

CA Workload Automation DE

Monitor Perspective Help

r11.3



CA Technologies Product References

This document references the following CA Technologies products:

- CA Workload Automation DE
- CA Workload Automation Desktop Client (CA WA Desktop Client)
- CA Workload Automation DE Web Client
- CA Workload Automation High Availability DE (CA WA High Availability)
- CA Workload Automation Web Services (CA WA Web Services)
- CA Workload Automation Agent for UNIX (CA WA Agent for UNIX)
- CA Workload Automation Agent for Linux (CA WA Agent for Linux)
- CA Workload Automation Agent for Windows (CA WA Agent for Windows)
- CA Workload Automation Agent for i5/OS (CA WA Agent for i5/OS)
- CA Workload Automation Agent for z/OS (CA WA Agent for z/OS)
- CA Workload Automation Agent for Application Services (CA WA Agent for Application Services)
- CA Workload Automation Agent for Web Services (CA WA Agent for Web Services)
- CA Workload Automation Agent for Micro Focus (CA WA Agent for Micro Focus)
- CA Workload Automation Agent for Databases (CA WA Agent for Databases)
- CA Workload Automation Agent for SAP (CA WA Agent for SAP)
- CA Workload Automation Agent for PeopleSoft (CA WA Agent for PeopleSoft)
- CA Workload Automation Agent for Oracle E-Business Suite (CA WA Agent for Oracle E-Business Suite)
- CA Workload Automation Restart Option EE (CA WA Restart Option)
- CA Spectrum® Service Assurance (CA Spectrum SA)

This Documentation, which includes embedded help systems and electronically distributed materials, (hereinafter referred to as the "Documentation") is for your informational purposes only and is subject to change or withdrawal by CA at any time.

This Documentation may not be copied, transferred, reproduced, disclosed, modified or duplicated, in whole or in part, without the prior written consent of CA. This Documentation is confidential and proprietary information of CA and may not be disclosed by you or used for any purpose other than as may be permitted in (i) a separate agreement between you and CA governing your use of the CA software to which the Documentation relates; or (ii) a separate confidentiality agreement between you and CA.

Notwithstanding the foregoing, if you are a licensed user of the software product(s) addressed in the Documentation, you may print or otherwise make available a reasonable number of copies of the Documentation for internal use by you and your employees in connection with that software, provided that all CA copyright notices and legends are affixed to each reproduced copy.

The right to print or otherwise make available copies of the Documentation is limited to the period during which the applicable license for such software remains in full force and effect. Should the license terminate for any reason, it is your responsibility to certify in writing to CA that all copies and partial copies of the Documentation have been returned to CA or destroyed.

TO THE EXTENT PERMITTED BY APPLICABLE LAW, CA PROVIDES THIS DOCUMENTATION "AS IS" WITHOUT WARRANTY OF ANY KIND, INCLUDING WITHOUT LIMITATION, ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NONINFRINGEMENT. IN NO EVENT WILL CA BE LIABLE TO YOU OR ANY THIRD PARTY FOR ANY LOSS OR DAMAGE, DIRECT OR INDIRECT, FROM THE USE OF THIS DOCUMENTATION, INCLUDING WITHOUT LIMITATION, LOST PROFITS, LOST INVESTMENT, BUSINESS INTERRUPTION, GOODWILL, OR LOST DATA, EVEN IF CA IS EXPRESSLY ADVISED IN ADVANCE OF THE POSSIBILITY OF SUCH LOSS OR DAMAGE.

The use of any software product referenced in the Documentation is governed by the applicable license agreement and such license agreement is not modified in any way by the terms of this notice.

The manufacturer of this Documentation is CA.

Provided with "Restricted Rights." Use, duplication or disclosure by the United States Government is subject to the restrictions set forth in FAR Sections 12.212, 52.227-14, and 52.227-19(c)(1) - (2) and DFARS Section 252.227-7014(b)(3), as applicable, or their successors.

Copyright © 2011 CA. All rights reserved. All trademarks, trade names, service marks, and logos referenced herein belong to their respective companies.

CA Technologies Product References

This document references the following CA Technologies products:

- CA Workload Automation DE
- CA Workload Automation Desktop Client (CA WA Desktop Client)
- CA Workload Automation DE Web Client
- CA Workload Automation High Availability DE (CA WA High Availability)
- CA Workload Automation Web Services (CA WA Web Services)
- CA Workload Automation Agent for UNIX (CA WA Agent for UNIX)
- CA Workload Automation Agent for Linux (CA WA Agent for Linux)
- CA Workload Automation Agent for Windows (CA WA Agent for Windows)
- CA Workload Automation Agent for i5/OS (CA WA Agent for i5/OS)
- CA Workload Automation Agent for z/OS (CA WA Agent for z/OS)
- CA Workload Automation Agent for Application Services (CA WA Agent for Application Services)
- CA Workload Automation Agent for Web Services (CA WA Agent for Web Services)
- CA Workload Automation Agent for Micro Focus (CA WA Agent for Micro Focus)
- CA Workload Automation Agent for Databases (CA WA Agent for Databases)
- CA Workload Automation Agent for SAP (CA WA Agent for SAP)
- CA Workload Automation Agent for PeopleSoft (CA WA Agent for PeopleSoft)
- CA Workload Automation Agent for Oracle E-Business Suite (CA WA Agent for Oracle E-Business Suite)
- CA Workload Automation Restart Option EE (CA WA Restart Option)
- CA Spectrum® Service Assurance (CA Spectrum SA)

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>.

Contents

Chapter 1: Subscribing to Workload 11

Application Generations.....	11
Create a Subscription	12
Subscribe to Workload	13
View an Application's Jobs	14
Display the Critical Path	15
Display a Graphical Overview of an Application	15
Set the Default Zoom Settings for Application Graphs	16
Hide Completed Applications Automatically.....	17

Chapter 2: Monitoring and Controlling Applications 19

Active Applications	20
Application States	20
Application Completion	21
Application Hold and Release	22
Application Unwait.....	22
Application Quiesce and Unquiesce.....	23
Monitor the Status of a Running Application.....	24
Issue a Command Against Application Generations	25
View the Gantt Chart of an Application	26
Insert a Job into an Active Application.....	27
Find the Troubled Jobs within an Application	29
Complete an Application.....	29
Display Application Details.....	30
Display the Command Log of an Application Generation	30
Hold an Application	32
Release a Held Application.....	33
Refresh an Application's Anticipated End Times.....	34
View the Comments of an Active Application.....	35
Unwait an Application.....	35
Quiesce a Server.....	36
Unquiesce Specific Events or Application Generations.....	37
Unquiesce an Application Generation	38
Trigger a New Generation of the Application	39
Hide Completed Applications Manually	41

Chapter 3: Monitoring and Controlling SubApplications 43

SubApplications in Active Applications	43
SubApplication Bypass and Unbypass.....	44
SubApplication Request and Unrequest	44
SubApplication Ready	45
SubApplication Unwait.....	45
Display SubApplications in an Active Application	46
Complete a SubApplication	46
Hold a SubApplication	47
Release a Held SubApplication.....	48
Bypass a SubApplication.....	49
Unbypass a SubApplication	50
Ready a SubApplication for Submission	51
Request a SubApplication	52
Unrequest a SubApplication.....	53
Unwait a SubApplication	54

Chapter 4: Monitoring and Controlling Jobs 57

Jobs in Active Applications	58
Job States	59
Job Commands Usage	63
Job's Spool File	63
Job Definition Reset	64
Job Bypass and Unbypass.....	65
Job Cancellation	65
Job Completion	66
Job Hold and Release	67
Job Request and Unrequest	67
Job Resubmission	67
Job Unwait	68
Predecessor Dependencies	68
Time Dependencies.....	70
Global Variable Dependencies	72
How Jobs with Resource Dependencies are Submitted.....	73
Issue a Command Against Jobs	73
Display Job Details.....	74
Display a Job's Anticipated End Time	74
Display the Command Log of a Job	75
Display a Job's Spool File	76
Reset a Job Definition.....	79
Bypass a Job	79

Unbypass a Job.....	80
Cancel a Running Job.....	81
Complete a Job.....	82
Hold a Job.....	83
Release a Held Job.....	83
Request a Job	84
Unrequest a Job	85
Resubmit a Job	86
Ready a Job for Submission.....	88
View the Comments of a Job.....	88
Unwait a Job.....	89
Set a Job's User Status.....	90
Drop Predecessor Dependencies	90
Reset Time Dependencies.....	91
Drop Variable Dependencies.....	93
Evaluate Variable Dependencies.....	94
Drop Resource Dependencies	95
Display Resource Usage	96
Verify Whether a Job is Running or Has Executed	97
Set a Windows Job Object's Properties.....	98
Set Process Priority for a Windows or a UNIX Job.....	101
Send a Signal to a UNIX Job	102
Retrieve the Spool File for an i5/OS Job.....	104
Respond to Suspended Jobs that Require Manual Intervention.....	106
Retrieve the JES Spool File of a Micro Focus Job.....	107
Browse JCL.....	109
Browse Copy JCL.....	109
Edit JCL	110
Edit Copy JCL	110
z/OS Jobs Restart Using CA WA Restart Option	111
Select the Steps to Restart	112
Let CA WA Restart Option Select the Steps to Restart	113
Specify Additional Restart Options	114
Specify CA WA Restart Option Statements	115
Issue a Database Command	116
SAP Commands	117
Issue an SAP Command Against an SAP-R3 or SAP-Job Copy Job	117
Issue an SAP Command Against an SAP-Batch Input Session Job.....	125
Issue an SAP Command Against an SAP-BW Info Package Job	127
Issue an SAP Command Against an SAP-BW Process Chain Job.....	129
Issue an SAP Command Against an SAP-Data Archiving Job	130
Issue an SAP Command Against an SAP-Process Monitor Job	130

Issue a PeopleSoft Command.....	131
Issue an Oracle E-Business Suite Command.....	132

Chapter 5: Working with Custom Views **135**

Custom Views.....	135
Custom View Filters	136
Open a Prepackaged Custom View	138
Create a New Text-Based Custom View	139
Copy an Existing Custom View	140
Create a Custom View Filter	141

Chapter 6: Working with Dashboard Views **145**

Dashboard Views.....	146
Open an Existing Status Message View	147
Create a New Status Message View	148
Copy an Existing Status Message View	151
Open an Existing System Monitor View	151
Create a New System Monitor View	152
Copy an Existing System Monitor View	154
Edit a Dashboard View	154
Delete a Dashboard View	155

Chapter 7: Resetting Resource Availability Counts **157**

Resource Availability Counts	157
Reset Resource Availability Count.....	158

Index **159**

Chapter 1: Subscribing to Workload

This section contains the following topics:

[Application Generations](#) (see page 11)

[Create a Subscription](#) (see page 12)

[Subscribe to Workload](#) (see page 13)

[View an Application's Jobs](#) (see page 14)

[Hide Completed Applications Automatically](#) (see page 17)

Application Generations

Each time an Application runs, the server creates an instance of that Application named generation and assigns it a number. The generation number increments with each Application generation. The generation number lets the server track each Application generation and identify the Application generation a particular job belongs to.

An Application generation may not complete within a 24-hour period. It can span many days or weeks. An Application may run many times in a day or have the same job running across multiple Applications. Multiple generations of the same Application may run at the same time.

The Application Monitor view organizes and displays the Applications and their generations in the Monitor perspective. A folder labeled with the Application name and the number of generations represents each Application. The generations appear under the Application folder. Each generation contains the Application name, the Application generation number, the Application state, and the number of jobs in the Application. The color of the generation represents the state of the generation.

More information:

[Trigger a New Generation of the Application](#) (see page 39)

Create a Subscription

You can create a subscription to display only the Applications and generations that you are interested in.

To create a subscription

1. Open the Monitor perspective.
2. Right-click the active server connection that contains the workload you want to view in the Application Monitor view.

A pop-up menu appears.

3. Select Subscribe with Filter.

The Subscription dialog opens.

4. Click New.

If you have multiple connections defined, the Select Servers page of the Subscription Definition dialog opens. Otherwise, you are taken directly to the Filter Definition page and you can skip the next step.

5. Select the servers that apply to this subscription and click Next, if applicable.

The Filter Definition pages opens.

6. Define a name for your subscription in the Filter Name field.
7. Choose from the following subscription options:

all Applications

Subscribes to all workload including completed Applications.

all active Applications

Subscribes to only Applications that are not complete.

Applications with names

Subscribes to the Applications that match the criteria in this field.

Note: You can enter a specific Application name or, to specify multiple Applications with similar names, enter a partial name using an asterisk (*) as a wildcard for any number of characters or a question mark (?) as a wildcard for a single character. This field is not case sensitive.

Last # of generations

Subscribes to the last specified number of generations of each Application that meets the criteria you specified above.

Default: All generations

8. (Optional) Select the Subscribe Automatically check box to subscribe automatically with the filter each time you connect to the applicable servers and open the Monitor perspective.

9. Click Save Subscription.

The new subscription is created.

Note: You can also create subscriptions when you are offline. From the main menu, click Windows, Preferences, Desktop Client, Monitor Perspective, Automatic Subscription, and then click New. From the Automatic Subscription page, you can also edit and remove subscriptions, and enable automatic subscription.

Subscribe to Workload

To view active workload, you must subscribe to receive the data from the server. You can subscribe to all of the workload data, subscribe to only active Applications, or subscribe with a filter that you create to limit the Applications in the display.

Note: You can only subscribe to one subscription filter for a server at a time. Each time you subscribe to a server, the server unsubscribes from the current subscription, if there is one, and subscribes to the new subscription.

To subscribe to workload

1. Open the Monitor perspective.
2. Right-click the active server connection that contains the workload you want to view in the Application Monitor view.

A pop-up menu appears.

3. Select one of the following subscription options:

Subscribe with Filter

Subscribes to workload using a defined subscription.

Subscribe All

Subscribes to all workload including completed Applications. The server creates a subscription named *server-alias_All* if one does not exist. You can enable the Subscribe Automatically option for this subscription to subscribe automatically to all workload on the server.

Subscribe Active

Subscribes to only Applications that are not complete. The server creates a subscription named *server-alias_Active* if one does not exist. You can enable the Subscribe Automatically option for this subscription to subscribe automatically to the active workload on the server.

If you subscribed with the Subscribe with Filter option, the Subscription dialog opens. Otherwise, the server subscribes to the Applications that match your subscription option and you can skip the next step.

4. Select the table row of the subscription you want to subscribe to and click Subscribe, if applicable.

The server subscribes to the Applications that match the subscription you subscribed to.

5. Double-click an Application folder.

The generations of the Application are displayed.

View an Application's Jobs

You can view the jobs that belong to an Application that you are monitoring.

To view an Application's jobs

1. Open the Monitor perspective.
2. Subscribe to workload on the server.

The Applications and their generations are displayed in the Application Monitor view.

3. Double-click the Application generation you want to view.

The jobs in the Application generation appear in the graphical view. Monitor displays the job state below each job and highlights the job by its state color.

Note: You can use the tabs at the top of the graphical view to navigate between the different views you have open and to close the individual views. You can use the two buttons in the top right corner to close multiple Application generations at a time.

Display the Critical Path

You can display an Application's critical path in the Monitor perspective. The server determines the critical path based on the longest path to the job that finishes last in an Application or another job designated as critical.

To display the critical path

1. Open the Monitor perspective.
2. Double-click the Application generation in the Application Monitor view.
The Application's graphical view appears.
3. Click Critical Path at the bottom of the workspace.
Monitor highlights the critical path in the Application and displays the lines between jobs on the critical path in green.

Display a Graphical Overview of an Application

You can display an overview of the graphical view and use it to zoom in and navigate an Application. This feature is a useful navigation tool for viewing large Applications.

To display a graphical overview

1. Open the Monitor perspective.
2. Double-click the Application generation in the Application Monitor view.
The Application's graphical view appears.
3. Click Overview above the workspace.
The graphical overview appears as a secondary window.
4. Move and resize the box within the secondary window to zoom in on a portion of the Application.
The Application view changes as indicated.

Set the Default Zoom Settings for Application Graphs

You can set the default zoom settings for Application graphs. When you open a graphical view of the Application, the zoom defaults are used. You can override the zoom defaults for a specific Application graph by zooming in and zooming out or scrolling through the graph by using the mouse wheel or the cursor.

To set the default zoom settings for Application graphs

1. Select Window, Preferences from the main menu.

The Preferences dialog opens.

2. Click Desktop Client, Monitor Perspective, Application Graphical Display in the left pane.

The Application Graphical Display page opens in the right pane.

3. Do *one* of the following in the Graph Fitting section:

- To fit the graph to the window size, select the Fit to Window check box.
- To specify a specific zoom factor, clear the Fit to Window check box, and select a value in the Use zoom factor field.

4. Click OK.

The zoom defaults are set for the Application graphs.

Note: Changes to the zoom defaults do not affect graphs that are open. You must close and reopen the graphs for the changes to apply.

Hide Completed Applications Automatically

You can automatically hide completed Applications from view to restrict the number of Applications that appear in the Monitor views.

Note: The administrator may schedule a job to purge completed Applications periodically.

To hide completed Applications automatically

1. Select Window, Preferences from the main menu.

The Preferences dialog opens.

2. Click Desktop Client, Monitor Perspective, Automatic Hide in the left pane.

The Automatic Hide page opens in the right pane.

3. Complete the following fields as appropriate:

Every

Specifies the frequency in minutes after which completed Applications are removed from view.

Default: 120 minutes (Completed Applications are removed 120 minutes after they complete.)

At

Sets the automatic hide feature to a specific hour each day.

Default: midnight (Completed Applications are removed each day at midnight.)

Rounded at

Specifies the time interval frequency at which completed Applications are removed from view, starting with the next occurrence of the selected unit of time. For example, if you select every 2 hours and the current time is 13:37:00, the completed Applications are automatically hidden at 14:00:00, 16:00:00, and so on.

Default: 30 minutes (Completed Applications are removed every 30 minutes.)

4. Click OK.

The automatic hide feature is set.

Note: Hiding Applications does not permanently delete them from the server. To restore Applications, you can right-click the server connection and resubscribe using the same filter you used before.

More information:

[Hide Completed Applications Manually](#) (see page 41)

Chapter 2: Monitoring and Controlling Applications

This section contains the following topics:

- [Active Applications](#) (see page 20)
- [Monitor the Status of a Running Application](#) (see page 24)
- [Issue a Command Against Application Generations](#) (see page 25)
- [View the Gantt Chart of an Application](#) (see page 26)
- [Insert a Job into an Active Application](#) (see page 27)
- [Find the Troubled Jobs within an Application](#) (see page 29)
- [Complete an Application](#) (see page 29)
- [Display Application Details](#) (see page 30)
- [Hold an Application](#) (see page 32)
- [Release a Held Application](#) (see page 33)
- [Refresh an Application's Anticipated End Times](#) (see page 34)
- [View the Comments of an Active Application](#) (see page 35)
- [Unwait an Application](#) (see page 35)
- [Quiesce a Server](#) (see page 36)
- [Unquiesce Specific Events or Application Generations](#) (see page 37)
- [Unquiesce an Application Generation](#) (see page 38)
- [Trigger a New Generation of the Application](#) (see page 39)
- [Hide Completed Applications Manually](#) (see page 41)

Active Applications

After an Event triggers, the Application becomes active. You can monitor an active Application's state, view its details and jobs, and issue commands against the Application to do the following:

- Complete the Application
- Hold or release the Application
- Refresh the estimated execution time of the Application
- Show the Application's comments
- Unwait the Application
- Trigger a new generation of the Application
- Insert jobs into the active Application
- View Gantt chart
- Find troubled job
- Hide selected Application(s)

Note: To run commands against Applications, you require the appropriate security permissions. For more information, contact your administrator.

Application States

When the server manages an active Application, the Application passes through multiple processing stages named states. The Application state is determined by the job in the Application with the most severe job state. For example, if an Application contains a job with a FAILED state, the server assigns the Trouble state to the Application. Some Application states such as Trouble require an action.

If you are using a third-party solution for CA Workload Automation DE, you can query the server for the Application state, expected end time, and percentage of completion.

Note: For more information about using web services to return the status of Application generations, see the *Programming Guide*.

The following table displays the states that Applications can fall into and the associated job states:

Application State	Associated Job States
Unknown (most severe)	UNKNOWN

Application State	Associated Job States
Trouble	AGENTDOWN, EXTSCHDOWN, FAILED, INACTIVE, Overdue condition, SUBERROR
Manual Intervention	MANHOLD, MANSUB, TASKWAIT
Applquiesce	APPLQUIESCE
Applhold	APPLHOLD
Applwait	APPLWAIT
Waiting	EXTWAIT, DEFINED, JANCWAIT, MONITOR, PREDWAIT, RESWAIT, SUBAPPLWAIT, SUBDELAY, SUBMIT, VARWAIT, WAITING
Processing	EXEC, MONITOR, READIED, READY
Complete (least severe)	BYPASSED, COMPLETE

Note: By default, CA WA Desktop Client displays Application generations in the Trouble state in red.

Application Completion

When you complete an Application, the server completes all jobs in the Application. The server deems completed jobs to have executed even if they have not been submitted, have failed, or are still running.

Note: You can monitor for forced completions by specifying a notification in the Application or job definition. The notification can send an email, trigger an Event, run a JavaScript script, or send an SNMP trap. For more information about setting up notifications, see the *Define Perspective Help*.

Similarly, when you complete a subApplication, the server completes all jobs in the subApplication.

Note: Completing a subApplication is different than bypassing a subApplication. If you bypass a subApplication, the server does not bypass the jobs in the subApplication until their dependencies are met. If you mistakenly bypass a subApplication, you can unbypass it before the jobs in the subApplication are eligible to run. If you complete a subApplication, it is completed immediately and cannot be undone.

