   Note: The name field provides a suggestion drop-down of all the existing searches that you have created. You can select one of the names from the list to overwrite that search, or you can provide a new name.
4. Personalize the search based on the preferences available in the dialog, such as Private, Public, and (optional) Favorite.

   Note: All users can view the Public searches, but only the author of the query can delete it.
5. Click Save.

   The search is saved. To view the saved search, in the Saved Investigations pane, and click Searches.

Note: If you have customized the tabular view in the Investigator after running a search, the customization is also saved with the search. When you run such a search, the Investigator displays the customized view. The Layout drop-down list lets you change the view. To view data without customization, select Current Layout from the drop-down list.

If you search for data after performing a navigation action, the search occurs on the data that the navigation action retrieved. The search panel header identifies the data.

Example:

As a system administrator, you are reviewing all user account data. To see the roles of a specific user, select that user from the Data pane, and click Show Roles in the Actions pane. The Data pane displays the roles for the selected user and the header indicates the same.

If you search for a specific role, the Data pane displays roles that are based on your search criteria. If you click the X icon after filtering the roles, the Data pane displays roles for all users that match the search criteria. If you click the X icon without searching anything, the Data pane displays roles for all the users.

Manage Saved Searches

The saved searches are displayed using the Search link in the Saved Investigations pane. You can select the Start a new search icon ( ) in the Investigator toolbar to see the searches related to a selected object. This drop-down list displays links to eight saved searches, all searches, and all policies. The eight searches are displayed in the following order:

- Private favorite
- Public favorite
- Private unfavorite
- Public unfavorite

Follow these steps:

1. Launch the Investigator (see page 28).
2. Select Saved Investigations at the bottom of the Discipline pane, and click Searches.

   The Data pane lists all public searches and your private searches that are related to the selected discipline.
3. Click the search that you want to load, and click Run Search in the Actions pane.

   The Data pane displays the result of the selected search. The search panel displays the search criteria.

Note the following behaviors:

- To edit or delete a search, select the search from the Data pane, select the Show/Edit Search Parameters or delete action from the Actions pane, and follow the prompts. After changing the criteria, you can run the search. Save the search with a new name if you do not want to overwrite the existing search.

  Note: The Delete Search link, however, does not appear for the searches saved by other users, thus blocks you from deleting them.

- Instead of the search criteria if you just want to change the options like public, private, or favorite, use the edit icon ( ) on the header of the Details pane.

Example: Editing a Search

You may have created and saved a complex search to retrieve your critical information. Your search has seven lines of filter criteria. If you want to remove the last two lines to retrieve a new result, you can edit and run the query.

Example: Investigator Searches

The search option helps you filter data based on customizable searches. This example helps you formulate searches. This example is based on a table, which contains the following columns:

Note: The search is case-sensitive.

System ID

Identifies a unique system identification number.

System Name

Identifies the systems to which CA Chorus is connected.

System ID	System Name
01	System1
11	System3
12	system1
27	System38
33	System385

Use the following examples to help formulate queries:

Column Name	Operator	Search Value	Meaning	Result
System Name	=	System1	Columns having a value *System1*.	■ System1
System Name	<>	System1	Columns having a value other than *System1*.	■ system1 ■ System3 ■ System38 ■ System385
System Name	like **Note:** To search for a specified pattern in a column, use the *like* operator. To define missing characters (character masking) in the pattern, use the percent (%) sign.	%38%	Columns having a value that follows the pattern, *38*.	■ System38 ■ System385
System Name	contains	3	Columns that contain the value *3*.	■ System3 ■ System38 ■ System385
System ID	starts with	1	Columns that start with the value *1*.	■ 11 ■ 12

Example: Grouping Criteria with Parenthesis

To search user IDs that start with A or B, specify the following search criteria:

- User ID starts with A OR User ID starts with B

If your search criteria is complex, the AND or OR operators can cause confusion. For example, if you are searching for administration users whose User IDs start with A or B, OR all managers whose user IDs start with C or D, group each line in the search criteria using parenthesis:

- (((User ID starts with A OR User ID starts with B) AND (User Role ='Admin')) OR ((User ID starts with C OR User ID starts with D) AND (UserRole = 'Manager')))

Save a Policy

After you performed a search, you can save it as a Policy.

Follow these steps:

1. Launch the Investigator (see page 28).
2. Select an option from the drop-down list in the discipline pane.
3. (Optional) Filter the objects by typing in the text box below the drop-down list.

   The tree displays filtered information.
4. Select a node from the tree to list down corresponding objects in the Data pane.
5. Filter the objects by adding some search criteria in the Search panel.
6. Click the Save icon ( ) in the toolbar, and select Save Policy to open the Save Policy dialog.
7. Provide the policy name, description, and information required for policy settings, and click Save.

   Note: For optimal use, set policy recurrence 30 or more, and specify a search criterion to produce a small number of objects (within a few hundreds). A different value may consume more resources, and make the system slow.

   The policy becomes available in CA Chorus Personal Policies.

Note: You can delete or overwrite policies - public and private - saved by you, but not those saved by other users. You can, however, save your policy as a private policy with the name of a public policy saved by a different user.

Run a Policy

You can manually run a policy any time without depending on its recurrence.

Follow these steps:

1. Launch the Investigator (see page 28).
2. Select Chorus from the drop-down list in the discipline pane.
3. (Optional) Filter the objects by typing in the text box below the drop-down list.

   The tree displays the filtered information.
4. Select Chorus Personal Policies from the tree to list the saved policies in the Data pane.
5. Select the appropriate policy row, and then click the Run Policy link in the Actions pane.

   The policy result appears in the Details pane.

Note: When running a policy, if an error occurs repetitively, the policy is paused after a certain number of attempts. You receive error details through WTO. You also receive an email notification if your email address is configured in the policy.

Change Policy Status

Each new policy is active by default. When saving, you can change its status to inactive or keep it as is. CA Chorus lets you change the policy status whenever required.

Follow these steps:

1. Run a policy (see page 35).
2. Select the appropriate object row.
3. Click the appropriate link in the Actions pane to activate or deactivate the policy.

   A message appears to confirm your action.
4. Click OK.

   CA Chorus updates the new status in the Policy Status column.

Edit Policy Parameters

CA Chorus lets you edit policy parameters with various features such as converting a policy to private or public, changing its activation state, or changing policy light settings.

Note: You can edit those public and private policies that only you, and *not* other users, have created. You can, however, save with different names the public policies created by other users.

Follow these steps:

1. Launch the Investigator (see page 28).
2. Select Chorus from the drop-down list in the discipline pane.
3. (Optional) Filter the objects by typing in the text box below the drop-down list.

   The tree displays filtered information.
4. Select CA Chorus Personal Policies from the tree to list the saved policies in the Data pane.
5. Select the appropriate policy from the object row in the data pane.
6. Click the Show/Edit Policy Parameters link under Policy Actions in the Actions pane.

   The search panel populates filter criteria of the policy.
7. Click the Save icon and select Save Policy.

   The Save Policy dialog displays the policy name, description, and existing parameters.
8. Make the required changes and click Save.

   The policy parameters are edited.

Go to the Chorus Personal Policies to see the saved edits. To go there, type Chorus Personal Policies in the text box below the drop-down list in the discipline pane.

Configure Parameters for Alert Notifications

CA Chorus lets the system administrator configure alert text to receive as email or write-to-operator (WTO) notifications. You receive WTO notification on the Mainframe screen and a notification to the configured email address.

Follow these steps:

1. Launch the Investigator (see page 28).
2. Select a discipline and a node from the object tree.
3. (Optional) Filter the objects by typing in the text box below the discipline drop-down list.

   The tree displays filtered information.
4. Filter the objects by adding search criteria in the Search panel.
5. Click the Save icon ( ) in the toolbar, and select Save Policy.

   The Save Policy window opens.
6. Expand Alert Settings, select the Send Alerts to Alert Module check box, and enter your alert text and the severity level.

   Note: The variable $ALERT_DETAIL translates this alert text into notification message.
7. Expand Email and enter the email Id, subject line, and email body.

   Note: Use the variable $ALERT_DETAIL in subject line and/or email body where you want details of the configured alert to appear.
8. Expand write-to-operator, and enter $ALERT_DETAIL in the WTO message box.
9. Click Save.

The parameter is configured for alert notification. The next time your alert criterion is met, you will receive a notification per these settings.

View Notification Messages

The Investigator toolbar includes a Notification icon. This icon changes colors to notify users that messages are available from the back-end products. These messages may display information such as product settings and product availability.

When messages are available, they appear by default. If you do not want to display the messages automatically, clear the Auto expand check box in the message box. You can also set this preference in the Preferences dialog.

The Notification icon displays a red flag when messages are available from the back-end products. To view the messages, click this icon.

Examples:

Storage Administration

Your site has three storage engines. One of the storage engines becomes unavailable for maintenance, upgrade, and so on. A notification message appears explaining this status change.

Database Administration

A notification message displays information that table results are limited to 5,000 rows in the Investigator. This notification is important because more data exists than what appears in the table.

Customize Your Tabular View

You can customize the tabular view (contents and order of appearance) to suit your preferences. You can include related table columns to construct a more useful query and display the results. By customizing your data, you see data based on your site-specific needs. The customizations apply only to the view where you changed them. CA Chorus saves these changes so they appear each time you log in using the same user ID and password.

Note: This procedure applies to various modules, for example, Alerts and Investigator. Continue with Step 2 to customize the Alerts module.

Follow these steps:

1. Launch the Investigator or configure the Alerts module.

   Note: If you are doing a new investigation, select an object from the Discipline pane to see its data in the Data pane.

2. Click the Customize Your Tabular Data View () icon on the toolbar.

   The Select and Reorder Columns dialog appears.

3. Select an object from the Objects list.

   The available columns for the selected object appear.

4. Select an available column.

   The column description appears.

