

IBM Workload Scheduler Troubleshooting Guide

Note

Before using this information and the product it supports, read the information in [Notices on page ccix](#).

This edition applies to version 10, release 1, modification level 0 of IBM Workload Automation (program number 5698-T09) and to all subsequent releases and modifications until otherwise indicated in new editions.

Contents

Note.....	ii	Configuration file syntax.....	57
.....	iii	xcli command.....	60
List of Figures.....	ix	Selecting programs, segments, and products.....	60
List of Tables.....	x	xcli command syntax.....	61
About this guide.....	xi	xcli messages.....	68
What is new in this release.....	xi	Chapter 5. Troubleshooting performance.....	72
Accessibility	xi	Chapter 6. Troubleshooting networks.....	73
Technical training.....	xi	Network recovery.....	73
Support information.....	xi	Initialization problems.....	73
Chapter 1. Getting started.....	13	Network link problems.....	74
Where products and components are installed.....	13	Replacement of a domain manager.....	75
Finding out what is installed.....	17	Replacement of a master domain manager.....	76
Built-in troubleshooting features.....	19	Other common network problems.....	76
Keeping up to date.....	20	Using SSL, no connection between fault-tolerant agent and its domain manager.....	76
Upgrading your whole environment.....	20	Timeout during Symphony download.....	77
Chapter 2. Logs and traces.....	21	After changing SSL mode, a workstation cannot link.....	77
Modifying log and trace levels: quick reference.....	21	Start and stop remote commands not working with firewall.....	78
Logs and traces: the difference.....	24	Remote command job fails to connect to remote computer.....	78
IBM Workload Scheduler logging and tracing.....	26	The domain manager cannot link to a fault-tolerant agent.....	78
Locations.....	26	Agents not linking after first JnextPlan on HP-UX.....	79
File switching.....	27	Fault-tolerant agents not linking to master domain manager.....	79
Customization.....	27	dynamic agent not found from console.....	81
Performance.....	31	Submitted job not running on a dynamic agent.....	81
Dynamic Workload Console log and trace files.....	31	Job status of submitted job continually running on dynamic agent.....	81
Activating traces.....	32	Network performance.....	81
IBM Workload Scheduler agent trace files.....	33	AWSITA245E- Agent is down but jobmanager is running.....	82
Retrieving IBM Workload Scheduler agent traces from the Dynamic Workload Console.....	33	Chapter 7. Troubleshooting engine problems.....	83
Self-Service Dashboard application log file.....	34	Composer problems.....	83
Dynamic workload scheduling log and trace files.....	34	Composer dependency error with interdependent object definitions.....	83
Dynamic agent log and trace files.....	35	The display cpu=@ command fails on UNIX.....	84
Trace configuration for the agent.....	37	Composer not authorized to access server error.....	84
Application server: log and trace files.....	40	Composer deletion of a workstation fails with the AWSJOM179E error.....	85
Setting application server traces.....	40	Job stream synchronicity problems.....	85
Chapter 3. Capturing product data.....	44	Exit from command line takes a long amount of time.....	85
Data capture utility.....	44	JnextPlan problems.....	86
When to run.....	44	JnextPlan fails to start.....	86
Prerequisites.....	45		
Command and parameters.....	46		
Tasks.....	48		
Data collection.....	49		
First failure data capture (ffdc).....	52		
Creating application server dumps.....	52		
Chapter 4. In-Flight Trace facility.....	54		
In-Flight Trace configuration file.....	56		
Changing the configuration	57		

JnextPlan fails because transaction log for the database is full.....	86	Corrupted characters appear in the command shell when executing cli commands.....	103
JnextPlan fails: Java out-of-memory error.....	87	Extended agent problems.....	104
JnextPlan fails with DB2 error: nullDSRA0010E.....	87	The extended agent for MVS returns an error....	104
JnextPlan fails - AWSJPL017E message.....	87	Planner problems.....	104
On windows operating systems JnextPlan fails with cscript error.....	88	Mismatch between Job Scheduler instances and preproduction plan.....	104
JnextPlan is slow.....	88	Planman deploy error when deploying a plugin.....	105
Remote workstation not initializing after JnextPlan.....	89	Insufficient space error while deploying rules....	105
Workstation not linking after JnextPlan.....	89	AWSJCO084E UpdateStats fails if job run time exceeds two hours.....	105
Job remains in "exec" status after JnextPlan but not running.....	89	Planman showinfo command displays inconsistent times.....	106
CreatePostReports.cmd, or Makeplan.cmd, or Updatestats.cmd, or rep8.cmd hang on Windows operating systems.....	91	Job stream duration might be calculated incorrectly as well as other time-related calculations.....	106
Possible performance impact on large production plans after JnextPlan.....	91	Bound z/OS shadow job carried forward indefinitely.....	107
Conman problems.....	92	Problems with DB2.....	107
AWSDEQ024E message is received on Windows.....	92	Timeout occurs with DB2.....	107
Duplicate ad-hoc prompt number.....	94	JnextPlan fails because DB2 transaction log is full.....	108
During Jnextplan fault-tolerant agents cannot be linked.....	94	DB2 might lock while making schedule changes.....	108
Submitting job streams with a wildcard loses dependencies.....	95	Problems with Oracle.....	109
Exit from command line takes a long amount of time.....	95	JnextPlan fails because the database transaction log is full.....	109
Job log not displayed.....	96	Oracle maintenance on UNIX not possible after installation.....	109
Cannot retrieve objects defined in folders on a workstation set to ignore.....	96	Dynamic workload broker fails to start after switching DB2 to Oracle.....	109
Fault-tolerant agent problems.....	97	Problems with MSSQL.....	110
Job fails in heavy workload conditions.....	97	An error is returned when deleting records with cascade option in MSSQL.....	110
Batchman fails on a fault-tolerant agent AWSDEC002E message.....	97	Problems with Informix.....	110
Fault-tolerant agents unlink from mailman.....	98	Solving deadlocks when using composer in Informix.....	110
Symphony file on the master domain manager not updated with fault-tolerant agent job status.....	98	Application server problems.....	111
Dynamic agent problems.....	99	Application server not starting after changes to SSL keystore password.....	111
Connection problems for the dynamic agent.....	99	Timeout occurs with the application server.....	111
JBDC driver problem with MSSQL database.....	99	Master Domain Manager may stop when trying to retrieve a big job log.....	112
Error AWSITA104E Unable to perform the system resources scan	100	AWKTSA050E error issued during submission.....	112
Event condition on dynamic agent does not generate any action.....	100	Event management problems.....	113
Job manager encounters a core dump.....	100	Troubleshooting an event rule that does not trigger the required action.....	113
Problems on Windows.....	101	After replying to a prompt, the triggered action is not performed.....	122
Interactive jobs not interactive using Terminal Services.....	101	Actions involving automatic sending of email fail.....	122
Services fail to start after workstation restart....	101	An event is lost.....	123
Batchup service fails to start.....	101		
Error relating to impersonation level.....	103		

Action not triggered.....	123	Deleting stdlist or one of its files while processes are still running.....	135
Event rules not deployed after switching event processor.....	124	Job with "rerun" recovery job remains in "running" state.....	136
Event LogMessageWritten is not triggered.....	125	Job statistics are not updated daily.....	136
Deploy (D) flag not set after ResetPlan.....	125	A Job Scheduler dependency is not added.....	136
Missing or empty event monitoring configuration file.....	125	Incorrect time-related status when time zone not enabled.....	136
Events not processed in correct order.....	126	Completed job or job stream not found.....	137
Stopeventprocessor or switcheventprocessor commands not working.....	126	Variables not resolved after upgrade.....	137
Event rules not deployed with large numbers of rules.....	127	Default variable table not accessible after upgrade.....	137
Problem prevention in IBM Workload Scheduler environment.....	127	Local parameters not being resolved correctly.....	137
On AIX operating systems the SSM agent crashes if you have a very large number of files to be managed using event-driven workload automation.....	127	Inconsistent time and date in job streams and jobs.....	138
File creation and deletion actions not triggered.....	128	Deleting leftover files after uninstallation is too slow.....	138
Managing concurrent accesses to Symphony.....	128	Special characters are corrupted in script job output.....	139
Scenario 1: Access to Symphony file locked by other processes.....	128	Failover Cluster Command Interface deprecated.....	139
Scenario 2: Access to Symphony file locked by stageman.....	128	Job stream creation failure due to missing rights on calendars.....	139
StartApp Server problems.....	128	Suppressing connector error messages	139
MakePlan problems.....	128	StartUp shows an error after upgrade.....	140
MakePlan fails to start.....	129	Chapter 8. Troubleshooting dynamic workload scheduling.....	142
Unable to establish communication with the server on host - AWSBEH023E.....	129	How to tune the rate of job processing.....	142
The user "twuser" is not authorized to access the server on host - AWSBEH021E.....	130	Remote command job fails.....	145
Database is already locked - AWSJPL018E.....	130	Monitoring and canceling jobs.....	145
An internal error has occurred - AWSJPL006E.....	130	On Windows 2012 the user interfaces for the interactive jobs are not visible on dynamic agents.....	148
The production plan cannot be created - AWSJPL017E	130	MASTERAGENTS workstation not updated in optman.....	149
An internal error has occurred - AWSJPL704E	130	Troubleshooting common problems.....	149
SwitchPlan problems.....	131	On AIX operating systems the concurrent submission of one hundred or more jobs on the same agent can result in a core dump or in a resource temporarily unavailable message.....	150
When SwitchPlan fails to start.....	131	Dynamic workload broker cannot run after the database is stopped.....	150
Previous Symphony file and Symnew file have the same run number - AWSBHV082E.....	131	OutofMemory exception when submitting a job.....	150
Create Post Reports.....	132	Jobsubmiterror exception when submitting a job.....	151
Update Stats.....	132	Chapter 9. Troubleshooting when automatically adding dynamic agent workstations to the plan.....	152
Miscellaneous problems.....	132	The dynamic agent workstation automatically added to the plan is not initialized.....	152
Error message for a locked database table or object in a table.....	133	Chapter 10. Troubleshooting Dynamic Workload Console.....	153
Command line programs error "user not authorized to access server".....	134	Troubleshooting connection problems.....	153
Rmstdlist command gives different results.....	134	The engine connection does not work.....	153
Rmstdlist command fails on AIX with exit code of 126.....	134		
Question marks found in the stdlist.....	135		

Test connection takes several minutes before returning failure.....	155	Web page error with Internet Explorer	168
Engine version and connection status not displayed.....	155	Panels frozen with Internet Explorer developer tools.....	169
Failure in testing connection or running reports if Oracle database in use.....	156	Simplified Chinese characters missing or corrupted when using Google Chrome or Apple Safari.....	169
Connection problem with the engine.....	156	Troubleshooting problems with graphical views.....	169
Engine connection failure when connecting to z/OS connector (v8.3x and v8.5x).....	156	Language-specific characters not correctly displayed in graphical views.....	169
Engine connection failure when connecting to z/OS connector v8.3x or distributed engine v8.3x.....	158	Plan View limit.....	169
Engine connection settings not checked for validity.....	159	JAWSITA122E or AWKRAA209E error in the Workload Designer.....	170
LDAP account locked after one wrong authentication attempt.....	159	Troubleshooting problems with database.....	170
Troubleshooting performance problems.....	160	Import settings to DB2 repository fails.....	170
With a distributed engine responsiveness decreases overtime.....	160	Troubleshooting other problems.....	171
Running production details reports overloading distributed engine.....	160	Jobs in READY status do not start.....	171
Troubleshooting user access problems.....	161	Composer deletion of a workstation fails with the AWSJOM179E error.....	172
Wrong user logged in when using multiple accesses.....	161	Data not updated after running actions.....	172
Unexpected user login request when using Single Sign-On.....	161	"Session has become invalid" message received.....	172
Authentication problem when opening the Workload Designer.....	162	Actions against scheduling objects return empty tables.....	173
Troubleshooting problems with reports	162	Default tasks not converted into the browser language.....	174
Validate command on a custom SQL query returns AWSWUI0331E error.....	162	"Access Error" when launching a task.....	174
Reports not displayed in a browser with a toolbar installed.....	163	Processing threads continue in the background.....	174
WSWUI0331E error when running reports.....	163	List of Available Groups is empty in Enter Task Information window.....	174
CSV report looks corrupted on Microsoft Excel not supporting UTF8.....	163	Blank page displayed (in High availability disaster recovery configuration).....	175
Insufficient space when running production details reports.....	164	Some exceptions might not be displayed in Language-specific in the Dynamic Workload Console.....	175
Report fields show default values after upgrade from v8.3 to v8.5.....	164	Extraneous exception logged in SystemOut.....	175
the specified run period exceeds the historical data time frame.....	165	Filtering task results might not work as expected.....	176
Troubleshooting problems with browsers.....	165	Sorting task results might not work as expected.....	178
Default tasks not converted into the browser language.....	166	Monitoring job streams on multiple engines does not respect scheduled time range on z/OS.....	178
"Access Error" when launching a task.....	166	Dynamic Workload Console 9.x login or graphical view pages do not display.....	178
Processing threads continue in the background.....	166	Java exception when performing a query on job streams in plan.....	179
Unresponsive script warning with Firefox.....	167	Import a property file bigger than 100MB.....	179
Plan View panel seems to freeze with Internet Explorer v7.....	167	Query results not displayed.....	180
Blank page displayed (in High availability disaster recovery configuration).....	167	Chapter 11. Troubleshooting workload service assurance.....	181
Workload Designer not showing on foreground with Firefox.....	168	Components involved in workload service assurance.....	181
Panels not displayed in Internet Explorer	168	Exchange of information.....	182
		Common problems with workload service assurance.....	182

Critical start times not aligned.....	182
Critical start times inconsistent.....	183
Critical network timings change unexpectedly.....	183
A critical job is consistently late.....	183
A high risk critical job has an empty hot list.....	184
Chapter 12. Troubleshooting switch manager.....	185
Event counter.....	185
Ftbox.....	186
Troubleshooting link problems.....	186
Common problems with the backup domain manager.....	190
The Symphony file on the backup domain manager is corrupted.....	191
Processes not killed on previous UNIX domain manager after running switchmgr.....	191
Agent cannot relink.....	191
Chapter 13. Synchronizing the database with the Symphony file.....	193
Chapter 14. Corrupt Symphony recovery.....	194
Recovery procedure on a master domain manager.....	194
Recovering using the backup master domain manager.....	194
Recover using the logman and ResetPlan commands.....	196
Recovering the plan from the archived plan.....	197
Recovery procedure on a fault-tolerant agent or lower domain manager.....	200
Recovery procedure on a fault-tolerant agent with resetFTA.....	201
Appendix A. Support information.....	203
Searching knowledge bases.....	203
Search the online product documentation.....	203
Search the Internet.....	203
Obtaining fixes.....	204
Receiving support updates.....	205
Appendix B. Date and time format reference - strftime.....	206
Notices.....	ccix
Index.....	213

List of Figures

Figure 1: ACCT_FS has not linked.....	187
Figure 2: Example output for conman sc @!@ run on the master domain manager.....	188
Figure 3: Example output for conman sc run on the domain manager.....	189
Figure 4: Example output for conman sc run on the unlinked workstation.....	189
Figure 5: Example output for conman sc @!@ run on the unlinked workstation.....	190
Figure 6: Example output for ps -ef grep writer run on the unlinked workstation.....	190

List of Tables

Table 1: Where to find other troubleshooting material.....	13
Table 2: Difference between logs and traces.....	25
Table 3: Locations of log files and trace files.....	32
Table 4: Locations of log and trace files.....	34
Table 5: Job processing status to queue jobs for dispatching.....	143
Table 6: Status mapping between Dynamic workload broker and IBM Workload Scheduler.....	145
Table 7: Default settings for new job run statistic reports.....	164
Table 8: Default settings for new job run history reports...	165
Table 9: strftime date and time format parameters.....	206

About this guide

Gives useful information about the guide, such as what it contains, who should read it, what has changed since the last release, and how to obtain training and support.

Troubleshooting Guide provides information about troubleshooting IBM Workload Scheduler and its components.

What is new in this release

Learn what is new in this release.

For information about the new or changed functions in this release, see *IBM Workload Automation: Overview*, section *Summary of enhancements*.

For information about the APARs that this release addresses, see the IBM Workload Scheduler Release Notes and the Dynamic Workload Console Release Notes. For information about the APARs addressed in a fix pack, refer to the readme file for the fix pack.

New or changed content is marked with revision bars.