More information:

[Complete an Application](#) (see page 29)

[Complete a SubApplication](#) (see page 46)

Application Hold and Release

When you place an Application on hold, the server holds all the jobs in the Application until you manually release the Application. You can hold an Application before its jobs start to run or while some of the jobs are running. Holding the Application does not affect jobs currently running, but does prevent the server from submitting new jobs. When you place an Application on hold, its state changes to APPLHOLD. Releasing the Application removes the APPLHOLD state and releases the remaining jobs once their dependencies are met.

If an Application is on hold, the server does not consider it complete even if all of its jobs have completed.

Similarly, when you place a subApplication on hold, the server holds all jobs in the subApplication until you release the subApplication.

Note: If your Application is in APPLWAIT state, unwait the Application to release the jobs.

More information:

[Hold an Application](#) (see page 32)

[Release a Held Application](#) (see page 33)

[Hold a SubApplication](#) (see page 47)

[Release a Held SubApplication](#) (see page 48)

Application Unwait

In an Application, you can specify that the Application waits until the jobs in the previous generation complete before running. The server marks an Application that waits for a previous generation in an APPLWAIT state. You can remove the APPLWAIT state by unwaiting the Application. When you unwait an Application, the selected Application no longer waits for its previous generation to complete before it runs.

Note: If your Application is in APPLHOLD state, you should release the Application to release the jobs.

More information:

[Unwait an Application](#) (see page 35)

Application Quiesce and Unquiesce

Quiescing the server quiesces all active Applications. Jobs that are in the EXEC state continue to execute, but no new jobs are submitted. All the Applications in the server are quiesced and the state of each Application changes to APPLQUIESCE.

You can unquiesce the server to unquiesce all the Applications at once, or you can unquiesce specific Applications on the server. After an active Application generation is unquiesced, it continues to run.

Note: For more information about quiescing the server, see the *Admin Perspective Help*.

More information:

[Quiesce a Server](#) (see page 36)

[Unquiesce Specific Events or Application Generations](#) (see page 37)

[Unquiesce an Application Generation](#) (see page 38)

Monitor the Status of a Running Application

You can view the following details for an Application to monitor its progress:

- Number of jobs
- State
- Start time
- Estimated end time
- Elapsed time
- Completion percentage

To monitor the status of a running Application

1. Open the Monitor perspective.
2. Subscribe to workload on the server.

The Applications and their generations are displayed in the Application Monitor view.

3. Right-click a running Application generation, and select Details from the drop-down menu

The Details dialog opens, displaying details such as Application state, start time, estimated end time, and completion percentage.

4. (Optional) Double-click the Application generation in the Application Monitor view to see its graphical view.

At the top of the graphical view, a progress bar shows the Application's start time, estimated end time, and completion percentage. You can also see the number of completed jobs and elapsed time.

Notes:

- The completion percentage is calculated based on the number of completed jobs or elapsed time depending on the option selected.
- By default, the completed jobs option button is selected. The completion percentage is calculated as $(\text{completed jobs} + \text{bypassed jobs}) / \text{total number of jobs}$.
- If you select the elapsed time option button, the completion percentage is calculated as $(\text{currentTime} - \text{startTime}) / (\text{anticipatedEndTime} - \text{startTime})$. The elapsed time option is enabled only if you enabled anticipated (estimated) end times in the Application. For more information about enabling anticipated end times in an Application, see the *Define Perspective Help*.
- When the Application is complete, estimated end time changes to end time.

- The information on the progress bar is updated in real time. You can turn off the progress bar by clicking the  icon.

Issue a Command Against Application Generations

You can issue a command against a single Application generation or multiple Application generations. For example, you can issue a command to complete, hold, release, or unwait Application generations.

Note: You cannot issue commands against Applications controlled on different servers.

To issue a command against Application generations

1. Open the Monitor perspective.
2. Subscribe to workload on the server.
3. Expand the generations of the Applications that you want to issue a command against in the Application Monitor view.
4. Select the Application generations that you want to issue a commands against.
Note: You can press the Ctrl key while making your selections to select multiple Application generations at once.
5. Right-click one of the selected Application generations, and select a command from the drop-down menu.

The server issues the command against the selected Application generations.

View the Gantt Chart of an Application

You can view the Gantt chart of an Application to monitor the progress of the Application as it runs relative to the current time. From the Gantt chart, you can see when the jobs in the Application will start and end, and how long the Application will run, based on historical run times (last 10 runs of the job) or a duration specified in the job definition. You can also view a Gantt chart of the jobs on the critical path in the Application.

Note: To view jobs in a Gantt chart, the jobs in the Application must have a start time and either an end time (completed jobs) or an anticipated end time (incomplete jobs). If the job has never run before and no duration is specified in the job definition, the server cannot calculate the job's end time.

Important! When the Application is defined, the Estimate end time option must be selected in the Application properties so the server can calculate start and end times. For more information about enabling anticipated end times in an Application, see the *Define Perspective Help*.

To view the Gantt chart of an Application

1. Open the Monitor perspective.
2. Subscribe to workload on the server.

The Applications and their generations are displayed in the Application Monitor view.

3. Right-click the Application generation and select View Gantt Chart from the drop-down menu.

The Gantt chart of the Application opens in a graphical view. On the left side of the view, a table displays the start time, end time or anticipated end time, duration, and completion percentage of each job in the Application. On the right side of the view, bars representing the Application and each of the jobs in the Application show the progress of the Application and its jobs relative to the current time. Arrows between the bars indicate predecessor dependencies. The color of the bars matches the color of the Application and job states.

Note: The Gantt chart is refreshed every 5 seconds by default or as the state of one of the jobs changes.

4. (Optional) Click Critical Path at the bottom of the view.

The Gantt chart displays the jobs that are on the critical path in the Application.

5. (Optional) Enter a filter in the Enter the Name filter field to limit the jobs displayed.

Note: You can use the asterisk (*) as a wildcard for zero or more characters and the question mark (?) as a wildcard for a single character. For example, to limit the Gantt chart to jobs that start with linux, enter **linux***. The filter is not case sensitive.

The Gantt chart displays the jobs that match your filter.

6. (Optional) Do the following to issue a command against the Application or one of its jobs:
 - a. Right-click the Application or job from the table or bar chart.
A list of commands appears.
 - b. Select the command from the pop-up menu.
The server issues the command against the Application or job.

Insert a Job into an Active Application

You can insert a job into an active Application. The inserted job runs immediately unless you define predecessor, time, variable, or resource dependencies, or insert the job on hold or on request.

Note: You cannot insert a job into an Application that is complete.

To insert a job into an active Application

1. Open the Monitor perspective.
2. Subscribe to workload on the server.
The Applications and their generations are displayed in the Application Monitor view.
3. Right-click the Application generation you want to insert the job into, select Insert Job from the first drop-down menu, and select the workload object you want to insert from the second drop-down menu.
The Basic page of the job definition opens.
4. Complete all required fields indicated by asterisks.
Note: When you insert a job into an active Application, you cannot use symbolic variables for the job name and qualifier. You can use symbolic variables in fields that get resolved at run time, such as the command or script name, if and only if a JavaScript script runs at run time to resolve those symbolic variables.
5. Select the Dependencies option from the left pane.
The Dependencies page opens.

6. Select and add the job's predecessor and successor dependencies by doing the following:
 - To add predecessor dependencies, select the predecessor jobs in the Current jobs list box, and click Add under the Predecessors list box.
 - To add successor dependencies, select the successor jobs in the Current jobs list box, and click Add under the Successors list box.

Note: You can press the Ctrl key while making your selections to select multiple jobs at once.

7. Complete the remaining pages for the job as appropriate.

Notes:

- For more information about the fields in the remaining pages, see the *Define Perspective Help*.
 - If you want to insert the job on hold, select On hold in the General page. The server holds the job until you release it.
 - If you want to insert the job on request, select Only on request in the General page. The server does not submit the job unless it is requested.
 - If you want the inserted job to wait for a job with the same name to complete in a previous Application generation, complete the Wait for previous generation section in the General page.
 - You can also provide the reason for inserting the job in the Reason field in the General page.
8. (Optional) Do the following if you want the job to wait for the previous generations:
 - a. Click General in the left pane.
The General page opens in the right pane.
 - b. Do *one* of the following:
 - a. Select the The last generation option in the Wait if the job with the same name is in field if you want the job to wait only for the job in the previous generation of the Application to complete.
 - b. Select the Any previous generation option in the Wait if the job with the same name is in field if you want the job to wait for the job in all the previous generations of the Application to complete.
 9. Click OK.

The inserted job appears in the Application's graphical view.

Find the Troubled Jobs within an Application

When an Application appears in a troubled state, you can locate the trouble jobs in the Application.

To find the troubled jobs within an Application

1. Open the Monitor perspective.
2. Subscribe to workload on the server.

The Applications and their generations are displayed in the Application Monitor view.

3. Right-click the Application generation, and select Find Troubled Job from the drop-down menu.

Note: You can press Ctrl+T to find more troubled jobs.

The troubled job is highlighted in the graphical view.

Complete an Application

You can complete an Application to complete all the jobs in that Application.

To complete an application

1. Open the Monitor perspective.
2. Subscribe to workload on the server.

The Applications and their generations are displayed in the Application Monitor view.

3. Right-click the Application generation you want to complete, and select Complete from the drop-down menu.

A confirmation dialog opens.

4. (Optional) Enter the command reason in the Reason field.

Limit: 1024 characters (including the timestamp and the user name)

Note: We recommend that the reason text not exceed 250 characters. The number of commands displayed in the command log is limited to the size of the command log buffer, which is 1024 characters. For example, if the reason text averages 10 characters per command, the command log displays up to 50 commands; if the reason text averages 250 characters per command, the command log displays up to 4 commands.

5. Click OK.

The server completes the Application. The state of each job in the Application generation and the Application generation change to COMPLETE.

More information:

[Application Completion](#) (see page 21)

Display Application Details

You can display the details of an Application that you are monitoring. For example, you can display the name of the Event that triggered the Application.

To display Application details

1. Open the Monitor perspective.
2. Subscribe to workload on the server.

The Applications and their generations are displayed in the Application Monitor view.

3. Right-click the Application generation and select Details from the drop-down menu.

The details of the Application are displayed.

Display the Command Log of an Application Generation

You can view the command log of an Application generation. The command log displays the most recent commands that were issued against the Application generation.

Note: The command log includes commands that were invoked programmatically by a JavaScript script using the `execCommand` function. For more information about the `execCommand` function, see the *Programming Guide*.

To display the command log of an Application generation

1. Open the Monitor perspective.
2. Subscribe to workload on the server.

The Applications and their generations are displayed in the Application Monitor view.

3. Right-click the Application generation and select Details from the drop-down menu.

The details of the Application generation are displayed.

4. Click View Application Command Log.

The Application Commands Log dialog opens. The Commands section displays a table with the following columns:

Date and Time

Displays the date and time the command was issued against the Application generation.

Command

Displays the name of the command issued against the Application generation.

User

Displays the name of the user that issued the command.

Programmatic

Specifies whether the command was invoked programmatically by a JavaScript script.

Reason

Displays the command reason, if any

Note: To display the full command reason in the Reason section, select the row in the table.

Notes:

- The number of commands displayed in the command log is limited to the size of the command log buffer, which is 1028 characters. For example, if the reason text averages 10 characters per command, the command log displays up to 50 commands; if the reason text averages 250 characters per command, the command log displays up to 4 commands.
- You can click the table headings to change the column sorting order.

5. (Optional) Click Export as Text .

The command log is saved as a text file.

6. (Optional) Click Export as CSV.

The command log is saved as a .csv file. You can open the .csv file in Microsoft Excel.

Hold an Application

You can hold an Application to hold all of the jobs in the Application. The server does not submit any new jobs in the Application.

To hold an Application

1. Open the Monitor perspective.
2. Subscribe to workload on the server.

The Applications and their generations are displayed in the Application Monitor view.

3. Right-click the Application generation you want to place on hold and select Hold from the drop-down menu.

A confirmation dialog opens.

4. (Optional) Enter the command reason in the Reason field.

Limit: 1024 characters (including the timestamp and the user name)

Note: We recommend that the reason text not exceed 250 characters. The number of commands displayed in the command log is limited to the size of the command log buffer, which is 1024 characters. For example, if the reason text averages 10 characters per command, the command log displays up to 50 commands; if the reason text averages 250 characters per command, the command log displays up to 4 commands.

5. Click OK.

The server holds the Application until you manually release it. The state of the Application changes to APPLHOLD.

More information:

[Application Hold and Release](#) (see page 22)

Release a Held Application

You can release a held Application to release all of the jobs in the Application. The server submits the jobs in the Application once their dependencies are met.

To release a held Application

1. Open the Monitor perspective.
2. Subscribe to workload on the server.

The Applications and their generations are displayed in the Application Monitor view.

3. Right-click the Application generation you want to release from hold, and select Release from the drop-down menu.

A confirmation dialog opens.

4. (Optional) Enter the command reason in the Reason field.

Limit: 1024 characters (including the timestamp and the user name)

Note: We recommend that the reason text not exceed 250 characters. The number of commands displayed in the command log is limited to the size of the command log buffer, which is 1024 characters. For example, if the reason text averages 10 characters per command, the command log displays up to 50 commands; if the reason text averages 250 characters per command, the command log displays up to 4 commands.

5. Click OK.

The server removes the APPLHOLD state and releases the Application. Each of the jobs in the Application run when all of its dependencies, if any, are met.

More information:

[Application Hold and Release](#) (see page 22)

Refresh an Application's Anticipated End Times

The server recalculates anticipated end times every time a job state changes in the Application. If you want to obtain a job's most current anticipated end time, you can refresh the Application's anticipated end times. For example, you may want to refresh an Application's anticipated end times if a job has exceeded its anticipated end time and the states of other jobs have not changed.

To refresh an Application's anticipated end times

1. Open the Monitor perspective.
2. Subscribe to workload on the server.

The Applications and their generations are displayed in the Application Monitor view.

3. Right-click the Application generation you want to refresh the anticipated end times of, and select Refresh Estimated Runtime from the drop-down menu.

The Refresh Estimated Runtime dialog opens.

4. (Optional) Enter the command reason in the Reason field.

Limit: 1024 characters (including the timestamp and the user name)

Note: We recommend that the reason text not exceed 250 characters. The number of commands displayed in the command log is limited to the size of the command log buffer, which is 1024 characters. For example, if the reason text averages 10 characters per command, the command log displays up to 50 commands; if the reason text averages 250 characters per command, the command log displays up to 4 commands.

5. Click OK.

The server recalculates the Application's anticipated end times.

View the Comments of an Active Application

You can view the comments of an active Application to view additional information stored about the Application.

To view the comments on an active Application

1. Open the Monitor perspective.
2. Subscribe to workload on the server.

The Applications and their generations are displayed in the Application Monitor view.

3. Right-click the Application generation you want to view the comments of, and select Show Application Comments from the drop-down menu.

The comments of the active Application are displayed.

Unwait an Application

You can unwait an Application so that the Application no longer waits for its previous generation to complete before it runs.

To unwait an Application

1. Open the Monitor perspective.
2. Subscribe to workload on the server.

The Applications and their generations are displayed in the Application Monitor view.

3. Right-click the Application generation you want to unwait, and select Unwait from the drop-down menu.

A confirmation dialog opens.

4. (Optional) Enter the command reason in the Reason field.

Limit: 1024 characters (including the timestamp and the user name)

Note: We recommend that the reason text not exceed 250 characters. The number of commands displayed in the command log is limited to the size of the command log buffer, which is 1024 characters. For example, if the reason text averages 10 characters per command, the command log displays up to 50 commands; if the reason text averages 250 characters per command, the command log displays up to 4 commands.

5. Click OK.

The server removes the APPLWAIT state and submits the Application.

More information:

[Application Unwait](#) (see page 22)

Quiesce a Server

You can quiesce (pause) the server to prevent the server from temporarily submitting jobs or triggering workload while the server remains active. Jobs that are running when the server is quiesced continue to run to completion. After you quiesce a server, you can selectively unquiesce certain Application generations and Events, giving you more control over what runs.

To quiesce a server

1. Open the Monitor perspective.
2. Subscribe to workload on the server.

The Applications and their generations are displayed in the Application Monitor view.

3. Right-click the server you want to unquiesce, and select Quiesce from the drop-down menu.

A confirmation dialog appears.

4. Click OK.

The server is quiesced. All the quiesced Application generations appear in the APPLQUIESCE state and all the quiesced Events appear in the QUIESCED state.

More information:

[Application Quiesce and Unquiesce](#) (see page 23)

Unquiesce Specific Events or Application Generations

You can unquiesce specific Application generations and Events. After an Event is unquiesced, it can be triggered to run its Application. After an active Application generation is unquiesced, it continues to run.

Note: The missed executions of the quiesced Event are not triggered.

To unquiesce specific Events or Application generations

1. Right-click the server node you want to unquiesce, and select Unquiesce from the pop-up menu.

The Unquiesce dialog appears.

2. Select *one* of the following options in the Application generations section:

ALL

Unquiesces all Application generations.

Application generation name

Unquiesces the specified Application generations in the format *application.generation*, where *application* specifies the Application name and *generation* specifies the generation number.

Example: appl.1,appl.2

Note: You can also use a wildcard for a partial name. For example, APPL* unquiesces all generations of Applications with names that start with APPL. You can use commas to separate multiple Application generations.

None

Does not unquiesce any Application generations.

3. Select *one* of the following options in the Event section:

ALL

Unquiesces all Events.

Event name

Unquiesces the specified Events in the format *eventprefix.eventname*, where *eventprefix* specifies the Event prefix and *eventname* specifies the Event name.

Example: CYBER.VERIFY

Note: You can use a wildcard for a partial name. For example, CYB* unquiesces all Events with prefixes that start with CYB. You can use commas to separate multiple Events.

None

Does not unquiesce any Events.

The specified Application and Events are unquiesced.

More information:

[Application Quiesce and Unquiesce](#) (see page 23)

Unquiesce an Application Generation

You can unquiesce an Application generation so that the server continues to run the jobs in the quiesced Application when their dependencies are met.

To unquiesce an Application generation

1. Open the Monitor perspective.
2. Subscribe to workload on the server.

The Applications and their generations are displayed in the Application Monitor view. Quiesced Application generations appear in the APPLQUIESCE state.

3. Right-click the quiesced Application generation you want to unquiesce, and select Unquiesce from the drop-down menu.

Note: The Unquiesce command is only listed in the drop-down menu for Applications that are in the quiesced state.

A confirmation dialog appears.

4. Click OK.

The server removes the APPLQUIESCE state and unquiesces the Application generation. Each of the jobs in the Application run when all of its dependencies, if any, are met.

More information:

[Application Quiesce and Unquiesce](#) (see page 23)

Trigger a New Generation of the Application

You can trigger a new generation of the Application while monitoring the workload. The server triggers the Event that triggered the Application generation.

To trigger a new generation of the Application

1. Open the Monitor perspective.
2. Subscribe to workload on the server.

The Applications and their generations are displayed in the Application Monitor view.

3. Right-click the Application generation you want to trigger a new generation of, and select Trigger from the drop-down menu.

The Trigger Event dialog opens.

4. Complete the following fields as appropriate:

Schedule criteria

(Optional) Specifies the time and the date when the Event triggers.

Default: Trigger the Event immediately

Note: This field does not apply to monitor Events.

Root jobs to run

(Optional) Specifies the jobs that are to be submitted when the Event is triggered. You can use this field to run, or rerun, a subset of jobs within the Application. Separate the list of jobs to run with commas. For example, enter **A,B,C**. To select a job with all of its successors (descendants), append a plus sign (+). For example, enter **D+**.

Note: This field does not apply to Forecast and Report Events.

Add new scheduled Event/Replace next scheduled Event

Indicates whether to add a new scheduled Event execution to the Event's usual schedule or replace the Event's next scheduled or expected execution with a new time.

Default: Add new scheduled Event

Notes:

- This field does not apply to monitor Events.
- You can override the default value for this option. From the main menu, select Window, Preferences, Desktop Client, Services Perspective, Event.

Submit Application on hold

(Optional) Indicates whether to submit the Application on hold when the Event is triggered.

Default: Not on hold

Notes:

- This field does not apply to Forecast and Report Events.
- You can override the default value for this option. From the main menu select Window, Preferences, Desktop Client, Services Perspective, Event.

User Parameters to pass to Event

(Optional) Specifies the parameter fields to pass user information to the Application that the Event triggers.

Note: This field does not apply to Forecast and Report Events.

5. (Optional) Specify variables to pass to the Event in the variables section.

Note: This step applies to monitor Events only. When the Event triggers, the server substitutes the values of these variables in the Application triggered by the Event. For example, when you trigger a File Trigger Event, you can specify a file name, and the server substitutes this file name in the APPL._filename symbolic variable wherever it occurs in the Application.

6. Click OK.

The server triggers a new generation of the Application.

More information:

[Application Generations](#) (see page 11)

Hide Completed Applications Manually

You can hide selected completed Applications from view to restrict the number of Applications that appear in the Monitor views.

To hide completed Applications manually

1. Open the Monitor perspective.
2. Subscribe to workload on the server.

The Applications and their generations are displayed in the Application Monitor view.

3. Select the completed Application generations that you want to hide.

Note: You can press the Ctrl key while making your selections to select multiple Application generations at once.

4. Right-click one of the Application generations you want to hide, and select Hide Selected Application(s) from the drop-down menu.

The selected completed Application generations are removed from the Application Monitor view.