5. Click the right arrow to add the selected column to the current view.

   The selected column appears in the Selected Columns list, for example, System ID. If a related object also shares this column then the related object name prefixes this column name, for example, Systems: System ID.

   Note: The left arrow and the cross button move the selected column back to Available Columns.

   Note: To display the current health state of the objects in the Data pane, select the Status column.

6. (Optional) Use the up and down arrow buttons to rearrange the column order.
7. (Optional) Select the column names from the Group By drop-down list to group the data based on the selected columns. This customization lets you see the count of data for the grouped columns in the Data pane. To expand or collapse the grouped data, click the arrow.
8. Click Save.

   Your column settings are saved for this view.

Export Investigator Table Data and Relations

Use this procedure to perform the following tasks from the Investigator toolbar:

- Export data from the visible tabular-columns of the currently selected object to a comma-separated value (CSV) file.
- Export relations and the visible tabular-columns of the currently selected object to an XML file.

Note: The program where you are prompted to save the file depends on your computer settings.

Follow these steps:

1. Launch the Investigator.

   Note: If you are doing a new investigation, select an object from the Discipline pane to see its data in the Data pane.

2. (Optional) Filter the data in the table.
3. Click the Export icon on the Investigator toolbar.

4. Perform *one* of the following tasks:

   a. Click Export to CSV to export tabular data to a CSV file.

   b. Click Export Metadata to export the visible tabular-columns of the currently selected object to an XML file.

      Note: You cannot export relations if no primary key columns are configured in the view.

5. Save the file to your desired location.

   Note: The Status column indicates the health state of the object. The following list helps you understand the health states of your objects:

   - 1—Danger
   - 2—Warning
   - 3—Normal

Generate a Report

When you display data in the Investigator, it creates search queries to find the requested data. You can save these search queries to a z/OS PDS member. When you save the queries, the Investigator adds JCL statements to create a batch job that you can submit to generate a report. The report displays the same data that is shown in the Investigator.

Creating the batch job lets you generate an updated report as needed, for monitoring purposes. You can submit the job manually, or you can automate the job submission using existing scheduling and output management products at your site. For example, your site may use CA Workload Automation or CA Scheduler to schedule jobs, and CA View or CA Deliver to manage output.

Follow these steps:

1. Launch the Investigator.
2. Click the Save icon, and select Save JCL from the pop-up menu.

   The Save JCL dialog appears.

   Note: You cannot save the JCL if the search query that you are saving contains only the Notes and Instance columns in the grid. If you attempt to do so, an Alert message appears. Add a column other than the Notes and Instance columns, and then save the JCL again.

3. Enter the following information:
   - (Optional) The name of the data set and member containing the JCL template to apply to your job. Do not perform this step unless more than one template is available. The default is *chorus_runtime_hlq*.CETJEZTR(EZTMPL01).
   - The name of the data set and member in which to save your JCL batch job. If the data set does not exist, it is created.
   - A description of the batch job being created. This description is inserted as a comment in the JCL, after the JOB statement information.

   Click Save.

   The Investigator saves the batch job to the specified data set. This job consists of search queries and JCL statements.
4. Edit the batch job:
   - Insert the account code.
   - Review the JCL comments and make other changes as needed.
5. Submit the batch job manually, or submit it using a job scheduler or workload automation task.

   Note: When you submit the job, CA Chorus must be running so that data can be retrieved from all sources.

   The report is generated as an output file on the mainframe.

View Data in a Table

The Investigator lets you view data in table rows. A grid contains a maximum of 300 rows, each representing an object or a group of objects. Each group is expandable and can further include up to another 300 objects.

Example:

You have a page that displays 10 grouped rows. On expanding the largest possible group, the page displays a total of 310 rows. You can expand only one group at a time.

View Data in a Chart

The Investigator lets you view column data in a pie chart. These charts provide a visual, customizable means of evaluating your data.

Follow these steps:

1. Launch the Investigator.
2. (Optional) Filter your data, if necessary.

3. Click the View Charts icon in the Investigator toolbar.
4. Use the drop-down lists to filter and customize your pie chart, and then click Add.
   a. Select the field to group your data.
   b. (Optional) Select the Aggregate format:

      Count

      Initiates the SQL function to display the number of rows that matches the specified criteria. This format is the aggregation default.

      Sum

      Initiates the SQL function to display the total sum of the numeric columns.

      Average

      Initiates the SQL function to display the average value of a numeric column.

   c. (Optional) Select the value that you are summing. This step does not apply when you select Count as the aggregation format.

   The customized pie chart appears in the center pane. If necessary, scroll down to see the table view.

   Note: If the number of slices is more than 10, the chart groups the rest of the slices under a slice named *Others*. The chart groups all columns having null values under a slice named *NULL*. The chart groups all columns having empty spaces as a value under a slice named *EMPTY*.

Example Pie Chart: Connected Systems

The example chart is based on a table that contains a column named System. This column identifies the systems to which CA Chorus is connected. System1 to System7 are the values that are listed under this column. CA Chorus has one connection to each system. If you select System as the field to group your data while charting column data, the following pie chart appears:

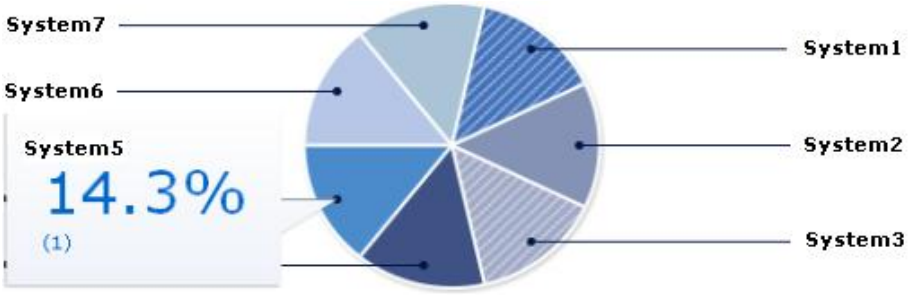

To see the label and percent of the column value, hover over the slices. For example, if you hover over the slice that represents System5, the callout in the figure appears.

Example Pie Chart: Aggregating Data

This example shows how to use data aggregation. Assume that a bank has two types of accounts: checking and savings. Each account has a dollar balance. This list shows the various options based on the same group by value. Each pie has two slices—one for checking and one for savings.

Note: The following italicized values are intentionally generic to illustrate the usage for all disciplines.

Group by: *account_type* and Aggregate: Count

Shows a pie chart where the size of the slice corresponds to the number of accounts in the type.

Group by: *account_type* and Aggregate: Sum Of: *balance*

Shows a pie chart where the size of the slice corresponds to the total dollar value of the accounts in the type.

Group by: *account_type* and Aggregate: Average Of: *balance*

Shows a pie chart where the size of the slice corresponds to the average size of the accounts in the account type.

Launch a Custom Link from the Investigator

The Investigator lets you launch a custom user interface or a Web link to view data related to CA Chorus objects. You can access both internal and external links. A custom action button in the Investigator toolbar connects you to the custom links. This button lists the linked user interfaces and Web pages customized for specific object types.

The custom action button is available in the Investigator and in the Custom Investigation modules.

Note: The custom links are configurable through metadata. For more information about configuring custom links, see the *CA Chorus Software Development Kit User Guide*.

Follow these steps:

1. Launch the Investigator (see page 28).
2. Select a discipline from the Discipline pane.
3. Select an object type from the object tree.
4. Click the drop-down arrow of the custom action button.

   A drop-down list displays the custom links corresponding to the selected object type. The first link in the list is the default link.
5. Click the appropriate link.

   The corresponding link opens in a new browser tab.

   Note: To launch the default link directly, click the plus (+) sign of the custom action button.

Visualize Data

The Topology Viewer provides a pictorial view of data in your system and their relationships. This view can simplify your ability to identify relationships as you manage your data. You can also drill down to isolate data within your system. The following color-codes help you assess the health state of each node:

- Gray—Normal or unknown
- Yellow—Warning
- Red—Danger

The Topology Viewer can help you perform the following tasks:

- Troubleshooting
- Taking inventory of your system

You can launch the Topology Viewer for one or more rows of data in the Investigator and can view the data pictorially. Each data node can be a member in a complex relationship with many data nodes. To see this relationship pictorially, right-click an object in the pictorial view and select an action from the context menu. When you display the child objects of an object, a group object appears in the Topology Viewer. The name of this group object is a combination of the relation name and the count of the child objects. Select this group object to launch the search panel.

The search panel helps you filter the pictorial view of data based on customizable search criteria. By default, the Query Name field in Search Panel displays the name of the selected relation; you can edit it to provide a custom name. The name of the group object changes to the custom name after you apply the search criteria.

Follow these steps:

1. Launch the Investigator (see page 28).
2. Select a discipline from the drop-down list.
3. Navigate in the tree to the data that you want to view.
4. Select one or more rows of data or highlight a row of data.

   The Actions pane opens.
5. Select Add to Topology Viewer under Navigation in the Actions pane.

   The Topology Viewer displays a pictorial view of your data. The default view is hierarchical. The label of the object displays the system name to which the selected object belongs.

   Note: The Add to Topology Viewer link, however, does not appear for the objects without any relations.
6. (Optional) Manipulate the view or print the visualization using the toolbar at the top of the Topology Viewer pane.
7. (Optional) Select the Replace Topology Viewer link under Navigation in the Actions pane to replace the current view of the object with a new view.

   Note: The Replace Topology Viewer link is disabled for the objects without any relations.

Note: For conceptual and procedural information and real-world examples, see the applicable discipline *User Guide*.

Manage Perspectives

After customizing the pictorial view of your data, you can save the view as a perspective. Saving a perspective helps you or others reuse the view and apply it to a different object of the same type.

Example:

You have customized the pictorial view of your data while working on an LPAR. Your colleagues are working one another LPAR. You want your colleagues to apply the customizations to a different object of the same type residing on that LPAR. To do so, you can save the customized pictorial-view as a public-perspective.