Accessibility

Accessibility features help users with a physical disability, such as restricted mobility or limited vision, to use software products successfully.

With this product, you can use assistive technologies to hear and navigate the interface. You can also use the keyboard instead of the mouse to operate all features of the graphical user interface.

For full information, see the Accessibility Appendix in the *IBM Workload Scheduler User's Guide and Reference*.

Technical training

Cloud & Smarter Infrastructure provides technical training.

For Cloud & Smarter Infrastructure technical training information, see: <http://www.ibm.com/software/tivoli/education>

Support information

IBM provides several ways for you to obtain support when you encounter a problem.

If you have a problem with your IBM software, you want to resolve it quickly. IBM provides the following ways for you to obtain the support you need:

- Searching knowledge bases: You can search across a large collection of known problems and workarounds, Technotes, and other information.
- Obtaining fixes: You can locate the latest fixes that are already available for your product.
- Contacting IBM Software Support: If you still cannot solve your problem, and you need to work with someone from IBM, you can use a variety of ways to contact IBM Software Support.

For more information about these three ways of resolving problems, see the appendix about support information in *IBM Workload Scheduler: Troubleshooting Guide*.

Chapter 1. Getting started with troubleshooting

Gives an overview of what troubleshooting information is contained in this publication, and where to find troubleshooting information which is *not* included.

This publication gives troubleshooting information about the IBM Workload Scheduler *engine*. The engine comprises the components of IBM Workload Scheduler that perform the workload scheduling activities, together with the command line by which they can be controlled.

Troubleshooting for other IBM Workload Scheduler activities, products, and components can be found in their relevant publications, as follows:

Table 1. Where to find other troubleshooting material

Activity, Product, or Component	Publication
Installation, upgrade, and uninstallation of IBM Workload Scheduler components and the Dynamic Workload Console	<i>Planning and Installation Guide</i>
Limited fault-tolerant agents for IBM i	<i>Limited Fault-tolerant Agent for IBM i</i>
IBM® Z Workload Scheduler	<i>Diagnosis Guide and Reference</i> <i>Messages and Codes</i>
IBM Workload Scheduler methods and plug-ins	<i>IBM Workload Automation: Scheduling Job Integrations with IBM Workload Automation</i>

Many of the procedures described in this publication require you to identify a file in the installation path of the product and its components. However, they can have more than one installation path, as described in [Where products and components are installed on page 13](#).

Where products and components are installed

Describes where the IBM Workload Scheduler products and components are installed.

This section commences by briefly introducing IBM Workload Automation and explaining how this concept impacts the installed structure of IBM Workload Scheduler.

IBM Workload Automation

IBM Workload Automation is the name of a family of products and components, which includes the following components:

- IBM Workload Scheduler
- IBM® Z Workload Scheduler

- IBM Workload Scheduler for Applications
- Dynamic Workload Console

Many IBM Workload Scheduler components are installed in what is called a *IBM Workload Automation instance*.

Installation paths

***TWA_home* installation path**

Many of the components are installed in an IBM Workload Automation instance. Although this is a notional structure it is represented on the computer where you install IBM Workload Automation components by a common directory referred to in the documentation as *TWA_home*. The path of this directory is determined when you install an IBM Workload Scheduler component for the first time on a computer. You have the opportunity to choose the path when you make that first-time installation, but if you accept the default path, it is as follows:

On UNIX™ operating systems

```
/opt/wa/server_<wauser><n>
```

On Windows™ operating systems

```
%Program Files%\wa\server<n>
```

where *<n>* is an integer value ranging from 0 for the first instance installed, 1 for the second, and so on.

This path is called, in the publications, *TWA_home*. For details about the directories created outside of *TWA_home*, see the section about directories created outside of *TWA_home* in *Planning and Installation Guide*.

***TWA_DATA_DIR* and *DWC_DATA_dir* configuration directories**

To simplify administration, configuration, and backup and recovery on UNIX systems, a new default behavior has been implemented with regard to the storage of product data and data generated by IBM® Workload Scheduler, such as logs and configuration information. These files are now stored by default in the *<data_dir>* directory, which you can optionally customize at installation time.

By default, this directory is *TWA_home/TWSDATA* for the server and agent components, and *DWC_home/DWC_DATA* for the Dynamic Workload Console. The product binaries are stored instead, in the installation directory.

You can optionally customize the *<data_dir>* directory at installation time by setting the **--data_dir** argument when you install using the command-line installation. If you want to maintain the previous behavior, you can set the **--data_dir** argument to the IBM® Workload Scheduler installation directory.

If you deploy the product components using Docker containers, the *<data_dir>* is set to the default directory name and location, and it cannot be modified.

To retrieve the `TWA_DATA_DIR` and `DWC_DATA_dir` location in case you have modified the default path, check the values for the `TWS_datadir` and `DWC_datadir` properties stored in the `twainstance<instance_number>.TWA.properties` file. The file is located in `/etc/TWA`.

Alternatively, you can also proceed as follows:

1. Browse to `<TWA_home>/TWS` path.
2. Source the `./twc_env.sh` shell script.
3. Type `echo $UNISONWORK`. As a result, the path to the `TWA_DATA_DIR` is returned.

IBM Workload Scheduler installation path

You can install more than one IBM Workload Scheduler component (master domain manager, backup master domain manager, domain manager, or backup domain manager) on a system, but each is installed in a separate instance of IBM Workload Automation, as described above.

The installation path of IBM Workload Scheduler is:

```
<TWA_home>/TWS
```

DWC_home installation path

The Dynamic Workload Console can be installed in the path of your choice, but the default installation path is as follows:

On Windows™ operating systems

```
%ProgramFiles%\wa\DWC
```

On UNIX™ operating systems

```
/opt/wa/DWC
```

On z/OS operating system

```
/opt/wa/DWC
```

IBM Workload Scheduler agent installation path

The agent also uses the same default path structure, but has its own separate installation directory:

```
<TWA_home>/TWS/ITA/cpa
```



Note: The agent also installs some files outside this path. If you have to share, map, or copy the agent files (for example when configuring support for clustering) share, map, or copy these files, as well:

On UNIX™ operating systems

```
/etc/teb/teb_tws_cpa_agent_<tws_user>.ini
/opt/IBM/CAP/EMICPA_default.xml
/etc/init.d/tebctl-tws_cpa_agent_<tws_user>
(on Linux and Solaris)
/etc/rc.d/init.d/tebctl-tws_cpa_agent_<tws_user>
```



```
(on AIX)
/sbin/init.d/tebctl-tws_cpa_agent_<tws_user>
(on HP-UX)
```

On Windows™ operating systems

```
%windir%\teb\teb_tws_cpa_agent_<tws_user>.ini
%ALLUSERSPROFILE%\IBM\CAP\EMICPA_default.xml
```

The agent uses the following configuration files which you might need to modify:

JobManager.ini

This file contains the parameters that tell the agent how to run jobs. You should only change the parameters if advised to do so in the IBM Workload Scheduler documentation or requested to do so by IBM Software Support. Its path is:

On UNIX™ operating systems

```
TWA_DATA_DIR/ITA/cpa/config/JobManager.ini
```

On Windows™ operating systems

```
TWA_home\TWS\ITA\cpa\config\JobManager.ini
```

JobManagerGW.ini

When a dynamic agent is installed and **-gateway** `local|remote` is specified, then this file contains the same parameters as the `JobManager.ini` file except for the following differences:

- The **ResourceAdvisorUrl** parameter points to the dynamic workload broker, and not the master domain manager.

The `JobManagerGW.ini` file is installed in the following location:

On UNIX™ operating systems

```
TWA_DATA_DIR/ITA/cpa/config/JobManagerGW.ini
```

On Windows™ operating systems

```
TWA_home\TWS\ITA\cpa\config\JobManagerGW.ini
```

ita.ini

This file contains parameters which determine how the agent behaves. Changing these parameters may compromise the agent functionality and require it to be reinstalled. You should only change the parameters if advised to do so in the IBM Workload Scheduler documentation or requested to do so by IBM Software Support. Its path is:

On UNIX™ operating systems

```
TWA_DATA_DIR/ITA/cpa/ita/ita.ini
```

On Windows™ operating systems

```
TWA_homeTWS\ITA\cpa\config\ita.ini
```

Installation path for files giving the dynamic scheduling capability

The files that give the dynamic scheduling capability are installed in the following path:

```
<TWA_home>/TDWB
```

The command line client installation path

The command line client is installed outside all *IBM Workload Automation instances*. Its default path is:

```
TWA_home/TWS/CLI
```

However, the information above supplies only the **default** paths. To determine the actual paths of products and components installed in your IBM Workload Automation instances, see [Finding out what has been installed in which IBM Workload Automation instances on page 17](#)

Finding out what has been installed in which IBM Workload Automation instances

How to identify which IBM Workload Scheduler components are installed on a computer.

About this task

If you are not the installer of IBM Workload Scheduler and its components, you might not know what components have been installed, and in which instances of IBM Workload Automation. Follow this procedure to find out:

1. Access the following directory:

UNIX™ and Linux™ operating systems

```
/etc/TWA
```

Windows™ operating systems

```
%windir%\TWA
```

2. List the contents of the directory. Each IBM Workload Automation instance is represented by a file called:

`twainstance<instance_number>.TWA.properties`. These files are deleted when all the products or components in an instance are uninstalled, so the number of files present indicates the number of valid instances currently in use.

3. Open a file in a text viewer.



Attention: Do not edit the contents of this file, unless directed to do so by IBM Software Support. Doing so might invalidate your IBM Workload Scheduler environment.

The contents are similar to this on a master domain manager :

```
#TWAInstance registry
#Tue Feb 26 09:28:08 EST 2021
TWA_path=/opt/wa/server_twsuser
TWA_componentList=TWS
TWS_version=10.1 Fix Pack 6
TWS_counter=1
TWS_instance_type=MDM
```

```
TWS_basePath=/opt/wa/server_twsuser/TWS
TWS_user_name=twsuser
TWS_wlpdir=/opt/wa/wlpEngine/wlp
TWS_datadir=/opt/wa/server_twsuser/TWSDATA
TWS_jdbcdir=/opt/wa/server_twsuser/TWS/jdbcdrivers/db2
```

The contents are similar to this on the Dynamic Workload Console:

```
#TWAInstance registry
#Tue Feb 26 09:42:10 EST 2021
TWA_path=/opt/wa/DWC
TWA_componentList=DWC
DWC_version=10.1 Fix Pack 6
DWC_counter=1
DWC_instance_type=DWC
DWC_basePath=/opt/wa/DWC
DWC_user_name=dwcadmin
DWC_wlpdir=/opt/wa/wlpDWC/wlp
DWC_datadir=/opt/wa/DWC/DWC_DATA
DWC_jdbcdir=/opt/wa/DWC/jdbcdrivers/derby
```

The important keys to interpret in this file are:

TWA_path

This is the base path, to which the installation added one or more of the following directories, depending on what was installed:

TWS

Where the IBM Workload Scheduler component is installed

DWC

Where the Dynamic Workload Console is installed

ssm

Where the Netcool® SSM monitoring agent is installed (used in event management)

TWA_componentList

Lists the components installed in the instance of IBM Workload Automation.

TWS_counter

Indicates if an IBM Workload Scheduler component is installed in this instance of IBM Workload Automation (when the value=1).

TWS_instance_type

Indicates which component of IBM Workload Scheduler is installed in this instance:

MDM

Master domain manager

BKM

Backup master domain manager

DDM

dynamic domain manager

FTA

Fault-tolerant agent or domain manager

TWS_user_name

The ID of the <<TWS_user>> of the IBM Workload Scheduler component.

TWS_wlplibdir

The installation directory of the WebSphere Application Server Liberty Base instance used by IBM Workload Scheduler.

TWS_datadir

The directory containing product data and data generated by IBM Workload Scheduler, such as logs and configuration information.

DWC_counter

Indicates if an instance of Dynamic Workload Console is installed in this instance of IBM Workload Automation (when the value=1)

DWC_user_name

The ID of the Dynamic Workload Console user.

DWC_wlplibdir

The installation directory of the WebSphere Application Server Liberty Base instance used by Dynamic Workload Console.

DWC_datadir

The directory containing product data and data generated by Dynamic Workload Console, such as logs and configuration information.

Built-in troubleshooting features

A list, brief description and links to more information on the tools and facilities which are built in to the product to facilitate troubleshooting.

IBM Workload Scheduler is supplied with the following features that assist you with troubleshooting:

- Informational messages that inform you of expected events.
- Error and warning messages that inform you of unexpected events.
- Message helps for the most commonly-occurring messages. See *IBM® Workload Automation: Messages and Codes*.
- A logging facility that writes all types of messages to log files, which you use to monitor the progress of IBM Workload Scheduler activities. See [IBM Workload Scheduler logging and tracing using CCLog on page 26](#).

- Various tracing facilities which record at varying levels of details the IBM Workload Scheduler processes for troubleshooting by IBM® Software Support. See [Difference between logs and traces on page 24](#) for more details.
- A facility to save a configurable level of log and tracing information in memory and then save all or part of this information to a single fully integrated file for troubleshooting by IBM® Software Support. See [In-Flight Trace facility for engine on page 54](#) for more details.
- An auditing facility that provides an audit trail of changes to the IBM Workload Scheduler database and plan for use in both monitoring and troubleshooting. For more details, see the section about Auditing in the Administration Guide.
- A configuration snapshot facility that you can use for backup, and also which provides IBM® Software Support with configuration information when unexpected events occur. See [Data capture utility on page 44](#).
- A facility that automatically creates a First Failure Data Capture (ffdc) configuration snapshot if the failure of any of the key components can be detected by its parent component. See [First failure data capture \(ffdc\) on page 52](#).
- An automatic backup mechanism of the `Symphony` file whereby each fault-tolerant agent and domain manager that receives a new `Symphony` file, automatically archives the previous `Symphony` to `Symphony.last` in the path `<TWA_home/TWS/`, so that a backup copy is always maintained. This permits viewing of the previous `Symphony` data in case there were any message updates on the job and job stream states that were lost between the agent and its master domain manager.
- A problem determination capability available from the Dynamic Workload Console to determine why jobs that are ready to start, do not start and the solution. See [Jobs in READY status do not start on page 171](#).

Keeping up-to-date with the latest fix packs

Reminds you that the best way to avoid problems is to apply fix packs

IBM Workload Scheduler fix packs contain fixes to problems that IBM®, you, or other customers have identified. Install the latest fix pack when it becomes available, to keep the product up to date.

Upgrading your whole environment

When upgrading, although compatibility with previous version components is a feature of IBM Workload Scheduler, potential problems can be avoided by upgrading all components to the new level as quickly as possible.

To avoid problems, ensure that when you upgrade to a new version of IBM Workload Scheduler you do so across your whole environment.

The components of this version of IBM Workload Scheduler are compatible with components of many previous versions (see *IBM® Workload Automation: Overview* for full details). However, running IBM Workload Scheduler in a mixed network increases the possibility of problems arising, because each new release of IBM Workload Scheduler not only adds functions, but also improves the stability and reliability of the various components. Try not to run in a mixed network for extended periods.

Chapter 2. Logging and tracing

Provides detailed information about logs and traces, and how to customize them and set the logging and tracing levels.

Information on the logging and tracing facilities of IBM Workload Scheduler, Dynamic Workload Console, and the WebSphere® Application Server is described in these topics:

- [Quick reference: how to modify log and trace levels on page 21](#)
- [Difference between logs and traces on page 24](#)
- [IBM Workload Scheduler logging and tracing using CCLog on page 26](#)
- [Dynamic Workload Console log and trace files on page 31](#)
- [Log file for the Self-Service Dashboard application on page 34](#)
- [Dynamic workload scheduling log and trace files on page 34](#)
- [Dynamic agent log and trace files on page 35](#)
- [Log and trace files for the application server on page 40](#)
- [Collect trace information on page 39](#)
- [Retrieving IBM Workload Scheduler agent traces from the Dynamic Workload Console on page 33](#)

Quick reference: how to modify log and trace levels

Quick reference information about how to modify log and tracing levels for all components.

Modifying IBM Workload Scheduler logging level

1. Edit the `TWSCCLog.properties` file, located in the following path:

On Windows operating systems

`<TWA_home>\TWS\`

On UNIX operating systems

`<TWA_DATA_DIR>/TWS/`

2. Modify the `tws.loggers.msgLogger.level` property.

This determines the type of messages that are logged. Change this value to log more or fewer messages, as appropriate, or on request from IBM® Software Support. Valid values are:

INFO

All log messages are displayed in the log. The default value.