Chapter 3: Monitoring and Controlling SubApplications

This section contains the following topics:

[SubApplications in Active Applications](#) (see page 43)
[Display SubApplications in an Active Application](#) (see page 46)
[Complete a SubApplication](#) (see page 46)
[Hold a SubApplication](#) (see page 47)
[Release a Held SubApplication](#) (see page 48)
[Bypass a SubApplication](#) (see page 49)
[Unbypass a SubApplication](#) (see page 50)
[Ready a SubApplication for Submission](#) (see page 51)
[Request a SubApplication](#) (see page 52)
[Unrequest a SubApplication](#) (see page 53)
[Unwait a SubApplication](#) (see page 54)

SubApplications in Active Applications

An active Application may contain subApplications. You can issue commands against subApplications to do the following:

- Complete the subApplication
- Hold or release the subApplication
- Bypass or unbypass a subApplication
- Ready a subApplication for submission
- Request or unrequest a subApplication
- Unwait the subApplication

Note: To run commands against the jobs in subApplications, you require the appropriate security permissions. For more information, contact your administrator.

SubApplication Bypass and Unbypass

When you bypass a subApplication, the server bypasses all jobs in the subApplication, preventing the jobs in the subApplication from running. When the dependencies of a job in the subApplication are met, the server bypasses the job and releases its successor jobs.

You can unbypass the subApplication before the server bypasses the jobs in the subApplication. For example, if a subApplication is bypassed by mistake, you can undo the action by unbypassing the subApplication. When the dependencies of a job in the subApplication are met, the server runs the job as usual.

Note: You can also prevent the jobs in a subApplication from running by completing the subApplication. When you complete a subApplication, the server immediately releases all successor jobs of jobs in the subApplication, even if their predecessors are still running. Unlike bypassing a subApplication, you cannot undo completing a subApplication.

More information:

[Bypass a SubApplication](#) (see page 49)

[Unbypass a SubApplication](#) (see page 50)

SubApplication Request and Unrequest

In an Application, you can specify that certain jobs only run when requested by a user. These jobs are named on-request jobs. When you request a subApplication, the server requests all on-request jobs in the subApplication. If requested, the on-request jobs in the subApplication run when their dependences have been met. If you do not request a subApplication, the server bypasses the on-request jobs in the subApplication when they are eligible to run and releases their successors.

To request a subApplication, some jobs in the subApplication must have been defined as on-request jobs. You can request a subApplication at any time before its on-request jobs are eligible to run in the schedule. Prior to its on-request jobs being submitted, you can unrequest the subApplication.

More information:

[Request a SubApplication](#) (see page 52)

[Unrequest a SubApplication](#) (see page 53)

SubApplication Ready

A job is not ready for submission until all of its predecessor and variable dependencies are satisfied and its submission time has been reached. To drop these dependencies for each job in a subApplication, you can ready the subApplication for submission. When you ready a subApplication for submission, the server does the following for each job in the subApplication:

- Drops the job's predecessor dependencies
- Drops the job's variable dependencies
- Abandons the job's submission time, if it has one
- Releases the job from manual hold, if it is being held

If a job has resource dependencies, you must remove those dependencies separately if you want the job to run immediately.

More information:

[Ready a SubApplication for Submission](#) (see page 51)

SubApplication Unwait

In a subApplication, you can specify that the subApplication waits until its jobs in a previous generation of the Application complete before running. The server marks the jobs in a subApplication that wait for a previous generation in a SUBAPPLWAIT state. You can remove the SUBAPPLWAIT state by unwaiting the subApplication. When you unwait a subApplication, the subApplication no longer waits for its jobs in a previous generation of the Application to complete before it runs.

More information:

[Unwait a SubApplication](#) (see page 54)

Display SubApplications in an Active Application

You can display subApplications when monitoring an active Application.

To display subApplications

1. Open the Monitor perspective.
2. Double-click the Application in the Application Monitor view.

The Application's graphical view appears.

3. Click subApplication at the bottom of the Application workspace.

The jobs are grouped in each subApplication and displayed in a graphical representation.

Note: To issue a command against a subApplication, right-click the subApplication and select a command from the pop-up menu.

Complete a SubApplication

You can complete a subApplication to complete all of the jobs in the subApplication.

To complete a subApplication

1. Open the Monitor perspective.
2. Subscribe to workload on the server.

The Applications and their generations are displayed in the Application Monitor view.

3. Click the plus sign (+) beside the Application generation that contains the subApplication.

Monitor displays the Application's subApplications.

4. Right-click the subApplication you want to complete, and select Complete from the drop-down menu.

A confirmation dialog opens.

5. (Optional) Enter the command reason in the Reason field.

Limit: 1024 characters (including the timestamp and the user name)

Note: We recommend that the reason text not exceed 250 characters. The number of commands displayed in the command log is limited to the size of the command log buffer, which is 1024 characters. For example, if the reason text averages 10 characters per command, the command log displays up to 50 commands; if the reason text averages 250 characters per command, the command log displays up to 4 commands.

6. Click OK.

The server completes each of the jobs in the subApplication.

More information:

[Application Completion](#) (see page 21)

Hold a SubApplication

You can hold a subApplication to hold all of the jobs in the subApplication. The server does not submit any new jobs in the subApplication.

To hold a subApplication

1. Open the Monitor perspective.
2. Subscribe to workload on the server.

The Applications and their generations are displayed in the Application Monitor view.

3. Click the plus sign (+) beside the Application generation that contains the subApplication.

Monitor displays the Application's subApplications.

4. Right-click the subApplication you want to hold, and select Hold from the drop-down menu.

A confirmation dialog opens.

5. (Optional) Enter the command reason in the Reason field.

Limit: 1024 characters (including the timestamp and the user name)

Note: We recommend that the reason text not exceed 250 characters. The number of commands displayed in the command log is limited to the size of the command log buffer, which is 1024 characters. For example, if the reason text averages 10 characters per command, the command log displays up to 50 commands; if the reason text averages 250 characters per command, the command log displays up to 4 commands.

6. Click OK.

The server holds each of the jobs in the subApplication.

More information:

[Application Hold and Release](#) (see page 22)

Release a Held SubApplication

You can release a held subApplication to release all of the jobs in the subApplication. The server submits the jobs in the subApplication once their dependencies are met.

To release a held subApplication

1. Open the Monitor perspective.
2. Subscribe to workload on the server.

The Applications and their generations are displayed in the Application Monitor view.

3. Click the plus sign (+) beside the Application generation that contains the subApplication.

Monitor displays the Application's subApplications.

4. Right-click the subApplication you want to release, and select Release from the drop-down menu.

A confirmation dialog opens.

5. (Optional) Enter the command reason in the Reason field.

Limit: 1024 characters (including the timestamp and the user name)

Note: We recommend that the reason text not exceed 250 characters. The number of commands displayed in the command log is limited to the size of the command log buffer, which is 1024 characters. For example, if the reason text averages 10 characters per command, the command log displays up to 50 commands; if the reason text averages 250 characters per command, the command log displays up to 4 commands.

6. Click OK.

The server releases each of the jobs from hold in the subApplication. Each of the jobs in the subApplication run when all of its dependencies, if any, are met.

More information:

[Application Hold and Release](#) (see page 22)

Bypass a SubApplication

You can bypass a subApplication to bypass all of the jobs in the subApplication, preventing the jobs in the subApplication from running.

To bypass a subApplication

1. Open the Monitor perspective.
2. Subscribe to workload on the server.

The Applications and their generations are displayed in the Application Monitor view.

3. Click the plus sign (+) beside the Application generation that contains the subApplication.

Monitor displays the Application's subApplications.

4. Right-click the subApplication you want to bypass, and select Bypass from the drop-down menu.

A confirmation dialog opens.

5. (Optional) Enter the command reason in the Reason field.

Limit: 1024 characters (including the timestamp and the user name)

Note: We recommend that the reason text not exceed 250 characters. The number of commands displayed in the command log is limited to the size of the command log buffer, which is 1024 characters. For example, if the reason text averages 10 characters per command, the command log displays up to 50 commands; if the reason text averages 250 characters per command, the command log displays up to 4 commands.

6. Click OK.

The server bypasses each of the jobs in the subApplication.

More information:

[SubApplication Bypass and Unbypass](#) (see page 44)

Unbypass a SubApplication

You can cancel your request to bypass all of the jobs in the subApplication by unbypassing the subApplication.

To unbypass a subApplication

1. Open the Monitor perspective.
2. Subscribe to workload on the server.

The Applications and their generations are displayed in the Application Monitor view.

3. Click the plus sign (+) beside the Application generation that contains the subApplication.

Monitor displays the Application's subApplications.

4. Right-click the subApplication you want to unbypass, and select Unbypass from the drop-down menu.

A confirmation dialog opens.

5. (Optional) Enter the command reason in the Reason field.

Limit: 1024 characters (including the timestamp and the user name)

Note: We recommend that the reason text not exceed 250 characters. The number of commands displayed in the command log is limited to the size of the command log buffer, which is 1024 characters. For example, if the reason text averages 10 characters per command, the command log displays up to 50 commands; if the reason text averages 250 characters per command, the command log displays up to 4 commands.

6. Click OK.

The server unbypasses each of the jobs in the subApplication.

More information:

[SubApplication Bypass and Unbypass](#) (see page 44)

Ready a SubApplication for Submission

You can ready a subApplication to mark all of the jobs in the subApplication ready for submission.

To ready a SubApplication for submission

1. Open the Monitor perspective.
2. Subscribe to workload on the server.

The Applications and their generations are displayed in the Application Monitor view.

3. Click the plus sign (+) beside the Application generation that contains the subApplication.

Monitor displays the Application's subApplications.

4. Right-click the subApplication you want to mark ready, and select Ready from the drop-down menu.

A confirmation dialog opens.

5. (Optional) Enter the command reason in the Reason field.

Limit: 1024 characters (including the timestamp and the user name)

Note: We recommend that the reason text not exceed 250 characters. The number of commands displayed in the command log is limited to the size of the command log buffer, which is 1024 characters. For example, if the reason text averages 10 characters per command, the command log displays up to 50 commands; if the reason text averages 250 characters per command, the command log displays up to 4 commands.

6. Click OK.

The server readies each of the jobs in the subApplication for submission. The jobs run when their resource dependencies, if any, have been met.

More information:

[SubApplication Ready](#) (see page 45)

Request a SubApplication

Before the jobs in a subApplication run, you can request the subApplication to request all of the on-request jobs in the subApplication.

To request a subApplication

1. Open the Monitor perspective.
2. Subscribe to workload on the server.

The Applications and their generations are displayed in the Application Monitor view.

3. Click the plus sign (+) beside the Application generation that contains the subApplication.

Monitor displays the Application's subApplications.

4. Right-click the subApplication you want to request, and select Request from the drop-down menu.

A confirmation dialog opens.

5. (Optional) Enter the command reason in the Reason field.

Limit: 1024 characters (including the timestamp and the user name)

Note: We recommend that the reason text not exceed 250 characters. The number of commands displayed in the command log is limited to the size of the command log buffer, which is 1024 characters. For example, if the reason text averages 10 characters per command, the command log displays up to 50 commands; if the reason text averages 250 characters per command, the command log displays up to 4 commands.

6. Click OK.

The server requests each of the on-request jobs in the subApplication.

More information:

[SubApplication Request and Unrequest](#) (see page 44)

Unrequest a SubApplication

You can cancel your request to submit all of the on-request jobs in a subApplication by unrequesting the subApplication.

To unrequest a subApplication

1. Open the Monitor perspective.
2. Subscribe to workload on the server.

The Applications and their generations are displayed in the Application Monitor view.

3. Click the plus sign (+) beside the Application generation that contains the subApplication.

Monitor displays the Application's subApplications.

4. Right-click the subApplication you want to unrequest, and select Unrequest from the drop-down menu.

A confirmation dialog opens.

5. (Optional) Enter the command reason in the Reason field.

Limit: 1024 characters (including the timestamp and the user name)

Note: We recommend that the reason text not exceed 250 characters. The number of commands displayed in the command log is limited to the size of the command log buffer, which is 1024 characters. For example, if the reason text averages 10 characters per command, the command log displays up to 50 commands; if the reason text averages 250 characters per command, the command log displays up to 4 commands.

6. Click OK.

The server unrequests each of the on-request jobs in the subApplication.

More information:

[SubApplication Request and Unrequest](#) (see page 44)

Unwait a SubApplication

You can unwait a subApplication so that the subApplication no longer waits for its jobs in a previous generation of the Application to complete before it runs.

Note: You can only unwait a subApplication if its jobs are in the SUBAPPLWAIT state.

To unwait a subApplication

1. Open the Monitor perspective.
2. Subscribe to workload on the server.

The Applications and their generations are displayed in the Application Monitor view.

3. Click the plus sign (+) beside the Application generation that contains the subApplication.

Monitor displays the Application's subApplications.

4. Right-click the subApplication you want to unwait, and select Unwait from the drop-down menu.

A confirmation dialog opens.

5. (Optional) Enter the command reason in the Reason field.

Limit: 1024 characters (including the timestamp and the user name)

Note: We recommend that the reason text not exceed 250 characters. The number of commands displayed in the command log is limited to the size of the command log buffer, which is 1024 characters. For example, if the reason text averages 10 characters per command, the command log displays up to 50 commands; if the reason text averages 250 characters per command, the command log displays up to 4 commands.

6. Click OK.

The server removes the SUBAPPLWAIT state and submits the jobs in the subApplication.

More information:

[SubApplication Unwait](#) (see page 45)

Chapter 4: Monitoring and Controlling Jobs

This section contains the following topics:

- [Jobs in Active Applications](#) (see page 58)
- [Issue a Command Against Jobs](#) (see page 73)
- [Display Job Details](#) (see page 74)
- [Display a Job's Spool File](#) (see page 76)
- [Reset a Job Definition](#) (see page 79)
- [Bypass a Job](#) (see page 79)
- [Unbypass a Job](#) (see page 80)
- [Cancel a Running Job](#) (see page 81)
- [Complete a Job](#) (see page 82)
- [Hold a Job](#) (see page 83)
- [Release a Held Job](#) (see page 83)
- [Request a Job](#) (see page 84)
- [Unrequest a Job](#) (see page 85)
- [Resubmit a Job](#) (see page 86)
- [Ready a Job for Submission](#) (see page 88)
- [View the Comments of a Job](#) (see page 88)
- [Unwait a Job](#) (see page 89)
- [Set a Job's User Status](#) (see page 90)
- [Drop Predecessor Dependencies](#) (see page 90)
- [Reset Time Dependencies](#) (see page 91)
- [Drop Variable Dependencies](#) (see page 93)
- [Evaluate Variable Dependencies](#) (see page 94)
- [Drop Resource Dependencies](#) (see page 95)
- [Display Resource Usage](#) (see page 96)
- [Verify Whether a Job is Running or Has Executed](#) (see page 97)
- [Set a Windows Job Object's Properties](#) (see page 98)
- [Set Process Priority for a Windows or a UNIX Job](#) (see page 101)
- [Send a Signal to a UNIX Job](#) (see page 102)
- [Retrieve the Spool File for an i5/OS Job](#) (see page 104)
- [Respond to Suspended Jobs that Require Manual Intervention](#) (see page 106)
- [Retrieve the JES Spool File of a Micro Focus Job](#) (see page 107)
- [Browse JCL](#) (see page 109)
- [Browse Copy JCL](#) (see page 109)
- [Edit JCL](#) (see page 110)
- [Edit Copy JCL](#) (see page 110)
- [z/OS Jobs Restart Using CA WA Restart Option](#) (see page 111)
- [Issue a Database Command](#) (see page 116)
- [SAP Commands](#) (see page 117)
- [Issue a PeopleSoft Command](#) (see page 131)
- [Issue an Oracle E-Business Suite Command](#) (see page 132)

Jobs in Active Applications

After the server selects a job to run in an Application, you can monitor and control the job in the Monitor perspective. You can monitor the job's state, view its details and output, and issue commands against the job to do the following:

- Change the job's definition
- Bypass or unbypass the job
- Cancel the job (does not apply to all workload objects)
- Complete the job
- Hold or release the job
- Request or unrequest the job
- Resubmit the job
- Ready the job
- Show the job's comments, if any
- Set the job's user status

If the job has predecessor, time, variable, or resource dependencies, you can issue additional commands to do the following:

- Drop the job's predecessor dependencies
- Reset the job's time dependencies
- Drop the job's variable dependencies
- Evaluate the job's variable dependencies
- Drop the job's resource dependencies
- List the job's resource usage

Some job types support additional commands. For example, if you run Windows, UNIX, Micro Focus, z/OS, SAP, PeopleSoft, or Oracle Applications workload, you can issue additional commands against these jobs.

Note: To run commands against jobs, you require the appropriate security permissions. For more information, contact your administrator.

Job States

When the server manages a job, the job passes through multiple processing stages named *states*. Some states, such as job failure, require action.

At any time during the course of its execution, a job has one state and zero or more conditions. A job's state and conditions depend on the job's properties, user interaction, and the success or failure of the workload being executed. The job's details displays the job's state and conditions.

Some states can appear as both a state and a condition. When multiple states apply to a job, state priorities determine the job's state. For example, suppose that a job waiting for a predecessor (PREDWAIT) is held at the Application level (APPLHOLD). Since the APPLHOLD state has a higher priority than the PREDWAIT state, the job's state becomes APPLHOLD with a condition of PREDWAIT. If the Application is released from hold, the job's state changes to PREDWAIT.

A job can pass through the following states, from highest priority to lowest priority:

BYPASSED

Indicates that the job has been bypassed. It is a final state. Before the job's predecessor dependencies have been met, the job has a Bypassreq condition, indicating that the job has been marked for bypassing.

COMPLETE

Indicates that the job has completed successfully. It is a final state if the Application that contains the job completes. If the job was completed manually, it has a Forced condition.

FAILED

Indicates that the job has failed. The job must be resubmitted or manually completed.

Note: Failed jobs that have an AutoResubmit condition are resubmitted automatically. By default, CA WA Desktop Client displays jobs with an AutoResubmit condition in yellow.

SUBERROR

Indicates that the submission of the job resulted in an error. The job must be resubmitted or manually completed. Common causes of submission errors include the following:

- Specifying an undefined agent in the job definition
- Specifying an invalid argument in the job definition
- Executing an invalid JavaScript script at run time
- Using an agent feature that is not supported on your agent version

Note: For information about what caused the submission error, see the job's details.

EXTSCHDOWN

Indicates that the external scheduler is down. The workload object must be completed manually.

EXEC

Indicates that the job is currently running.

DEFINED

Indicates that the SAP job is defined on the SAP system.

SUBMIT

Indicates that the SAP job has been released on the SAP system.

INACTIVE

Indicates that the Oracle Applications job is inactive because the concurrent request is either disabled, suspended, or on hold.

MONITOR

Indicates that the job is monitoring for a condition to occur. For continuous monitoring, the server triggers an Alert each time the monitor condition is met. For non-continuous monitoring, the job completes when the monitor condition is met.

MANSUB

Indicates that the z/OS job is to be submitted outside CA Workload Automation DE, such as a job that is submitted manually by a user.

AGENTDOWN

Indicates that the agent referenced in the job definition is down.

APPLHOLD

Indicates that the Application that contains the job is in a hold state. The job cannot run until the Application is released. Holding an Application does not affect jobs currently running but does prevent the server from submitting new jobs.

APPLWAIT

Indicates that the Application that contains the job is waiting for a previous generation of the same Application to complete. The job cannot run until the previous generation of the Application completes.

SUBAPPLWAIT

Indicates that the subApplication that contains the job is waiting for a previous generation of the subApplication to complete. The job cannot run until the previous generation of the subApplication completes.

JANCWAIT

Indicates that the job is waiting for the same job to complete in a previous generation of the Application. The job cannot run until the same job completes in a previous generation of the Application.

MANHOLD

Indicates that the job is manually held. The job cannot run until the manual hold is released.

PREDWAIT

Indicates that the job is waiting for a predecessor. The job cannot run until the release condition is met or dropped.

WAITING

Indicates that the job is waiting until a specific time for submission. This state is always paired with the Timewait condition. The job cannot run until the submission time is met or reset to the current time.

EXTWAIT

Indicates that the external job is waiting for its home job to complete in the home Application.

SUBDELAY

Indicates that the submission of this job has been delayed because the agent is quiesced.

VARWAIT

Indicates that the job is waiting for one or more variable dependencies to be met. The job cannot run until all the required variable dependencies are met, dropped, or abandoned.

RESWAIT

Indicates that the job is waiting to acquire resources. The job cannot run until all the required resource dependencies are met, dropped, or abandoned.

TASKWAIT

Indicates that the task requires manual completion.

READY

Indicates that the job is ready to be submitted (all of its dependencies have been met). A workload object will not remain in this state.

UNKNOWN

Indicates that the server cannot determine the state of the workload object.

Job Conditions

During its execution, a job can have one or more conditions in addition to its state. Some job states can also appear as a condition if another applicable state has a higher priority. A state cannot appear as both a state and a condition simultaneously.

In addition to states that can appear as conditions, a job can have the following conditions:

AutoResubmit

Indicates that the job will be resubmitted automatically. The job's state is FAILED.

Note: By default, CA WA Desktop Client displays failed jobs with an AutoResubmit condition in yellow.

Bypassreq

Indicates that the job has been requested for bypassing. When the job's time and predecessor dependencies are met, the server bypasses the job and the job's state becomes BYPASSED.

Forced

Indicates that the job has been manually completed. The job's state is always COMPLETE.

Note: You can monitor for forced completions by specifying a notification in the Application or job definition. The notification can send an email, trigger an Event, run a JavaScript script, or send an SNMP trap. For more information about setting up notifications, see the *Define Perspective Help*.

Overdue

Indicates that the job is overdue because it is late starting, it is late completing, or it exceeded a maximum execution time.

Note: By default, Desktop Client displays jobs with an Overdue condition in red.

Premend

Indicates that the job completed prematurely.