To save a perspective, click the Perspectives and Maps icon, click Save Perspective, fill out the appropriate fields, and click Save. You can save a perspective with only one root node.

To apply a perspective, select an object from the Topology Viewer, click the Perspectives and Maps icon, click Apply Perspective, click the + sign on the perspective, and select Apply. The owners of the perspectives can edit, duplicate, and delete their entries.

When you modify an applied public-perspective, your edits are saved to a copy. The original perspective stays unchanged.

If you change a perspective (for example, applying a filter, adding nodes, and so on) after it has been applied to an object, you cannot remove the perspective.

Manage Maps

After customizing the pictorial view of your data, you can save the view as a map. Saving a map helps you or others load the view.

Example:

You have customized the pictorial view of your data while working on an LPAR. Your colleagues are working one another LPAR. You want to share the customizations with your colleagues. To do so, you can save the customized pictorial-view as a public map.

To save a map, click the Perspectives and Maps icon, click Save Map, fill out the appropriate fields, and click Save. Unlike the perspectives, you can save a map with any number of root nodes.

To load a map, click the Perspectives and Maps icon, click Load Map, click the + sign on the map, and select Apply. The owners of the maps can edit, duplicate, and delete their entries.

Note: The maps that you have saved are stored under the *Chorus\Topology Maps* section of the Discipline drop-down list in the Investigator tree view. After selecting a map, follow the actions available in the Actions pane to continue.

When you modify a public map, your edits are saved to a copy. The original map stays unchanged.

View 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.

TSF lets you quickly compare performance data, which can help you complete the following tasks:

- Troubleshoot an issue
- Identify an area that is approaching a questionable threshold
- Compare data from a different or similar time period

Each selected metric produces a chart and each selected entity produces a line on the chart. You can produce up to four charts with up to four entities on each chart. You can set time ranges, start dates, and end dates.

Data presented on the charts appears in local time.

More information:

Base Entities and Contributors (see page 51)

Display an Entity in a Time Series Facility Chart

Performance objects added to the Time Series Facility (TSF) appear in the TSF as *entities*. An entity is a virtual object that holds one or more performance metrics that can be charted on a graph. These metrics are sent from a registered CA product and then collected and stored in TSF. The entity can be any object type that the CA product supports. The CA product generates the metrics and the discipline for view access to its metrics.

You can use the TSF to chart the performance information within a specific time frame. These charts help you see performance issues, identify trends, and compare current and historical data.

Follow these steps:

1. Launch the Investigator (see page 28).
2. Navigate the Investigator tree to the applicable performance folder, which varies by discipline.

   Example:

   The Security and Storage disciplines have TSF entities under Statistics.

   The Infrastructure Management discipline has TSF entities under:

   - CICS, IMS, WebSphere MQ and z/OS Statistics (CA SYSVIEW)
   - IP Stacks, Network Interfaces, Network Summary

   A list of available performance objects appears.
3. Select or highlight the required object in the table, and click Add Entity to Time Series.

   The Investigator switches to the TSF view, and the entity is added to the TSF.

   Note: You can add as many entities as required to the TSF. Use the Table View icon in the toolbar to switch your view.
4. Select an entity from the Chart Tools section.

   The list of metrics for the selected entity appears.

   Note: Different entities can have different metrics. The metrics that are presented are from the entity that is bold. If you expect to see a metric, but it is not presented in the metric list, click the entity to make it bold. A new metric list for that entity appears.
5. Select the applicable metric.

6. (Optional) Repeat Steps 2 through 5 to select more entities and metrics. You can select up to four entities and up to four metrics with each entity.

   The entity and metrics are highlighted.

7. Click Perform Charting.

   TSF produces a chart for each selected metric. Each entity becomes a line on the chart. The end time on the chart is the last collected data point. To change the section of the chart that appears, use the slider control.

   Note: If no data points exist for the selected metrics, click Latest Recorded (see page 49) to see the most recent recorded data. For more information about the TSF charts, see the *Troubleshooting Guide.*

8. (Optional) Specify the period for the chart using one of the available methods in the upper-left corner.

   Note: If you change the charting period using one of these methods, the amount of data that is being aggregated changes. Therefore, the values for the Y axis of the chart change to reflect the increase or decrease in the aggregate data.

   a. Select a ZOOM option.

   b. Use the slider control to see data for the remaining data in the chart.

   c. Specify the start date, start time, end date, and end time for the chart using the time-based drop-downs.

The TSF chart displays performance information for the selected entity.

Compare Time Series Facility Chart Time Ranges

After you create a TSF chart, you can duplicate the chart with the Duplicate Chart icon (that is, the icon next to the Contributors button). The duplicated chart appears beneath the original chart with independent end time and slider controls. The duplicated charts show the same duration as the original chart. You can compare multiple charts, each with a different time range. The changes that you make to the timeslider of one chart appear in your other charts. The duplicated charts show the same duration as the original chart.

Example:

If you have a spike on your chart, duplicate the chart and compare it with the same time on the previous day using the time slider.

Latest Recorded Data

You can view the last recorded data for the charts that are currently displayed in the TSF, for the selected entities and metrics. When no current data exists for a selected entity, this option helps you see the most recent recorded data.

Click Latest Recorded in the ZOOM options. The link is enabled when the entities and metrics that you select match the current charts. If your selection does not match the current charts, the link is disabled.

Note: If you select multiple entities, TSF takes the most recent date and time at which the data was last collected for the entities. The chart is displayed according to this date and time.

One of the following results appears based on the last date and time the data was recorded:

- If the end date and time at which the data was last recorded falls within the slider time range, you can view the updated chart. The End Date/Time displays the corresponding values.
- If the end date and time at which the data was last recorded falls outside the slider time range, a new query is made. Also, a chart is displayed with the last recorded data. This behavior is similar to charting with a new end date and time.
- When data does not exist in the database, a message appears indicating this fact. You can see this message when you hover on the Latest Recorded link, and also when you click the link.

Note: TSF determines the latest recorded details at the time the original chart was displayed. If new data samples occur after this time, the samples are not included in the chart. Request another chart to view the new samples.

Analyze Contributors

After you have created a TSF chart, use Contributors to perform the following actions:

- Break down a base entity to show the contributors to metric values.
- Compare and chart different or new contributors.

Use the TSF contributors feature to analyze what contributed to the metric values shown for a base entity. This action lets you drill deeper into performance data. For example, if your base entity shows high CPU, you can identify which contributor is having the most impact.

Follow these steps:

1. Access an existing TSF chart.
2. Select an entity from the Contributors drop-down list, and click Contributors.

   The Entities panel becomes the Base Entity panel and shows the component parts that contribute to the selected entity.
3. Click the Contributor Type drop-down list, and select a type to analyze.

The list of available contributors for the selected type and metric appears. The contributors are sorted in descending order based on the metric being charted. When multiple metrics are charted, the contributor is sorted and based on the first metric selected and charted.

Note: If the Base Entity is fully qualified (does not have any available contributors to the metric), the Contributor Type list is empty. No further analysis is possible.

4. Select the applicable contributors, and click Perform Charting.

   The chart shows the base entity plotted with the new contributor plots.

Base Entities and Contributors

An understanding of *base entities* and *contributors* is useful before you begin analyzing data. To understand this relationship, visualize these terms in a parent (base entity) and child (contributor) relationship. Parents in a family are a base entity and each child is a contributor to the family. Depending upon the metric, each child could contribute differently. Some children do not have contributors. Parents are also children (contributors) in an extended family.

When you view a performance metric for a base entity in TSF, you see the total view, for example, total CPU. When you plot contributors for a base entity, you see the contributors that are part of the total view, for example, each component contributing to the total CPU. You can add or remove contributors to isolate performance areas. Additionally, you can use TSF to change a contributor into a base entity.

Example:

For CA Chorus for Security and Compliance Management, a system could be a base entity with each Command Propagation Facility (CPF) node as a contributor.

Change the Base Entity of a Time Series Contributor

Use TSF contributors to change the base entity and chart different combinations of contributor types. This action lets you analyze the contributors to metric values using different starting points or base entities. For example, you can use this feature if you have five contributors to a metric value, but you want to isolate one contributor.

A base entity value added from the contributors list cannot be removed from the base entity list. To remove the base entity, exit and reenter the contributors display.

Follow these steps:

1. Access an existing TSF chart.
2. Select an entity from the contributors drop-down list and click Contributors.

   The Entities panel becomes the Base Entity panel and shows the component parts that contribute to the selected entity.

3. Click the Contributor Type drop-down list, and select a type so you can add or update the base entity.
4. Update the base entity by performing *one* of the following steps:
   - Add the contributor to the existing Base Entity criteria by selecting the Set in Base Entity option from the Actions button for that contributor. This action effectively increases the filtering criteria of available contributors to the Base Entity.
   - Expand or replace a base entity component criteria by performing the following steps:
     - Uncheck a base entity component. This action effectively expands the Base Entity contributor filtering criteria allowing a bigger range of contributors from which to select.
     - Uncheck and replace the base entity component with a different specified value.

   Note: You can replace a base entity component value by unchecking an existing base entity. Doing so lets you expand the contributors that match that Base Entity Criteria. You can then select the same type from the Contributor type drop-down list, and select a new Contributor value to Set into Base Entity using the Actions button.

   The base entity is changed.
5. Click Perform Charting.

   The TSF chart is recalculated and displays the new base entity and any selected contributors.

Change the Criteria of a TSF Contributor

Use the TSF contributors feature to change the criteria for an existing TSF chart. This option lets you manipulate data based on your changing needs.

Follow these steps:

1. Access an existing TSF chart.
2. 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 passed by the Investigator for the selected entity.
3. Click the Contributors drop-down list.

   The drop-down list shows the valid contributors.

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

4. Select a contributor type.

   A list of all the contributors for the selected entity and metric combination appears.