WARNING

All messages except *informational* messages are displayed.

ERROR

Only *error* and *fatal* messages are displayed.

FATAL

Only messages which cause IBM Workload Scheduler to stop are displayed.

3. Save the file. The change is immediately effective.
4. Modify the `Conman.ini` to customize logging levels for conman commands.

The `Conman.ini` file is located in the following path:

On UNIX™ operating systems

`TWA_DATA_DIR/ITA/cpa/config`

On Windows™ operating systems

`TWA_home\TWS\ITA\cpa\config`

5. Save the file. The change is immediately effective.

See [Engine log and trace customization on page 27](#) for more details.

Modifying IBM Workload Scheduler tracing level

1. Edit the `TWSCCLog.properties` file, located in the following path:

On Windows operating systems

`<TWA_home>\TWS\`

On UNIX operating systems

`<TWA_DATA_DIR>`

2. Modify `twc.loggers.trc<component>.level`.

This determines the type of trace messages that are logged. Change this value to trace more or fewer events, as appropriate, or on request from IBM® Software Support. Valid values are:

DEBUG_MAX

Maximum tracing. Every trace message in the code is written to the trace logs.

DEBUG_MID

Medium tracing. A medium number of trace messages in the code is written to the trace logs.

DEBUG_MIN

Minimum tracing. A minimum number of trace messages in the code is written to the trace logs.

INFO

All *informational*, *warning*, *error* and *critical* trace messages are written to the trace. The default value.

WARNING

All *warning*, *error* and *critical* trace messages are written to the trace.

ERROR

Only *error* and *critical* messages are written to the trace.

CRITICAL

Only messages which cause IBM Workload Scheduler to stop are written to the trace.

3. Save the file. The change is immediately effective.
4. Modify the `Conman.ini` to customize logging levels for conman commands.

The `Conman.ini` file is located in the following path:

On UNIX™ operating systems

`TWA_DATA_DIR/ITA/cpa/config`

On Windows™ operating systems

`TWA_home\TWS\ITA\cpa\config`

5. Save the file. The change is immediately effective.

See [Engine log and trace customization on page 27](#) for more details.

Modifying Dynamic Workload Console tracing level

To activate the Dynamic Workload Console traces at run time, locate and edit the `trace.xml` template.

To modify the `trace.xml` template, follow this procedure:

1. Copy the template file from the `templates` folder to a working folder.
2. Edit the template file in the working folder with the desired configuration.
3. Optionally, create a backup copy of the relevant configuration file present in the `overrides` directory in a different directory. Ensure you do not copy the backup file in the path where the template files are located.
4. Copy the updated template file to the `overrides` folder. Maintaining the original folder structure is not required.
5. Changes are effective immediately.

See: [Activating and deactivating traces in Dynamic Workload Console on page 32](#) for more details.

Modifying WebSphere Application Server Liberty Base tracing level

To modify the trace level on the WebSphere Application Server Liberty Base, edit the `trace.xml` file as necessary.

1. Copy the template file from the `templates` folder to a working folder.
2. Edit the template file in the working folder with the desired configuration.
3. Optionally, create a backup copy of the relevant configuration file present in the `overrides` directory in a different directory. Ensure you do not copy the backup file in the path where the template files are located.
4. Copy the updated template file to the `overrides` folder. Maintaining the original folder structure is not required.
5. Changes are effective immediately.

See [Setting the traces on the application server for the major IBM Workload Scheduler processes on page 40](#) for more details.

Managing dynamic agent tracing level

To manage traces for the dynamic agent, refer to the following sections:

- [See command usage and verify version on page 37](#)
- [Enable or disable trace on page 37](#)
- [Set trace information on page 38](#)
- [Show trace information on page 38](#)
- [Collect trace information on page 39](#)

You can also configure the traces when the agent is not running by editing the [JobManager.Logging] section in the `JobManager.ini` file as described in [Configuring the agent](#) section. This procedure requires that you stop and restart the agent.

Difference between logs and traces

Describes the difference between log and trace messages, and indicates in which languages they are available.

IBM Workload Scheduler and the Dynamic Workload Console create both log and trace messages:

Log messages

These are messages that provide you with information, give you warning of potential problems, and inform you of errors. Most log messages are described in *Messages and Codes*. Log messages are translated into the following languages:

- Chinese - simplified
- Chinese - traditional
- French
- German
- Italian
- Japanese
- Korean
- Portuguese - Brazilian
- Spanish

Messages are written to the log file in the language of the locale set on the computer where they were generated, at the moment when they were generated.

Trace messages

These are messages for IBM® Software Support that provide in depth information about IBM Workload Scheduler processes. In most cases they are in English. Whereas log messages are written so that you can

understand them in relation to the activity you were performing, trace messages might not be. There is no guarantee that you can diagnose any error situations from the information they contain.

The traces are provided at several different levels and in several different forms:

Messages for IBM® Software Support

These are similar to log messages, and while not intended for customer use, can be sometimes helpful to experienced customers who know the product well. The information they contain is used by IBM® Software Support to understand problems better.

Specific software traces

These are traces written directly by the program code normally indicating the values of variables being used in complex processes. They are not for use by the customer.

Automatic software traces

These are traces issued automatically by the code when it enters and exits code modules. They are not for use by the customer.

The following table gives more detailed information:

Table 2. Difference between logs and traces

Characteristics	Log Messages	Messages for IBM® Software Support	Specific software traces	Automatic software traces
Translated	✓			
Documented in IBM Documentation	✓	Some		
Written to	✓			
Windows operating systems				
<code><TWA_home>\TWS\stdlist\logs</code>				
UNIX operating systems				
<code>TWA_DATA_DIR/stdlist/logs/</code>				
Written to	✓	✓	✓	
Windows operating systems				
<code><TWA_home>\TWS\stdlist\traces</code>				
UNIX operating systems				
<code>TWA_DATA_DIR/stdlist/traces/</code>				

Table 2. Difference between logs and traces (continued)

Characteristics	Log Messages	Messages for IBM® Software Support	Specific software traces	Automatic software traces
Logging level, format etc. controlled by <code>TWSCCLog.properties</code>	✓	✓	✓	
Logging level, format etc. controlled by TWSFullTrace				✓
Optionally written to memory by TWSFullTrace and written to disc by that utility when requested.	✓	✓	✓	✓

If you want to merge the logs and traces controlled by `TWSCCLog.properties` into one file, set the **localopts** option `merge stdlist` to `yes`.

IBM Workload Scheduler logging and tracing using CCLog

Describes the log and trace files created by the CCLog logging engine, and how they are configured.

CCLog is a logging engine that creates log files in a defined structure. It can be used to monitor many products from a variety of software suppliers. The configuration supplied with IBM Workload Scheduler uses it uniquely for the processes of IBM Workload Scheduler.

The CCLog engine is used wherever any of the following components are installed:

- Master domain manager
- Backup master domain manager
- Fault-tolerant agent

The contents of this section are as follows:

- [Engine log and trace file locations on page 26](#)
- [Engine log and trace file switching on page 27](#)
- [Engine log and trace customization on page 27](#)
- [Engine log and trace performance on page 31](#)

Engine log and trace file locations

Describes where to find the engine log and trace files produced by CCLog.

All log and trace files produced by IBM Workload Scheduler are stored in:

On Windows operating systems

`TWA_home\TWS\stdlist`

On UNIX operating systems

`TWA_DATA_DIR/stdlist`

The files have different names, depending on the settings in the `localopts` file:

merge stdlists = yes

- `<yyyymmdd>_NETMAN.log`

This is the log file for netman.

- `<yyyymmdd>_TWSMERGE.log`

This is the log file for all other processes.

merge stdlists = no

`<yyyymmdd>_<process_name>.log`

where `<process_name>` can be one of the following values:

```
APPSRVMAN
BATCHMAN
CONNECTR
JOBMAN
JOBMON
MAILMAN
NETMAN
WRITER
```

Low-level traces, and open source library messages that do not conform to the current standard IBM Workload Scheduler message format (for instance, some SSL stdout and stderr messages), are found in the following files: `<yyyy.mm.dd>/<process_name>`, where `<process_name>` is as above. For more information, see *User's Guide and Reference*.



Note: You can add a local option `restricted stdlists` to your `localopts` file to limit access to the `stdlist` directory on your UNIX™ workstation. For details, see *Administration Guide*.

Engine log and trace file switching


Describes when new log and trace files with the next day's datestamp are created.

The IBM Workload Scheduler log files are switched every day, creating new log files with the new datestamp, at the time set in the `startOfDay` global options (`optman`).

Engine log and trace customization

Describes how you can customize the CCLog logging and tracing facility. You can modify the appearance of the log and the logging and tracing levels.

You can customize the information written to the log files by modifying selected parameters in its properties file. The changes you can make affect the format of the log or trace file and the logging level or trace level.

 **Attention:** Do not change any parameters in this file other than those detailed here, otherwise you might compromise the logging facility.

The CCLog properties file is located in the following path:

On Windows operating systems

`<TWA_home>/TWS`

On UNIX operating systems

`TWA_DATA_DIR`

Parameters

The parameters that can be modified are as follows:

Logging level

`twsl.loggers.msgLogger.level`

This determines the type of messages that are logged. Change this value to log more or fewer messages, as appropriate, or on request from IBM® Software Support. Valid values are:

INFO

All log messages are displayed in the log. The default value.

WARNING

All messages except *informational* messages are displayed.

ERROR

Only *error* and *fatal* messages are displayed.

FATAL

Only messages which cause IBM Workload Scheduler to stop are displayed.

You can customize the logging and tracing level for conman commands by setting the relevant values for the **`twsl.loggers.trcConman.level`** property, available in the `Conman.ini`.

The `Conman.ini` file is located in the following path:

On UNIX™ operating systems

`TWA_DATA_DIR/ITA/cpa/config`

On Windows™ operating systems

`TWA_home\TWS\ITA\cpa\config`

Tracing level**twc.loggers.trc<component>.level**

This determines the type of trace messages that are logged. Change this value to trace more or fewer events, as appropriate, or on request from IBM® Software Support. Valid values are:

DEBUG_MAX

Maximum tracing. Every trace message in the code is written to the trace logs.

DEBUG_MID

Medium tracing. A medium number of trace messages in the code is written to the trace logs.

DEBUG_MIN

Minimum tracing. A minimum number of trace messages in the code is written to the trace logs.

INFO

All *informational*, *warning*, *error* and *critical* trace messages are written to the trace. The default value.

WARNING

All *warning*, *error* and *critical* trace messages are written to the trace.

ERROR

Only *error* and *critical* messages are written to the trace.

CRITICAL

Only messages which cause IBM Workload Scheduler to stop are written to the trace.

Log format parameters

fomatters.basicFmt.dateTimeFormat

This contains a specification of the date and time format used by CCLog when adding the date and time stamp to the message header. The format uses the standard *strftime* format convention, used by many programming libraries. The full format details can be found by searching the Internet, but a synthesis of the commonly used definitions is included in [Date and time format reference - strftime on page 206](#).

fomatters.basicFmt.separator

This defaults to the pipe symbol "|", and is used to separate the header of each log message, which contains information such as the date and time stamp and the process that issued the error, from the body, which contains the process-specific information such as the issuing process, the message number and the message text. You can change the separator to another character or characters, or set it to null.

twshnd.logFile.className

This indicates if CCLog uses semaphore memory to write to the log file. The default setting (`ccg_filehandler`) tells CCLog to write each line of a multiline message separately. Each process interleaves each line of its multiline messages with messages from other processes, if necessary, improving performance. While this approach could potentially make the log files more difficult to read, this interleaving only occurs in extreme situations of very high use, for example when many jobs are running concurrently.

The setting `ccg_multiproc_filehandler`, defines that each process completes writing any log message, including multiline messages, before freeing the log file for another process to use. This can have an impact on performance when many processes are running concurrently.

twshnd.loggers.className

This indicates the type of log layout you want to use, determining the number of fields in the log record header. The default setting (`ccg_basiclogger`) tells CCLog to put just the date/time stamp and the process name in the header. The alternative setting is `ccg_pdlogger`, which contains more information in the header, thus reducing the length of the log records available for the message text.

twshnd.loggers.organization

This defaults to *IBM®* and is used to differentiate between log entries from applications from different suppliers when the same instance of CCLog is being used by more than one software supplier. IBM Workload Scheduler is supplied with a unique instance, and thus unique log files, so if this value is prefixed to your log messages, you can set the value of this parameter to null to avoid it being displayed.

twsl.loggers.product

This defaults to *TWS* and is used to differentiate when the same log files are used by more than one product. IBM Workload Scheduler is supplied with unique log files, so if this value is prefixed to your log messages, you can set the value of this parameter to null to avoid it being displayed.

Other parameters

No other parameters must be modified. To do so risks compromising the logging or tracing facility, or both.

Making changes effective

Making your changes effective depends on the type of change:

Changes to log or trace levels

If you change the `twsl.loggers.msgLogger.level` or the `twsl.loggers.trc<component>.level`, the change is immediately effective after the file has been saved.

All other changes

Restart IBM Workload Scheduler to make overall changes effective; restart a process to make process-specific changes effective.

Engine log and trace performance

Describes what impact logging and tracing has on the product's performance.

If you use the default configuration, CCLog does not normally have a significant impact on performance. If you believe that it is impacting performance, check that the default values for the parameters `twslHnd.logFile.className` and `twslloggers.className` are as described in [Engine log and trace customization on page 27](#), and have not been set to other values.

However, even if the default parameters are in use, you might find that in situations of very heavy workload, such as when you have many jobs running simultaneously on the same workstation, multiline log messages become interleaved with messages from other processes. The length of log messages has been increased to offset this risk, but if you find it becoming a problem, contact IBM® Software Support for advice on how to reset the previous settings, which avoided the interleaved messages, but had an impact on performance at busy times.

Dynamic Workload Console log and trace files

This section describes the Dynamic Workload Console log and trace files, where to find them, and how to modify log and tracing levels.

[Table 3: Locations of log files and trace files on page 32](#) lists the log and trace files created by the Dynamic Workload Console:

Table 3. Locations of log files and trace files

Path	Files	Content
<p>Dynamic Workload Console:</p> <p>On Windows</p> <p><DWC_home> \stdlist\appserver\dwcServer\logs</p> <p>On UNIX:</p> <p><DWC_DATA_dir>/stdlist/appserver/dwcServer/logs</p>	<p>console.log, messages.log trace.log</p>	<p>The Dynamic Workload Console run time logs and traces.</p>
<p>On Windows</p> <p><DWC_home>\logs</p> <p>On UNIX</p> <p><DWC_DATA_dir>/installation/logs</p>	<p>dwcinst_<version_number>.log</p>	<p>The Dynamic Workload Console installation log.</p>

Activating and deactivating traces in Dynamic Workload Console

Describes how to activate or deactivate the Dynamic Workload Console traces.

Traces activation and deactivation

About this task

This task activates Dynamic Workload Console traces.

To activate the Dynamic Workload Console traces at run time, locate and edit the `trace.xml` template as described below.

Templates for the Dynamic Workload Console are stored in the following paths:

On UNIX operating systems

`DWC_home/usr/servers/dwcServer/configDropins/templates`

On Windows operating systems

`DWC_home\usr\servers\dwcServer\configDropins\templates`

When you edit the file with your customized settings for the Dynamic Workload Console, move it to the following paths:

On UNIX operating systems

`DWC_DATA_dir/usr/servers/dwcServer/configDropins/overrides`

On Windows operating systems

`DWC_home\usr\servers\dwcServer\configDropins\overrides`

To modify the `trace.xml` template, follow this procedure:

1. Copy the template file from the `templates` folder to a working folder.
2. Edit the template file in the working folder with the desired configuration.
3. Optionally, create a backup copy of the relevant configuration file present in the `overrides` directory in a different directory. Ensure you do not copy the backup file in the path where the template files are located.
4. Copy the updated template file to the `overrides` folder. Maintaining the original folder structure is not required.
5. Changes are effective immediately.

IBM Workload Scheduler agent trace files

Describes how to collect trace files for the IBM Workload Scheduler agent.

You can collect log and trace files for the IBM Workload Scheduler agent, by performing the following actions:

When the agent is running:

By using the `twstrace` command as described in [twstrace command on page 37](#).

When the agent is stopped:

By configuring the `[JobManager.Logging]` section in the `JobManager.ini` file.

For more information, see the section about Configuring log message properties `[JobManager.Logging.cilog]` in *IBM Workload Scheduler: Administration Guide*.

This procedure requires that you stop and restart the dynamic agent.