Request

Indicates that the job runs only if it is requested.

Requested

Indicates that the job was requested. The job runs when its dependencies are met.

Timewait

Indicates that the job is waiting for a time dependency to be satisfied. The job may be in a WAITING state.

Job Commands Usage

You can issue commands against a single job or multiple jobs. The commands you can issue depend on the workload object and its state. For example, you cannot complete a job that is already complete or bypassed.

You can issue a command against multiple jobs with the following limitations:

- If you select jobs with different states, you can only issue commands that apply to the selected states. For example, you can resubmit jobs that have terminated abnormally (such as FAIL or SUBERROR), but you cannot resubmit bypassed jobs.
- If you select different types of jobs, you can only issue commands that are applicable for every selected job. For example, if you select two UNIX jobs and three SAP R/3 jobs, you can only apply commands that are applicable to both jobs, such as Bypass, but not commands applicable to only SAP jobs.
- You cannot cancel multiple jobs with a single command.
- You cannot issue commands against jobs controlled on different servers.

More information:

[Issue a Command Against Jobs](#) (see page 73)

Job's Spool File

You can display a job's output named a spool file. For example, the spool file may contain the standard output for a UNIX job or the results of an SQL query for a Database job. You can retrieve the spool file for any executing, failed, or completed job.

Depending on the type and version of the agent the job is running on, you can view the entire spool file, selected lines, or the number of lines you specify.

More information:

[Display a Job's Spool File](#) (see page 76)

FTP Job Spool File

To verify that the transfer completed successfully without file corruption, you can check the job's spool file. The spool file may contain the following:

- If the data was transferred using SSL FTP, the spool file contains a response as follows:

```
AUTH TLS
234 AUTH command OK. Initializing SSL connection.
```

- If the file was compressed and transferred without file corruption, the spool file contains a response as follows:

```
Downloaded 81920/26119 bytes (original/compressed) in 0.161 seconds, 496.89
Kbytes/sec.
```

- If the file was downloaded successfully, the spool file contains the following response:

```
Download successful
```

Job Definition Reset

You can change, add, or remove some properties before a job's execution or resubmission. These changes are temporary and are associated only with a particular instance of a job.

You cannot reset the job definition of a bypassed job, or a job in AGENTDOWN state. If you reset the job definition for a job in a FAIL or SUBERROR state, the changes do not occur until you resubmit the job.

For example, you may have incorrectly typed a script name when you defined a job and, as a result, the job failed. You can reset the job definition with the correct script name and resubmit the job.

You can use any symbolic variables that the server sets at run time or trigger time. When the job runs, the server automatically resolves all the variables.

More information:

[Reset a Job Definition](#) (see page 79)

Job Bypass and Unbypass

Before the server submits a job to run in the schedule, you can prevent the job from running by bypassing it. When the job's dependencies are met, the server bypasses the job and releases its successor jobs. For example, job A releases job B, which releases job C. While job A runs, you request that job B be bypassed. When job A completes, the server bypasses job B, changes its state to `BYPASSED`, and releases job C.

You can unbypass the job before the server bypasses the job. For example, if a job is bypassed by mistake, you can undo the action by unbypassing the job. When the job's dependencies are met, the server runs the job as usual.

Note: You can also prevent a job from running by completing it. When you complete a job, the server immediately releases the job's successor jobs, even if its predecessors are still running. Unlike bypassing a job, you cannot undo completing a job.

More information:

[Bypass a Job](#) (see page 79)

[Unbypass a Job](#) (see page 80)

Job Cancellation

When you cancel a running job, the server flags the job as failed and does not release the job's successors. To release the job's successors, you must complete the job.

If you cancel a Windows or UNIX job, the agent cancels the running process, and the server flags the job as failed.

Note the following about job cancellation:

- Some workload objects do not support job cancellation.
- To cancel a Micro Focus job, you may have to issue the cancel command twice. After issuing the first command, if the job continues to execute, issue the cancel command again.
- In most cases, you cannot cancel an SQL or DB Stored Procedure job.

More information:

[Cancel a Running Job](#) (see page 81)

Job Completion

When you force a job to complete, the server completes the job immediately. The server deems the completed job to have executed, even if the job has not been submitted, has failed, or is still running. If you force a job to complete, its successors can run.

Note: You can monitor for forced completions by specifying a notification in the Application or job definition. The notification can send an email, trigger an Event, run a JavaScript script, or send an SNMP trap. For more information about setting up notifications, see the *Define Perspective Help*.

If you mistakenly complete a job, you can insert another occurrence of it with the required dependencies. Because job names must be unique, you could use the same job name but add a different qualifier. However, the job's successors may already have been released.

Note: Completing a job is different than bypassing a job. If you bypass a job, the server does not bypass the job until its dependencies are met. If you mistakenly bypass a job, you can unbypass it before the job is eligible to run. If you complete a job, it is completed immediately and cannot be undone.

More information:

[Complete a Job](#) (see page 82)

[Job Conditions](#) (see page 62)

External Job Completion

If you complete a job in the home Application, the server flags its corresponding External jobs as complete.

If you complete an External job in the distant Application, only that job is flagged as complete.

Example: External Job Completion

Suppose that the home Application for job X is APPL1 and the distant Application for job X is APPL2. If you complete job X in the home Application (APPL1), the server completes the External job X in APPL2 and releases its successors. If you complete job X in the distant Application (APPL2), job X in the home Application (APPL1) is unaffected.

Job Hold and Release

Prior to job submission, you can manually hold a job until you want it to run. The server does not release the job until you manually release it and all of its dependencies have been met.

When you place a job on hold, you can include a reason for the hold.

For example, suppose that a programmer needs to make a last-minute change before a job can run. You place the job on hold with the reason “last-minute change”. When the change has been made, you release the held job. The job is now eligible to run.

Note: You cannot manually hold an External job.

More information:

[Hold a Job](#) (see page 83)

[Release a Held Job](#) (see page 83)

Job Request and Unrequest

In an Application, you can specify that certain jobs only run when requested by a user. These jobs are named on-request jobs. If you request an on-request job, the job runs when all of its dependencies are met. If an on-request job is not requested, the server bypasses the job when it is eligible to run and releases its successors.

To request a job, the job must be defined as an on-request job. You can request a job any time before it is eligible to run in the schedule. Prior to being submitted, you can unrequest the job.

More information:

[Request a Job](#) (see page 84)

[Unrequest a Job](#) (see page 85)

Job Resubmission

If a job terminates abnormally (such as FAIL or SUBERROR), you can resubmit it. You cannot resubmit bypassed jobs, jobs in an AGENTDOWN state, or jobs in completed Applications. You can resubmit completed jobs in active Applications.

Prior to job resubmission, you can reset the job's definition.

Note: You cannot resubmit an External job.

More information:

[Resubmit a Job](#) (see page 86)

Job Unwait

In a job definition, you can specify that the job waits until the same job completes only in the previous Application generation or in all previous Application generations before running. The server marks a job that waits to complete in a previous Application generation in a JANCWAIT state. You can remove the JANCWAIT state by unwaiting the job. When you unwait a job, the selected job no longer waits to complete in a previous Application generation before it runs.

More information:

[Unwait a Job](#) (see page 89)


Predecessor Dependencies

Many jobs can run simultaneously in an Application. Predecessor dependencies describe the relationships between jobs in an Application. By default, a job runs after all its predecessors complete successfully. A job that waits for other jobs to complete is in the PREDWAIT state.

Example: Dependencies between jobs in an Application

Every day, jobs A, B, C, and D run. On Fridays, job E runs. On the last workday of the month, job F also runs.

The following illustration displays the run frequency of all the jobs.



The server submits a job after all of the job's predecessors complete successfully. In this Application, the jobs run in the following order:

- Job B and job C run after job A completes successfully
- Job D runs after job B and job C complete successfully
- Job E runs on Fridays after job D completes successfully
- On the last workday of the month, there are two possibilities:
 - If the last workday of the month is a Friday, job F runs after job E completes successfully
 - If the last workday of the month is not a Friday, job F runs after job D completes successfully

More information:

[Drop Predecessor Dependencies](#) (see page 90)

Time Dependencies

A job can contain several time dependencies. When a job reaches its job submission time, the server checks for predecessor, variable, and resource dependencies and submits the job if it meets all the dependencies.

You can use time dependencies to do the following at scheduled times:

- Delay job submission
- Mark a job overdue (Overdue condition)
- Mark a job that ends prematurely (Premend condition)
- Bypass a job
- Abandon a job's predecessor dependencies
- Abandon a job's variable dependencies
- Abandon a job's resource dependencies

You can specify time dependencies by specifying absolute or relative times such as 9PM, 13:00, NOW PLUS 10 MINUTES, or 10PM TODAY PLUS 2 WORKDAYS.

The server resolves your scheduling criteria to a single date and time as follows:

- If you specify a time without a date (for example, 9PM), the server resolves the date to the next occurrence of that time.
- If you specify a date without a time (for example, 1ST MONDAY OF MONTH), the server resolves the time to midnight (00:00:00).
- If you specify a frequency (for example, MONDAYS 10PM), the server resolves the date and time to the next occurrence of that date and time.

More information:

[Reset Time Dependencies](#) (see page 91)

Overdue Jobs Monitoring

A job can become overdue if the job is late starting, the job is late completing, or the job exceeds its maximum execution time. If a job becomes overdue, the server adds the Overdue condition to the job's conditions. By default, CA WA Desktop Client displays jobs with an Overdue condition in red.

You can monitor overdue jobs in the following ways:

- View the job details to check if a job has an Overdue condition. The Status field displays whether the job is overdue because of a late start (the job is late starting) or late completion (the job is late completing or it exceeds its maximum execution time).
- Use the custom view Overdue to monitor jobs with Overdue conditions.
- Create a notification for the job so that the server automatically takes action when a job is overdue. The notification can alert you and other users by email or SNMP trap, trigger an Event, or run a JavaScript script.

The Overdue condition is cleared when the condition is met. For example, if a job is late starting, the Overdue condition is cleared when the job starts. If a job is late completing, the Overdue condition is cleared when the job completes successfully.

If you reset a job's dueout (overdue) time, the server compares the new overdue time to the current time and resets the time as follows:

- If the new dueout time is in the past, the server adds the Overdue condition to the job's conditions and sets the job's dueout time to the new time.
- If the new dueout time is in the future, the server sets the job's dueout time to the new time.

Note: If the job is currently marked with an overdue condition, you cannot clear the overdue condition by resetting a dueout time. If the job is late starting, you cannot reset the late start time. If the job is late completing or has exceeded its maximum execution time, you cannot reset either the late completion time or the maximum execution time.

More information:

[Reset Time Dependencies](#) (see page 91)

Global Variable Dependencies

Jobs can have global variable dependencies. A global variable dependency is a variable expression that must be satisfied before a job is submitted. Such jobs are submitted when all of the global variable dependencies (and the time, predecessor, and resource dependencies) are met, dropped, or abandoned.

For example, you can define a job that only runs when a global variable named quota is assigned a value greater than or equal to 1000. If the global variable dependency is not met at job submission time, the job goes into a VARWAIT state and waits for the dependency to be met, dropped, or abandoned.

There are several ways to manage global variable dependencies. You can set up, edit, and remove variable dependencies by using the Define perspective of CA WA Desktop Client. While monitoring an Application generation in the Monitor perspective, you can set up, edit, remove, and drop variable dependencies and reset the time to abandon variable dependencies. You can also drop variable dependencies by using the CLI. To check the status of jobs in a VARWAIT state, you can evaluate the variable dependencies by using the Monitor perspective or the CLI.

Variable dependencies are stored in the relational database for the server. These dependencies are kept and restored after a warm start of the server or after the server starts as a Primary in a CA WA High Availability configuration. In a cold start of the server, these dependencies are deleted from the relational database.

You can set up global variable dependencies in all job types except for External-Other Scheduler and External-Same Scheduler.

More information:

[Drop Variable Dependencies](#) (see page 93)

[Evaluate Variable Dependencies](#) (see page 94)

How Jobs with Global Variable Dependencies are Submitted

After a job meets all its time and predecessor dependencies, the server does the following:

- Checks if the job has any global variable dependencies.
 - If the global variable dependencies are not met, the job goes into a VARWAIT state.
 - When all of the global variable dependencies are met, dropped, or abandoned, the server removes the VARWAIT state.
- Checks if the job has any resource dependencies. When the resource dependencies are met, the server submits the job.

How Jobs with Resource Dependencies are Submitted

After a job meets all its time and predecessor dependencies, the server does the following:

- Readies the job for submission.
- Checks if the job has requested any resources.
- Checks if resource units are available if the job has requested resources.
 - If all the requested resource units are not available, the job goes into a RESWAIT state.
 - When all the requested resource units are available, the server removes the RESWAIT state and submits the job.

More information:

[Display Resource Usage](#) (see page 96)

[Drop Resource Dependencies](#) (see page 95)

Issue a Command Against Jobs

You can issue a command against a single job or multiple jobs. For example, you can issue a command to complete, bypass, cancel, hold, request, ready, or resubmit jobs.

To issue a command against jobs

1. Select the jobs that you want to issue a command against in the Monitor perspective.

Note: You can press the Ctrl key while making your selections to select multiple jobs at once.

2. Right-click one of the selected jobs, and select a command from the pop-up menu.
The server issues the command against the selected jobs.

More information:

[Job Commands Usage](#) (see page 63)

Display Job Details

You can display a job's details for more information about the job. For example, you may want to view a job's details to see why it failed or to view its conditions.

To display a job's details

1. Right-click the job in the Monitor perspective.

A list of job commands appears.

2. Select Details.

The details of the job are displayed.

Note: The details you see depend on the job definition and the state of the job. Monitor only displays details that have values for that job.

Display a Job's Anticipated End Time

If an Application has Estimate end time enabled, you can display the anticipated end times of its jobs in the Monitor perspective.

To display a job's anticipated end times

1. Right-click the job in the Monitor perspective.

A list of job commands appears.

2. Select Details.

The Details dialog displays the job's anticipated end time in the Anticipated End Time field.

Display the Command Log of a Job

You can view the command log of a job that displays the most recent commands issued for that selected job.

Notes:

- If the job belongs to a subApplication, the command log includes commands issues against the subApplication.
- The command log includes commands that were invoked programmatically by a JavaScript script using the `execCommand` function. For more information about the `execCommand` function, see the *Programming Guide*.

To display the command log of a job

1. Right-click the job in the Monitor perspective.

A list of job commands appears.

2. Select Details.

The job details are displayed.

3. Click View Jobs Command Log.

The Job Commands Log dialog opens. The Commands section displays a table with the following columns:

Date and Time

Displays the date and time the command was issued against the job.

Command

Displays the name of the command issued against the job.

Note: The table only displays commands that were successful. Failed commands do not appear.

User

Displays the name of the user that issued the command.

Programmatic

Specifies whether the command was invoked programmatically by a JavaScript script.

Reason

Displays the command reason, if any

Note: To display the full command reason in the Reason section, select the row in the table.

Notes:

- The number of commands displayed in the command log is limited to the size of the command log buffer, which is 1028 characters. For example, if the reason text averages 10 characters per command, the command log displays up to 50 commands; if the reason text averages 250 characters per command, the command log displays up to 4 commands.
- You can click the table headings to change the column sorting order.

4. (Optional) Click Export as Text .

The command log is saved as a text file.

5. (Optional) Click Export as CSV.

The command log is saved as a .csv file. You can open the .csv file in Microsoft Excel.

Display a Job's Spool File

You can display a spool file to view the output of the job.

To display a spool file of the job

1. Right-click the job in the Monitor perspective.

A list of job commands appears.

2. Select Retrieve Spool File.

The Retrieve Spool File dialog opens.

3. (Optional) Select the Selection Criteria option button and complete the following fields as appropriate:

Line Contains

Specifies a text string to search for.

Note: You cannot use wildcards in this field.

Start of range

Specifies the start point of the spool file portion to be retrieved in lines, bytes, kilobytes, or megabytes.

End of range

Specifies the end point of the spool file portion to be retrieved in lines, bytes, kilobytes, or megabytes.

Lines

Retrieves the specified spool file portion by lines.

Bytes

Retrieves the specified spool file portion by bytes.

Kbytes

Retrieves the specified spool file portion by kilobytes.

Mbytes

Retrieves the specified spool file portion by megabytes.

Last match

Retrieves the last match of the specified text string within the specified range.

Note: To retrieve all matches, clear this check box.

4. (Optional) Select the Line numbers option button and complete the following field:

Last number of lines to display

Specifies the number of lines to retrieve from the end of the spool file.

5. (Optional) Select the Append to existing display check box to add the results returned to the existing display.

Note: This option can be useful if you retrieve the spool file in segments.

6. Click one of the following options:

Retrieve Selected

Retrieves a selection of the spool based on criteria you specify. You can specify a text string to search for within a range or specify the last number of lines to display.

Retrieve Next

Retrieves the entire spool file in segments.

Note: Use this option if you are retrieving large spool files or you are unsure of the spool file size.

Retrieve All

Retrieves the entire spool file.

Note: If you retrieve a large spool file, this option can slow down the server.

The server displays the spool file of the job.

Note: If you are retrieving the spool file for an HTTP job and receive an error message stating that the line is too long, you can open the file manually to view the line.

Example: Retrieve Part of a Spool File

Suppose that you want to search lines 12 to 20 in a spool file for the string "error" and return the last line in the range that contains the string.

To retrieve a part of the spool file

1. Right-click the job in the Monitor perspective.
A list of job commands appears.
2. Select Retrieve Spool File.
The Retrieve Spool File dialog opens.
3. Select the Selection Criteria and the Lines option buttons.
4. Enter **error** in the Line contains field.
5. Enter **12** in the Start of range field and **20** in the End of range field.
6. Select the Last match check box.
7. Click Retrieve Selected.

The server searches lines 12 to 20 in the spool file for the string "error". The spool file displays the last line in the range that contains the string "error".

More information:

[Job's Spool File](#) (see page 63)

Reset a Job Definition

You can reset a job definition if you need to modify its details for a particular instance.

To reset a job definition

1. Right-click the job in the Monitor perspective.
A list of job commands appears.
2. Select Reset Definition.
The Reset Definition dialog opens.
3. Modify the details as appropriate and click OK.
The server resets the job definition for the current instance.

Note: To make the changes permanent, update the job's definition in the Define perspective.

More information:

[Job Definition Reset](#) (see page 64)

Bypass a Job

Before the server submits a job, you can bypass a job to prevent it from running.

To bypass a job

1. Right-click the job in the Monitor perspective.
A list of job commands appears.
2. Select Bypass.
A confirmation dialog opens.

3. (Optional) Enter the command reason in the Reason field.

Limit: 1024 characters (including the timestamp and the user name)

Note: We recommend that the reason text not exceed 250 characters. The number of commands displayed in the command log is limited to the size of the command log buffer, which is 1024 characters. For example, if the reason text averages 10 characters per command, the command log displays up to 50 commands; if the reason text averages 250 characters per command, the command log displays up to 4 commands.

4. Click OK.

If the job is held, or is waiting for time or predecessor dependencies, the server adds the Bypassreq condition to the job's conditions. After the job is released from hold, or the time and predecessor dependencies are satisfied or dropped, the server bypasses the job.

Note: If you bypass a job that is waiting for variable or resource dependencies to be satisfied, the server bypasses the job immediately.

More information:

[Job Bypass and Unbypass](#) (see page 65)

Unbypass a Job

You can cancel your request to bypass a job by unbypassing it.

To unbypass a job

1. Right-click the job in the Monitor perspective.
A list of job commands appears.
2. Select Unbypass.
A confirmation dialog opens.

3. (Optional) Enter the command reason in the Reason field.

Limit: 1024 characters (including the timestamp and the user name)

Note: We recommend that the reason text not exceed 250 characters. The number of commands displayed in the command log is limited to the size of the command log buffer, which is 1024 characters. For example, if the reason text averages 10 characters per command, the command log displays up to 50 commands; if the reason text averages 250 characters per command, the command log displays up to 4 commands.

4. Click OK.

The server removes the Bypassreq condition from the job's conditions. When the job's dependencies are met, the server runs the job as usual.

More information:

[Job Bypass and Unbypass](#) (see page 65)

Cancel a Running Job

You may want to cancel a job if it has been submitted in error, is running much longer than expected, or is looping.

To cancel a running job

1. Right-click the job in the Monitor perspective.

A list of job commands appears.

2. Select Cancel Active Job.

A confirmation dialog opens.

3. (Optional) Enter the command reason in the Reason field.

Limit: 1024 characters (including the timestamp and the user name)

Note: We recommend that the reason text not exceed 250 characters. The number of commands displayed in the command log is limited to the size of the command log buffer, which is 1024 characters. For example, if the reason text averages 10 characters per command, the command log displays up to 50 commands; if the reason text averages 250 characters per command, the command log displays up to 4 commands. You can provide a reason only for UNIX, Windows, i5/OS, and Tandem job types that are in the EXEC state.

4. Click OK.

The server cancels the running job.

More information:

[Job Cancellation](#) (see page 65)

Complete a Job

You can complete a job to let its successors run or to let the Application that contains the job complete.

To complete a job

1. Right-click the job in the Monitor perspective.

A list of job commands appears.

2. Select Complete.

A confirmation dialog opens.

3. (Optional) Enter the command reason in the Reason field.

Limit: 1024 characters (including the timestamp and the user name)

Note: We recommend that the reason text not exceed 250 characters. The number of commands displayed in the command log is limited to the size of the command log buffer, which is 1024 characters. For example, if the reason text averages 10 characters per command, the command log displays up to 50 commands; if the reason text averages 250 characters per command, the command log displays up to 4 commands.

4. Click OK.

The state of the job changes to COMPLETE with a Forced condition.