5. Click Perform Charting.

   The new chart is generated from the selected entities.

View Default TSF Chart

CA Chorus lets you launch with a single click the time series UI, displaying default TSF charts with predefined metrics.

Follow these steps:

1. Launch the Investigator (see page 28).
2. Select an appropriate statistics object from the objects tree in the discipline pane.
3. Select an appropriate object row from the objects grid in the data pane.
4. Click the Add Entry to Time Series link in the Actions pane.

   The TSF chart for the selected object appears displaying default entities and metric values.

Note: The Chart Tools pane at the right side of the window displays options such as to select and remove metrics. You can use these options to create a chart with different metric values other than the default ones. You can also remove the selected entity and search new entity to create an all new chart. Selection of irrelevant entity or metrics, however, does not generate a chart.

Investigator Paths

A *path* is a linear series of views (pages, locations, objects, charts) that you can save as an example for others or for your use. Paths help you quickly return to a specific area in the product for training, troubleshooting, or daily operations.

A path entry is automatically displayed in the History pane when you navigate in the Investigator. A breadcrumb represents each step in a path. When you click a breadcrumb to return to a point in the path, the corresponding historical data for the breadcrumb appears in the Investigator.

If you decide to complete the current path later, save it as an incomplete path. When you close the Investigator without saving a path, the Investigator saves it as an unsaved path. Incomplete paths and unsaved paths are stored as private paths with an incomplete status.

Note: The Investigator can store only one unsaved path per user. When a new unsaved path is present, the Investigator overwrites the old path.

Examples:

- Consider the series of tasks you complete to obtain a frequently used list of data. You can save these tasks as a path so that when you need the list of data, you can load the path from the Investigation Launcher. This action saves the time that is required to repeat the tasks.
- Use the paths when you train employees. You could ask an employee to locate specific data in the Investigator. They could then show you the path that they used to find the information.

Save a Path in the Investigator

You can save a path so that you or other users can investigate the path later. Saving a path lets you quickly return to the path to investigate the activity. Path permissions determine who can view or edit a path. Public permissions let all users view the path. Private permissions let only you edit or view the path.

Private and public paths are also saved in the Knowledge Center for reference, training, and similar purposes.

Save a path with a title convenient for you to remember. For example, *Path* followed by your role such as Security Administrator.

Follow these steps:

1. Click Save Path in the History pane.
2. Enter a name and description for the path, select the completion status, set path permissions, and click Save.

   Note: You can set permissions for a complete path only. Permissions for incomplete paths are set to private.

   The Investigator saves the path that you can view in the Investigation Launcher as well as on the Investigator window. To view a path on the Investigator window, click the Paths link in the Saved Investigations pane. To edit the path, click the path, click the edit icon ( ) in the Details pane, enter your changes, and click Save.

Manage Saved Paths

You can view, edit, or delete a saved path by using the Investigator.

Follow these steps:

1. Add the Investigation Launcher to the dashboard.
2. Select from the Investigation Launcher the path that you want to manage to the Investigator.
3. Select one of the following options:
   - To view path details, click Load or Resume in the Actions pane.

     A Complete path takes you to the first breadcrumb, and an Incomplete path takes you to the last breadcrumb.
   - To edit the path, click the edit icon in the Details pane, make necessary changes, and click Save.
   - To delete the path, click Delete Path in the Actions pane, and follow further instructions.

     Note: The Delete Path link, however, does not appear for the paths saved by other users, thus refrains you from deleting them.

How to Maintain Path Relevance

We recommend that you regularly clear unnecessary paths from the Investigator and the Investigation Launcher. Doing so lets you start a new path in the History pane in the Investigator and maintain a clean Investigation Launcher. Clearing and deleting paths from the Investigation Launcher also helps you keep the paths relevant and easily accessible.

To remove paths, perform one or both of the following:

- Click Clear Path in the History pane in the Investigator and follow the prompts.
- Click the path that you want to delete in the Investigation Launcher, click Delete, and follow the prompts. You can delete only paths that you saved.

Manage Breadcrumbs

Breadcrumbs represent steps of a path. You can click a breadcrumb to view the status and stored data of the path at that point in time.

You can add comments to a breadcrumb in the History pane. Comments can identify events for you and other users to view. Comments, the path, and time and date when the breadcrumb created appear when you hover over the breadcrumb.

Follow these steps:

1. Click the arrow on the bottom of the applicable breadcrumb and click Comment.
2. Type the comment that you want to add and click Save.

   The comments icon identifies the annotated breadcrumb.

Note: You can remove breadcrumbs from a path to make the path relevant. To remove a breadcrumb, click the arrow at the bottom of the applicable breadcrumb, click Remove Breadcrumb, and follow the prompts.

Start a Custom Investigation

Essentially, instead of reviewing data in the full Investigator, the *Custom Investigation* module lets you select a subset of data to view in this module. Doing so lets you narrow the data you are doing to simplify your view, which can lead to efficiencies. The module requires only one-time configuration, following which you can frequently monitor and perform actions on these subsets in the subsequent sessions.

Custom Investigation supports multiple module configurations on the same dashboard, enabling convenient comparison of objects. The module saves current status of the page with the most recent search and action. A new session restores the saved data for you to resume your task from the same point.

Follow these steps:

1. Add Custom Investigation to the dashboard from the Module Library.

   The home page populates the available disciplines for the logged in user.
2. Enter a name for the window, select appropriate discipline and object node.
3. (Optional) Enter a search keyword for the object node.
4. Click Save.

   A new window opens in your dashboard with the name you entered. This window displays the selected object node and corresponding object rows.
5. (Optional) Click the View In Investigator link in the Actions pane.

   The Investigator window opens with details about the selected object node.

Every time you perform an action, the Custom Investigation module saves it as your latest state. Therefore, your last state before a logout is always saved for the new session.

How to Share Knowledge Using Notes

As a CA Chorus administrator, use the Investigator to view and analyze information that is stored in discipline-specific data repositories. While analyzing object-specific information, you can create notes to record your observations, findings, and general comments. The Investigator lets you create and insert notes to a row of data representing a path. A note can be a private or public information that is related to an object.

Notes act as on-screen reminders and greatly increase communication among users and identify changes in the usage of system resources, interaction with a database, and similar activities. Notes also appear as results when you search the Knowledge Center.

For example, you have investigated and identified the system that causes service outage during the peak working hours. You raised a Service Desk ticket to resolve this issue. You can add a note that mentions the service desk ticket number to the identified system. This note helps your colleagues understand that the system is under maintenance. You can revisit and update the note when you get updates on your Service Desk ticket.

The following illustration shows how an administrator shares knowledge using notes:

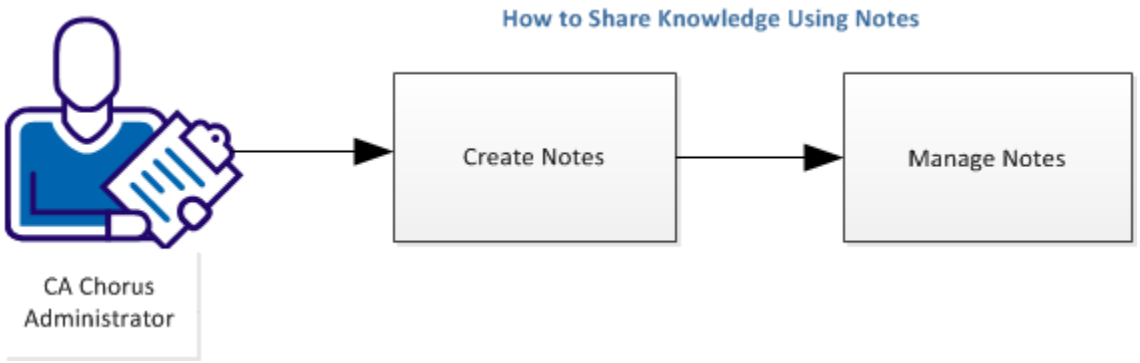

To share knowledge using notes, complete the following tasks:

1. Create Notes (see page 57)
2. Manage Notes

Create Notes

Notes can greatly increase communication among users, serve as logs for data, or identify errors or ongoing research projects. In the Investigator, you can add notes to a table row. In addition, the notes appear in the Notes module according to the permissions that you set. A note is saved with descriptive details, including a numeric ID, the author, and time and date of creation.

In this example, you created a Service Desk ticket to resolve the problem with one of your systems. You can create a note with the following details and attach the note to the row in the Data pane that identifies the system:

- A short description of a problem.
- The Service Desk ticket number.
- The status of the ticket.

Use the following procedure to add a note to a table row in the Investigator.

Follow these steps:

1. Launch the Investigator.
2. Highlight the table row where you want to add the note, and click the Add Notes icon () in the Investigator toolbar.
3. Enter the note text, select the permissions, and click Save.

   Note: A note can be 1024-character long and can include any text and special characters except < and >. Also, we recommend that you identify your discipline - such as security administrator - in the note title or body.

   The note is added to the selected table row and to the Notes module. Table rows that contain notes include a Notes icon () in the Notes column. To view notes for a table row, click the Notes icon. The All Notes window displays all notes for the table row. Public notes that you added appear as search results in the Knowledge Center.

Manage Notes

The Notes module houses all notes in its Private Notes and Public Notes tabs. A note is stored with the time and date of creation for viewing later. The Private Notes tab displays notes that only the author of the note can view and edit. The Public Notes tab contains notes that all users can view but only the author of the note can edit or delete. When you select a note by clicking the + button, more information appears for managing the note.

The auto-refresh option is enabled by default for authorized users. When this option is enabled, any new note created at any work space (or module) refreshes the Notes module once every 30 seconds. For example, when a note is created in the Alerts module, it appears in the Notes module after 30 seconds. Click Refresh () to disable or enable the auto-refresh option. If you are not an authorized user, click Refresh () to refresh the data manually.