Using the Dynamic Workload Console

Log in to the Dynamic Workload Console and proceed as described in [Retrieving IBM Workload Scheduler agent traces from the Dynamic Workload Console on page 33](#).

Retrieving IBM Workload Scheduler agent traces from the Dynamic Workload Console

Describes how to collect trace files for the IBM Workload Scheduler agent from the Dynamic Workload Console.

About this task

You can collect trace files for the IBM Workload Scheduler agent, by performing the following actions:

1. Log in to the Dynamic Workload Console.
2. From the **Monitoring and Reporting** menu, click **Monitor Workload**.
3. In the **Engine** field, select one or more engines.
4. In the **Object Type** field, Select **Workstation**.
5. in the **More Actions** menu, select **Collect Agent Logs**.

A .zip file is created containing all agent traces.



Note:

An error message might be displayed if the size of the .zip file is very large. In this case, try and reduce the size of the log files on the agent workstation.

Log file for the Self-Service Dashboard application

The log file for the Self-Service Dashboards mobile application can be configured in the Dynamic Workload Console global settings file. The log file is enabled by default.

For more information, see the section about auditing mobile app activity in *Dynamic Workload Console User's Guide*

The log file is written to the following path:

On UNIX operating systems

`<DWC_DATA_dir>/stdlist/appserver/dwcServer/audit_SSD.log`

On Windows operating systems

`<TWA_home>\stdlist\appserver\dwcServer\audit_SSD.log`

Dynamic workload scheduling log and trace files

The logs and traces produced by the dynamic workload scheduling processes are in most part included in the log and trace files of the IBM Workload Scheduler master domain manager. In addition, the files listed in [Table 4: Locations of log and trace files on page 34](#) also contain log and trace material from these processes.

Table 4. Locations of log and trace files

Component	Path	Trace files	Log files	Content
IBM Workload Scheduler master domain manager	On Windows systems <TWA_home> \stdlist\appserver\engineServer\logs	native_stderr.log	messages.log	Additional log files used by dynamic workload scheduling
		native_stdout.log		
		serverStatus.log		
	On UNIX systems <TWA_DATA_DIR>/stdlist/appserver/engineServer/logs	startServer.log		
		stopServer.log		
		SystemErr.log		
		trace.log		
IBM Workload Scheduler agent	TWA_home /TWS/stdlist/JM	JobManager_trace.log	JobManager_message.log	Log and trace files
		ita_trace.log	ita_message.log	
	TWA_home /TWS/stdlist/JM/JOBMANAGER-FFDC/yy-mm-dd/		JobManager_message.log	Processing error log file

Dynamic agent log and trace files

Describes how to collect log and trace files for the agent.

You can collect log and trace files for the agent, by performing the following actions:

When the agent is running:

By using the `twstrace` command as described in [twstrace command on page 37](#).

When the agent is stopped:

By configuring the [JobManager.Logging] section in the `JobManager.ini` file.

For more information, see the section about Configuring log message properties [JobManager.Logging.cclg] in *IBM Workload Scheduler: Administration Guide*.

This procedure requires that you stop and restart the dynamic agent.

The log messages are written in the following file:

On Windows operating systems:

`<TWA_home>\TWS\stdlist\JM\JobManager_message.log`

On UNIX and Linux operating systems:

`<TWA_DATA_DIR>/stdlist/JM/JobManager_message.log`

The trace messages are written in the following file:

On Windows operating systems:

- `<TWA_home>\TWS\stdlist\JM\ITA_trace.log`
- `<TWA_home>\TWS\stdlist\JM\JobManager_trace.log`
- `<TWA_home>\TWS\JavaExt\logs\javaExecutor0.log`

On UNIX and Linux operating systems:

- `<TWA_DATA_DIR>/stdlist/JM/ITA_trace.log`
- `<TWA_DATA_DIR>/stdlist/JM/JobManager_trace.log`
- `<TWA_DATA_DIR>/JavaExt/logs/javaExecutor0.log`

Logging information about job types with advanced options

You can use the `logging.properties` file to configure the logging process for job types with advanced options, with the exception of the Executable and Access Method job types.

The `logging.properties` file is located on the IBM® Z Workload Scheduler Agent, located in the following path:

On Windows operating systems:

`<TWA_home>/TWS/JavaExt/cfg/logging.properties`

On UNIX and Linux operating systems:

<TWA_DATA_DIR>/JavaExt/cfg/logging.properties

After installation, this file is as follows:

```
# Specify the handlers to create in the root logger
# (all loggers are children of the root logger)
# The following creates two handlers
handlers = java.util.logging.ConsoleHandler,
           java.util.logging.FileHandler

# Set the default logging level for the root logger
.level = INFO

# Set the default logging level for new ConsoleHandler instances
java.util.logging.ConsoleHandler.level = INFO

# Set the default logging level for new FileHandler instances
java.util.logging.FileHandler.level
    = ALL
java.util.logging.FileHandler.pattern
    = C:\TWA_home\TWS\JavaExt\logs\javaExecutor%g.log
java.util.logging.FileHandler.limit
    = 1000000
java.util.logging.FileHandler.count
    = 10

# Set the default formatter for new ConsoleHandler instances
java.util.logging.ConsoleHandler.formatter =
    java.util.logging.SimpleFormatter
java.util.logging.FileHandler.formatter =
    java.util.logging.SimpleFormatter

# Set the default logging level for the logger named com.mycompany
com.ibm.scheduling = INFO
```

You can customize:

- The logging level (from INFO to WARNING, ERROR, or ALL) in the following keywords:

`.level`

Defines the logging level for the internal logger.

`com.ibm.scheduling`

Defines the logging level for the job types with advanced options. To log information about job types with advanced options, set this keyword to ALL.

- The path where the logs are written, specified by the following keyword:

`java.util.logging.FileHandler.pattern`

Configuring trace properties when the agent is running

Use the **twstrace** command to set the trace on the agent when it is running.

Using the **twstrace** command, you can perform the following actions on the agent when it is running:

- [See command usage and verify version on page 37.](#)
- [Enable or disable trace on page 37.](#)
- Set the traces to a specific level, specify the number of trace files you want to create, and the maximum size of each trace file. See [Set trace information on page 38.](#)
- [Show trace information on page 38.](#)
- Collect trace files, message files, and configuration files in a compressed file using the command line. See [Collect trace information on page 39.](#)
- Collect trace files, message files, and configuration files in a compressed file using the Dynamic Workload Console. See [Retrieving IBM Workload Scheduler agent traces from the Dynamic Workload Console on page 33.](#)

You can also configure the traces when the agent is not running by editing the [JobManager.Logging] section in the `JobManager.ini` file as described in [Configuring the agent](#) section. This procedure requires that you stop and restart the agent.

twstrace command

Use the **twstrace** command to configure traces, and collect logs, traces, and configuration files (ita.ini and jobManager.ini) for agents. You collect all the information in a compressed file when it is running without stopping and restarting it.

See command usage and verify version

To see the command usage and options, use the following syntax.

Syntax

```
twstrace -u | -v
```

Parameters

-u

Shows the command usage.

-v

Shows the command version.

Enable or disable trace

To set the trace to the maximum or minimum level, use the following syntax.

Syntax

```
twstrace -enable | -disable
```

Parameters

-enable

Sets the trace to the maximum level. The maximum level is **1000**.

-disable

Sets the trace to the minimum level. The minimum level is **3000**.

Set trace information

To set the trace to a specific level, specify the number of trace files you want to create, and the maximum size the trace files can reach, use the following syntax.

Syntax

```
twstrace [ -level <level_number> ] [ -maxFiles <files_number> ] [ -maxFileBytes <bytes_number> ]
```

Parameters

-level <level_number>

Sets the trace level. Specify a value in the range from 1000 to 3000, which is also the default value. Note that if you set this parameter to 3000, you have the lowest verbosity level and the fewest trace messages. To have a better trace level, with the most verbose trace messages and the maximum trace level, set it to **1000**.

-maxFiles <files_number>

Specify the number of trace files you want to create.

-maxFileBytes <bytes_number>

Set the maximum size in bytes that the trace files can reach. The default is **1024000** bytes.

Show trace information

To display the current trace level, the number of trace files, and the maximum size the trace files can reach, use the following syntax.

Syntax

```
twstrace -level | -maxFiles | -maxFileBytes
```

Parameters

-level

See the trace level you set.

-maxFiles

See the number of trace files you create.

-maxFileBytes

See the maximum size you set for each trace file

Example

Sample

The example shows the information you receive when you run the following command:

```
twstrace -level -maxFiles -maxFileBytes
AWSITA176I The trace properties are: level="1000",
max files="3", file size="1024000".
```

Collect trace information

To collect the trace files, the message files, and the configuration files in a compressed file, use the following syntax.

Syntax

```
twstrace -getLogs [ -zipFile <compressed_file_name> ] [ -host <host_name> ] [ -protocol {http | https} ] [ -port <port_number> ] [ -iniFile <ini_file_name> ]
```

Parameters

-zipFile <compressed_file_name>

Specify the name of the compressed file that contains all the information, that is logs, traces, and configuration files (ita.ini and jobManager.ini) for the agent. The default is **logs.zip**.

-host <host_name>

Specify the host name or the IP address of the agent for which you want to collect the trace. The default is **localhost**.

-protocol http|https

Specify the protocol of the agent for which you are collecting the trace. The default is the protocol specified in the **.ini** file of the agent.

-port <port_number>

Specify the port of the agent. The default is the port number of the agent where you are running the command line.

-iniFile <ini_file_name>

Specify the name of the **.ini** file that contains the SSL configuration of the agent for which you want to collect the traces. If you are collecting the traces for a remote agent for which you customized the security certificates, you must import the certificate on the local agent and specify the name of the **.ini** file that contains this configuration. To do this, perform the following actions:

1. Extract the certificate from the keystore of the remote agent.
2. Import the certificate in a local agent keystore. You can create an ad hoc keystore whose name must be **TWSClientKeyStore.kdb**.
3. Create an **.ini** file in which you specify:

- **0** in the **tcp_port** property as follows:

```
tcp_port=0
```

- The port of the remote agent in the **ssl_port** property as follows:

```
ssl_port=<ssl_port>
```

- The path to the keystore you created in Step 2 on page 39 in the **key_repository_path** property as follows:

```
key_repository_path=<local_agent_keystore_path>
```

Log and trace files for the application server

The log and trace files for WebSphere Application Server Liberty Base can be found in:

Application server run time log and trace files

On UNIX:

- `<TWA_DATA_DIR>/stdlist/appserver/engineServer/logs/messages.log`
- `<TWA_DATA_DIR>/stdlist/appserver/engineServer/logs/trace.log`

On Windows:

- `<TWA_home>\TWS\stdlist\appserver\engineServer\logs\messages.log`
- `<TWA_home>\TWS\stdlist\appserver\engineServer\logs\trace.log`

Trace files containing messages related to the plan replication in the database

On UNIX:

- `<TWA_DATA_DIR>/stdlist/appserver/engineServer/logs/PlanEventMonitor.log.0`
- `<TWA_DATA_DIR>/stdlist/appserver/engineServer/logs/PlanEventMonitor.log.1`

On Windows:

- `<TWA_home>\TWS\stdlist\appserver\engineServer\logs\PlanEventMonitor.log.0`
- `<TWA_home>\TWS\stdlist\appserver\engineServer\logs\PlanEventMonitor.log.1`

Setting the traces on the application server for the major IBM Workload Scheduler processes

About this task

The application server handles all communications between the IBM Workload Scheduler processes. The trace for these communications is set to `tw_s_info` by default (information messages only). The application server can be set to trace `all` communications, either for the whole product or for these specific groups of processes:

- Command line
- Connector
- Database
- Planner
- Utilities
- Dynamic workload broker

Significant impact on performance: Activating traces for the WebSphere Application Server Liberty Base leads to a significant impact on performance, especially if you set the tracing to `all`. Thus you are strongly advised to identify the process group where the problem that you want to trace is occurring, and only set the trace to that group.

To modify the trace level on the WebSphere Application Server Liberty Base, edit the `trace.xml` file as necessary.

Templates for the master domain manager are stored in the following paths:

On UNIX operating systems

`TWA_home/usr/servers/engineServer/configDropins/templates`

On Windows operating systems

`TWA_home\usr\servers\engineServer\configDropins\templates`

Templates for the Dynamic Workload Console are stored in the following paths:

On UNIX operating systems

`DWC_home/usr/servers/dwcServer/configDropins/templates`

On Windows operating systems

`DWC_home\usr\servers\dwcServer\configDropins\templates`

When you edit the file with your customized settings for the master domain manager, move it to the following paths:

On UNIX operating systems

`TWA_DATA_DIR/usr/servers/engineServer/configDropins/overrides`

On Windows operating systems

`TWA_home\usr\servers\engineServer\configDropins\overrides`

When you edit the file with your customized settings for the Dynamic Workload Console, move it to the following paths:

On UNIX operating systems

`DWC_DATA_dir/usr/servers/dwcServer/configDropins/overrides`

On Windows operating systems

`DWC_home\usr\servers\dwcServer\configDropins\overrides`

1. Copy the template file from the `templates` folder to a working folder.
2. Edit the template file in the working folder with the desired configuration.

3. Optionally, create a backup copy of the relevant configuration file present in the `overrides` directory in a different directory. Ensure you do not copy the backup file in the path where the template files are located.
4. Copy the updated template file to the `overrides` folder. Maintaining the original folder structure is not required.
5. Changes are effective immediately.

For example, to modify the trace level on WebSphere Application Server Liberty Base, perform the following steps:

1. Copy the `trace.xml` file from the `TWA_home/usr/servers/engineServer/configDropins/templates` folder to a working folder.
2. Edit the template file in the working folder by changing the following string:

```
<variable name="trace.specification" value="*info"/>
```

into

```
<variable name="trace.specification" value="com.ibm.tws.dao.model.
=all:com.ibm.tws.dao.rdbms.=all"/>
```

3. Copy the updated template file to `TWA_DATA_DIR/usr/servers/engineServer/configDropins/overrides`. Changes are effective immediately.

Traces are stored in `DWC_DATA_dir/appserver/engineServer/logs`.

The **trace.specification** can be found in `trace.xml` files and can refer to a specific component (`tws_XXX`) or to the whole product, as follows:

tws_all

```
"com.ibm.tws.=all:org.apache.wink.server.=all:com.hcl.tws.*=all"
```

tws_alldefault

```
"com.ibm.tws.*=error=enabled"
```

tws_broker_all

```
"com.ibm.scheduling.*=all:TWSAgent=all"
```

tws_rest

```
"com.ibm.tws.conn.=all:com.ibm.tws.twsd.rest.=all:org.apache.wink.server.*=all"
```

tws_cli

```
"com.ibm.tws.cli.=all:com.ibm.tws.objects.=all"
```

tws_utils

```
"com.ibm.tws.util.*=all"
```

tws_conn

```
"com.ibm.tws.conn.=all:com.ibm.tws.objects.=all:com.ibm.tws.updatemanager.=all:com.ibm.tws.dao.plan.=a
ll"
```

tws_db

```
"com.ibm.tws.dao.model.=all:com.ibm.tws.dao.rdbms.=all"
```

tws_planner

```
"com.ibm.tws.planner.=all:com.tivoli.icalendar.=all:com.ibm.tws.runcycles.=all:com.ibm.tws.conn.planner.=all  
:com.ibm.tws.cli.planner.*=all"
```

tws_secjni

```
"com.ibm.tws.audit.=all:com.ibm.tws.security.=all"
```

tws_engine_broker_all

```
"com.ibm.tws.=all:com.ibm.scheduling.=all:TWSAgent=all"
```

Editing the logging element above with the **traceSpecification** value to `tws_all`, enables

```
"com.ibm.tws.=all:org.apache.wink.server.=all:com.hcl.tws.*=all".
```

Other values are reported in variable tags above. You can also replace the value of the **trace.specification** parameter with a custom string.

Chapter 3. Capturing data in the event of problems

Describes the facilities available for data capture in the event of problems occurring. It provides full details of the Data capture utility and the provisions for first failure data capture.

In the event of any problems occurring while you are using IBM Workload Scheduler, you might be asked by the IBM® Customer Support to supply information about your system that might throw a light on why the problem occurred. The following utilities are available:

- A general data capture utility command that extracts information about IBM Workload Scheduler and related workstations, system-specific information, and data related to WebSphere Application Server Liberty Base; see [Data capture utility on page 44](#).
- A first failure data capture (ffdc) facility built into **batchman** and **mailman** that automatically runs the data capture utility when failures occur in **jobman**, **mailman**, or **batchman**; see [First failure data capture \(ffdc\) on page 52](#).
- WebSphere Application Server Liberty Base javadump command to create the heap dump for WebSphere Application Server Liberty Base that runs on the Dynamic Workload Console and the master domain manager; see [Creating application server dumps on page 52](#).