More information:

[Job Completion](#) (see page 66)

Hold a Job

You can hold a job to prevent the server from submitting that job until you manually release it.

To hold a job

1. Right-click the job in the Monitor perspective.

A list of job commands appears.

2. Select Hold.

3. (Optional) Enter the command reason in the Reason field.

Limit: 1024 characters (including the timestamp and the user name)

Note: We recommend that the reason text not exceed 250 characters. The number of commands displayed in the command log is limited to the size of the command log buffer, which is 1024 characters. For example, if the reason text averages 10 characters per command, the command log displays up to 50 commands; if the reason text averages 250 characters per command, the command log displays up to 4 commands.

4. Click OK.

The server adds the MANHOLD condition to the job's conditions. The server holds the job until you manually release it.

More information:

[Job Hold and Release](#) (see page 67)

Release a Held Job

You can release a held job so that the server will submit the job when its dependencies have been met.

To release a held job

1. Right-click the job in the Monitor perspective.

A list of job commands appears.

2. Select Release.

A confirmation dialog opens.

3. (Optional) Enter the command reason in the Reason field.

Limit: 1024 characters (including the timestamp and the user name)

Note: We recommend that the reason text not exceed 250 characters. The number of commands displayed in the command log is limited to the size of the command log buffer, which is 1024 characters. For example, if the reason text averages 10 characters per command, the command log displays up to 50 commands; if the reason text averages 250 characters per command, the command log displays up to 4 commands.

4. Click OK.

The server removes the MANHOLD condition from the job's conditions. The job runs when all of its dependencies are met.

More information:

[Job Hold and Release](#) (see page 67)

Request a Job

If you want an on-request job to run, you must request it before it is eligible to run.

Note: The job must be defined as an on-request job.

To request a job

1. Right-click the job in the Monitor perspective.

A list of job commands appears.

2. Select Request.

A confirmation dialog opens.

3. (Optional) Enter the command reason in the Reason field.

Limit: 1024 characters (including the timestamp and the user name)

Note: We recommend that the reason text not exceed 250 characters. The number of commands displayed in the command log is limited to the size of the command log buffer, which is 1024 characters. For example, if the reason text averages 10 characters per command, the command log displays up to 50 commands; if the reason text averages 250 characters per command, the command log displays up to 4 commands.

4. Click OK.

Monitor adds the Requested condition to the job's conditions. When the job's dependencies are met, the server runs the job.

More information:

[Job Request and Unrequest](#) (see page 67)

Unrequest a Job

You can cancel your request to submit an on-request job by unrequesting the job.

Note: You must define the job as an on-request job and request it.

To unrequest a job

1. Right-click the job in the Monitor perspective.

A list of job commands appears.

2. Select Unrequest.

A confirmation dialog opens.

3. (Optional) Enter the command reason in the Reason field.

Limit: 1024 characters (including the timestamp and the user name)

Note: We recommend that the reason text not exceed 250 characters. The number of commands displayed in the command log is limited to the size of the command log buffer, which is 1024 characters. For example, if the reason text averages 10 characters per command, the command log displays up to 50 commands; if the reason text averages 250 characters per command, the command log displays up to 4 commands.

4. Click OK.

Monitor removes the Requested condition from the job's conditions. When the job's dependencies are met, the server bypasses the job.

More information:

[Job Request and Unrequest](#) (see page 67)

Resubmit a Job

You can resubmit a job to rerun a failed job or a completed job in an active Application.

To resubmit a job

1. Right-click the job in the Monitor perspective.

A list of job commands appears.

2. Select Resubmit.

A confirmation dialog opens.

3. (Optional) Enter the command reason in the Reason field.

Limit: 1024 characters (including the timestamp and the user name)

Note: We recommend that the reason text not exceed 250 characters. The number of commands displayed in the command log is limited to the size of the command log buffer, which is 1024 characters. For example, if the reason text averages 10 characters per command, the command log displays up to 50 commands; if the reason text averages 250 characters per command, the command log displays up to 4 commands. You cannot provide the command reason for a Micro Focus job.

4. Click OK.

The server resubmits the job. If you made changes to the job's properties in the Monitor perspective, the server resubmits the job using these changes.

Note: If you dropped the job's variable dependencies at any time before resubmission, the variable dependencies are also dropped from the resubmitted job.

Example: Resubmitting a MicroFocus Job from a Particular Step

You can restart a Micro Focus job from a particular step. This example restarts a Micro Focus job from step name STEP3 of the following JCL:

```
/INCSTEP JOB 'INC STEP PROGRAM',CLASS=A,MSGCLASS=A

//STEP1 EXEC PGM=IEFBR14,PARM='11'
//*SYSOUT DD SYSOUT=*
//INCST2 EXEC PGM=INCORSMP,PARM='000030'
//SYSOUT DD SYSOUT=*
//STEP3 EXEC PGM=SLEEPSMP,PARM='000030'
//SYSOUT DD SYSOUT=*
//
```

Note: To restart a Micro Focus job from a particular step, verify that JCL content is selected in the job definition. This command is not supported if you have JCL reference selected.

To resubmit a MicroFocus job from a particular step

1. Right-click the job in the Monitor perspective.
A list of job commands appears.
2. Select Resubmit.
3. Enter **STEP3** in the Restart from step field.

Note: To resubmit the job from the first step, leave this field blank.

The server resubmits the job from the specified step. If you made changes to the job's JCL or properties in the Monitor perspective, the server resubmits the job using these changes.

More information:

[Job Resubmission](#) (see page 67)

Ready a Job for Submission

You can ready a job for submission to drop the job's variable and predecessor dependencies, abandon the job's submission time (if it has one), and release the job from manual hold (if it is being held).

Note: If a job has resource dependencies, you must remove these dependencies separately if you want the job to run immediately.

To ready a job for submission

1. Right-click the job in the Monitor perspective.

A list of job commands appears.

2. Select Ready.

A confirmation dialog opens.

3. (Optional) Enter the command reason in the Reason field.

Limit: 1024 characters (including the timestamp and the user name)

Note: We recommend that the reason text not exceed 250 characters. The number of commands displayed in the command log is limited to the size of the command log buffer, which is 1024 characters. For example, if the reason text averages 10 characters per command, the command log displays up to 50 commands; if the reason text averages 250 characters per command, the command log displays up to 4 commands.

4. Click OK.

The server readies the job for submission. The job runs when its resource dependencies have been met, if any.

Note: When you ready a job, the graphical view of the Application changes in the Monitor perspective. The dependency lines between the job and its predecessors are removed from the view and the position of the job may also change.

View the Comments of a Job

You can view the comments of a job in an active Application to view additional information stored about the job.

To view the comments of a job

1. Right-click the job in the Monitor perspective.

A list of job commands appears.

2. Select Show Comments.

The comments of the active job are displayed.

Unwait a Job

You can unwait a job so that the job no longer waits for the same job to complete in a previous Application generation before it runs.

Note: You can only unwait jobs that are in the JANCWAIT state.

To unwait a job

1. Right-click the job in the Monitor perspective.

A list of job commands appears.

2. Select Unwait.

A confirmation dialog opens.

3. (Optional) Enter the command reason in the Reason field.

Limit: 1024 characters (including the timestamp and the user name)

Note: We recommend that the reason text not exceed 250 characters. The number of commands displayed in the command log is limited to the size of the command log buffer, which is 1024 characters. For example, if the reason text averages 10 characters per command, the command log displays up to 50 commands; if the reason text averages 250 characters per command, the command log displays up to 4 commands.

4. Click OK.

The server removes the JANCWAIT state. The job runs when all of its dependencies are met.

More information:

[Job Unwait](#) (see page 68)

Set a Job's User Status

You can set the user status for a job to communicate information to other users, notifying them of the reason for an action. For example, if you bypass a job, you can use the User status field to notify others of the reason for bypassing the job.

Note: You can access a job's user status using the WOB._userstatus built-in symbolic variable.

To set a job's user status

1. Right-click the job in the Monitor perspective.
A list of job commands appears.
2. Select Update User Status.
3. Enter the job's user status in the User status field.
The job's user status is set.

Drop Predecessor Dependencies

You can drop some or all of a job's predecessor dependencies before a job becomes eligible for submission if the job no longer needs to wait for its predecessors to complete, or you want the job to become eligible for submission sooner.

Note: You can drop a dependency only for a running Application.

To drop predecessor dependencies

1. Right-click the job in the Monitor perspective.
A list of job commands appears.
2. Select Drop Predecessors.
The Drop Predecessors dialog opens.
3. Do one of the following:
 - Select a predecessor dependency or hold the Ctrl key and select multiple predecessor dependencies in the Predecessors list box. Click Drop to drop the selected predecessor dependencies.
 - Click Drop All to drop all predecessor dependencies at once.

4. (Optional) Enter the command reason in the Reason field.

Limit: 1024 characters (including the timestamp and the user name)

Note: We recommend that the reason text not exceed 250 characters. The number of commands displayed in the command log is limited to the size of the command log buffer, which is 1024 characters. For example, if the reason text averages 10 characters per command, the command log displays up to 50 commands; if the reason text averages 250 characters per command, the command log displays up to 4 commands.

The appropriate predecessor dependencies are dropped.

Note: When you drop a job's predecessor dependencies, the graphical view of the Application changes in the Monitor perspective. The dependency lines between the job and its predecessors are removed from the view and the position of the job may also change.

More information:

[Predecessor Dependencies](#) (see page 68)

Reset Time Dependencies

Before the job runs, you can use the Reset Times command to display, remove, change, or add time dependencies for the current run.

To reset time dependencies

1. Right-click the job in the Monitor perspective.

A list of job commands appears.

2. Select Reset Times.

The Reset Times dialog opens.

3. Add or modify the date or time in the following fields:

Do not submit before

Specifies the job's submission time. The submission of the job is delayed until this time.

Overdue if not started by

Specifies the latest acceptable time by which the job must start before it is considered overdue (late start time).

Note: If the job is currently overdue because it is late starting, you cannot reset the late start time.

Overdue if not completed by

Specifies the time by which a job should complete successfully before it is considered overdue (late completion time).

Note: If the job is currently overdue because it is late completing or has exceeded its maximum execution time, you cannot reset the late completion time.

Abandon predecessor dependencies at

Submits a job without its predecessor dependencies once it meets a specified time.

Note: This field does not override a manual hold or submission time dependency.

Abandon submission at

Specifies the latest time this job can be submitted. If the job is not submitted by this time, it is bypassed.

Abandon resources at

Specifies the time at which you no longer want the job to wait for its resource dependencies.

Abandon variables at

Specifies the time at which you no longer want the job to wait for its variable dependencies.

Overdue if execution time exceeds

Specifies the maximum acceptable execution time in minutes. If the job exceeds this elapsed execution time, it is considered overdue.

Note: If the job is currently overdue because it is late completing or has exceeded its maximum execution time, you cannot reset the maximum execution time.

Premature if execution time is less than

Specifies the minimum acceptable execution time in minutes. If the job completes before this elapsed execution time, it is considered premend (premature end).

Reason

(Optional) Specifies the command reason.

Limit: 1024 characters (including the timestamp and the user name)

Note: We recommend that the reason text not exceed 250 characters. The number of commands displayed in the command log is limited to the size of the command log buffer, which is 1024 characters. For example, if the reason text averages 10 characters per command, the command log displays up to 50 commands; if the reason text averages 250 characters per command, the command log displays up to 4 commands.

4. Click OK.

The server resets the job's time dependencies.

More information:

[Time Dependencies](#) (see page 70)

[Overdue Jobs Monitoring](#) (see page 71)

Drop Variable Dependencies

You can drop some or all of a job's variable dependencies before a job becomes eligible for submission if the job no longer needs to wait for variable dependencies, or you want the job to become eligible for submission sooner.

Note: You can drop a dependency only for a running Application.

To drop variable dependencies

1. Right-click the job in the Monitor perspective.

A list of job commands appears.

2. Select Drop Variable Dependencies.

The Drop Variable Dependencies dialog opens.

3. Do one of the following:

- Select a variable dependency, or hold the Ctrl key and select multiple variable dependencies in the Variable Dependency Expressions table. Click Drop Selected to drop the selected variable dependencies.
- Click Drop All to drop all variable dependencies at once.

4. (Optional) Enter the command reason in the Reason field.

Limit: 1024 characters (including the timestamp and the user name)

Note: We recommend that the reason text not exceed 250 characters. The number of commands displayed in the command log is limited to the size of the command log buffer, which is 1024 characters. For example, if the reason text averages 10 characters per command, the command log displays up to 50 commands; if the reason text averages 250 characters per command, the command log displays up to 4 commands.

The appropriate variable dependencies are dropped.

More information:

[Global Variable Dependencies](#) (see page 72)

Evaluate Variable Dependencies

When a job is in a VARWAIT state, you can evaluate the job's variable dependencies. Evaluating the dependencies is helpful when you want to check the status of the job and the expressions that are waiting to be met before the job can be submitted.

To evaluate variable dependencies

1. Right-click the job that is in a VARWAIT state in the Monitor perspective.

A list of job commands appears.

2. Select Evaluate Variable Dependencies.

The Evaluate Variable Dependencies dialog opens. The job's variable dependencies are displayed in the Variable Expressions Evaluation Result table. A False value in the Evaluation column indicates that the variable expression in that row is waiting to be met.

3. Click Close.

The dialog closes.

More information:

[Global Variable Dependencies](#) (see page 72)

Drop Resource Dependencies

You can drop some or all of a job's resource dependencies when the job no longer must wait for all its resource dependencies to be met.

To drop resource dependencies

1. Right-click the job in the Monitor perspective.

A list of job commands appears.

2. Select Drop Resources.

The Drop Resources dialog opens.

3. Do *one* of the following:

- Select a single resource dependency or press the Ctrl key and select multiple resource dependencies in the Resources list box. Click Drop to drop the selected resource dependencies.
- Click Drop All to drop all resource dependencies at once.

4. (Optional) Enter the command reason in the Reason field.

Limit: 1024 characters (including the timestamp and the user name)

Note: We recommend that the reason text not exceed 250 characters. The number of commands displayed in the command log is limited to the size of the command log buffer, which is 1024 characters. For example, if the reason text averages 10 characters per command, the command log displays up to 50 commands; if the reason text averages 250 characters per command, the command log displays up to 4 commands.

The server drops the appropriate resource dependencies.

More information:

[How Jobs with Resource Dependencies are Submitted](#) (see page 73)

Display Resource Usage

You can view the resources the job requires and the jobs that currently use the required resources in the Monitor perspective.

Note: To view resource usage status, a job must be in RESWAIT state.

To view resource usage

1. Right-click the job in the Monitor perspective.
A list of job commands appears.
2. Select List Resource Usage.
The job's resource usage status is displayed.

More information:

[How Jobs with Resource Dependencies are Submitted](#) (see page 73)

Verify Whether a Job is Running or Has Executed

You can verify whether a job is running, has executed and completed, or has executed and failed.

To verify whether a job is running or has executed

1. Right-click the job in the Monitor perspective.

A list of job commands appears.

2. Select Process Verify.

The Process Verify dialog opens. The Process ID field displays the Job ID (this is the process ID for the job on the computer where the job runs).

Note: If you cannot see a Job ID, the job has not run through an EXEC state, so you cannot verify the selected process.

3. Leave the number that is displayed in the Process ID field to verify the selected job.

Note: You can also enter the Job ID of any other job on the agent computer to verify if that process is running or has executed. In the Monitor perspective, you can hover over the job you want to verify to see the Job ID.

4. Click OK.

One of the following Process status messages is displayed:

Unable to catch or not child

Indicates that the job has executed and completed or the job has executed and failed.

Process is child and running

Indicates that the job is running.

Note: To display server responses, from the main menu click Window, Preferences, Desktop Client, Monitor Perspective, Show server response, and click OK. The process status is also recorded in the tracelog.txt file and the agent transmitter.log file.

Set a Windows Job Object's Properties

You can modify the properties for a Windows job object that a running Windows job is associated with. The job must be in the EXEC state.

Note: After all of the processes associated with a job object complete, the job object no longer exists.

To set a Windows job object's properties

1. Right-click the job in the Monitor perspective.

A list of job commands appears.

2. Select Set Job Object.

The Set Job Object dialog opens.

3. Complete the following required field in the dialog as appropriate:

Job object name

Specifies the name of the existing Windows job object that you want to modify.

Note: The name you specify does not have to be the job object that the selected job is associated with. You can modify any existing Windows job object properties in the Set Job Object dialog by entering that Windows job object name.

4. (Optional) Specify the following additional information:

Job memory

Defines the maximum virtual memory in bytes allocated to *all* processes associated with the job object. If the total memory used for all processes associated with the job object exceeds this limit, the job that is trying to use memory fails.

Note: To specify memory in megabytes, add M after the memory value. To specify memory in kilobytes, add K after the memory value.

Examples: 50M, 500K

Process memory

Defines the maximum virtual memory in bytes allocated to *each* process associated with the job object. If the memory used for a single process exceeds this limit, the job fails.

Note: To specify memory in megabytes, add M after the memory value. To specify memory in kilobytes, add K after the memory value.

Examples: 50M, 500K

Job time

Defines the maximum CPU time in milliseconds allocated to *all* processes associated with the job object. If the total CPU time for all processes associated with the job object exceeds this limit, all jobs associated with the job object fail.

Process time

Defines the maximum CPU time in milliseconds allocated to *each* process associated with the job object. If the CPU time used for a single process exceeds this limit, the job fails.

Priority class

Indicates the process priority for *all* processes in the job object as follows:

- High—Indicates processes that must be executed immediately. These processes can use nearly all available CPU time.
- Above normal—Indicates processes that have priority above the Normal level, but below the High level.
- Normal—Indicates processes without special scheduling needs.
- Below normal—Indicates processes that have priority above the Idle level, but below the Normal level.
- Idle—Indicates processes that will run only when the system is idle.

Active process limit

Defines the maximum number of simultaneously active processes allowed in the job object.

Note: Changing the Active process limit does not affect jobs that are already running. For example, suppose that a Windows job object has its Active process limit set to three, and two jobs associated with the job object are already running. If you change the limit to

one, those two jobs will continue to run. However, if you insert a new Windows job into an Application and associate the new job with that Windows job object, the new job will not run. The job object has already reached its Active process limit of one.

Reason

(Optional) Specifies the command reason.

Limit: 1024 characters (including the timestamp and the user name)

Note: We recommend that the reason text not exceed 250 characters. The number of commands displayed in the command log is limited to the size of the command log buffer, which is 1024 characters. For example, if the reason text averages 10 characters per command, the command log displays up to 50 commands; if the reason text averages 250 characters per command, the command log displays up to 4 commands.

The Windows job object's properties are set.

Example: Modify a Windows Job Object's Process Memory

Suppose that a job named PROCJOB is executing and using 8 MB of memory. PROCJOB is associated with a Windows job object named ProcJobObject. ProcJobObject has its process memory set to 20 MB and its priority set to High.

As PROCJOB is executing, ProcJobObject's process memory is changed to 10 MB and its priority is changed to Idle. The process runs with a low priority, and when the process reaches 10 MB, the process ends.

To modify a Windows job object's process memory

1. Enter the following information in the Set Job Object dialog:
 - Job object name—ProcJobObject
 - Process Memory—10 M
2. Select Idle from the Priority drop-down list.
3. Click OK.

Example: Set the Process Priority for a Windows Job Associated with a Windows Job Object

Suppose that a Windows job is associated with the Windows job object named newJO, and is in the EXEC state. The priority for all processes in the newJO job object is increased to Above normal, but the other job object properties are not changed.

To set the process priority for a Windows job associated with a Windows job object

1. Enter **newJO** in the Job object name field of the Set Job Object dialog.
2. Select Above normal from the Priority drop-down list.
3. Click OK.

Set Process Priority for a Windows or a UNIX Job

You can change the process priority of a Windows or UNIX job. Process priority determines the order in which processes are scheduled on the processor. Depending on the priority level, process priority can speed up or slow down a process.

Note: If the Windows job is associated with a Windows job object, you can only change the process priority for that job using the Set Job Object command.

To set process priority for a Windows or a UNIX job

1. Right-click the Windows or UNIX job in the Monitor perspective.
A list of job commands appears.
2. Select Set Priority.
The Process Priority dialog opens.
3. Select one of the following process priority levels in the Priority field:

High

Indicates processes that must be executed immediately. These processes can use nearly all available CPU time.

Above normal

Indicates processes that have priority above the Normal level, but below the High level.

Normal

Indicates processes without special scheduling needs.

Below normal

Indicates processes that have priority above the Idle level, but below the Normal level.

Idle

Indicates processes that will run only when the system is idle.

Note: You can increase a UNIX job's process priority only if the job runs on a computer with the agent started by the root account. If the agent is not started by root and you select a higher process priority, the job runs with the current process priority (the priority is not increased) and an error message is recorded in the job's spool file.

4. (Optional) Enter the command reason in the Reason field.

Limit: 1024 characters (including the timestamp and the user name)

Note: We recommend that the reason text not exceed 250 characters. The number of commands displayed in the command log is limited to the size of the command log buffer, which is 1024 characters. For example, if the reason text averages 10 characters per command, the command log displays up to 50 commands; if the reason text averages 250 characters per command, the command log displays up to 4 commands.

5. Click OK.

The process priority for the Windows or the UNIX job is set.

Note: The Set Priority command does not set the process priority for subprocesses on Windows.

Send a Signal to a UNIX Job

You can send a signal to a UNIX job while it is executing. The signal can perform actions on the running process such as kill, suspend, or release.

To send a signal to a UNIX job

1. Right-click the UNIX job in the Monitor perspective.

A list of job commands appears.

Note: You can hold the Ctrl key and select multiple UNIX jobs.

2. Select Send Signal.

The Send Signal dialog opens.