You can customize the tabular view of the dialog data by clicking the wrench icon.

In this example, you created a note to communicate the details of the Service Desk ticket with your colleagues. As your Service Desk ticket goes through various phases, communicate the status to keep your colleagues updated. Also, when the issue is resolved, you can delete the note.

The following procedure helps you edit and delete notes:

Follow these steps:

1. Add the Notes module to the Dashboard from the Module Library.
2. Select the applicable note from the Notes module by clicking the + button.

   The note entry expands to display the note.

   Note: For longer notes, only the first part of the note appears. Click View Detail to view the complete note.
3. Review the note contents. For example, a note may explain the root cause and resolution of a recent service outage.
4. Select one of the following options:

   Note: The Edit and Delete options appear only for the note author.

   - Click Edit, make your changes, and click Save. You can also change the permissions in this window to provide or limit note access. For example, you have uncovered new details about the recent outage, and you want to add them so that your colleagues have the best information.
   - Click Delete, and confirm the deletion.

   The note is edited or deleted.

You created a note to store your object-specific observations, findings, and general comments. You enabled knowledge transfer by sharing the note.

Chapter 3: Using the Quick Links Module

How to Add a Discipline Hyperlink

This scenario shows how an application developer adds a discipline hyperlink to access data and tools quickly and respond to requests to troubleshoot issues. Each hyperlink in the Quick Links module provides shortcut access to various interfaces. Many companies create their own tools to manage various tasks or house data. Adding a link to these components in this module can help you efficiently manage tasks and your time from CA Chorus.

The following illustration shows how an application developer adds a discipline hyperlink:

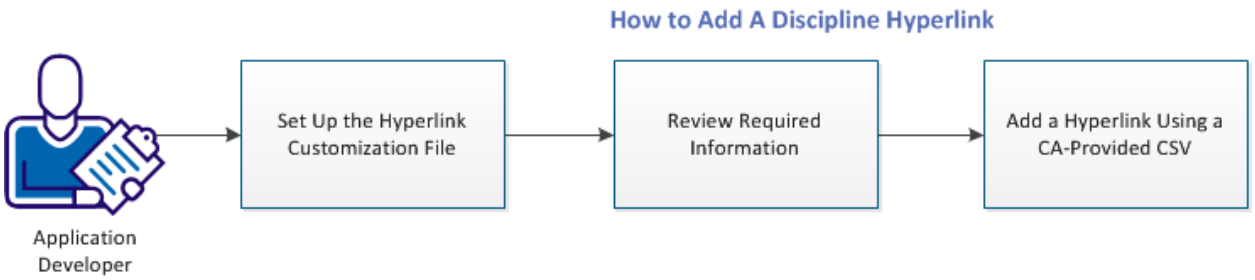

The default hyperlinks display by discipline. However, the links that you add are not necessarily specific to a discipline. The following diagram shows the module layout and the default links for each discipline:

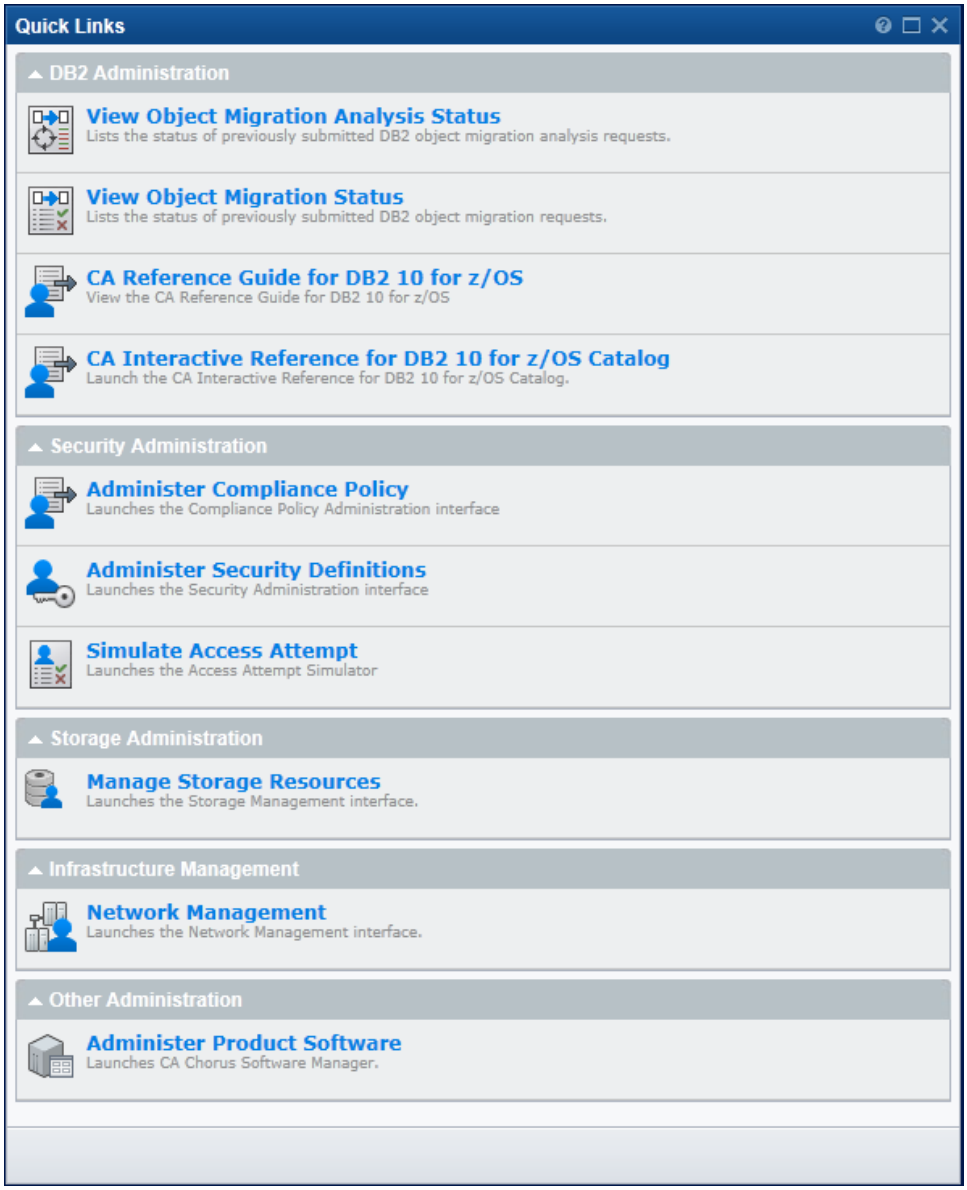

Note: Disable pop-up blockers for your browser before you access hyperlinks through the Quick Links module.

Use the following topics to add a hyperlink:

1. Set Up the Hyperlink Customization File (see page 63)
2. Review Required Information (see page 63)

3. Add a Hyperlink Using a CA-Provided CSV (see page 65)

Set Up the Hyperlink Customization File

Use this procedure to set up the hyperlinks customization file.

Follow these steps:

1. Determine if you have run ETJI0130 in *chorus_runtime_hlq*.CETJJCL.

   Important! You need only run this job once. If you run it extra times, you erase your custom hyperlinks.

   - If it has not been run, go to the next step.
   - If it has been run, go to the next topic.

2. Edit member ETJI0130 in *chorus_runtime_hlq.*CETJJCL as described in the member, and submit for execution.

   This job creates a default custom hyperlink file on the configuration file system. Edit this file to add custom hyperlinks in the Quick Links module.

3. Restart the CA Chorus Application Server for these changes to take effect.

Review Required Information

The following information is required:

Quick Link Schema:

URL

Unique resource address that launches a browser-based interface when a user clicks the hyperlink. This hyperlink can point to any accessible URL. For example, an external web site, an internal web application within your organization running elsewhere, a web application running under the same CA Chorus Application Server instance as CA Chorus.

Note: If you want to pass data, for example: passticket and username, as part of the URL, use a URL that supports POST requests.

When a hyperlink requires dynamic information, use variables (for example, server and port) and a velocity script so you do not have to enter hard-coded values. Note the following velocity script rules:

- Embed the entire URL expression in double quotes.
- Escape each double quote included inside the velocity expression with another double quote. For example, see the $USERID expression.
- Place any SQL provided inside $db.oneResult() in double quotes. Escape these quotes with another set of double quotes.

The following list details current variables and how to use them in a velocity script:

Note: For the following $db.oneResult references, this function can take any valid SQL expression and can return exactly one value, which is the first column of the result set. The SQL expression can be anything that Teiid can execute.

$SERVER

```
#set ($server = $db.oneResult(""CALL
chorus_platform_config.config_map_get('server')""))
```

This chorus_platform_config stored procedure takes a key to the global map in CA Chorus, which returns the corresponding value.

$PORT

```
#set ($port = $db.oneResult(""CALL
chorus_platform_config.config_map_get('port')""))
```

This chorus_platform_config stored procedure takes a key to the global map in CA Chorus, which returns the corresponding value.

$USERID

```
#set ($userid = $db.oneResult(""CALL
chorus_platform_config.get_chorus_user()""))
```

This chorus_platform_config stored procedure identifies the user ID.

$PASSTICKET

```
#set ($passticket = $db.oneResult(""CALL
chorus_platform_config.generate_passticket('$userid','applid')""))
```

This chorus_platform_config stored procedure takes user ID and applid and returns the passticket.

Label

Clickable text to launch the URL.

Description

Text that appears below the hyperlink label in the Quick Links module.

Icon URL

Absolute URL of the icon that appears next to the hyperlink label and description in the module. This link can point to any accessible URL, which may or may not be on the same CA Chorus Application server as CA Chorus.

If you do not have an absolute URL, use the CA icon at *chorus_install_directory*/resources/resources/images/icons/Support-Icon_32_A.png

Section Name

Text that is used to group related hyperlinks and display them under a tab. This text is the tab header label.