Data capture utility

The data capture utility is a script named `wa_pull_info` which extracts information about a product instance of IBM Workload Scheduler.

The data capture utility script is located in the following path:

On Windows operating systems

`%windir%`

On UNIX operating systems

`TWA_home/TWS/bin`

and can be run from the UNIX™ or DOS prompt on the master domain manager, the backup master domain manager, or agents.

On HP and Solaris agents, use the `tw_inst_pull_info` command, located in the `TWA_home/TWS/bin` directory.

When to run the utility

Describes the circumstances in which you would use the data capture utility.

Use the data capture utility in these circumstances:

- An IBM Workload Scheduler process has failed, but the automatic ffdc facility has not detected the failure and run the script for you (see [First failure data capture \(ffdc\) on page 52](#))
- IBM Workload Scheduler is very slow or is behaving in any other abnormal way
- You are requested to do so by IBM® Software Support

Using the utility when you need to switch to the backup master domain manager

If the master domain manager fails you might decide that you want to switch to the backup master domain manager to keep your scheduling activities running. If you also want to run the data capture utility you have two choices:

Data capture first

Run the data capture utility first to ensure that the information extracted is as fresh as possible. Then run **switchmgr**.

Switchmgr first

In an emergency situation, where you must continue scheduling activities, run **switchmgr** immediately and then run the data capture utility on both the new master domain manager and the new backup master domain manager as soon as **switchmgr** has completed.

Prerequisites

Describes the prerequisites for running the `wa_pull_info` data capture utility.

Where the utility can be run

The utility can be run on the master domain manager, the backup master domain manager or a standard or fault-tolerant agent.

Who can run it

The utility must be run by one of the following users:

- Any IBM Workload Scheduler user
- Root (recommended on UNIX™ or Linux™ systems)
- Administrator (on Windows™ systems)

To determine the best user to run the script, make the following considerations:

Troubleshooting any type of problem

- On UNIX™ operating systems, the user running the script must have read access to the `/etc` and `/etc/TWS` directories and read access to the `/etc/TWS/TWSRegistry.dat` file

Troubleshooting installation problems

- On UNIX™ operating systems, run the script as root to ensure to gather all installation information.

Troubleshooting problems when the product is running

- The script will only extract database object descriptions to which the user running it has EXTRACT permission in the Security file. The *<TWS_User>* (the user who performed the installation) normally has full access to all database objects, so this is the best user to run the script.
- The IBM Workload Scheduler instance must have a Symphony file otherwise some information will not be extracted.

Command and parameters

Describes the command syntax and parameters of the data capture utility.

Command syntax

Run the data capture utility with the following command:

```

wa_pull_info -?

wa_pull_info
-user userid
[-component DWC|TWS]
[-date yyyy-mm-dd]
[-output output_path]
[-workdir working_directory]
    
```

This is the syntax for UNIX® operating systems; on Windows™ use `wa_pull_info.exe`

Parameters

-user

The product user that you specify when you install the component. This user must exist in the registry file located on the master domain manager in the following path:

On Windows operating systems

C:\Windows

On UNIX operating systems

/etc/TWS

or on the Dynamic Workload Console in the following path:

On Windows operating systems

C:\Windows\TWA

On UNIX operating systems

/etc/TWA

This parameter is mandatory.

-component

The component whose data you want to capture. Supported values are `DWC` and `TWS`. The default value is `TWS`. This parameter is optional.

-date

Used as the base date for collected data logs. If not specified, the script uses the current date by default and the day before the current date. Run the data capture utility as soon as a problem occurs, to collect the data specific to the date and time of the problem. Thus, if the problem occurs on the current date, this option is not required. If the problem occurred earlier, then the date on which the problem occurred must be specified in the `yyyy-mm-dd` format. Either the current date or the specified date is used to identify which files and logs are extracted. This parameter is optional.

-output

The base directory location where the collected data is stored. Ensure you have write access to the specified directory. This parameter is optional. The default value is as follows:

On Windows operating systems

```
\tmp\wadatagather\output
```

On UNIX operating systems

```
/tmp/wadatagather/output
```

Ensure you do not enter a path that contains the `stdlist` folder located in `TWA_DATA_DIR`.

-workdir

The working directory used by the command for storing data while running. When the command stops running, the working directory is deleted. Ensure you have write access to the specified directory and enough space is available. This parameter is optional. The default value is as follows:

On Windows operating systems

```
\tmp\wadatagather\workdir
```

On UNIX operating systems

```
/tmp/wadatagather/workdir
```

Ensure you do not enter a path that contains the `stdlist` folder located in `TWA_DATA_DIR`.

-?

Displays the usage of the command.

Comments

Regardless of the parameters you specify, the command always returns at least the following information:

`datagather_summary.log`

A summary of the data retrieved by the command.

system_info

A folder containing several files which store system info.

registry_info

A folder containing IBM® Workload Scheduler registry files, if available.

Examples

Consider the following examples:

```
<twS_inst_path>/TWS/bin/wa_pull_info -component TWS -user <twS_user>
```

```
<dwC_inst_path>/DWC/tools/wa_pull_info -component DWC -user <dwC_user>
```

Tasks

Describes the tasks performed by the data capture utility.

Check that the user exists

The script verifies if the specified user exists in the `TWSRegistry.dat` file. If it does, the `<TWS_HOME>` directory used for data collection is extracted from the `TWSRegistry.dat` file. (UNIX™ only) If the specified user does not exist, the script verifies if the user exists in the `/etc/passwd` file. If no user exists, the script terminates.

Check the user permissions

The commands that are used during the data collection try to retain the original ownership of the files; when the script is run on Solaris platforms, the ownership of the files might change. If the script is run by a IBM Workload Scheduler user (for example, not the root user) the script collects the available instance data.



Note:

Some Windows™ security policies can affect which data is extracted.

Create the directories in which to store the collected data

The script first creates the `<log_dir_base>` directory, where `<log_dir_base>` is the value provided for the `-log_dir_base` option. Within the `<log_dir_base>` directory, the script creates the `twS_info` directory and its subdirectories `twS_yyyymmdd_hhmmss`, where yyyy=year, mm=month, dd=day, hh=hour, mm=minute and ss=seconds.

Collect data

The script collects system and product-specific data, as described in [Data collection on page 49](#).

Create the TAR file

UNIX™

The script creates the TAR file `TWS_yyyyymmdd_hhmmss.tar` and compresses it to `TWS_yyyyymmdd_hhmmss.tar.Z`, or if the operating system is Linux_i386, `TWS_yyyyymmdd_hhmmss.tar.gz`.

Windows™

On Windows™ operating systems there is no built-in compression program, so the script does not create a compressed file. If you intend to send the data to IBM® Software Support you should use your own compression utility to create the compressed archive.

Data collection

Describes the data collected by the data capture utility.

System-specific data

For system-specific data, the script performs the following operations:

- Extracts local CPU node information
- Extracts the environment for the current IBM Workload Scheduler instance
- Extracts `nslookup` information for local CPU
- Extracts `netstat` information for local CPU
- Extracts Services information
- Extracts the current running processes
- (UNIX® only) Extracts a list of the files and directories under `/usr/Tivoli/TWS`
- Extracts the current available disk space for `%TWA_HOME%`
- Extracts the current available disk space for the `tmp` directory
- (UNIX® only) Extracts the current system disk space
- (UNIX® only) Extracts the current disk space of root `filesystem`
- (Solaris 10.x or above) Extracts `zonecfg` information
- (AIX® only) Copies `netsvc.conf`
- (UNIX® only, except AIX®) Copies the `nsswitch.*` files
- Copies the host and services files

IBM Workload Scheduler-specific data

For IBM Workload Scheduler-specific data, the script performs the following operations:

Collects IBM Workload Scheduler messages, as follows:

- Generates a list of the `.msg` files
- Extracts a list of the files in the following path:

On Windows™ operating systems

`<TWA_home>\TWS\ftbox`

On UNIX® operating systems

`<TWA_DATA_DIR>/ftbox`

Collects IBM Workload Scheduler information, as follows:

- Extracts information about the IBM Workload Scheduler instance installation
- Extracts the IBM Workload Scheduler Security file
- Extracts a list of all files under the `%TWA_HOME%` directory
- Extracts the database definitions to `flatfiles`
- (UNIX® only) Extracts the `optman` output
- (UNIX® only) Extracts `planman` "showinfo" output
- Copies `jobmanrc.cmd` and `jobmanrc` (if it exists)
- Copies the `schedlog` files of the previous day (the option `-date` is not used)
- Copies the `schedlog` files of the day on which the problem occurred, day - 1 and day + 1 (the option `-date` is used)
- Copies files located in the following path:

On Windows™ operating systems

`<TWA_home>TWS\audit\database\plan\${today} & ${yesterday}`

On UNIX® operating systems

`<TWA_DATA_DIR>\audit\database\plan\${today} & ${yesterday}`

- Copies the `BmEvents.conf` file and the event log (if `%TWA_HOME%\BmEvents.conf` exists)
- Copies the content of the `BmEvents` log file (if `%TWA_HOME%\BmEvents.conf` exists)
- Copies the `TWSRegistry.dat` file
- (UNIX® only) Copies all files from `/etc/TWA`, `/tmp/TWA*`, `/tmp/twsinst*`, `/tmp/tws9*`, and `TWA_DATA_DIR/stdlist/logs`.
- Copies the content of the `%TWA_HOME%\version` directory
- Copies the files of the local workstation (the master domain manager and the backup master domain manager are also workstations on which jobs can be scheduled)
- (Windows™ only) If the z/OS® connector is installed locally, copies the `TWSZOSConnRegistry.dat` file

Collects IBM Workload Scheduler logs, as follows:

- Copies the `TWSUser`, `BATCHUP`, `NETMAN`, `TWSMERGE`, and `joblog stdlist` files for current and previous date
- Copies the `TWSMERGE` and `NETMAN` log files from the `stdlist\logs` directory for current and previous date
- Copies the `TWSMERGE BATCHUP` and `NETMAN stdlist` files from the `stdlist\traces` directory for current and previous date
- Collects output of various conman commands: `sc`, `sj`, `ss`



Note: The `NETMAN` log files also contain information about the `mailman` process.

Collects IBM Workload Scheduler files, as follows:

- Extracts a list of the files in the `%TWA_HOME%\ITA` directory
- Extracts a list of the files in the `%TWA_HOME%\stdlist\JM` directory
- Extracts a list of the files in the `%TWA_HOME%\jmJobTableDir` directory
- Copies all the files in the `%TWA_HOME%\stdlist\JM` directory
- Copies all the files in `%TWA_HOME%\jmJobTableDir`

Collect xtrace information from IBM Workload Scheduler processes as follows:

- Generates snapshot files for IBM Workload Scheduler processes in raw format
- Generates snapshot files in XML format from the raw format

If IBM Workload Scheduler for Applications is installed on the workstation, collects data on the methods, as follows:

- Copies the content of the `%TWA_HOME%\methods` directory (if it exists)
- (Windows™ only) Collects information about the `Peoplesoft` method
- Collects information about the `r3batch` method
- (UNIX® only) Collects the `r3batch` picklist results

WebSphere Application Server Liberty Base-specific data

For WebSphere Application Server Liberty Base-specific data, the script performs the following operations:

- Extracts a list of the IBM Workload Scheduler server files specific to WebSphere Application Server Liberty Base
- Collects the IBM Workload Scheduler server Liberty profile ()
- Copies select IBM Workload Scheduler application files specific to WebSphere®
- Collects all configuration files:
 - `datasource.xml`
 - `ssl_variables.xml`
 - `wauser_variables.xml`
 - `ports_variables.xml`
 - `host_variables.xml`
 - `jvm.options`
 - `authentication_config.xml`
- Collects the `javacore*.txt` files from the path `<PROFILE_HOME>`
- Collects the data source properties
- Collects the host properties
- Collects the security properties

First failure data capture (ffdc)

Describes how the data capture tool is used automatically by components of the product to create a *first failure data capture* of the products logs, traces and configuration files.

To assist in troubleshooting, several modules of the product have been enabled to create a *first failure data capture* in the event of failure. This facility uses the data capture tool `wa_pull_info` (see [Data capture utility on page 44](#)) to copy logs, traces, configuration files and the database contents (if the database is on DB2®) and create a compressed file that you can send to IBM® Software Support.

This tool is run in the following circumstances:

Jobman fails

If batchman detects that jobman has failed, it runs the script, placing the output in
`<TWA_home>/stdlist/yyyy.mm.dd/collector/JOBMAN`

Batchman fails

If mailman detects that batchman has failed, it runs the script, placing the output in
`<TWA_home>/stdlist/yyyy.mm.dd/collector/BATCHMAN`

Mailman fails

If mailman detects that it itself has failed with a terminal error, it runs the script, placing the output in
`<TWA_home>/stdlist/yyyy.mm.dd/collector/MAILMAN`. Note that process hard stops, for example, segmentation violations, are not tracked by mailman itself.

Netman child process fails

If netman detects that one of its child processes has failed, it runs the script, placing the output in
`<TWA_home>/stdlist/yyyy.mm.dd/collector/NETMAN`

Only one data collection is kept for each day. Each day a new data collection overwrites the previous day's collection.

Within each of the target output directories, the output file is stored in the `/tws_info/TWS_YYYYMMDD_HHMMSS` directory.

To perform `ffdc`, the `wa_pull_info` script is run by a script called `collector.sh (.cmd)`. You can customize this script (located in `<TWA_home>/TWS/bin`) to apply different parameters to the `wa_pull_info` script for any of the enabled modules (jobman, mailman, batchman and netman)

Creating application server dumps

About this task

See the related WebSphere Application Server Liberty Base documentation at [How to generate javacores/thread dumps, heapdumps and system cores for the WebSphere Application Server Liberty profile](#).

The following is an example of how to create the heap dump for both the Dynamic Workload Console and the master domain manager:

1. Set the environment by running the `setEnv.sh` script located in the path:

Dynamic Workload Console

On Windows™ operating systems:

```
DWC_INST_DIR\DWC\appservertools\setEnv.bat
```

On UNIX™ operating systems:

```
DWC_DATA_dir/DWC/appservertools/setEnv.sh
```

master domain manager

On Windows™ operating systems:

```
TWA_home\server_<wauser>\appservertools\setEnv.bat
```

On UNIX™ operating systems:

```
TWA_DATA_DIR/server_<wauser>/appservertools/setEnv.sh
```

2. Submit the following command to create the heap dump:

Dynamic Workload Console and master domain manager

On Windows™ operating systems:

```
WLP_INST_DIR\bin\server javadump <dwc_Server_Name> --include=heap
```

On UNIX™ operating systems:

```
WLP_INST_DIR//bin/server javadump <dwc_Server_Name> --include=heap
```

Results

You can locate the dump file in the following path:

Dynamic Workload Console

On Windows™ operating systems:

```
DWC_INST_DIR\stdlist\appserver\dwcServer
```

On UNIX™ operating systems:

```
DWC_DATA_dir/stdlist/appserver/dwcServer
```

master domain manager

On Windows™ operating systems:

```
TWA_home\stdlist\appserver\engineServer
```

On UNIX™ operating systems:

```
TWA_DATA_DIR/stdlist/appserver/engineServer
```

Chapter 4. In-Flight Trace facility for engine

Describes the tracing facility for troubleshooting the IBM Workload Scheduler engine. This facility is called In-Flight Trace.

This document describes the IBM Workload Scheduler server tracing facility that replaced Autotrace from version 8.6. The facility is designed to be used by IBM® Software Support, but is fully described here so that you understand how to use it if requested to do so by IBM® Software Support.

The IBM Workload Scheduler server tracing facility (hereafter called In-Flight Trace) is a facility used by IBM® Software Support to help solve problems in IBM Workload Scheduler. At maximum capacity it can trace the entry into and exit from every IBM Workload Scheduler function, plus many other events, and includes all log and trace messages currently issued by the CCLog facility.

In-Flight Trace has been conceived as a multi-product tool, although this description concentrates on its use for IBM Workload Scheduler.

It works as follows:

Existing trace calls

In-Flight Trace uses the logging and tracing facilities still used by the CCLog logging and tracing mechanism, and which were used by the Autotrace facility in releases before 8.6.