3. Complete the following fields in the dialog as appropriate:

Signal

Specifies the signal number to send to the UNIX job.

Notes:

- Signal numbers can vary across platforms. For more information about the signal numbers, see the documentation for your UNIX platform.
- On most systems, signal 2 interrupts the process and signal 3 kills the process.

Reason

(Optional) Specifies the command reason.

Limit: 1024 characters (including the timestamp and the user name)

Note: We recommend that the reason text not exceed 250 characters. The number of commands displayed in the command log is limited to the size of the command log buffer, which is 1024 characters. For example, if the reason text averages 10 characters per command, the command log displays up to 50 commands; if the reason text averages 250 characters per command, the command log displays up to 4 commands.

4. Click OK.

The specified signal is sent to the UNIX job.

Note: You can verify if the signal was sent using the Job Command log.

More information:

[Display the Command Log of a Job](#) (see page 75)

Retrieve the Spool File for an i5/OS Job

When an i5/OS job completes, the i5/OS system stores the output in a spool file that is created from the job's job log. These spool files are native objects on the i5/OS system.

You can use the Retrieve Spool File command in the Monitor perspective to do the following tasks:

- View an i5/OS job's spool file list.
- Retrieve an entire spool file.
- Retrieve the last lines of a spool file.
- Retrieve part of a spool file that meets some filtering criteria.

To retrieve the spool file for an i5/OS job

1. Right-click the job in the Monitor perspective.
A list of job commands appear.
2. Select Retrieve Spool File.
The Retrieve Spool File dialog opens.
3. Select one of the available spool files, and click Retrieve Selected.
A second Retrieve Spool File dialog opens.
4. (Optional) Select the Selection Criteria option button and complete the following fields as appropriate:

Line Contains

Specifies a text string to search for.

Note: You cannot use wildcards in this field.

Last match

Retrieves the last match of the specified text string within the specified range.

Note: To retrieve all matches, clear this check box.

Start of range

Specifies the start point of the spool file portion to be retrieved in lines, bytes, kilobytes, or megabytes.

End of range

Specifies the end point of the spool file portion to be retrieved in lines, bytes, kilobytes, or megabytes.

5. (Optional) Select the Line numbers option button and complete the following field:

Last number of lines to display

Specifies the number of lines to retrieve from the end of the spool file.

6. (Optional) Select the Append to existing display check box to add the results returned to the existing display.

Note: This option can be useful if you retrieve the spool file in segments.

7. Click one of the following options:

Retrieve Selected

Retrieves a selection of the spool based on criteria you specify. You can specify a text string to search for within a range or specify the last number of lines to display.

Retrieve Next

Retrieves the entire spool file in segments.

Note: Use this option if you are retrieving large spool files or you are unsure of the spool file size.

Retrieve All

Retrieves the entire spool file.

Note: If you retrieve a large spool file, this option can slow down the server.

The server displays the spool file of the job.

Respond to Suspended Jobs that Require Manual Intervention

An i5/OS job that appears in an INTVRQ state requires manual intervention. In this state, the job is temporarily suspended and waiting for a user reply to continue its execution. For example, suppose that you schedule a job to save data in a SAV file. If the SAV file already contains data, the job prompts you to confirm that the data in the SAV file can be overwritten.

To respond to suspended jobs that require manual intervention

1. Right-click the job in the Monitor perspective.

A list of job commands appear.

2. Select Reply.

The Reply dialog opens.

3. Check the Message field for the possible responses to enter.

Note: The Message field displays message text from the i5/OS system.

4. Enter a text string in the Reply field.

5. (Optional) Enter the command reason in the Reason field.

Limit: 1024 characters (including the timestamp and the user name)

Note: We recommend that the reason text not exceed 250 characters. The number of commands displayed in the command log is limited to the size of the command log buffer, which is 1024 characters. For example, if the reason text averages 10 characters per command, the command log displays up to 50 commands; if the reason text averages 250 characters per command, the command log displays up to 4 commands.

6. Click OK.

The job resumes execution.

Retrieve the JES Spool File of a Micro Focus Job

You can display the JES spool file of a Micro Focus job to view the output of the job. You can view the entire spool file, selected lines, or a number of lines you specify.

Note: To retrieve the JES spool file of a Micro Focus job, your system requires the CA WA Agent for UNIX, Linux, or Windows and the CA WA Agent for Micro Focus.

To retrieve the JES spool file of a Micro Focus job

1. Right-click the Micro Focus job in the Monitor perspective.
A list of job commands appears.
2. Select Retrieve JES Spool File.
The Retrieve JES Spool File dialog opens.
3. (Optional) Select the Selection Criteria option button and complete the following fields as appropriate:

Line Contains

Specifies a text string to search for.

Note: You cannot use wildcards in this field.

Start of range

Specifies the start point of the spool file portion to be retrieved in lines, bytes, kilobytes, or megabytes.

End of range

Specifies the end point of the spool file portion to be retrieved in lines, bytes, kilobytes, or megabytes.

Lines

Retrieves the specified spool file portion by lines.

Kbytes

Retrieves the specified spool file portion by kilobytes.

Mbytes

Retrieves the specified spool file portion by megabytes.

Last match

Retrieves the last match of the specified text string within the specified range.

Note: To retrieve all matches, clear this check box.

4. (Optional) Select the Line numbers option button and complete the following field:

Last number of lines to display

Specifies the number of lines to retrieve from the end of the spool file.

5. (Optional) Select the Append to existing display check box to add the results returned to the existing display.

Note: This option can be useful if you retrieve the spool file in segments.

6. (Optional) Enter the sequence number of the data definition to retrieve the output of a particular data definition (DD).

Note: To display a list of available data definitions, leave this field blank.

7. Click one of the following options:

Retrieve Selected

Retrieves a selection of the spool based on criteria you specify. You can specify a text string to search for within a range or specify the last number of lines to display.

Retrieve Next

Retrieves the entire spool file in segments.

Note: Use this option if you are retrieving large spool files or you are unsure of the spool file size.

Retrieve All

Retrieves the entire spool file.

Note: If you retrieve a large spool file, this option can slow down the server.

Example: Data Definitions for a JCL

The following lists the data definitions associated with a particular JCL:

SEQ#	DDNAME	STEPNAME	PROCSTEP	NUM	RECS	T	RCFM	LRECL
00001	JESYSMSG			00000246	A	LSEQ	00132	
00002	SYSPRINT	STEP00		00000034	A	F	00121	
00003	SYSPRINT	STEP01		00000001	A	F	00133	
00004	SYSPRINT	STEP10		00000010	A	F	00121	
00005	SYSPRINT	STEP11		00000011	A	F	00121	
00006	SYSPRINT	STEP12		00000022	A	F	00121	
00007	SYSPRINT	STEP20		00000010	A	F	00121	
00008	SYSPRINT	STEP30		00000010	A	F	00121	
00009	SYSPRINT	STEP40		00000006	A	F	00121	
00010	SYSPRINT	STEP50		00000006	A	F	00121	
00011	SYSPRINT	STEP60		00000006	A	F	00121	

Browse JCL

After JES processes a z/OS job, you can browse JCL to view the JCL of a z/OS job.

Note: To browse JCL, your system requires CA WA Agent for z/OS.

To browse JCL

1. Right-click the z/OS job in the Monitor perspective.

A list of job commands appears.

2. Select Browse JCL.

The details of the JCL are displayed.

Note: You can right-click in the text display area to display a popup menu that has Copy, Select All, and Find commands. CA WA Desktop Client also supports the standard keyboard shortcuts such as Ctrl+C.

Browse Copy JCL

After JES processes a z/OS job, you can browse copy JCL to view the copy JCL of a z/OS job.

Note: To browse copy JCL, your system requires CA WA Agent for z/OS.

To browse copy JCL

1. Right-click the z/OS job in the Monitor perspective.

A list of job commands appears.

2. Select Copy JCL.

The details of the copy JCL are displayed.

Note: You can right-click in the text display area to display a popup menu that has Copy, Select All, and Find commands. CA WA Desktop Client also supports the standard keyboard shortcuts such as Ctrl+C.

Edit JCL

After JES processes a z/OS job, you can edit JCL to update the JCL of a z/OS job.

Note: To edit JCL, your system requires CA WA Agent for z/OS.

To edit JCL

1. Right-click the z/OS job in the Monitor perspective.

A list of job commands appears.

2. Select Edit JCL.

The Edit JCL dialog opens.

3. Modify the JCL as required.

Note: You can right-click in the text editing area to display a popup menu that has Undo, Redo, Copy, Cut, Paste, Select All, and Find commands. CA WA Desktop Client also supports the standard keyboard shortcuts such as Ctrl+C.

4. Click Upload to save your changes to the mainframe.

The JCL is edited.

Edit Copy JCL

After JES processes a z/OS job, you can edit copy JCL to update the copy JCL of a z/OS job. For example, in case of a failure, you can edit the copy JCL, update the JCL library name to the copy JCL library name in the job definition, and resubmit the job without affecting the JCL source.

Note: To edit copy JCL, your system requires CA WA Agent for z/OS.

To edit copy JCL

1. Right-click the z/OS job in the Monitor perspective.

A list of job commands appears.

2. Select Edit Copy JCL.

The Edit Copy JCL dialog opens.

3. Modify the copy JCL as required.

Note: You can right-click in the text editing area to display a popup menu that has Undo, Redo, Copy, Cut, Paste, Select All, and Find commands. CA WA Desktop Client also supports the standard keyboard shortcuts such as Ctrl+C.

4. Click Upload to save your changes to the mainframe.

The copy JCL is edited.

z/OS Jobs Restart Using CA WA Restart Option

CA WA Restart Option is an advanced rerun and restart product that works with CA Workload Automation DE to restart batch jobs. CA WA Restart Option recommends the restart point in failed jobs, can make the necessary adjustments to batch JCL, and can perform the necessary data set cleanups for an error-free restart.

Note: To restart z/OS jobs using CA WA Restart Option, your system requires CA WA Agent for z/OS and CA WA Restart Option.

When restarting z/OS-Regular jobs using CA WA Restart Option, you can select the steps to restart or let CA WA Restart Option choose which steps to run on restart.

You can restart a job from the first step, restart a single step or multiple steps, restart all steps starting from a selected step, or exclude a range of steps from running on restart. If you select a restart step, CA WA Restart Option can report on the restart. CA WA Restart Option predicts errors for the restart and identifies if bypassing a step is acceptable.

If you do not select a restart step, CA WA Restart Option can recommend the steps to run on restart and list the cleanup actions it will perform on the data sets to ensure an error-free restart. If CA WA Restart Option detects a severe error, it explains the nature of the error and indicates how you can proceed to correct the error.

Select the Steps to Restart

When restarting z/OS-Regular jobs using CA WA Restart Option, you can select the steps to restart and ask CA WA Restart Option to report on the restart.

To select the steps to restart

1. Right-click the z/OS-Regular job you want to restart in the Monitor perspective.

A list of job commands appears.

2. Select Restart.

The Step Summary page of the Restart Option EE dialog displays the steps in the job and the data sets used by those steps.

Note: To display the data definitions (DD) associated with the step name, select a step. To display the data sets associated with the DD name, select the DD name.

3. Do *one* of the following:
 - To restart all steps, select the first step, and select the all steps starting with the step selected option button.
 - To restart a single step, select the step you want to restart, and select the only the steps selected option button.
 - To restart multiple steps, hold down the Ctrl key, select the steps you want to restart, and select the only the steps selected option button.
 - To restart all steps starting from a selected step, select the step you want to resubmit the job from, and select the all steps starting with the step selected option button.
4. Click Restart Analysis.

The Restart Option EE Action Summary dialog displays a report on the selected restart. The Restart Option EE has predicted the following errors section displays any errors that CA WA Restart Option predicts for the restart. The Actions to be taken by Restart Option EE section lists the actions CA WA Restart Option takes on the restart. If CA WA Restart Option needs to run an additional step to recreate a required data set, it would be indicated.

5. Click Resubmit.

CA WA Restart Option resubmits the job based on the steps you selected.

Let CA WA Restart Option Select the Steps to Restart

When restarting z/OS-Regular jobs using CA WA Restart Option, you can let CA WA Restart Option select the steps to restart.

To let CA WA Restart Option select the steps to restart

1. Right-click the z/OS-Regular job you want to restart in the Monitor perspective.
A list of job commands appears.

2. Select Restart.

The Step Summary page of the Restart Option EE dialog displays the steps in the job and the data sets used by those steps.

3. Select the let Restart Option EE choose option button.
4. Click Restart Analysis.

The Restart Option EE Restart Action Summary dialog displays the recommendations of CA WA Restart Option. The Actions to be taken by Restart Option EE section recommends a restart point and lists the actions it will perform to ensure an error-free restart.

5. (Optional) View a summary of recommendations by step by doing the following:

- a. Click Restart Step Summary.

The Restart Option EE Step Summary dialog lists the steps in the job, and reveals the steps CA WA Restart Option restarts and the steps it bypasses. For each step being restarted, CA WA Restart Option supplies a reason.

Note: For more information about a particular step, you can double-click the step.

- b. Click Exit after reading the summary.

The Restart Option EE Step Summary dialog closes.

6. (Optional) View a list of actions that CA WA Restart Option must perform on data sets by doing the following:

- a. Click Cleanup.

The Restart Option EE Cleanup Action Summary dialog opens.

- b. Click Cleanup Summary.

- c. Click Exit after reading the summary.

The Restart Option EE Cleanup Summary dialog closes.

7. (Optional) View the error log of CA WA Restart Option by doing the following:

- a. Click Restart Option EE Log.

The Restart Option EE Log dialog displays information when CA WA Restart Option detects a severe error. If CA WA Restart Option detects an error, the dialog indicates how you can proceed to correct the error.

- b. Click Exit after reading the log.

The Restart Option EE Log dialog closes.

8. Click Resubmit.

CA WA Restart Option resubmits the job based on its recommendations.

Specify Additional Restart Options

When restarting z/OS-Regular jobs using CA WA Restart Option, you can specify the name of the data set and member containing the JCL you want to submit, a range of steps to run on restart, or a list of steps to exclude on restart.

To specify additional restart options

1. Right-click the z/OS-Regular job you want to restart in the Monitor perspective.

A list of job commands appears.

2. Select Restart.

The Step Summary page of the Restart Option EE dialog displays the steps in the job and the data sets used by those steps.

3. Click Restart in the left pane.

The Restart page opens in the right pane.

4. Complete the fields as required.

5. (Optional) Click Restart Analysis to display the restart analysis.

6. Click Resubmit to resubmit the job.

CA WA Restart Option resubmits the job using the additional restart options you specified.

Specify CA WA Restart Option Statements

You can specify predefined and custom CA WA Restart Option statements to specify special processing options.

For example, you can instruct CA WA Restart Option to do the following:

- Back out data sets created by the job.
- Perform data set cleanup.
- Run the job despite errors predicted by CA WA Restart Option.
- Check for errors prior to a job's submission.
- Perform automatic recovery of missing data sets.
- Honor previously coded condition codes and bypass steps accordingly.

To specify CA WA Restart Option statements

1. Right-click the z/OS-Regular job you want to restart in the Monitor perspective.
A list of job commands appears.
2. Select Restart.
The Step Summary page of the Restart Option EE dialog displays the steps in the job and the data sets used by those steps.
3. Click Restart Option EE in the left pane.
The Restart Option EE page opens in the right pane.
4. Select the predefined statements you want CA WA Restart Option to send and the parameter to pass with each.
Note: These statements depend on the defaults set in the initialization parameters on the CA WA Agent for z/OS. For more information about the initialization parameters, contact your agent administrator.
5. (Optional) Enter your own CA WA Restart Option statements starting at line 7.
Note: For each statement, select the line in the Send column and enter the statement. Do not prefix the statement with ENCPARM. You can include up to 12 statements in total.
6. (Optional) Click Restart Analysis to display the restart analysis.
7. Click Resubmit.
CA WA Restart Option resubmits the job using the specify processing options you specified.

Issue a Database Command

You can issue a Database command against a Database job to monitor or control the job.

To issue a Database command

1. Right-click the Database job in the Monitor perspective.

A list of job commands appears.

2. Select one of the following Database commands:

Cancel Request

Sends a request to the database to cancel an SQL or DB Stored Procedure job.

Note: The job must be in the EXEC state. The agent sends the cancel request to the database, but the database may not cancel the job if the job is already processing.

Delete Monitor

Stops monitoring the database table for the changes specified in the DB Monitor job.

Note: The job must be in the EXEC or MONITOR state.

Delete Trigger

Stops monitoring the database table for the changes specified in the DB Trigger job.

Note: The job must be in the EXEC or MONITOR state.

A confirmation dialog opens.

3. (Optional) Enter the command reason in the Reason field.

Limit: 1024 characters (including the timestamp and the user name)

Note: We recommend that the reason text not exceed 250 characters. The number of commands displayed in the command log is limited to the size of the command log buffer, which is 1024 characters. For example, if the reason text averages 10 characters per command, the command log displays up to 50 commands; if the reason text averages 250 characters per command, the command log displays up to 4 commands.

The server issues the command against the job.

SAP Commands

You can issue SAP commands against the following SAP job types:

SAP-R3

Schedules an SAP R/3 job on your SAP system.

SAP-Job Copy

Copies an existing SAP R/3 job.

SAP-Batch Input Session

Imports data from external systems to the SAP system.

SAP-BW Info Package

Transfers data from a data source to an SAP Business Warehouse system.

SAP-BW Process Chain

Creates Process Chains on the SAP system.

SAP-Data Archiving

Stores information in an SAP Archiving Object.

SAP-Process Monitor

Monitors for a specific SAP process status.

Issue an SAP Command Against an SAP-R3 or SAP-Job Copy Job

You can issue an SAP command against an SAP-R3 job or SAP-Job Copy job to monitor or control the job.

To issue an SAP command against an SAP-R3 or SAP-Job Copy job

1. Right-click the SAP-R3 or SAP-Job Copy job in the Monitor perspective.
A list of job commands appears.
2. Select SAP commands.
The SAP commands sub-menu appears.
3. Select one of the following commands:

JobCancel

Cancels a running SAP-R3 or SAP-Job Copy job. If the job has not completed successfully, the server marks the job as failed.

JobCopy

Creates a copy of an SAP-R3 or SAP-Job Copy job on the SAP system.

JobDelete

Deletes an SAP-R3 job from the SAP system.

JobModify

Lets you modify an ABAP step parameter. You can modify a job only if it is scheduled in the SAP system, and has not started. The parameters displayed on the JobModify dialog are based on the original job definition. You can only modify a step. You cannot add a step.

Note: This command does not apply to SAP-Job Copy jobs.

Job Start

Starts an SAP-R3 or SAP-Job Copy job that is waiting for a manual start. This command has the following options:

- Start ASAP—Releases the job as soon as possible after other scheduled jobs complete. This is the default.
- Start immediately—Releases the job immediately.

Set Job Class

Lets you to modify the SAP class of an SAP-R3 or SAP-Job Copy job.

Job Children Get

Lists the details of the SAP-R3 or SAP-Job Copy job's children. Children are jobs spawned by a parent job.

JobStatusCheck

Checks the SAP-R3 or SAP-Job Copy job status recorded in the SAP system database and the actual job status. If the server finds any discrepancies between the job status recorded in the SAP system and the actual job status, it corrects the status and updates the database.

Note: The agent has an automatic monitoring component that periodically monitors the status of an SAP job. The JobStatusCheck command returns the same status as the automatic monitoring component.

JobStatusGet

Displays the SAP-R3 or SAP-Job Copy job status recorded in the SAP system database.

Monitor Children

Monitors an SAP-R3 or SAP-Job Copy job's children. Children are jobs spawned by a parent job. If the job has children, the Monitor perspective lists the following information for each child job:

- Job name
- Job count
- Status
- Parent-child relation

Note: A child job can spawn children jobs and become a parent job as well.

Job Get Spool List

Displays an image of output that the selected ABAP produced, according to the print parameters of the SAP-R3 or SAP-Job Copy job. The Step number field is used to display a particular step.

Note: CA WA Desktop Client may need to connect to an FTP site, depending on the amount of data retrieved. To do that, CA WA Desktop Client needs a user ID and password for the FTP site. If the information has changed, you may be prompted for it.

Job Get Dump

Displays the dump of the SAP-R3 or SAP-Job Copy job that is produced on the SAP system when the job fails. If the selected job has a dump available, the dump appears in the Job Get Dump dialog. You can scroll through the dump or copy and paste it into a text document.

Note: To use this command, the agent administrator must set the `sap.job.dump_ref.enabled` parameter in the `agentparm.txt` file to true.

Job Header Read

Displays the header of the SAP-R3 or SAP-Job Copy job.

Job Read

Displays the header and step of the SAP-R3 or SAP-Job Copy job.

JobDefGet

Displays the SAP-R3 or SAP-Job Copy job definition.

Note: The command does not return values for the OperSys cover page (PRUNX) and print priority (PRIOT) parameters.

Job Log Read

Displays the SAP system job log.

Get Global Audit

Initiates a query of the current global audit-level settings of all SAP-R3 or SAP-Job Copy jobs.

Get Job Count

Retrieves the number of SAP-R3 or SAP-Job Copy jobs with the same job name. The Filter field is used to set the criteria you want to use to filter the jobs that run on the SAP system, for example, AFTER*.

Job Select

Selects a set of SAP-R3 or SAP-Job Copy jobs in the SAP system that match a specific selection criteria.

Set Global Audit

Sets the global audit level for an SAP-R3 or SAP-Job Copy job. The Level(0,1,2,3,D) field is used to set the level.

Variant Read

Displays variant values for all ABAP steps in an SAP-R3 or SAP-Job Copy job. The ABAP name field is used to set the ABAP name.