Values: Icnfrastructure Management, Security Administration, Storage Administration

Note: You can also enter your own section name, such as General.

Role

Column that determines to which discipline this hyperlink belongs.

Values: CHORUS.ROLE.INFRASTRUCTURE, CHORUS.ROLE.SECURITY, CHORUS.ROLE.STORAGE, CHORUS.ROLE.NONE

Note: By specifying CHORUS.ROLE.NONE, you indicate that the link is to appear in the module regardless of the disciplines to which the user has access.

Add a Hyperlink Using a CA-Provided CSV

CA Chorus includes a facility to add hyperlinks as comma-separated value (CSV) file entries. Use this procedure to add a discipline hyperlink using a CA-provided CSV. When adding a hyperlink to the Quick Links module using a CSV file, you can use the custom-quicklinks.txt file at *chorus_install_directory/*config.

Follow these steps:

1. Navigate to *chorus_install_directory*/config.
2. Open the custom-quicklinks.txt file.
3. Add the required schema information, and save the file.

   Note: In both examples, to view each hyperlink, you must have access to applicable resource (bold). To add this resource name, contact your security administrator. For the steps to add this user authorization, see the *Site Preparation Guide*.

   Example: Variables Not Required

   This example adds a link named *CA Support Online.* Use all values noted, except the icon file name. Place the icon of your choice in the appropriate folder and use the name here.

   ```
   http://support.ca.com,View CA Support Online,Launches the home page.,resources/resources/images/icons/CA_32b.png,SDK,CHORUS.ROLE.NONE
   ```

   Example: Variables Required

   This example shows how the Security Administration Interface is added to the Quick Links module using variables and a velocity script. In the first two lines, you specify the schema information, including the server and port variables. The remainder of the text is a velocity script to simplify how you add or modify hyperlinks that require dynamic information.

   Note: *CALDAP* indicates where you specify the APPLID, which allows single signon.

   ```
   http://$server:$port/SecAdmin,Administer Security Definitions,Launches the Security Administration interface,resources/resources/images/icons/admin_securitydefintn_32b.png,Security Administration,CHORUS.ROLE.SECURITY

   "#set ($server = $db.oneResult(""CALL chorus_platform_config.config_map_get('server')""))

   #set ($port = $db.oneResult(""CALL chorus_platform_config.config_map_get('port')""))

   #set ($userid = $db.oneResult(""CALL chorus_platform_config.get_chorus_user()""))

   #set ($passticket = $db.oneResult(""CALL chorus_platform_config.generate_passticket('$userid','CALDAP')""))

   http://$server:$port/SecAdmin?var1=$userid&var2=$passticket"
   ```

4. Restart the CA Chorus Application Server for these changes to take effect.
5. Log back in to CA Chorus.
6. Launch the Quick Links module.

   The new hyperlink appears under the applicable discipline section.
7. Click the new hyperlink.

   Your link opens.

You have configured the Quick Links module to include one of your UI tools.

Chapter 4: Using the Policy Status Light Module

How to Configure the Policy Status Light Module

As a system administrator, you monitor system performance to ensure optimal productivity for your users. CA Chorus offers various tools to help you monitor performance, including the Alerts module, Metrics panel, and the Policy Status Light module.

This scenario shows how an administrator configures the Policy Status Light module to monitor the memory usage of a server. A low-memory condition can cause unexpected program behaviors. By monitoring the usage, you can intervene before a low-memory condition causes problems.

To monitor performance, you need policies that identify the area to monitor and thresholds that indicate an important change in your data. A *policy* is the rule to evaluate a data point in an object against a user-specified value. Policy defines the actions to be executed on the successful evaluation of the rule.

Example:

Your server object has memory usage as the data point. You can define policies to change the colors in the Policy Status Light module when the memory usage reaches 85 percent and 95 percent. Monitoring memory usage at 85 percent and 95 percent requires two different policies. If these policies are not available, you must create them. You link these policies to the Policy Status Light module, which can help you monitor the health state of your server by displaying green, yellow, or red backgrounds to alert you to perform changes. From this module, you can see issues as they arise, investigate them, and identify root cause.

The following illustration shows how a system administrator configures the Policy Status Light module. This scenario uses this memory usage example to illustrate how to configure this module.

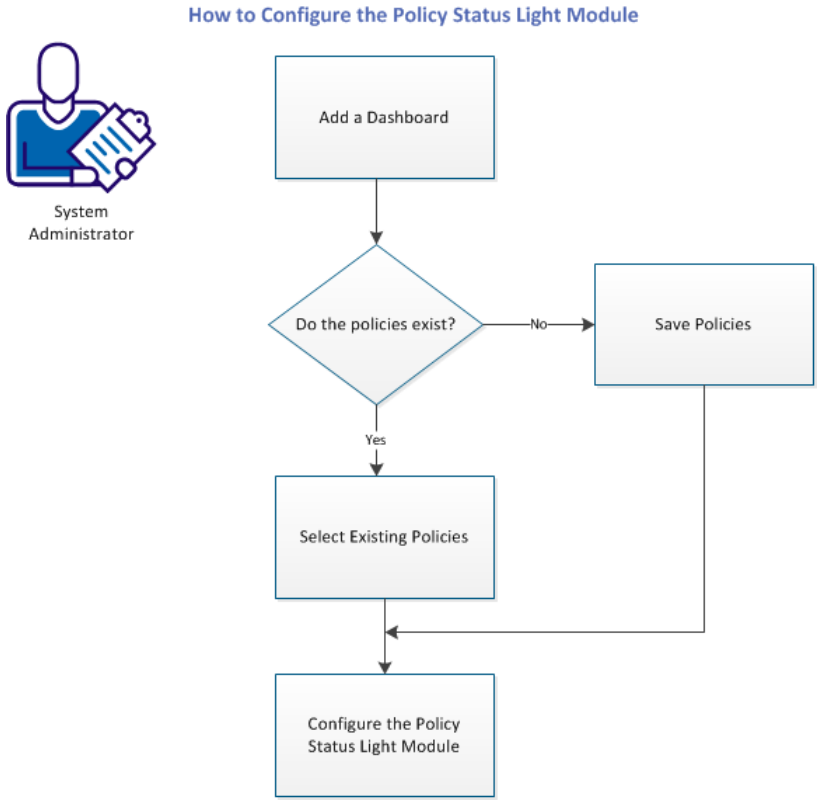

To monitor the memory usage, complete the following tasks:

1. Add a Dashboard (see page 71)
2. Choose one of the following options:
   - Select Existing Policies (see page 71)
   - Save Policies (see page 34)
3. Configure the Policy Status Light Module (see page 72)

Add a Dashboard

A *dashboard* is a customizable area that contains modules that are necessary for your tasks and projects. For example, you can create a dashboard to monitor the memory usage for your critical server.

Follow these steps:

1. Log in to CA Chorus.
2. Click the plus sign dashboard tab.
3. Select the option to create a dashboard, enter a name for your dashboard, for example: Memory Usage at Product Lab, and click Add Dashboard.

   The dashboard is added. You can add modules, metrics, and text boxes related to monitoring performance of this site.

   Note: The dashboard name can be up to 21 characters. Numbers (0 through 9), characters (a through z and A through Z), underscores, and blank spaces are allowed.

Select Existing Policies

To monitor memory usage at 85 percent and 95 percent, you need two different policies. The Policies Object Picker lets you create, select, and manage the policy objects. Before selecting the policies, review the objects, rules, and actions that are set for the policies.

If the policies do not exist, skip this procedure and go to Save a Policy (see page 34).

Follow these steps:

1. Open the Module Library.
2. Add the Policy Status Light module to your Dashboard.
3. Click the link to configure the module.

   The Policies Object Picker window displays Chorus Personal Policies and predefined policies.
4. Select Chorus Personal Policies under the Policies node from the left pane.
5. Select the policies that are defined for memory usage, and click Add To Cart in the Actions pane under Policy Actions.

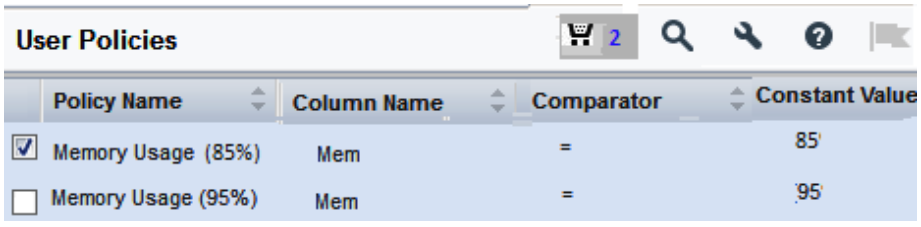

Note: You can repeat step 5 across different tree nodes and disciplines. The Shopping Cart icon () displays the total count of policies that you have selected. Click this icon to see the list of selected policies. To remove a policy from the cart, select the check box, click Remove, and then click Save.

6. (Optional) Review the policy details in the Details pane.
7. Click Save.

You have selected the policies to monitor the memory usage of your server. To configure the module immediately after selecting the policies, continue with Step 5 of the procedure Configure Policy Status Light (see page 72).

Configure the Policy Status Light Module

You can configure the Policy Status Light module to monitor predefined and user-customized policies in a color-coded window. The colors indicate if the system has exceeded the thresholds in your policies, which can help you respond quickly to performance changes. For example, Green color indicates that the performance change does not meet the defined criteria, and does not fetch any results. On the other hand, red or yellow color appears with appropriate result when performance change meets the defined criteria.

Note: If you have already configured the module, skip steps 1 through 4 of the following procedure.

Follow these steps:

1. Open the Module Library.
2. Add the Policy Status Light module to your Dashboard.
3. Click the link to configure the Policy Status Light.

   The Policies Object Picker window displays the titles of personal and predefined policy lists in the left pane.

4. Click the applicable titles in the left pane, and select those policies in the middle pane that you need to monitor the memory usage of your system.