Function entry and exit

In addition, the IBM Workload Scheduler engine product build now inserts trace calls in the code to record the entry to and exit from every function and assigns a sequential numeric function ID to each function. The trace calls use these IDs to identify the functions.

Building the xdb.dat symbols database

During the same process, the build creates the **xdb.data** symbols database associating the name of each function with the function ID. In this way, the trace writes the minimum information possible to the trace record (the function ID), which can then be expanded to give the function name later for viewing.

The build also stores in the database the source file and line number of each function.

Further, it stores the name of the component which "owns" the function. One program contains many components, each of which contains many functions.

The symbols database is the key to managing the activation/deactivation and filtering of the traces. The information it contains is encrypted.

Tracing in shared memory

The traces are written to shared memory. This is divided into segments, and the traces chosen to be written to each segment are written in an endless loop. At maximum capacity (tracing all events on all functions) the traces might loop every few seconds, while at minimum capacity (tracing just one little-used function), the trace might not loop for months.

Segments

You can choose to use any number of segments (each is identified by a unique number) and for each segment can determine how much shared memory the segment is to use. More and bigger segments consume more memory, with all the normal consequences that entails.

Programs

Any number of IBM Workload Scheduler programs can be configured to be saved to the same segment. You decide which programs are to be traced to which segments, and whether those segments are to be enabled for tracing, by modifying the basic configuration. Any of the IBM Workload Scheduler programs and utilities can be configured for tracing.

Basic configuration

The basic configuration determines which segments are enabled for tracing, and makes an initial determination of whether the tracing for a specific program is activated. It is achieved by editing a configuration file with a text editor. The IBM Workload Scheduler engine (the product) must be restarted to make the changes take effect. The configuration is divided into the following sections:

Global

This section not only includes general information like the product code and the segment size, but also acts as a "catch-all", where traces from programs not specifically configured are configured.

<program>

If a program is not to be traced under the "global" section, a specific program section must be configured, defining which segment the program is to be traced in, and other basic information. The information in a program section overrides that in the global section, but just for that program.

Activating and deactivating traces

For segments which are enabled, traces for specific programs can be activated and deactivated on-the-fly, from the command line, as these flags are held in memory.

Trace levels

Events in the code have been assigned trace levels. The lower the level, the more drastic the event. The levels range from reporting only unrecoverable errors, through recoverable errors, warnings, and informational messages and three debug levels to the maximum reporting level, where even function entry and exit events are recorded.

Trace levels can also be changed on-the-fly, from the command line, without restarting the engine.

Snapshots

In-Flight Trace lets you take a snapshot of the current contents of the traces for a program or segment and save it to a file. You can optionally clear the memory in the segment after taking the snapshot. The snapshot file is in the internal format, containing function IDs, etc., and is not easily readable. It must be formatted to make it readable.

Formatting the snapshot

A command-line option lets you format a snapshot file for the standard output. The output can be in CSV or XML format, and information about the source data (file name and line number) is automatically included. Or you can select the standard trace format (one line per trace record) and choose whether to include the source information. And finally you can choose whether to include the header information (ideal for a printed output) or not (ideal for the creation of a file you are going to analyze programmatically).

Filtering

The tooling-up of the code is a fully automatic process and you might find that your traces include frequently used components or functions that are not causing any problems. You would like to exclude them from the trace and you do this by using the command line to create a filter file, in which you can specify to include all and then exclude any combinations of specific components, functions, and source files. Alternatively, you can exclude all and then include any combinations of specific components, functions, and source files. Functions can also be included or excluded by specifying a range of function IDs.

Once created, a filter file is declared either in the global section of the configuration file or one of the program sections. You can have more than one filter file which you use with different programs, however, note that the filter is applied at segment level. This means that if you have two programs writing to the same segment, the filter is applied to both even if it is only specified for one of them.

Existing filter files can be modified from the command line.

Products

In-Flight Trace is conceived as a multi-product facility. Each product has its own separate configuration file. Multiple instances of the facility can be run on the same system, completely independently of each other. However, you can also control one product from the tracing facility of another, by identifying the product to which to apply the commands. For example, if you had two versions of IBM Workload Scheduler running on the same system, you could control the In-Flight Trace facility for both of them from one place, inserting the appropriate product code when required by the command syntax.

In-Flight Trace configuration file

Describes the In-Flight Trace configuration file, `xtrace.ini`.

The In-Flight Trace configuration file is used to initialize the shared memory at product startup. The information in shared memory determines which traces are saved at which level. All function trace calls are parsed by the trace facility to determine if they should be saved.

The In-Flight Trace configuration file is found in the following path:

```
<TWA_home>/TWS/xtrace/xtrace.ini
```

An example of the file is as follows:

```
[ _GLOBAL_ ]
Product      = <PRODUCT>
Enabled      = y
```

```

Active      = y
SegNum      = 1
FilterFile  = $(install_dir)/bin/xfull.xtrace
SegSize     = 10240
Level       = 80
SegPerUser  = n

[netman]
Enabled     = y
Active      = y
SegNum      = 2
Level       = 80

```

Changing the configuration

Describes how to modify the configuration file.

About this task

To permanently change the configuration in shared memory that controls the tracing, edit the file, save it, and restart the product. On UNIX™ platforms you must also clean up the memory by running the tracing command with the **-clean** parameter between stopping and restarting the product. Thus, the procedure to change the configuration file is as follows:

UNIX™

1. Modify the configuration file
2. Save the configuration file
3. Stop the product
4. Run `xcli -clean`
5. Restart the product

Windows™

1. Modify the configuration file
2. Save the configuration file
3. Stop the product
4. Restart the product

You can change much of the configuration in shared memory that controls the tracing by using the `xcli` command (see [xcli command syntax on page 61](#)). However, any changes made in this way are not updated in the configuration file, so at the next initialization, unless you have specifically edited the file, the parameters used are those that were in the file last time you restarted the product.

Configuration file syntax

The configuration file is divided into sections.

Each section begins with a header in one of the two following formats (the square brackets are required, they are not command syntax indicators):

```
[ _GLOBAL_ ]
[<program>]
```

[_GLOBAL_]

There must be only one [_GLOBAL_] section containing general information about the product and the tracing configuration for all programs that do not have a specific section.

[<program>]

You can define a separate section for each of the IBM Workload Scheduler programs ([<program>]). The following programs are most likely to require tracing:

```
APPSRVMAN
BATCHMAN
JOBMAN
JOBMON
MAILMAN
NETMAN
WRITER
JAVA (the connector)
```

However, you can trace any executable program such as COMPOSER, CONMAN and all the utilities stored in the IBM Workload Scheduler `/bin` directory.

You cannot have more than one instance of a section for the same program.

If a program has no specific section, its trace configuration uses the defaults in the [_GLOBAL_] section. Details defined in the program sections in almost all cases override the corresponding values in the [_GLOBAL_] section (the exception is Product).

The program name is not case-sensitive. For example, you can write `Netman`, `NetMan`, `netman` OR `NETMAN`.



Note: On UNIX™ operating systems, `JOBMAN` and `jobman` are two separate programs performing different functions. This means that on UNIX™ operating systems, because of the case-insensitivity, if you set a trace configuration section for either `JOBMAN` or `jobman`, both programs will be traced using that section and therefore to the same segment. This is a limitation that cannot be avoided at present.

Available keys (each key can be defined only once in each section):

Product

[_GLOBAL_] only. Product identification string. Required.

Enabled

Specifies if the segment is enabled. If you change the enablement of a segment by changing this value and saving the configuration file, you must restart the product to make the change effective. If the segment in the [_GLOBAL_] section is not enabled, the entire tracing facility is disabled. Enter "y" or "n".

Active

Specifies whether tracing for the specific program is active. If the [`_GLOBAL_`] section is not activated, the tracing for all programs without a specific section is not activated. This value can be changed without restarting the product by using the tracing command. Enter "y" or "n".

SegNum

Determines the segment number to use for tracing for a specific section. More than one program can be defined for the same segment in different sections. The `SegNum` specified in the [`_GLOBAL_`] section is used by any program that does not have a specific section defined. If you change the segment number of a program by changing this value and saving the configuration file, you must restart the product to make the change effective. Enter any numeric value.

FilterFile

Specifies the file that contains the criteria for filtering components, functions, or source files. The file is applied at segment level, so you cannot specify different filter files for different programs that use the same segment. This value can be changed without restarting the product by using the tracing command. Enter the fully qualified file path.

The default filter file supplied with the product does not trace the top 5% most-used routines (by being most-used they are less likely to exhibit problems).

SegSize

Specifies the segment size (Kb). If this value is supplied more than once in different sections for the same segment, the trace facility uses the highest of the supplied values. If you change the size of a segment by changing this value and saving the configuration file, you must restart the product. Enter a numeric value.

The full shared memory usage is the sum of all enabled segments, plus several Kbytes for the control data.

On UNIX™, ensure that you do not exceed the configurable kernel parameter which determines the maximum size of shared memory.

Level

Specifies the maximum level to be traced for the program. Enter one of the following numeric values:

Level	Description
10	Unrecoverable
20	Error
30	Warning
40	Informational
50	Debug minimum
60	Debug medium
70	Debug maximum

Level	Description
80	Function entry and exit

If this value is supplied more than once in different sections for different programs that trace in the same segment, the trace facility uses the appropriate values for each program. Thus, the segment might contain traces for one program at level 10 and for another at level 80.

This value can be changed without restarting the product by using the tracing command.

SegPerUser

Specifies if the segment provides access to the owner of the segment only (y) or all users (n). Enter "y" or "n"

In-Flight Trace command: xcli

This topic describes the command used to control all aspects of the run time behavior of In-Flight Trace.

The `xcli` command modifies the information in shared memory. Shared memory is initialized from the information in the configuration file, but any changes to shared memory that are made using the options of this command are *not* saved in the configuration file.

Selecting programs, segments, and products

Describes how to select programs, segments, and products in the `xcli` command.

In many of the parameters of the `xcli` command, you are required to select a program or a segment, and optionally a product. To avoid repeating the same information, details of how to do this are supplied here:

Program

Select a program for a specific action by identifying the global section (`[_GLOBAL_]`) or any of the configuration file sections containing IBM Workload Scheduler programs (`[<program>]`).

Segment

Select one of the segment numbers that were defined in the configuration file when the shared memory was initialized. If you need to use extra segments or redistribute the programs within the segments, you must edit and save the configuration file and then stop and restart the IBM Workload Scheduler engine.

Product

The tracing facility is multi-product. However, if you run the `xcli` command from the same directory as a configuration file, you automatically run it on the product defined in that configuration file, without having to define the product in the command.

But if you are using In-Flight Trace to trace more than one product, and you want to use the command supplied with product A to modify the tracing of product B, you must supply the product code for product B as a parameter to the command, by adding the `-P <product>` parameter to the command string. This parameter is only applicable to the `-snap`, `-query`, `-active`, `-level`, and `-filter` subcommands.

xcli command syntax

Gives the full syntax of the xcli command.

Controls all of the runtime aspects of the In-Flight Trace facility. It modifies the information in shared memory. Shared memory is initialized from the information in the configuration file, but any changes to shared memory made using the options of this command are *not* saved in the configuration file.

You must be the **TWS_user** to run the command.

Syntax

xcli

-snap <snap_file>

{ -p <program> | -s <segment> }

[-descr <description>]

[-clean]

[-P <product>]

-format <snap_file>

-d <symbols_database>

[-full]

[-noHeader]

[-standard [-source] | -xml | -csv]

-query [-p <program> | -s <segment>] [-P <product>]

-active { y | n }

{ -p <program> | -s <segment> | -all }

[-P <product>]

-level <level>

{ -p <program> | -s <segment> | -all }

[-P <product>]

-filter <filter_file>

{ -p <program> | -s <segment> }

[-P <product>]

-createFilter <filter_file> **-d**<symbols_database>

[-add_all |

-add_comp <component> | **-remove_comp** <component> |

-add_func <function_name> | -remove_func <function_name> |
 -add_func_id <function_ID> | -remove_func_id <function_ID> |
 -add_func_id_range <from> <to> | -remove_func_id_range <from> <to> |
 -add_filter <filter_file> | -remove_filter <filter_file>] ...

-modifyFilter <filter_file> -d <symbols_database>
 [-add_all | -remove_all
 -add_comp <component> | -remove_comp <component> |
 -add_func <function_name> | -remove_func <function_name> |
 -add_func_id <function_ID> | -remove_func_id <function_ID> |
 -add_func_id_range <from> <to> | -remove_func_id_range <from> <to> |
 -add_filter <filter_file> | -remove_filter <filter_file>] ...

-clean

-config [<config_file>]

Arguments

-snap <snap_file>

Saves a snapshot of part of the shared memory to the indicated file. For the snapshot, you can use the following parameters:

{ -p <program> | -s <segment> }

Define if the snapshot is for either a program or a segment. If it is made for a program which shares a segment with other programs, the whole segment is snapped, but the header information shows which program it was snapped for. See also [Selecting programs, segments, and products on page 60](#).

[-descr <description>]

Supply a description for the snapshot. Surround it with double quotation marks if it contains spaces.

[-clean]

Optionally clear the entire segment memory after taking the snapshot. If any process is still using the memory, the clean operation cannot be performed and a warning message is given.



Note: If your snapshot is of a program, this option clears the memory for all traces in the segment for which the program is configured, including those of any other programs that have been configured to write to it.

[-P <product >]

See [Selecting programs, segments, and products on page 60](#).

The snap file header information is as follows:

```
"Snap information:
" Product:          <product>
" Description:      <description>
" Snap platform:    <platform>
" Snap time (GMT):  <time>
" Snap program:     <program>
" Snap segment:     <segment>
"   Segment size:   <size>(Kb)
"   Segment use:    <percent_used>
```

-format <snap_file>

Formats the supplied snapshot file for the standard output. The formatting options are:

-d <symbols_database>

Supply the name of the symbols database to use for the formatting. The database must be either the same version as the instance of IBM Workload Scheduler from which the snap was captured (ideally), or a later version. The default symbols database is **xdb.dat**.

[-full]

If the snap was taken of a single program in a multi-program segment, use this option to send the full set of traces (all programs) to the standard output, rather than that of the single program as determined by the header information in the snap file.

[-noHeader]

Use this to suppress the output of the header information. The standard output then just consists of trace messages, which is more acceptable as input to an analysis program.

[-standard [-source] | -xml | -csv]

Define the formatting of the traces. If you have selected **-standard**, use the optional parameter **-source** to add information about the source file and line number. This source information is automatically included in the **-xml** and **-csv** options. If you supply none of these, the format defaults to **-standard**.

-query

Outputs the enablement or activation state of a program or segment. Without parameters, this option displays information about the entire configuration to the standard output. The parameters are:

[-p <program> | -s <segment>]

Optionally define whether the query is for a specific program or a specific segment. See also [Selecting programs, segments, and products on page 60](#).

[-P <product>]

See [Selecting programs, segments, and products on page 60](#).

-active { y | n }

Activates (y) or deactivates (n) a program or segment in memory, or all programs and segments. The parameters are as follows:

{ -p <program> | -s <segment> | -all }

Activate either a specific program or a specific segment, or all programs and segments. See also [Selecting programs, segments, and products on page 60](#).

[-P <product>]

See [Selecting programs, segments, and products on page 60](#).

-level <level>

Sets the tracing level for a program or segment in memory. Specify one of the following level codes:

Level	Description
10	Unrecoverable
20	Error
30	Warning
40	Informational
50	Debug minimum
60	Debug medium
70	Debug maximum
80	Function entry and exit

For example, to trace only unrecoverable failures and errors, supply "20".

The parameters are as follows:

{ -p<program> | -s<segment> | -all }

Set the level for either a specific program or a specific segment, or all programs and segments. See also [Selecting programs, segments, and products on page 60](#).

[-P<product>]

See [Selecting programs, segments, and products on page 60](#).

-filter <filter_file>

Applies a new filter file for a program or segment in shared memory. The parameters are as follows:

{ -p <program> | -s <segment> }

Determine the filter file to be used for either a program or a segment. See also [Selecting programs, segments, and products on page 60](#).

[-P <product>]

See [Selecting programs, segments, and products on page 60](#).

default

The filter file is created using the **-createFilter** option. In this option (and the associated **-modifyFilter** option) you specify any components and functions you want to include or exclude from the tracing (see below for more details). This information is written in the filter file as a list of all functions in the symbols database (by ID) with a bit set to indicate whether they are to be included or excluded. The default symbols database is **xdb.dat**.