4. Complete the fields in the dialog as appropriate, and click OK.

The server issues the command against the job.

JobCancel Command

You can use the JobCancel command to cancel a running SAP job. If the job has not completed successfully, the server marks the job as failed.

Note: The agent parameter file (agentparm.txt) determines if the server also cancels children jobs for a parent job. Children are jobs spawned by a parent job.

For more information about the following agent properties, contact your agent administrator:

- sap.job.children.cancel
- sap.job.children.monitor
- sap.job.children.recursive

JobCopy Command

You can use the JobCopy command to create a copy of an SAP job on the SAP system. The server creates a copy of the job on the SAP system with a new job count (job number). You can copy an entire SAP job or copy the job starting from a certain step until the end of the job. You can also use the JobCopy command to rerun an entire job or, if an ABAP failed, rerun the job from the step following the failed step.

If you leave the Step number field blank, the Monitor perspective copies the entire job. If you leave the Target job name field blank, the copy job has the same job name as the original job, and the server gives the copy job a new job count (job number).

Note: The new job is in the SCHEDULED state on the SAP system. You can start the job in SAP GUI.

JobDelete Command

You can use the JobDelete command to delete an SAP-R3 or SAP-Job Copy job from the SAP system.

You can delete a job in the following SAP states:

- Canceled
- Completed
- Finished
- Released
- Scheduled

If the agent is configured to delete children jobs, you must ensure that all children are inactive before you issue the JobDelete command. If you delete a parent job while there are still outstanding children jobs, you cannot access the children from the Monitor perspective. To ensure that all children are canceled before issuing the delete command, you must issue the JobCancel command for each child. You can only cancel active jobs.

The agent parameter file (agentparm.txt) determines if the server also deletes children for a parent job. You can contact your agent administrator for details about the following agent properties:

- sap.job.children.delete
- sap.job.children.monitor
- sap.job.children.recursive

If the job you are deleting has children, ensure all of the children are in one of the listed states before you delete the job. If you delete a job parent when there are still outstanding children jobs, the children cannot be accessed from the Monitor perspective. In this case, you need to delete the outstanding children from the SAP system.

If the agent is configured to monitor the children, a job with children is marked FAILED under the following conditions:

- A parent job is deleted while its children are still active.
- A parent job is finished and any child job was terminated.
- A parent job is finished and any child job was deleted.

Note: When an SAP job is in the COMPLETE state, the parent job and all its children are complete.

JobStatusGet Command

You can use the JobStatusGet command to display the job status recorded in the SAP system database.

The following table displays a list of SAP job statuses and the equivalent CA Workload Automation DE job states:

SAP Job Status	CA Workload Automation DE Job State
ACTIVE	EXEC
FINISHED	COMPLETE
READY	SUBMIT
RELEASED	SUBMIT
SCHEDULED	DEFINED
TERMINATED	FAILED

Job Select Command

You can use the Job Select command to select a set of jobs in the SAP system that matches a specific selection criteria.

You can use the following selection criteria in this command:

Aborted

Indicates jobs in the ABORTED state.

Client number

(Optional) Specifies the client number used for the Remote Function Call (RFC) connection.

Limit: 3 digits

Example: 800

Event ID

(Optional) Specifies the Event ID.

Limit: 32 characters

Event param

(Optional) Specifies the Event parameter.

Limit: 65 characters

Finished

Indicates jobs in the FINISHED state.

From

Specifies the planned start date and time.

Job count

(Optional) Specifies the batch job number.

Limit: 8 characters

Job group

(Optional) Specifies the name of a Job Group.

Limit: 12 characters

Job name

Specifies the job name.

Limit: 32 characters

Note: You can use wildcards to select multiple job names containing the defined string. For example, *string* will select all job names containing the characters defined as string.

Job user

Specifies the owner of the job. This is the user on the job tab.

Limit: 12 characters

No date

Indicates jobs without a start date.

Ready

Indicates jobs in the READY state.

Released

Indicates jobs in the RELEASED state.

RFC dest

(Optional) Specifies the three-character SAP destination (SAP system ID).

Default: Destination specified during the agent installation

Example: cyb

Running

Indicates jobs in the RUNNING state.

SAP user

(Optional) Specifies the user ID for the RFC connection.

Limit: 12 characters

Scheduled

Indicates jobs in the SCHEDULED state.

To

(Optional) Specifies the planned end date and time.

With predecessors

Indicates jobs that start after predecessors.

Note: CA WA Desktop Client may need to connect to an FTP site, depending on the amount of data retrieved. To do this, CA WA Desktop Client needs a user ID and password for the FTP site. If the information has changed, you may be prompted for it.

Issue an SAP Command Against an SAP-Batch Input Session Job

You can issue an SAP command against an SAP-Batch Input Session job to monitor or control the job.

To issue an SAP command against an SAP-Batch Input Session job

1. Right-click the SAP-Batch Input Session job in the Monitor perspective.

A list of job commands appears.

2. Select SAP commands.

The SAP commands sub-menu appears.

3. Select one of the following commands:

Cancel

Cancels an active SAP-Batch Input Session job on the SAP system. If the job has not completed successfully, the server marks the job as failed.

Delete

Deletes an SAP-Batch Input Session job from the SAP system.

Details

Displays the status of an SAP-Batch Input Session job.

Start

Starts an SAP-Batch Input Session job that is waiting for a manual start. This command has the following options:

- Start immediately—Releases the job immediately.
- Start ASAP—Releases the job as soon as possible after other scheduled jobs complete. This is the default.

4. Complete the fields in the dialog as appropriate, and click OK.

The server issues the command against the job.

Cancel Command

You can use the Cancel command to cancel an active SAP-Batch Input Session job on the SAP system. If the job has not completed successfully, the server marks the job as failed.

SAP-Batch Input Session (BDC) jobs run in the following two phases:

- ABAP runs and creates a BDC session.
- The system processes the created BDC session.

Before confirming an SAP-Batch Input Session job cancellation, you must check the Cancel dialog to determine whether you are canceling the ABAP or the BDC session. The following situations may take place:

- If the Cancel dialog has no job number in the BDC Queue Id and no BDC name, the Cancel command cancels the ABAP.
- If the Cancel dialog has a job number in the BDC Queue Id field and a name in the BDC name field, the Cancel command cancels the BDC session.

Delete Command

You can use the Delete command to delete an SAP-Batch Input Session job from the SAP system.

Note: You can issue the Delete command only when the SAP-Batch Input Session job is in a TERMINATED or COMPLETE state. If the SAP-Batch Input Session job is active, you can use the Cancel command before proceeding with the delete.

SAP-Batch Input Session (BDC) jobs run in the following two phases:

- ABAP runs and creates a BDC session.
- The system processes the BDC session.

Before confirming an SAP-Batch Input Session job deletion, you must check the Delete dialog to determine whether you are deleting the job that ran the ABAP or the BDC session.

The following situations may take place:

- If the Delete dialog has a name in the Job name field and a job number in the Job ID field, the Delete command deletes the job that ran the ABAP specified in the SAP-Batch Input Session job definition.
- If the Job Delete dialog has a name in the BDC Name field and a job number in the BDC Queue Id field, the Delete command deletes the BDC session.

Details Command

You can use the Details command to display the status of a BDC job.

You can scroll down to the Status field to see the BDC job's overall status. The values in the Status field describe the following BDC job's progress:

- BDC Session Created
- BDC Submit Completed
- BDC Session Active
- BDC Session Finished
- BDC Session Terminated

If the final job status for a job appears with an asterisk (for example, BDC Session Terminated*), the Minimum Processed Rate and Maximum Error Rate values in the job definition details modified the job status. In this case, the values in the LStatus field describe the following transaction summary for the session:

- Transactions (total)
- ErrorRate (actual, acceptable value from job details)
- ProcessedRate (actual, acceptable value from job details)

Issue an SAP Command Against an SAP-BW Info Package Job

You can issue an SAP command against an SAP-BW Info Package job to monitor or control the job.

To issue an SAP command against an SAP-BW Info Package job

1. Right-click the SAP-BW Info Package job in the Monitor perspective.

A list of job commands appears.

2. Select SAP commands.

The SAP commands sub-menu appears.

3. Select one of the following commands:

Stop

Stops an SAP-BW Info Package job from running. If the job has not completed successfully, the server marks the job as failed.

Info Pack Status

Displays the status of an SAP-BW Info Package job.

Info Pack Details

Displays the details of a BW Info Package.

4. Complete the fields in the dialog as appropriate, and click OK.

The server issues the command against the job.

Info Pack Status Command

You can use the Info Pack Status command to display the status of a BW Info Package job.

The following table lists the Info Package statuses:

Status Parameter	Value
LOCATION	Current processing step
QUALINFO	Pertains to the QUALSTATUS status. Descriptive informational messages on manually set jobs.
QUALSTATUS	Manually set status of this Info Package run ■ G-green (completed) ■ Y-yellow (running) ■ R-red (failed) Note: If this field is blank, the Info Package status is the same as the status shown in the TECHSTATUS field.
TECHINFO	Informational messages pertaining to the TECHSTATUS, such as, "All Docs processed successfully" or "Correct data request; processing running"
TECHSTATUS	Technical status of this Info Package run ■ G-green (Completed) ■ Y-yellow (Running) ■ R-red (Failed)

Issue an SAP Command Against an SAP-BW Process Chain Job

You can issue an SAP command against an SAP-BW Process Chain job to monitor or control the job.

To issue an SAP command against an SAP-BW Process Chain job

1. Right-click the SAP-BW Process Chain job in the Monitor perspective.

A list of job commands appears.

2. Select SAP commands.

The SAP commands sub-menu appears.

3. Select one of the following commands:

Get Chain Status

Displays the status of a BW Process Chain.

List Chain Processes

Lists the processes of a BW Process Chain.

Get Process Log

Displays the process log of a BW Process Chain. In the List Chain Processes dialog, you can select the chain and click Get Process Logs.

Note: Before using this command, you must first issue the List Chain Processes command against the SAP-BW Process Chain job.

Get Message Details

Displays the message details of a BW Process Chain. In the Get Message Details dialog, you can specify the Message ID, Message number, and Message format fields.

Interrupt Chain

Interrupts a BW Process Chain.

Restart Chain

Restarts a BW Process Chain.

Get Chain Log

Displays the log of an SAP-BW Process Chain job.

Note: Because there may be numerous processes in an SAP-BW Process Chain job, the response may take several minutes.

4. Complete the fields in the dialog as appropriate, and click OK.

The server issues the command against the job.

Issue an SAP Command Against an SAP-Data Archiving Job

You can issue an SAP command against an SAP-Data Archiving job to monitor and control the job.

To issue an SAP command against an SAP-Data Archiving Job

1. Right-click the SAP-Data Archiving job in the Monitor perspective.

A list of job commands appears.

2. Select SAP commands.

The SAP commands sub-menu appears.

3. Select one of the following commands:

Cancel

Cancels an active SAP-Data Archiving job. If the job has not completed successfully, the server marks the job as failed.

Delete Job

Deletes an SAP-Data Archiving job from the SAP system. You can delete an SAP-Data Archiving job in the following SAP states: Scheduled, Released, Ready, Finished, and Cancelled.

Monitor Children

Monitors an SAP-Data Archiving job's children. Children are jobs spawned by a parent job.

4. (Optional) Complete the fields in the dialog as appropriate, and click OK.

The server issues the command against the job.

Issue an SAP Command Against an SAP-Process Monitor Job

You can issue an SAP command against an SAP-Process Monitor job to display the process that is being monitored.

To issue an SAP command against an SAP-Process Monitor job

1. Right-click the SAP-Process Monitor job in the Monitor perspective.

A list of job commands appears.

2. Select SAP commands.

The SAP commands sub-menu appears.

3. Select the Monitor Processes command to display the process an SAP-Process Monitor job is monitoring.

The server issues the command against the job.

Monitor Processes Command

You can use the Monitor Processes command to display the process an SAP-Process Monitor job is monitoring.

If the SAP-Process Monitor job is defined as a continuous monitor (an Alert was specified on the Basic dialog), the processes that meet the criteria specified on the Basic dialog appear in the Monitor Processes dialog.

Note: Only the most recent match appears in the Monitor Processes dialog.

SAP-Process Monitor jobs defined as continuous stay in the MONITOR state in the Application Monitor view until they are forced complete.

If the SAP-Process Monitor job is defined as non-continuous (no Alert was specified on the Basic dialog), the first process that matches the criteria on the Basic dialog appears in the Monitor Processes dialog. The Monitor Processes dialog remains empty until a match occurs.

SAP-Process Monitor jobs defined as non-continuous stay in the MONITOR state in the Application Monitor view until a match is detected. When the server detects a match, it marks the SAP-Process Monitor job complete.

Issue a PeopleSoft Command

You can issue a PeopleSoft command against a PeopleSoft job to monitor or control the job.

To issue a PeopleSoft command

1. Right-click the PeopleSoft job in the Monitor perspective.

A list of job commands appears.

2. Select one of the following PeopleSoft commands:

Cancel Job

Cancels a running PeopleSoft job. If the job has not completed successfully, the server marks the job as failed.

Delete Job

Deletes a PeopleSoft job.

Note: You cannot delete a PeopleSoft job in a SUBERROR state.

Hold Job

Places a submitted or running PeopleSoft job on hold.

Restart Job

Restarts a held or failed PeopleSoft job using PeopleSoft.

Retrieve Trace File

Retrieves the trace file that stores the output for a PeopleSoft job.

A confirmation dialog opens.

3. (Optional) Enter the command reason in the Reason field.

Limit: 1024 characters (including the timestamp and the user name)

Note: We recommend that the reason text not exceed 250 characters. The number of commands displayed in the command log is limited to the size of the command log buffer, which is 1024 characters. For example, if the reason text averages 10 characters per command, the command log displays up to 50 commands; if the reason text averages 250 characters per command, the command log displays up to 4 commands. The Retrieve Trace File command does not support the Reason field.

The server issues the command against the job.

Issue an Oracle E-Business Suite Command

You can issue an Oracle E-Business Suite command against an Oracle E-Business Suite job to monitor or control the job.

To issue an Oracle E-Business Suite command

1. Right-click the Oracle E-Business Suite job you want to issue the command against in the Monitor perspective.

A list of job commands appears.

2. Select one of the following Oracle E-Business Suite commands:

Cancel Request

Cancels an active Oracle E-Business Suite job. If the job has not completed successfully, the server marks the job as failed.

Hold Request

Holds an active Oracle E-Business Suite job.

Remove Hold Request

Removes a hold on an Oracle E-Business Suite job.

View Completion Details

Displays the completion details of an Oracle E-Business Suite job.

View Log File

Retrieves the log of an Oracle E-Business Suite job.

View Output File

Displays the output of an Oracle E-Business Suite job.

Note: Some Oracle E-Business Suite jobs generate empty output files or no output files at all.

The server issues the command against the job.

Chapter 5: Working with Custom Views

This section contains the following topics:

- [Custom Views](#) (see page 135)
- [Open a Prepackaged Custom View](#) (see page 138)
- [Create a New Text-Based Custom View](#) (see page 139)
- [Copy an Existing Custom View](#) (see page 140)
- [Create a Custom View Filter](#) (see page 141)

Custom Views

CA Workload Automation DE provides prepackaged custom views for monitoring and controlling your workload. You can also create your own text-based custom views to display only the information you want, in the format you want.

Custom views display fields you select in a column format. You can set the background and foreground colors, the type of font, and the order of the fields. You can limit the amount of information in your custom view by applying a filter. Depending on your needs, your filters can range from simple to complex.

You can use the following prepackaged custom views:

Completed

Displays completed jobs

Critical Path

Displays jobs on a critical path

Failed

Displays failed jobs

Held

Displays held jobs

No Filter

Displays all jobs that you are subscribed to

Not Completed

Displays jobs that have not completed

Overdue

Displays overdue jobs

Ready

Displays jobs ready to run

Running

Displays running jobs

Waiting

Displays waiting jobs

Note: The names of these prepackaged custom views can be changed either upon installation or by other users.

You can import and export custom views between installations of CA WA Desktop Client using the import/export wizards, which you can access from the File menu.

Custom View Filters

You can apply a filter to a custom view to limit the amount of information that appears for your workload. Each custom view contains filter statements you set. The filter statements can be changed at any time.

You can create custom view filters that are simple or complex. For example, a simple filter might display all Applications named PAYROLL, while a more complex filter might display jobs that are waiting or jobs that are in trouble.

In the Monitor perspective, you can design a custom view for jobs with a wide array of scheduling criteria. For example, you can create a display for jobs that have been waiting to run in an active Application since a certain time and are waiting for a resource before they can be submitted.

In the Filter page of the Custom View Configuration dialog, you can define multiple filter criteria and link them with 'and' and 'or' operands.

Example: Display All Jobs That Are Waiting

Suppose that you want to view all jobs that are waiting. You might create a filter using the exact state names as the values:

Predwait

Indicates the jobs waiting for predecessors to complete

Waiting

Indicates the jobs waiting for a specific time for submission

Reswait

Indicates the jobs waiting for resources

Subdelay

Indicates the jobs that have delayed submission

The following table shows how to create a filter to display all jobs that are in waiting:

Oper	Field Name	Relationship	Value	Oper	And/Or
	State	Is	Predwait		Or
	State	Is	Waiting		Or
	State	Is	Reswait		Or
	State	Is	Subdelay		Or

To simplify the filter, you can replace the first three filter criteria with the following single statement:

State Contains wait

Example: Display All Jobs That are in Trouble

The following table shows how to create a filter to display all jobs that are in trouble:

Oper	Field Name	Relationship	Value	Oper	And/Or
	State	Contains	fail		Or
	State	Is	suberror		Or
	State	Is	inactive		Or

More information:

[Create a Custom View Filter](#) (see page 141)

Open a Prepackaged Custom View

You can use the prepackaged custom views to show information about the jobs you subscribed to based on job state and other characteristics.

To open a prepackaged custom view

1. Open the Custom Views view in the Monitor perspective.
The Custom Views tree displays the predefined custom views.
2. Right-click the predefined custom view you want to use, and select Open from the pop-up menu.
Monitor displays the text-based custom view of the workload.
3. (Optional) Do the following to re-sort a custom view:
 - To re-sort a column in ascending order, click the column heading until an upward arrow appears in the column heading.
 - To re-sort a column in descending order, click the column heading until a downward arrow appears in the column heading.
 - To remove the sort order, click the column heading until no arrow appears in the column heading.

Note: You may have to re-sort multiple columns to achieve the results you want. For example, a custom view may be sorted by Application name, then by generation number, followed by job name. If you want to re-sort by generation number, you can remove the sort order on the Application name column and re-sort the generation number column in ascending or descending order.

Monitor displays the custom view in the sort order you specified.

Create a New Text-Based Custom View

You can create a new custom view to monitor and control the jobs you are interested in.

To create a custom view

1. Open the Custom Views view in the Monitor perspective.
The Custom Views tree lists the predefined custom views.
2. Click the Create a New Custom View icon below the Custom Views title bar.
The Presentation page of the Custom View Configuration dialog opens.
3. Enter a name for your custom view in the Custom view name field.
4. Do one of the following to select the fields you want to display in the custom view:
 - To move a field, select the field from the Available fields section, and click the right arrow (>).
 - To move multiple fields, hold down the Ctrl key, select the fields you want to move, and click the right arrow (>).
 - To move all the fields, click the double right arrow (>>).

Note: To remove a field from the Selected fields section, you can select the fields you want to remove and click the left arrow (<), or click the double left arrow (<<) to remove all fields.

The selected fields are moved to the Selected fields section from the Available fields section.

5. (Optional) Enter an alias for the field name in the Alias column if you want to shorten the field name.
Note: For example, you can shorten generation number to gen #.
6. (Optional) Click the Sort Order column for the field and select one of the following:
 - Ascending—Sorts from A to Z, lowest to highest number, or by earliest date.
 - Descending—Sorts from Z to A, highest to lowest number, or by latest date.

The Sort Priority column displays a sort priority for the selected field.

Note: When viewing workload using a custom view, you can re-sort the data by clicking a column heading.

7. (Optional) Click [+] and [-] buttons to adjust the sorting priority.
Note: The field with Sort Priority 1 is the field with the highest priority in the new custom view. For example, if the Agent Name field has Sort Priority 2, it is second on priority, click [-] to set the Agent Name field's Sort Priority to 1 to make it the highest priority field. The column positions are independent of the sort priority in the new custom view.
8. (Optional) Select a field from the Selected fields section, and click Move Up or Move Down to change the order of the fields.

9. Click OK.

The custom view appears in the tree view window.

Copy an Existing Custom View

You can copy a custom view and save it with a different name to create a new custom view.

To copy a custom view

1. Open the Custom Views view in the Monitor perspective.

The Custom Views tree lists the predefined custom views.

2. Right-click the custom view you want to copy, and select Copy from the pop-up menu.

The Copy Custom View Definition dialog opens.

3. Enter a name for the custom view and click OK.

The copy of the custom view appears in the tree view.

4. (Optional) Do the following to modify the custom view:

- a. Right-click the custom view copy you created, and select Edit from the pop-up menu.

The Presentation page of the Custom View Configuration dialog opens.

- b. Modify the custom view as required and click OK.

5. Right-click your custom view in the Custom Views view, and select Open from the pop-up menu.

The custom view opens and displays your workload.

Create a Custom View Filter

You can create custom view filters to limit the amount of information that appears for your workload.

To create a custom view filter

1. Open the Custom Views view in the Monitor perspective.
The Custom Views tree lists the predefined custom views.
2. Right-click the custom view you want to create a filter for, and select Edit from the pop-up menu.
The Presentation page of the Custom View Configuration dialog opens.
3. Click Filter in the left pane.
The Filter page opens in the right pane.
4. Click Add.
A new row is added to the Criteria section.
5. Select a value from the drop-down list in the Field Name field.
6. Select one of the following operators from the drop-down list in the Relationship field:

Contains

Displays partial and full matches of a value. For example, Application Name Contains pay displays all Applications that have names containing “pay” such as payroll and BCpayroll.