   Note: If you have a long list, use the View Filter option to select policies.

5. Click Add To Cart in the Actions pane under Policy Actions.

6. (Optional) Click the cart icon, and select the applicable added policies.

   Note: You can also remove policies from the shopping cart. Select the check box, and click Remove in the Actions pane under Navigation.

7. Click Save.

   The Policy Status Light displays the health states of policies on the dashboard.

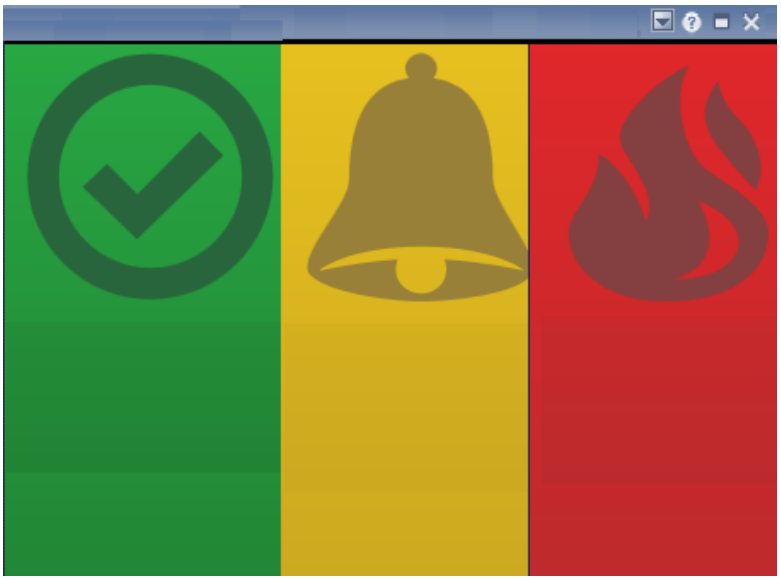

   Hover over a Policy Status Light to see the list of objects that are in the same condition. If you have only one policy and you delete it, the Policy Status Light turns green to indicate that the policy is no longer available.

   Example: Multiple Policy Usage

   Consider a situation where your policy contains other objects such as processor usage. When the processor usage and the memory usage are in the worst condition, the hover text displays memory usage and processor usage.

8. (Optional) Customize the Policy Status Light view using the drop-down list on the title bar.

   Example:

   You have four policies named P1, P2, P3, and P4. The following list details your policies:

   - P1 has one object of type T1; the health state of this object is red.
   - P2 has two objects of type T1; the health state of these objects is yellow.

- P3 has three objects of type T2; the health state of these objects is red.
- P4 has four objects of type T3; the health state of these objects is yellow.

The following table shows how to configure the Policy Status Light to display your objects:

Option	View
Show Overall (default view)	One red light and hover text for one P1 object and three P3 objects.
Group By Policy	Four lights: ■ One red light for P1 and hover text for one P1 object. ■ One red light for P3 and hover text for three P3 objects. ■ One yellow light for P2 and hover text for two P2 objects. ■ One yellow light for P4 and hover text for four P4 objects.
Group By Object Type	Three lights: ■ One red light for T1 and hover text for one P1 object. ■ One red light for T2 and hover text for three P3 objects. ■ One yellow light for T3 and hover text for four P4 objects.
Show Object Instances	Nine lights, which is the sum of the policy objects.

You have created a memory usage dashboard for your site Product Lab, created and selected policies, configured the Policy Status Light module, and started monitoring the memory usage of your server.

Note: If you receive a query canceled notification, reconfigure the Policy Status Light module.

In the future, if a threshold is crossed, right-click the policy on dashboard, and select the Investigate option to identify root cause. You can then use several tools in the Investigator to identify the root cause. For example, add the object of the policy that showed 95 percent of memory usage to the Time Series Facility, duplicate the chart, and compare it with the same time on the previous day using the time slider.

Chapter 5: Using the Knowledge Center

How the Knowledge Center Works

The Knowledge Center is the repository for all documentation in CA Chorus. Content can include:

- CA product documentation
- User-generated documentation, including:
  - Your paths and public paths that you have saved while working in the Investigator.
  - Your notes and public notes.
- Websites
- Links to third-party documentation
- Chicago-Soft MVS/QuickRef messages

Knowledge Center content can help you perform tasks and become familiar with the product.

When you click Help, the Knowledge Center displays help topics according to your task or your current user interface location. The Knowledge Center identifies your current task or user interface location by gathering information and keywords from the following sources:

- The module where you are working
- The task that you are performing
- Highlighted text
- Search terms that you enter in the Knowledge Center search field

The Knowledge Center compares this information with existing topics in the Knowledge Center repository. Links to topics that match the criteria are returned in the Knowledge Center window.

The Knowledge Center also returns links to documentation topics that you have added to the Knowledge Center repository. Adding documentation to the repository lets you build a set of documentation that applies more directly to your role or tasks you perform regularly.

Supported File Types

You can add many types of documentation to the Knowledge Center repository. The Knowledge Center repository supports the following file types:

Important! The Knowledge Center only supports the indexing of UTF-8 (ASCII) file formats. To index an EBCDIC file, create a copy and convert it to UTF-8 format, and then index the UTF-8 copy.

- HTML files (.html and .htm)

  Note: If you are including your HTML files in the repository, internally linked HTML files are not uploaded or indexed automatically. CA Chorus uploads and indexes only the first level of HTML files.

- Text files (.txt)
- PDF files (.pdf)
- Microsoft Word files (.doc and .docx)
- Microsoft PowerPoint files (.ppt and .pptx)
- Microsoft Excel files (.xls and .xlsx)

Note: Your browser or system settings determine whether files open directly when you click the file links in the Knowledge Center window.

Searching the Knowledge Center

The Knowledge Center searches for documentation that is based on keywords. Keywords include items such as product features, parts of the interface, tasks. If you want your uploaded documentation to appear in your search results, use the keywords from your document as the search terms. The Knowledge Center displays the first five results containing the keywords for each data source, including documents that you added to the repository.

Note: Using the drop-down list named Result per data source, you can change the number of results that the Knowledge Center displays for each data source selected.

Important! CA Chorus consolidates indexes periodically to improve search efficiency. If you do not receive search results immediately after you upload a file, you may have requested data during the consolidation process. In this case, retry the search after a few minutes.

Search Documentation

The Knowledge Center appears when you click the Help icon on any CA Chorus window. When the Knowledge Center opens, the window contains matches for the module where you were working. You can refine these results by entering more search keywords in the Search entry field. Search keywords override the original returned matches.

Follow these steps:

1. Type the search term in the Search entry field in the Knowledge Center browser and click the Search icon.
2. (Optional) Click Advanced Search to refine your search and get more relevant results. You can enter the search criteria and can select the search sources.

   Note: You can customize the data sources in the Knowledge Center to suit your preferences, and save the customized list. You can rearrange the data sources. The search results appear in the same order.

   Note: Administrators often have overlapping responsibilities so we recommend that you configure your search settings based on your role. Doing so can improve the relevance of your search results.

   A list of documentation files matching the search term appears in the Knowledge Center browser.

Search MVS/QuickRef

The Knowledge Center integrates with the MVS/QuickRef™ product by Chicago-Soft, Ltd. When you encounter an error, this feature lets you access the MVS/QuickRef messages directly in CA Chorus. All matches appear in the Knowledge Center search results.

Follow these steps:

1. Highlight the item that you want to search for and click the Help icon.

   Note: The item must be a message, message ID, source code, or similar item. MVS/QuickRef searches for the first highlighted term.

   The Knowledge Center displays the MVS/QuickRef matches that correspond to the highlighted term.
2. Click the match that you want to view.

   The complete MVS/QuickRef content for the term appears.

How to Add Your Files to the Knowledge Center

This scenario shows how a system administrator includes and indexes documentation in the Knowledge Center. Indexing is a process that reviews your document and creates a record of the content and location of your document. You can also index websites. This record speeds up the search. When you first install CA Chorus, the Knowledge Center is already populated with indexed information for CA Chorus and your back-end products.

Adding your own documentation offers the following benefits:

- Increases the number of resources that you can query from the Knowledge Center
- Shares information available to other users of the same system

Important! You may use the Knowledge Center to index content available only on the LPAR where CA Chorus is running. This content is not available when the product switches to a different LPAR. Switching can occur when you have configured the product for High Availability (HA). After the product reverts to the LPAR where the content resides, the content is available.

Important! The Knowledge Center only supports the indexing of UTF-8 (ASCII) file formats. To index an EBCDIC file, create a copy and convert it to UTF-8 format, and then index the UTF-8 copy.