Any filter files defined in the configuration file are loaded into shared memory at initialization. If you use this option, the shared memory area is overwritten with the new contents. If the new filter file has been created using a different symbols database than the original file, a warning is given, because it is advisable to use the same symbols database when creating the new filter file.

The default filter file supplied with the product is set to not trace the 5% most-used routines, on the basis that the most-used routines are less likely to create problems because they are well tried and tested.

-createFilter <filter_file>

Creates the filter file named in the parameter. The file must not already exist. There is no facility to view a filter file, so use meaningful names and maintain your own documentation of the contents of each filter file.

To populate the file supply one or more of the following parameters. If you *add* an item, its traces will be saved; if you *remove* an item, its traces will *not* be saved. By default, all components and functions are removed.

-d <symbols_database>

Identify the symbols database to use to verify the component names, and the function names and IDs.

-add_all

Add all components and functions to the filter file. Use this with one of the **-remove** options to create an exclusive "all except ..." filter.

-add_comp <component> | -remove_comp <component>

Add a component to the file or remove one that has already been added. For example, you could add all components using **-add_all** and then remove just one, which would be easier than adding all of the required components individually. Discover component names by viewing a formatted snapshot.

-add_func <function_name> | -remove_func <function_name>

Add a function to the file or remove one that has already been added. For example, you could add a component using **-add_comp** and then remove one of its functions, which would be easier than adding all of the required functions individually. Discover function names by viewing a formatted snapshot.

-add_func_id <function_ID> | -remove_func_id <function_ID>

Adds a function to the file by ID, or removes one that has already been added. For example, you could add a component using **-add_comp** and then remove one of its functions, which would be easier than adding all of the required functions individually. A function ID is a sequential number allocated to a function when the product was built, and stored in the symbols database. Discover function IDs by viewing a formatted snapshot.

-add_func_id_range <from> <to> | -remove_func_id_range <from> <to> |

Adds a range of functions to the file by ID or removes a range of functions that have been already added. Discover function IDs by viewing a formatted snapshot.

-add_filter <filter_file> | -remove_filter <filter_file>

Adds or removes the contents of an existing (different) filter file, as follows:

Adding a filter file

If you add a filter file, the items in that filter file which are set to be filtered (traced) are added to whatever other filter criteria you might have set.

Removing a filter file

If you remove a filter file, the items in the filter file which are set to be filtered (traced) are removed from whatever other filter criteria you might have set.

For example, you might create a filter file that configures the tracing of the communications functions. You could then add this set of functions to your filter set in one command, or remove them, depending on whether you think the communications are part of the problem you are trying to solve.

The add and remove actions are processed in the order you submit them. Thus if you add a function ID, and then remove a range that includes that ID it is removed from the criteria. But if you remove the range, and then add the ID, it is added to the criteria.

-modifyFilter <filter_file>

Modifies the existing filter file named in the parameter.

This subcommand takes all of the parameters used in the **-addFilter** subcommand in the same way, with the addition of the following action:

-remove_all

Removes all components and functions from the filter file. Use this with one of the **-add** options to create an inclusive "all of the following" filter.

-clean

On UNIX™ operating systems only, use this to delete the shared memory segments after you have modified and saved the configuration file, and stopped the product. If a segment is in use, it is marked for deletion and will be automatically deleted when no longer in use.

-config [<config_file>]

This initializes the memory. It is run automatically when the IBM Workload Scheduler engine is restarted, using the default configuration file `./xtrace.ini`. In normal circumstances, you never need to run this manually. If you believe that the shared memory is corrupted, it is better to restart the product, which automatically re-initializes the memory.

Example**Examples**

The following examples are a scenario for using the trace to troubleshoot an instance of IBM Workload Scheduler which is hanging for 5 minutes when you run a particular utility command, without giving any log messages to indicate why.

The presupposition is that you have the following configuration file:

```
[ _GLOBAL_ ]
Product      = 9.3.0-IBM-IWS-00
Enabled      = y
Active       = n
SegNum       = 1
FilterFile   = $(install_dir)/bin/xfull.xtrace
SegSize      = 10240
Level        = 80
SegPerUser   = n

[netman]
Enabled      = y
Active       = n
SegNum       = 2
Level        = 80

[batchman]
Enabled      = y
Active       = n
SegNum       = 3
Level        = 80
```

1. Start the tracing

Tracing is enabled but inactive for three segments. You think the problem is not network related, so netman is not involved. To activate the other two segments run the following commands:

```
xcli -active y -s 1
xcli -active y -s 3
```

2. Adjust the levels for minimum debug

You want to trace as much activity as possible so that you understand what is happening. So you adjust the tracing levels to minimum debug:

```
xcli -level 50 -s 1
xcli -level 50 -s 3
```

3. Take a snapshot when the product hangs

Restart IBM Workload Scheduler and run the utility again. When the product hangs immediately take a snapshot of each segment. You include the option to clean the memory after the snapshot:

```
xcli -snap main_snap -s 1
      -descr "Snap of segment 1 when TWS hangs after using utility" -clean
xcli -snap batchman_snap -s 3
      -descr "Snap of batchman when TWS hangs after using utility" -clean
```

4. Format the trace to view it

Run the following command for a standard format for each file, and save it to a text file:

```
xcli -format main_snap -d xdb.dat > main_snap.txt
xcli -format batchman_snap -d xdb.dat > batchman_snap.txt
```

5. The problem seems to be with batchman, but you need more detail

After examining the two snap files it seems as though the problem is occurring in batchman, but you need more detail:

```
xcli -level 80 -s 3
```

6. Take another snapshot of batchman when the product hangs

Restart IBM Workload Scheduler and run the utility again. When the product hangs immediately take another snapshot of batchman's segment:

```
xcli -snap batchman2_snap -s 3 -descr "Second snap of batchman (level 80)"
```

7. Format the trace again to view it

Run the following command to save the snap file in XML format to a file:

```
xcli -format batchman2_snap -d TWS86SymDB -xml > batchman2_snap.xml
```

You now have a well-formatted XML file of the traces to examine in details and determine where the problem is occurring.

xcli messages

Lists all the messages that might be issued by the xcli command.

This section details the messages that might be produced by xcli and explains what they mean.

Incorrect syntax in configuration file.

In-Flight Trace has found syntax that it cannot parse in the configuration file. Check the syntax carefully with the information in this manual. Correct the error and rerun the command.

Cannot create the semaphore '%d', error %ld.

The error message is from the operating system. There might be a memory usage problem which requires an operating system reboot.

Cannot lock the semaphore, error %ld.

The error message is from the operating system. There might be a memory usage problem which requires an operating system reboot.

Cannot create the shared memory, error %ld.

The error message is from the operating system. There might be a memory usage problem which requires an operating system reboot.

Cannot map the shared memory, error %ld.

The error message is from the operating system. There might be a memory usage problem which requires an operating system reboot.

Incorrect value for key %s in section %s'.

The syntax of the configuration file is correct but the indicated key in the indicated section has an incorrect value.

The tracing facility is not active.

If the "Enable" key in the [_GLOBAL_] section is set to "n", the tracing facility is disabled. To enable it, edit the configuration file, set the "Enable" key in the [_GLOBAL_] section to "y", save the file and restart IBM Workload Scheduler.

Unable to open file %s', error %d.

The error message is from the operating system. Check the error code. The file might be open in another process, or the user running the command might not have rights to open the file. Correct the problem and try the command again.

Not enough free memory to allocate %d bytes.

The message indicates the memory required by your configuration. Either reduce the amount of memory used by the configuration by editing the configuration file, changing the values, saving the file and restarting IBM Workload Scheduler, or free some memory by closing other applications. You might also be able to enlarge the memory paging file. Use the **-config** option to reinitialize the memory.

Unable to write to file %s', error %d.

The error message is from the operating system. Check the error code. The file might have been deleted by another process, or the user running the command might not have rights to write to the file. Correct the problem and try the command again.

Unable to read from file %s', error %d.

The error message is from the operating system. Check the error code. The file might have been deleted by another process, or the user running the command might not have rights to read from the file. Correct the problem and try the command again.

The selected file does not contain a valid snapshot.

You have identified a snapshot file to format, but either it is not a snapshot file or the snapshot file was not written correctly. Check the name you supplied. If it was not correct, reissue the command with the correct file name. If the file name is correct, rerun the snap to regenerate the file.

Memory not correctly initialized.

The shared memory has not been correctly created. Check that there is sufficient free memory to create the shared memory you have defined in the configuration file.

The tracing facility is not active for program %s.

You have requested to change tracing information for the indicated program which is not active. Activate the program first, using the **-active** option.

The tracing facility is not active for segment %s.

You have requested to change tracing information for the indicated segment which is not active. Activate the segment first, using the **-active** option.

Operation successful.

No-brainer! Whatever you were doing has worked!

Unable to remove the semaphore %x, error %d.

The error message is from the operating system. There might be a memory usage problem which requires an operating system reboot.

Unable to remove the shared memory %x, error %d.

The error message is from the operating system. There might be a memory usage problem which requires an operating system reboot.

The tracing facility is not active for product %s.

You have identified a product for which the "Enable" key in the [_GLOBAL_] section is set to "n", and so the tracing facility is disabled. To enable it, edit the appropriate configuration file, setting the "Enable" key in the [_GLOBAL_] section to "y", save the file and restart the product.

The maximum number of products (%d) has already been reached.

In-Flight Trace can only trace a limited number of products at one time, regardless of the amount of memory available. You have reached that limit!

The sections '%s' and '%s' have the same segment number but different %s.

Some of the keys in a section are "segment-based", in that if more than one section traces to the same segment, they must have the same values. For example, the filter file for programs that trace to the same segment must be the same. Either change the programs to trace to different segments or supply the same filter file for all programs that trace to the same segment.

Segment %d is too small for the filter file '%s'.

The space you allocate to a segment is used to store the filter file for more than 12,000 functions, in addition to the traces. In this case, you have not created sufficient space to store the filter file. You cannot change the size of the filter file, because it uses one entry for every function in the product, regardless of whether that function is or is not filtered for tracing. So you must increase the segment size by editing the configuration file, saving it and restarting IBM Workload Scheduler.

Cannot open the symbols database '%s'.

Either the symbols database does not exist with the name you supplied, the user running the command does not have the rights to open the file, or the file is corrupted. Check that the name is correct and ensure that you are the *TWS_user*.

Too many input parameters.

The syntax of the command you supplied is not correct. Check the syntax with what is documented in this publication and try the command again.

An error occurred while opening the symbols database '%s'.

The database file might be corrupted.

Warning: the function ID %d is not in the symbols database.

You have tried to add or remove a function which is not in the symbols database. Check the source from which you obtained the function name or ID. Check that you are using the correct symbols database. The default symbols database is **xdb.dat**. Correct the error and try the command again.

Duplicated section '%s' in the configuration file.

Each program section can only be present once in the configuration file. Perhaps you copied a section intending to change the name but did not. Edit the configuration file, save it and restart IBM Workload Scheduler.

Warning: There is a mismatch between the size of the new filter file and the previous one (new size = %d, previous size = %d).

You have used the **-filter** option to supply a new filter file, but the new filter file was generated using a symbols database different from that used when the filter file currently in use was created. Put a different way, you seem to have used different symbols databases to create two different filter files, and the two different databases have different numbers of functions. In-Flight Trace can continue tracing but the filtering might not be applied correctly. You are advised always to use only the symbols database generated when the version of the product you are tracing was built.

Cannot clean up the shared memory because some process is currently using it.

You have used the **-clean** option to clean the shared memory, but one or more processes is still using the shared memory, so it cannot be cleaned. Use your system resources to determine which process is using the shared memory, stop it, and retry the **-clean** option.

Chapter 5. Troubleshooting performance issues

This refers you to the *Administration Guide* for the resolution of performance problems.

The performance of IBM Workload Scheduler can depend on many factors. Preventing performance problems is at least as important as resolving problems that occur. For this reason, all discussion of performance issues has been placed together in the chapter on performance in the *IBM® Workload Scheduler: Administration Guide*.

Chapter 6. Troubleshooting networks

Describes how to recover from short-term and long-term network outages and offers solutions to a series of network problems.

This section describes how to resolve problems in the IBM Workload Scheduler network. It covers the following topics:

- [Network recovery on page 73](#)
- [Other common network problems on page 76](#)

Network recovery

Several types of problems might make it necessary to follow network recovery procedures. These include:

- Initialization problems that prevent agents and domain managers from starting properly at the start of a new production period. See [Initialization problems on page 73](#).
- Network link problems that prevent agents from communicating with their domain managers. See [Network link problems on page 74](#).
- Loss of a domain manager, which requires switching to a backup. See [Replacement of a domain manager on page 75](#).
- Loss of a master domain manager, which is more serious, and requires switching to a backup or other more involved recovery steps. See [Replacement of a master domain manager on page 76](#).



Note: In all cases, a problem with a domain manager affects all of its agents and subordinate domain managers.

Initialization problems

Initialization problems can occur when IBM Workload Scheduler is started for a new production period. This can be caused by having IBM Workload Scheduler processes running on an agent or domain manager from the previous production period or a previous IBM Workload Scheduler run. To initialize the agent or domain manager in this situation, perform the following steps:

1. For a domain manager, log into the parent domain manager or the master domain manager. For an agent, log into the agent domain manager, the parent domain manager, or the master domain manager.
2. Run the Console Manager and issue a stop command for the affected agent.
3. Run a link command for the affected agent. This initializes and starts the agent.

If these actions fail to work, check to see if netman is running on the affected agent. If not, issue the StartUp command locally and then issue a link command from its domain manager.

If there are severe network problems preventing the normal distribution of the new Symphony file, a fault-tolerant agent or subordinate domain manager can be run as a standalone system, provided the following conditions are met:

- The Sinfonia file was generated on the master domain manager after the network problem occurred, and so has never been transferred to the agent or domain manager
- You have some other method, such as a physical file transfer or FTP to transfer the new Sinfonia file from the master domain manager to the agent or subordinate domain manager.
- The master domain manager and the agent or subordinate domain manager have the same processor architecture.

If these conditions are met, perform the following steps:

1. Stop the agent or domain manager.
2. Delete the `<TWA_home>/TWS/Symphony` file on the agent or domain manager.
3. Copy the file `<TWA_home>/TWS/Sinfonia` from the master domain manager to the `<TWA_home>/TWS` directory on the agent or domain manager.
4. Rename the copied file `<TWA_home>/TWS/Symphony`
5. Run StartUp to start the agent or domain manager.

Any inter-workstation dependencies must be resolved locally using appropriate console manager commands, such as **Delete Dependency** and **Release**.

Network link problems

IBM Workload Scheduler has a high degree of fault tolerance in the event of a communications problem. Each fault-tolerant agent has its own copy of the `Symphony` file, containing the production period's processing. When link failures occur, they continue processing using their own copies of the `Symphony` file. Any inter-workstation dependencies, however, must be resolved locally using appropriate console manager commands: `deldep` and `release`, for example.

While a link is down, any messages destined for a non-communicating workstations are stored by the sending workstations in the `<TWA_home>/TWS/pobox` directory, in files named `<workstation>.msg`. When the links are restored, the workstations begin sending their stored messages. If the links to a domain manager are down for an extended period of time, it might be necessary to switch to a backup (see *IBM® Workload Scheduler: Administration Guide*).



Note:

1. The **conman submit job** and **submit schedule** commands can be issued on an agent that cannot communicate with its domain manager, provided that you configure (and they can make) a direct HTTP connection to the master domain manager. This is configured using the `conman connection` options in the `localopts` file, or the corresponding options in the `useropts` file for the user (see the *IBM® Workload Scheduler: Administration Guide* for details).



However, all events have to pass through the domain manager, so although jobs and job streams can be submitted, their progress can only be monitored locally, not at the master domain manager. It is thus always important to attempt to correct the link problem as soon as possible.

2. If the link to a standard agent workstation is lost, there is no temporary recovery option available, because standard agents are hosted by their domain managers. In networks with a large number of standard agents, you can choose to switch to a backup domain manager.

Troubleshooting a network link problem

About this task

When an agent link fails it is important to know if the problem is caused by your network or by IBM Workload Scheduler. The following procedure is run from the master domain manager to help you to determine which:

1. Try using telnet to access the agent: `telnet <node>:<port>`
2. Try using ping to access the agent: `ping <node>:<port>`
3. Run `nslookup` for the agent and the master domain manager from both, and check that the information on each system is the same from each system
4. Run `netstat -a |grep <port>` and check if any `FIN_WAIT_2` states exist
5. Verify that the port number of the master domain manager matches the entry for "`nm port`" in the `localopts` file of the master domain manager
6. Verify that the port number of the agent matches the entry for "`nm port`" in the `localopts` file of the agent
7. Check the `netman` and `TWSMerge` logs on both the master domain manager and the agent, for errors.