Does not contain

Displays matches that do not contain the value. For example, Application Name Does not contain pay displays all Applications with names that do not include “pay” anywhere in the name.

Is

Displays exact matches of a value. For example, Application Name Is payroll displays only Applications named “payroll.”

Is not

Displays everything that does not exactly match the value. For example, Application Name Is not payroll displays all Applications with names that are not “payroll.”

Begins with

Displays matches for the beginning of the value. For example, Application Name Begins with Pay displays all Applications that begin with “Pay” such as Pay1 and Payroll.

Ends with

Displays matches for the end of the value. For example, Application Name Ends with roll displays all Applications that end with “roll” such as Payroll.

Greater Than

Display matches that are greater than the value. For example, Average Run Time Greater Than 60 displays all jobs with average execution times greater than 60 minutes.

Less Than

Display matches that are less than the value. For example, Average Run Time Less Than 60 displays all jobs with average execution times less than 60 minutes.

7. Enter a value in the Value field.

Note: You can enter a portion of a word or an entire word.

8. Click Apply.

The filter displays in the Result text box.

9. Click OK.

The filter is applied to the custom view.

Example: Create a Custom View for the Payroll Application

You can create a custom view to display only the Payroll Application.

To create a custom view for the Payroll Application

1. Open the Custom Views view in the Monitor perspective.
The Custom Views tree lists the predefined custom views.
2. Right-click the custom view you want to create a filter for, and select Edit from the pop-up menu.
The Presentation page of the Custom View Configuration dialog opens.
3. Click Filter in the left pane.
The Filter page opens in the right pane.
4. Click Add.
A new row is added to the Criteria section.
5. Select Application name from the Field Name drop-down list.
6. Select Is from the Relationship drop-down list.
7. Enter Payroll in the Value field.
8. Click Apply.
The filter displays in the Result text box.
9. Click OK.
The filter is applied to the custom view.

More information:

[Custom View Filters](#) (see page 136)

Chapter 6: Working with Dashboard Views

This section contains the following topics:

[Dashboard Views](#) (see page 146)

[Open an Existing Status Message View](#) (see page 147)

[Create a New Status Message View](#) (see page 148)

[Copy an Existing Status Message View](#) (see page 151)

[Open an Existing System Monitor View](#) (see page 151)

[Create a New System Monitor View](#) (see page 152)

[Copy an Existing System Monitor View](#) (see page 154)

[Edit a Dashboard View](#) (see page 154)

[Delete a Dashboard View](#) (see page 155)

Dashboard Views

The Monitor Perspective contains a dashboard view that provides status and historical information in predefined status message and system monitor custom views. You can also create, edit, or delete views to display the information you want.

The Status Message View provides status information as a list of messages. The status messages represent state or historical information.

State messages indicate the current state of a component. Only one state message is displayed for each component, for example, the state of an agent. When the server receives a new state message, it deletes the older state message for the specified component. You can view state messages on the following:

- Agent status—up, down, quiesced
- Number of active Applications
- Details of the client connections (standalone CLI or CA WA Desktop Client)
- Memory usage including total memory, free memory, and maximum memory allocated
- Primary server start time and type
- Standby server start time and type
- License status of the server
- LDAP status—active, inactive
- Server quiesced status—quiesced, unquiesced

Note: The server node displays the status as Quiesced but the status of the server instances do not change.

Historical messages represent events. Multiple historical messages can be displayed, for example, SNMP messages, log messages, and so on. You can view historical messages on the following:

- SNMP traps
- Details of Login failures
- Events that failed to run due to insufficient user permissions of the user
- Error messages from the tracelog.txt file

The System Monitor View displays the following status information in a graphical dashboard:

- Number of active Applications
- Number of CA WA Desktop Client connections
- Server memory status

- Single server start time and type (in a stand-alone configuration)
- Primary and Standby server start time and type (in a CA WA High Availability configuration)
- Number of defined agents
- Number of agents in Active, Inactive and Quiesce states
- LDAP server status
- License status, the date of expiry, and the number of agents using the license

System Monitor views are updated in real time as information is received from the server.

Open an Existing Status Message View

You can open an existing status message view to display the messages in that view. The message views provide status information on the following as a list of messages:

- Servers
- Agents
- The LDAP server
- Number of active Applications
- Number of CA WA Desktop Client connections
- Memory used by the server (max, free, total)
- Additional historical information on changes over a certain period of time (for example, SNMP messages)
- Error log messages from the tracelog.txt file

Note: You must have Allow access to the STATUSMSG* permission to open an existing status message view.

To open a prepackaged dashboard view

1. Open the Dashboard Views view in the Monitor perspective.

The Dashboard Views tree displays the status message and system monitor custom views.

2. Right-click the predefined status message view you want to open, and select Open from the pop-up menu.

The server connections are displayed in the left pane of the custom view. You can expand each server connection to view the list of the messages. The messages are sorted by the specified group.

3. Double-click a message to view the details of the message.

The message details are displayed in the right pane of the custom view.

Create a New Status Message View

You can create a new status message view to select the message types that you want to display in the new view. You can also specify the message sorting order of the new view.

Note: You must have Allow access to the STATUSMSG* permission to create a new status message view.

To create a new status message view

1. Open the Dashboard Views view in the Monitor perspective.

The Dashboard Views tree displays the status message and system monitor custom views.

2. Right-click Status Message Views, and select New from the pop-up menu.

The Edit Status Message View dialog opens.

3. Enter a name for your view in the Name field on the Types page.

4. (Optional) Enter a description for the new view in the Description field.

5. Do *one* of the following to select the message types you want to display in the new view:

- To add a message type, select the message type from the Available message types section, and click Add>.
- To add multiple message types, hold down the Ctrl key, select the message types from the Available message types section, and click Add>.
- To add all the message types, click Add All>>.

The selected message types are moved to the Selected message types section from the Available message types section.

Note: To remove a message type from the Selected message types section, select the message types you want to remove and click <Remove or click <<Remove All to remove all message types.

6. Click Columns in the left pane.

The Columns page opens in the right pane.

7. Select a value in the Group by drop-down list to sort the new view by message type, date, or severity.

8. Do *one* of the following to select the message columns you want to display in the new view:

- To add a message column, select the message column from the Available message columns section, and click Add>.
- To add multiple message columns, hold down the Ctrl key, select the message columns from the Available message columns section, and click Add>.
- To add all the message columns, click Add All>>.

The selected message columns are moved to the Selected message columns section from the Available message columns section.

Note: To remove a message column from the Selected message columns section, select the message columns you want to remove and click <Remove or click <<Remove All to remove all message columns.

9. Click OK.

The new Status Message view appears in the Dashboard Views tree under Status Message Views.

Example: Create a Status Message View to Display the Server Status

You can create a status message view to display only the server status.

To create a status message view to display the server status

1. Right-click Status Message View and select New from the pop-up menu.
2. Enter **Server Status** in the Name field on the Types page.
3. Select Server status and Quiesce status from the Available message types section, and click Add>.
4. Click Columns in the left pane.
5. Select Date in the Group by drop-down list to sort the new view by date.
6. Select Severity, Date, and Connection from the Available message columns section, and click Add>.
7. Click OK.

The Server Status Message view appears in the Dashboard View tree under Status Message Views.

8. Double-click the Server Status view to open the view.

All the messages are displayed under your server connections in the left pane.

9. Double-click a message to view the message details in the right pane.

Copy an Existing Status Message View

You can copy a status message view and save it with a different name to create a new status message view. You can also add or remove message types from the copied view.

Note: You must have Allow access to the STATUSMSG* permission to copy an existing status message view.

To copy an existing status message view

1. Open the Dashboard Views view in the Monitor perspective.

The Dashboard Views tree displays the status message and system monitor custom views.

2. Right-click the status message view you want to copy, and select Copy from the pop-up menu.

The Edit Status Message View dialog opens.

3. Enter a name for the new view in the Name field on the Types page.
4. Modify the view as required and click OK.

The copied status message view appears in the Dashboard Views tree under Status Message Views.

Open an Existing System Monitor View

You can open an existing system monitor view to monitor the graphical controls in that view. The system monitor views provide status information on the following in a graphical dashboard::

- Total number of active Applications, CA WA Desktop Client connections, defined agents, and their states
- Server memory status
- Single server start time and type (in a stand-alone configuration)
- Primary and Standby server start time and type (in a CA WA High Availability configuration)
- LDAP server status
- License status, the date of expiry, and the number of agents using the license

Note: You must have Allow access to the STATUSMSG* permission to open an existing system monitor view.

To open an existing system monitor view

1. Open the Dashboard Views view in the Monitor perspective.

The Dashboard Views tree displays the status message and system monitor custom views.

2. Right-click the predefined system monitor view you want to open, and select Open from the pop-up menu.

The graphical dashboard is displayed in the right pane of the custom view.

Create a New System Monitor View

You can create a new system monitor view to select the controls you want to monitor in the new graphical dashboard view.

Note: You must have Allow access to the STATUSMSG* permission to create a new system monitor view.

To create a new system monitor view

1. Open the Dashboard Views view in the Monitor perspective.

The Dashboard Views tree displays the status message and system monitor custom views.

2. Right-click System Monitor Views, and select New from the pop-up menu.

The Edit Dashboard View dialog opens.

3. Enter a name for your view in the Name field.
4. (Optional) Enter a description for the new view in the Description field.
5. (Optional) Complete the following fields as required:

Number of columns to render the controls

Specifies the number of columns to display in the new view.

Limits: 1-4

Default: 1

Horizontal spacing between controls (pixels)

Specifies the horizontal spacing between the columns in pixels.

Default: 5

Vertical spacing between controls (pixels)

Specifies the vertical spacing between the columns in pixels.

Default: 5

Use the same lengths for all columns

Indicates whether to use the same size for all the columns in the new view.

6. Do one of the following to select the controls you want to display in the new view:
 - To add a field, select the control from the Available controls section, and click Add>.
 - To add multiple controls, hold down the Ctrl key, select the controls from the Available controls section, and click Add>.
 - To add all the controls, click Add All>>.

The selected controls are moved to the Selected controls section from the Available controls section. The Preview section displays the layout of new view.

Note: To remove a control from the Selected controls section, select the control you want to remove and click <Remove or click <<Remove All to remove all controls.

7. Click OK.

The new system monitor view appears in the Dashboard Views tree under System Monitor Views.

Example: Create a System Monitor View to Display the Agent Status

You can create a system monitor view to display the agent status in a graphical format.

To create a system monitor view to display the agent status

1. Open the Dashboard Views view in the Monitor perspective.
2. Right-click the System Monitor Views, and select New from the pop-up menu.
3. Enter **Agent Status** in the Name field.
4. Select Agent Status – Pie Chart and Summary Monitor from the Available controls section, and click Add>.
5. Enter **3** in the Number of columns to render the controls field.
6. (Optional) View the layout of the new system monitor view in the Preview section.
7. Click OK.

The Agent Status system monitor view appears in the Dashboard Views tree under System Monitor Views.

8. Double-click the Agent Status view under System Monitor Views.

The Agent Status view displays the agent status pie chart in the left pane and the summary monitor in the right pane of the view. The connected servers appear as tabs at the bottom of the view.

Copy an Existing System Monitor View

You can copy a system monitor view and save it with a different name to create a new status monitor view. You can also add or remove controls to create a view that meets your requirements.

Note: You must have Allow access to the STATUSMSG* permission to copy an existing system monitor view.

To copy an existing system monitor view

1. Open the Dashboard Views view in the Monitor perspective.
The Dashboard Views tree displays the status message and system monitor custom views.
2. Right-click the existing system monitor view you want to copy, and select Copy from the pop-up menu.
The Edit Dashboard View dialog opens.
3. Enter a name for the new view in the Name field.
4. Modify the view as required and click OK.
The copied system monitor view appears in the Dashboard Views tree under System Monitor Views.

Edit a Dashboard View

You can edit a dashboard view to suit your business needs. For example, you had designed a view to see the messages related to the server, but now you also want to see the messages related to agents.

Note: You must have Allow access to the STATUSMSG* permission to edit a dashboard view.

To edit a dashboard view

1. Open the Dashboard Views view in the Monitor perspective.
The Dashboard Views tree displays the status message and system monitor custom views.
2. Right-click the dashboard view you want to edit, and select Edit from the pop-up menu.
The Edit Status Message View or Edit Dashboard View dialog opens.
3. Modify the dashboard view as required and click OK.
The modified dashboard view is saved.

Delete a Dashboard View

You can delete a dashboard view if you do not require it any more.

Note: You must have Allow access to the STATUSMSG* permission to delete a dashboard view.

To delete a dashboard view

1. Open the Dashboard Views view in the Monitor perspective.
The Dashboard Views tree displays the status message and system monitor custom views.
2. Right-click the dashboard view you want to delete, and select Delete from the pop-up menu.
A confirmation dialog opens.
3. Click Yes.
The view is deleted from the Dashboard Views tree in the left pane.

Chapter 7: Resetting Resource Availability Counts

This section contains the following topics:

[Resource Availability Counts](#) (see page 157)

[Reset Resource Availability Count](#) (see page 158)

Resource Availability Counts

You can reset the resource availability counts to increase or decrease the number of available units of a resource. You can also increase or decrease the maximum number of available units of a renewable resource.

For example, you may need to reset a resource's availability counts in the following cases:

- During periods of low activity, you can set a threshold resource from 0 to 1 to let low priority jobs run.
- Before shutting down an agent, you can set a threshold resource from 1 to 0 to quiesce the agent.
- To let more database jobs run concurrently, you can set a renewable resource from 5 to 6.

Reset Resource Availability Count

You can reset resource availability counts to increase or decrease the number of available units of a resource. For renewable resources, you can also increase or decrease the maximum number of available units.

To reset resource availability count

1. Open the Monitor perspective.
2. Right-click the server connection in the Application Monitor view and select Modify Resources.

The Modify Resources dialog opens.

3. Select a resource name in the Resource name drop-down list.
4. Select one of the following options from the Property name drop-down list.

Availability

Indicates the number of available units of the resource.

Maximum Availability

Indicates the maximum number of available units of the resource. Applies to renewable resources only. Changing the maximum availability count also changes the availability count by the same amount. For example, if you increase the maximum availability count from three to five, the availability count increases by two.

5. Select one of the following options from the drop-down list:

Reset value to

Sets the availability count of the resource to the specified value.

Increase value by

Increases the availability count of the resource by the specified value.

6. Specify the absolute or relative availability count in the text box.

Note: For renewable resources, if the availability count exceeds the maximum availability count, the availability count is set to the maximum availability count.

7. Click OK.

The availability count of the resource is reset.

Index

A

- active process limit • 98
- anticipated end times
 - refreshing • 34
- Application generations • 11
- Applications
 - active • 20
 - command log • 30
 - commands • 20
 - completing • 29
 - details • 30
 - generations • 11
 - graphical overview • 15
 - hiding completed • 17, 41
 - holding • 32
 - inserting jobs in active • 27
 - issue commands against • 25
 - monitoring • 24
 - refreshing anticipated end times • 34
 - relationships • 68
 - releasing • 33
 - states • 20
 - subscribing to • 13
 - troubled jobs in active • 29
 - unquiescing • 37, 38
 - unwaiting • 35
 - viewing active • 13
 - viewing comments • 35
 - viewing Gantt chart • 26
 - viewing jobs within active • 14
 - zoom settings • 16
- availability counts, resources • 157

B

- bypassing
 - jobs • 79
 - subApplications • 49

C

- CA WA Restart Option
 - defined • 111
 - restarting z/OS jobs • 111
- cancelling running jobs • 81
- command log • 30, 75

- commands
 - Application • 20
 - Database • 116
 - issuing against Applications • 25
 - issuing against jobs • 73
 - job • 58, 63
 - Oracle E-Business Suite • 132
 - PeopleSoft • 131
 - SAP • 117, 125, 127, 129, 130
 - subApplication • 43
- comments
 - viewing Application • 35
 - viewing job • 88
- completing
 - Applications • 29
 - jobs • 82
 - subApplications • 46
- conditions, job • 62
- copy JCL
 - browsing • 109
 - editing • 110
- counts
 - resource availability • 157
- critical paths
 - custom views • 135
 - displaying • 15
- custom view filters
 - creating • 141
 - defined • 136
 - examples • 136, 141
- custom views
 - copying • 140
 - creating • 139, 141
 - defined • 135
 - filters • 136
 - opening prepackaged • 138
 - prepackaged • 135

D

- dashboard views
 - defined • 146
 - deleting • 155
 - editing • 154
- data definition • 107
- Database jobs

- issuing commands against • 116
- deleting
 - dashboard views • 155
- dependencies
 - global variable • 72
 - predecessor • 68
 - resource • 73
 - time • 70
- dropping
 - predecessor dependencies • 90
 - resource dependencies • 95
 - variable dependencies • 93
- dueout times
 - reset • 71

E

- Events
 - unquiescing • 37
- external jobs
 - completion • 66

F

- FTP jobs
 - verifying status • 64

G

- Gantt charts
 - Application, active • 26
- generations
 - Application • 11
- global variable dependencies
 - defined • 72
 - dropping • 93
 - evaluating • 94
 - job submission • 72

H

- hiding
 - completed Applications • 17, 41
- holding
 - Applications • 32
 - jobs • 83
 - subApplications • 47

I

- i5/OS jobs
 - retrieving spool file • 104

- inserting jobs in active Applications • 27

J

- JCL
 - browsing • 109
 - editing • 110
- JES spool file, retrieving • 107
- job
 - memory • 98
 - time • 98
- job objects, Windows • 98
- jobs
 - bypassing • 79
 - cancelling running • 81
 - command log • 75
 - commands • 58, 63
 - completing • 82
 - conditions • 62
 - details • 74
 - holding • 83
 - ID • 97
 - inserting in active Application • 27
 - issuing commands against • 73
 - monitoring forced complete • 21, 62, 66
 - process priority • 101
 - readying for submission • 88
 - releasing • 83
 - requesting • 84
 - resetting definition • 79
 - resubmitting • 86
 - setting user status • 90
 - spool file • 63
 - states • 59
 - troubled, within an Application • 29
 - unbypassing • 80
 - unrequesting • 85
 - unwaiting • 89
 - verifying process • 97
 - viewing comments • 88
 - viewing in Application • 14

M

- MicroFocus jobs
 - cancellation • 65
 - JES spool file, retrieving • 107
 - resubmitting • 86
- monitoring
 - Application • 24

overdue jobs • 71

O

Oracle E-Business Suite jobs
 issuing commands against • 132
overdue jobs
 monitoring • 71

P

PeopleSoft jobs
 issuing commands against • 131
predecessor dependencies
 defined • 68
 dropping • 90
priority
 class • 98
 process • 101
process
 ID • 97
 limit • 98
 memory • 98
 status • 97
 time • 98
 verifying • 97

Q

quiescing
 server • 36

R

readying for submission
 jobs • 88
 subApplications • 51
refreshing anticipated end times • 34
releasing
 Applications • 33
 jobs • 83
 subApplications • 48
requesting
 jobs • 84
 subApplications • 52
resetting
 job definition • 79
 time dependencies • 91
resources
 availability counts • 157
 displaying usage status • 96
 dropping • 95

 setting availability count • 158
resubmitting jobs • 86
retrieving
 JES spool file • 107
 spool file • 76

S

SAP jobs
 controlling • 117
SAP-Batch Input Session jobs
 issuing commands against • 125
SAP-BW Info Package jobs
 issuing commands against • 127
SAP-BW Process Chain jobs
 issuing commands against • 129
SAP-Data Archiving jobs
 issuing commands against • 130
SAP-Job Copy jobs
 issuing commands against • 117
SAP-Process Monitor jobs
 issuing commands against • 130
SAP-R3 jobs
 issuing commands against • 117
schedule criteria
 time dependencies • 70
server
 quiescing • 36
 responses • 97
show server response • 97
signal, send • 102
spool file, retrieving • 76, 104, 107
state
 Application • 20
 job • 59
status message view
 copying • 151
 creating • 148
 opening • 147
subApplications
 bypassing • 49
 commands • 43
 completing • 46
 displaying in an active Application • 46
 holding • 47
 readying for submission • 51
 releasing • 48
 requesting • 52
 unbypassing • 50

- unrequesting • 53
- unwaiting • 54
- subscribing to Applications • 13
- subscriptions, creating • 12
- system monitor views
 - copying • 154
 - creating • 152
 - opening • 151

T

- time dependencies
 - defined • 70
 - resetting • 91

U

- unbypassing
 - jobs • 80
 - subApplications • 50
- UNIX jobs
 - cancelation • 65
 - sending a signal • 102
 - setting process priority • 101
 - verifying process • 97
- unquiescing
 - Application generations • 37, 38
 - Events • 37
- unrequesting
 - jobs • 85
 - subApplications • 53
- unwaiting
 - Applications • 35
 - jobs • 89
 - subApplications • 54
- user status, setting • 90

V

- variable dependencies
 - defined • 72
 - dropping • 93
 - evaluating • 94
 - job submission • 72
- verifying
 - zoom settings, Application graphs • 16

W

- Windows job objects
 - setting properties • 98
- Windows jobs

- cancelation • 65
- setting process priority • 101
- verifying process • 97
- workload
 - custom views • 135
 - subscribing to • 13
 - viewing active • 13

Z

- z/OS jobs
 - browsing copy JCL • 109
 - browsing JCL • 109
 - editing copy JCL • 110
 - editing JCL • 110
 - restarting using CA WA Restart Option • 111
- zoom settings, Application graphs • 16