The following illustration describes the steps to add files:

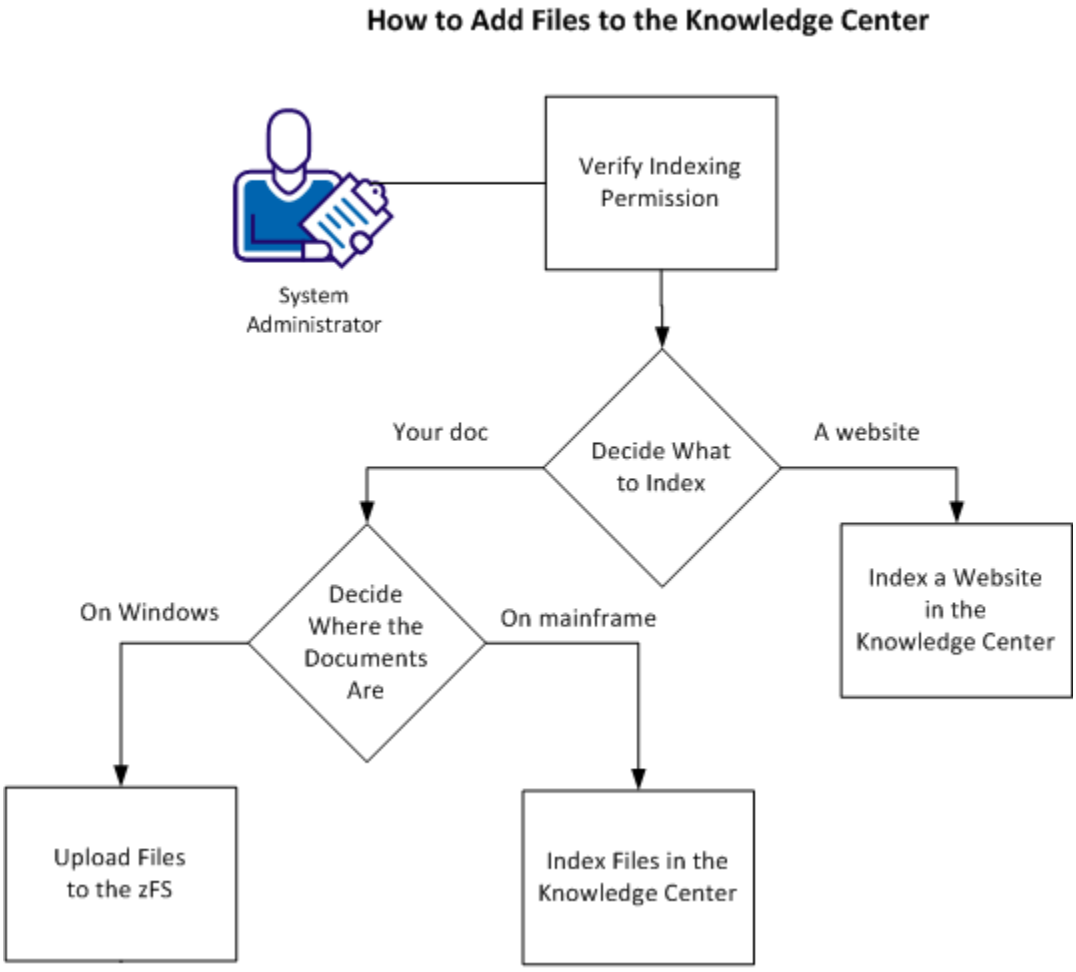

This scenario describes the following steps:

1. Verify Indexing Permission (see page 80)
2. Decide What to Index (see page 80)
3. Decide Where the Documents Are (see page 80)
4. Upload Files to the zFS (see page 81)
5. Index Files in the Knowledge Center (see page 81)
6. Index a Website in the Knowledge Center (see page 82)

Verify Indexing Permission

Only users with indexing permissions are allowed to index and upload user content to the Knowledge Center. This limitation prevents users from including unverified content in the Knowledge Center.

The wrench icon on the Knowledge Center window provides you access to the Knowledge Center Settings dialog. This dialog lets you upload and index documentation and websites. If you do not see the wrench icon on the Knowledge Center window, contact the security administrator for your external security product.

During product installation, your installer used an external security product such as CA ACF2™, CA Top Secret®, or IBM RACF. These security products define the CAMFC resource class that permits you to index content for the Knowledge Center. Only the security administrator can grant you the right to manage content in the Knowledge Center.

Decide What to Index

As a system administrator of the Knowledge Center, you can decide what type of information to include in the Knowledge Center. You can index documentation that you use every day in Microsoft Office, such as .doc, .xls, and .ppt files. In addition, you can index PDF, HTML, and text files that you have on your Windows system.

You can also index an entire website of information. Websites can include third-party websites, SharePoint sites, or wikis. CA Chorus crawls down to two levels in the website. Crawling to two levels means that, when a link refers to another link, CA Chorus indexes content at the first and second links only. Any content at a link referenced from the second-level link and further on is not indexed.

Decide Where the Documents Are

Understanding the location of your documents determines the steps to include them in the Knowledge Center. All files must be on the zFS before indexing.

If you have files on a Windows system, upload the files (see page 81) to the zFS. The files are automatically indexed.

If you already have files on the mainframe, you do not need to upload them to the zFS. You can begin indexing immediately.

Upload Files to the zFS

If you have files on a Windows system that you want to include in the Knowledge Center, upload the files to the zFS. The Knowledge Center indexes the uploaded files automatically. If your files exist on the mainframe, you do not need to upload them.

After you upload files, they are placed on the zFS in a folder that you specify during the installation of CA Chorus. If you change these source files on your Windows system, upload the files again. The files are indexed automatically.

Note: Links for uploaded files that contain nonalphanumeric characters in the file name are not supported when they appear in search results. To view these files, remove or replace the nonalphanumeric characters before uploading the files.

You must have Knowledge Center configuration and write authority to the zFS destination folder to upload your Windows files to the zFS.

Follow these steps:

1. Click the Help icon.
2. Click the wrench icon at the top right of the Knowledge Center window.

   The Knowledge Center Settings dialog opens.
3. Click the Upload Documents tab, browse to the location of the file that you want to add, select the file, and click Upload. Repeat this step if you want to add more files.

   A progress bar appears, which tracks the file upload progress. The progress bar disappears after the uploading completes. The destination directory appears at the top of the bottom pane and the uploaded files appear in the bottom pane. The file is now uploaded to the zFS and is automatically indexed.

   Note: If you are including your HTML files in the repository, internally linked HTML files are not uploaded or indexed automatically. CA Chorus uploads and indexes only the first level of HTML files.

Index Files in the Knowledge Center

The files you upload in the Knowledge Center are not available for search queries until they are indexed. If you upload files from Windows, the Knowledge Center indexes them automatically after the upload. If you include mainframe files in the Knowledge Center, you must index them manually following this procedure.

Follow these steps:

1. Click the Help icon.
2. Click the wrench icon at the top of the Knowledge Center window.

   The Knowledge Center Settings dialog opens.

3. Click the Index Documents tab and enter the path of the folder that contains the files that you want to index. You can also click Browse to locate the folder that contains your files.
4. Select the folders to index, and click the right arrow.

   The selected folders appear in the Folders to Index pane.

   Note: If you already uploaded files from a Windows file system, the Upload Documents tab indicates the location of the uploaded files in the zFS.
5. Click Index.

   The resulting index replaces the existing version of the index for the Knowledge Center repository.

   Important! Indexing many files can cause high CPU usage, which can cause the CA Chorus Application Server to become nonresponsive. We recommend indexing only a few files at a time.

Any search request in the Knowledge Center now yields links to information contained in the indexed files.

Index a Website in the Knowledge Center

You can index websites, SharePoint sites, and wikis in the Knowledge Center. Indexing a website increases the number of resources that you can access from the Knowledge Center. The Knowledge Center crawls only up to two levels of the website. So, if links refer to other links, only two levels are included in the index. After the website is indexed, content from the site appears in the Knowledge Center search results.

Note: Indexing large websites can result in extended indexing time, possibly hours.

Follow these steps:

1. Click the Help icon.
2. Click the wrench icon at the top of the Knowledge Center window.

   The Knowledge Center Settings dialog opens.
3. Click the Index URLs tab, enter the URL of the site that you want to index, and click Index.

   Note: If internet proxy settings or user credentials are required to access a website, you are prompted to configure the proxy server (see page 83).

The site is indexed. The URL and index status appear in the Index URLs pane of the dialog.

Any search request in the Knowledge Center now yields links to information contained on your website.

Configure a Proxy Server

The Proxy Settings window opens if the internet proxy settings or user credentials are required to access a website.

Your server must have access to the Internet before you can connect to websites. If you connect to the Internet through a proxy server, enter the proxy server information so you can load websites. If you do not know the URL of your proxy server, contact your Network Administrator.

Enter the applicable information in the Proxy Settings dialog, and click Submit. This information is saved the first time you enter it and remains with the session until the proxy server session is terminated.

The proxy server configuration is saved.

Manage Files in the Knowledge Center

You can modify or delete documentation that you have added to the Knowledge Center repository.

Note: Access to the Knowledge Center configuration is restricted. The *Administration Guide* contains information about defining this access permission. To gain access, contact your Security Administrator.

To modify your previously uploaded PC files, change the files locally and upload. The files are indexed automatically.

To modify your previously uploaded mainframe files, change the files and reindex.

You can also delete the previously uploaded files when necessary. In the Upload Documents tab, select the file that you want to delete, click the remove icon, and confirm deletion. Clear the index (see page 83) so that the documentation that you previously indexed no longer appears in the search.

Clear Indexed Documentation

You can clear an index from the Knowledge Center repository. Clearing an index is necessary to verify that previously indexed documentation no longer appears in the search. This task can be necessary when the information is outdated or incorrect. Clearing an index is required when you remove an existing index and create another index.

Note: Access to the Knowledge Center configuration is restricted. The *Administration Guide* contains information about defining this access permission. To gain access, contact your Security Administrator.

Follow these steps:

1. Click the wrench icon in the Knowledge Center window.

   The Knowledge Center Settings dialog opens.

2. Remove all folders from the Folders to Index pane and click Index.

   A dialog appears to confirm the deletion of the existing indexes.

3. Click Yes.

   Existing user indexes are removed from the Knowledge Center repository.

   Note: You cannot remove certain types of indexed documentation, such as CA or International DB2 Users Group (IDUG) documentation.

View the Index Log

The index log displays the most recent results of the documentation indexing process. Indexing is the process of marking the content of files with keywords and adding the files to the Knowledge Center repository. View the Index Log to complete the following tasks:

- Determine which files were indexed and added to the repository.
- Determine which files were skipped because the file type is not supported or the file was already indexed.
- Determine if the indexing process is complete.
- Determine the number of empty folders in the Folders to Index pane.

Follow these steps:

1. Click the wrench icon in the Knowledge Center browser.

   The Knowledge Center Settings dialog opens.

2. Click View Log on the Index Documents tab.

   The Index Log window opens.

   Note: If you detect a problem, index your selected folders again.