Note:

1. Any issues found in steps [1 on page 75](#) to [4 on page 75](#) suggest that there are problems with the network
2. Any issues found in steps [5 on page 75](#) to [7 on page 75](#) suggest that there are problems with the IBM Workload Scheduler configuration or installation

If this information does not provide the answer to the linking issue, call IBM® Software Support for further assistance.

The commands used in steps [1 on page 75](#) to [4 on page 75](#) are IP network management commands, information about which can be obtained in the Internet. The following technical note also provides useful information about their use:

<http://www.ibm.com/support/docview.wss?rs=0&uid=swg21156106>

Replacement of a domain manager

A domain manager might need to be changed as the result of network linking problems or the failure of the domain manager workstation itself. It can be temporarily replaced by switching any full status agent in its domain to become the new domain manager, while the failed domain manager is repaired or replaced.

The steps for performing this activity are as described for the planned replacement of a domain manager; see *IBM® Workload Scheduler: Administration Guide*.

Replacement of a master domain manager

If you lose a master domain manager, you have to perform all of the steps described in *IBM® Workload Scheduler: Administration Guide* for the planned replacement of a master domain manager.

Other common network problems

The following problems could be encountered:

- [Using SSL, no connection between a fault-tolerant agent and its domain manager on page 76](#)
- [After changing SSL mode, a workstation cannot link on page 77](#)
- [In a configuration with a firewall, the start and stop remote commands do not work on page 78](#)
- [The domain manager cannot link to a fault-tolerant agent on page 78](#)
- [Agents not linking to master domain manager after first JnextPlan on HP-UX on page 79](#)
- [Fault-tolerant agents not linking to master domain manager on page 79](#)
- [The dynamic agent cannot be found from Dynamic Workload Console on page 81](#)
- [Submitted job is not running on a dynamic agent on page 81](#)
- [Job status of a submitted job is continually shown as running on dynamic agent on page 81](#)
- [Network performance on page 81](#)
- [AWSITA245E - Agent is down but jobmanager is running on page 82](#)

Using SSL, no connection between a fault-tolerant agent and its domain manager

In a network using SSL authentication, no connection can be established between a fault-tolerant agent and its domain manager. The standard lists of the two workstations display messages like in the following examples.

- On the domain manager, mailman messages:

```
+ ++++++
+ AWSBCV082I Workstation FTAHP, Message: AWSDEB009E Data
+ transmission is not possible because the connection is broken.
+ The following gives more details of the error: Error 0.
+ ++++++
+ AWSBCV035W Mailman was unable to link to workstation: rsmith297;
+ the messages are written to the PO box.
+ ++++++
```

- On the fault-tolerant agent, writer messages:

```

/* *****
/* AWSBCW003E Writer cannot connect to the remote mailman. The
/* following gives more details of the error: "
/* AWSDEB046E An error has occurred during the SSL handshaking. The
/* following gives more details of the error: error:140890B2:SSL
/* routines:SSL3_GET_CLIENT_CERTIFICATE:no certificate returned
/* *****
/* AWSDEZ003E **ERROR**(cpu secs 0)

```

Cause and solution:

In the localopts file of either the domain manager or the fault-tolerant agent, the SSL port statement is set to 0.

Correct the problem by setting the SSL port number to the correct value in the localopts file. You then need to stop and restart netman on the workstation so that it can now listen on the correct port number.

Timeout during Symphony download - AWSDEB003I Writing socket Resource temporarily unavailable

A dynamic domain manager is installed. The installation completes successfully. The dynamic domain manager is not linked to the master domain manager.

In the TWSMERGE.log of the master domain manager, the following mailman message is displayed:

```

+ *****
+ AWSBCV082I Workstation DDM, Message: AWSDEB003I Writing socket:
+ Resource temporarily unavailable.
+ *****

```

The same problem occurs when installing a backup domain manager or a fault-tolerant agent.

Cause and solution:

In the localopts file of the master domain manager, the mm symphony download timeout is set to 0 minutes.

Correct the problem by setting the mm symphony download timeout to 1 minute.

After changing SSL mode, a workstation cannot link

You have changed the SSL mode between a workstation and its domain manager. However, you are unable to relink to the workstation from the domain manager.

Cause and solution:

The following Symphony file and message files at the workstation must be deleted after a change of SSL mode, otherwise the data does not match:

```
Symphony  
Sinfonia  
$HOME/*.msg  
$HOME/pobox/*.msg
```

In a configuration with a firewall, the start and stop remote commands do not work

In a configuration with a firewall between the master domain manager and one or more domain managers, the start and stop commands from the master domain manager to the fault-tolerant agents in the domains do not work. This is often the case when an "rs final" ends and the impacted fault-tolerant agents are not linked.

Cause and solution:

The fault-tolerant agents belonging to these domains do not have the *behind firewall* attribute set to *on* in the IBM Workload Scheduler database. When there is a firewall between the master domain manager and other domains, start and stop commands must go through the IBM Workload Scheduler hierarchy. This parameter tells the master domain manager that the stop request must be sent to the domain manager which then sends it to the fault-tolerant agents in its domain.

Use either the Dynamic Workload Console or the composer cpuname command to set to the *behind firewall* attribute *on* in the workstation definitions of these fault-tolerant agents.

Remote command job fails to connect to remote computer

After submitting a remote command job, error message AWKRCE012E indicates that an error has occurred establishing a connection to the remote computer.

Cause and solution:

There are several possible causes to this problem:

- The host name specified for the computer where the remote command instance is running does not exist.
- The port number is incorrect, for example, a port number which is different from the port number configured for a specific protocol.
- The protocol type specified is unable to establish a connection because the remote computer is not open to using that particular protocol.

See the topic about the remote command job definition in the *IBM Workload Scheduler: User's Guide and Reference* for information about the connection settings for the remote computer.

The domain manager cannot link to a fault-tolerant agent

The domain manager cannot link to a fault-tolerant agent. The stdlist records the following messages:

```
+ ++++++  
+ AWSEDW020E: Error opening IPC
```

```
+ AWSEDW001I: Getting a new socket: 9
+ ++++++
```

Cause and solution:

The fault-tolerant agent has two netman processes listening on the same port number. This is the case if you installed more than one IBM Workload Scheduler instance on the same workstation and failed to specify different netman port numbers.

Stop one of the two netman services and specify a unique port number using the `nm port` local option (`localopts` file).

Ensure that the workstation definition on the master domain manager is defined with the unique port number or it will not be able to connect.

Agents not linking to master domain manager after first JnextPlan on HP-UX

You have successfully installed the components of your network with the master domain manager on HP-UX. You perform all the necessary steps to create a plan and run your first JnextPlan, which appears to work correctly. The Symphony file is distributed to the agents but they cannot link to the master domain manager, even if you issue a specific link command for them. The conman error log shows that the agents cannot communicate with the master domain manager.

Cause and solution:

One possible cause for this problem is that while on HP-UX host names are normally limited to eight bytes, on some versions of this platform you can define larger host names. The problem occurs if you define the master domain manager's host name as more than eight bytes. When you install the master domain manager on this host a standard operating system routine obtains the host name from the operating system, but either truncates it to eight bytes before storing it in the database, or stores it as "unknown". When you install the agents, you supply the longer master domain manager host name. However, when the agents try to link to the master domain manager they cannot match the host name.

To resolve this problem, perform the following steps:

1. Change the workstation definition of the master domain manager to the correct host name.
2. Run `ResetPlan -scratch`.
3. Run `JnextPlan`.

The agents now link.

Fault-tolerant agents not linking to master domain manager

A fault-tolerant agent does not link to its master domain manager and any other link problem scenarios documented here do not apply.

Cause and solution:

The cause of this problem might not be easy to discover, but is almost certainly involved with a mismatch between the levels of the various files used on the fault-tolerant agent.

To resolve the problem, if all other attempts have failed, perform the following cleanup procedure. However, note that this procedure loses data (unless the fault-tolerant agent is not linking after a fresh installation), so should not be undertaken lightly.

Perform the following steps:

1. Using conman "unlink @;noask" or the Dynamic Workload Console, unlink the agent from the master domain manager
2. Stop IBM Workload Scheduler, in particular netman, as follows:
 - a. conman "stop;wait"
 - b. conman "shut;wait"
 - c. On Windows™ only; shutdown
 - d. Stop the SSM agent, as follows:
 - On Windows™, stop the Windows™ service: IBM Workload Scheduler SSM Agent (for <<TWS_user>>).
 - On UNIX™, run **stopmon**.



Note: If the conman commands do not work, enter the following command:

UNIX™

```
ps -ef |grep <<TWS_user>> & kill -9
```

Windows™

```
<TWA_home>\TWS\unsupported\listproc & killproc
```

3. **Risk of data loss:** Removing the following indicated files can cause significant loss of data. Further, if jobs have run on the fault-tolerant agent for the current plan, without additional interaction, the fault-tolerant agent will rerun those jobs.

Remove or rename the following files:

```
<TWS_home>\TWS\*.msg
    \Symphony
    \Sinfonia
    \Jobtable
    \pobox\*.msg
```



Note: See [Corrupt Symphony file recovery on page 194](#) for additional options.

4. Start netman with StartUp run as the <TWS_user>
5. Issue a "link" command from the master domain manager to the fault-tolerant agent
6. Issue a conman start command on the fault-tolerant agent.

The IBM® technical note describing this procedure also contains some advice about starting with a lossless version of this procedure (by omitting step 3 on page 80) and then looping through the procedure in increasingly more-aggressive ways, with the intention of minimizing data loss. See <http://www.ibm.com/support/docview.wss?uid=swg21296908>

The dynamic agent cannot be found from Dynamic Workload Console

You correctly installed a dynamic agent but cannot see it from the Dynamic Workload Console.

Cause and solution:

A possible cause for this problem might be that either the dynamic workload broker hostname, **-tdwbhostname**, or the dynamic workload broker port, or both, and which are both registered on the agent, are not known in the network of the master domain manager because the broker host is in a different DNS domain.

Edit the JobManager.ini configuration file (for its path, see [Where products and components are installed on page 13](#)). Edit the following parameter:

```
ResourceAdvisorUrl = https://<servername>:
                    31116/JobManagerRESTWeb/JobScheduler/resource
```

Submitted job is not running on a dynamic agent

From Dynamic Workload Console, you can see a dynamic agent, but the submitted job appears as "No resources available" or is dispatched to other agents.

Cause and solution:

A possible cause might be that the local hostname of a registered dynamic workload broker server on the agent is not known in the network of the master domain manager because it is in a different DNS domain.

Edit the JobManager.ini configuration file (for its path, see [Where products and components are installed on page 13](#)). Edit the following parameter:

```
FullyQualifiedHostname = <servername>
```

Job status of a submitted job is continually shown as running on dynamic agent

From Dynamic Workload Console, you can see a dynamic agent, but the job status of a submitted job is continually in the running state.

Cause and solution:

A possible cause might be that the master domain manager local hostname is not known in the network of the agent because it is in a different DNS domain.

Open the JobDispatcherConfig.properties file and edit the parameter JDURL=https://<localhostname>

See the *Administration Guide* for more details about editing this file.

Network performance

If your network shows performance problems, you can use the environment variable TWS_TRUSTED_ADDRESS to resolve these problems. You can set TWS_TRUSTED_ADDRESS to an IP address or leave it blank. If you set it to an IP address, all

the processes use this address for their connections. If you leave it blank, the processes use the address returned by the operating system. If you do not create this variable, the processes use the IP address returned by the Symphony file.

AWSITA245E - Agent is down but jobmanager is running

If while using **conman** or the Dynamic Workload Console you receive one of the following error messages stating that the agent is down, while jobmanager process is up and running, causes and solutions might be as follows:

Cause and solution

An error occurred getting the response of the HTTP request.

- If the detailed error reports problems about `CURL Error ##`, check the following link: <http://curl.haxx.se/libcurl/c/libcurl-errors.html>.
- If the error is 35, perform the following steps:
 1. Enable **gskit** traces, as described in the following section: <http://www-01.ibm.com/support/docview.wss?uid=swg21283690>.
 2. Edit the `JobManager.ini` file and add a the following line `GSK_TRACE_FILE = /tmp/gskit.log` in the `ITA` `Env` section.
 3. Reproduce the problem.
 4. Collect the `/tmp/gskit.log` file.

Chapter 7. Troubleshooting common engine problems

Gives solutions to problems that might occur with the modules and programs that comprise the basic scheduling "engine" on the master domain manager.

This section details commonly occurring problems and their solutions in components and activities not already discussed in previous chapters.

Other common problems are dealt with in other guides, or other chapters of this guide:

- For installation problems see the *IBM® Workload Scheduler: Planning and Installation Guide*.
- For network problems see [Troubleshooting networks on page 73](#)
- For problems with the fault-tolerant switch manager see [Troubleshooting the fault-tolerant switch manager on page 185](#).
- For problems with the Symphony file see [Corrupt Symphony file recovery on page 194](#).

The problems are grouped according to their typology:

Composer problems

The following problems could be encountered with composer:

- [Composer gives a dependency error with interdependent object definitions on page 83](#)
- [The display `cpu=@` command does not work on UNIX on page 84](#)
- [Composer gives the error "user is not authorized to access server" on page 84](#)
- [The deletion of a workstation fails with the "AWSJOM179E error on page 85](#)
- [When using the composer add and replace commands, a Job Scheduler has synchronicity problems on page 85](#)
- [Exiting the composer or conman command line takes an abnormally long amount of time on page 85](#)

Composer gives a dependency error with interdependent object definitions

You are running composer to add or modify a set of object definitions where one object is dependent on another in the same definition. An error is given for the dependency, even though the syntax of the definition is correct.

Cause and solution:

Composer validates objects in the order that they are presented in the command or the definition file. For example, you define two jobs, and the first-defined (*job_tom*) has a `follows` dependency on the second-defined (*job_harry*). The object validation tries to validate the `follows` dependency in *job_tom* but cannot find *job_harry* so gives an error and does not add the job to the database. However, it then reads the definition of *job_harry*, which is perfectly valid, and adds that to the database.

Similarly, this problem could arise if you define that a job needs a given resource or a Job Scheduler needs a given calendar, but you define the resource or calendar *after* defining the job or Job Scheduler that references them.

This problem applies to all composer commands that create or modify object definitions.

To resolve the problem, you can just simply repeat the operation. In the above example the following happens:

- The first job defined (*job_tom*) now finds the second job (*job_harry*) which was added to the database initially.
- You receive a "duplicate job" error for the second.

Alternatively, you can edit the object definition and retry the operation with just the object definition that gave the error initially.

To ensure that the problem does not reoccur, always remember to define objects in the order they are to be used. Define depending jobs and job streams before dependent ones. Define referred objects before referring objects.



Note: There is a special case of this error which impacts the use of the validate operation. Because validate does not add *any* job definitions to the database, correct or otherwise, *all* interdependent job definitions give an error.

In the example above, the problem would not have occurred when using add, new, create, or modify if the job definition of *job_harry* preceded that of *job_tom*. *job_harry* would have been added to the database, so the validation of *job_tom* would have been able to verify the existence of *job_harry*. Because the validate command does not add *job_harry* to the database, the validation of the `follows` dependency in *job_tom* fails.

There is no workaround for this problem when using validate. All you can do is to ensure that there are no interdependencies between objects in the object definition file.

The `display cpu=@` command does not work on UNIX™

In UNIX™, nothing happens when typing `display cpu=@` at the composer prompt.

Cause and solution:

The `@` (atsign) key is set up as the "kill" character.

Type `stty -a` at the UNIX™ prompt to determine the setting of the `@` key. If it is set as the "kill" character, then use the following command to change the setting to be "control/U" or something else:

```
stty kill ^U
```

where `^U` is "control/U", not caret U.

Composer gives the error "user is not authorized to access server"

Troubleshooting for a user authorization error in composer.

You successfully launch composer but when you try to run a command, the following error is given:

```
user is not authorized to access server
```

Cause and solution:

This is a problem that is common to several CLI programs; see [Command line programs \(like composer\) give the error "user is not authorized to access server" on page 134](#).

The deletion of a workstation fails with the "AWSJOM179E error

You want to delete a workstation either using Composer or the Dynamic Workload Console and the following error occurs:

```
AWSJOM179E An error occurred deleting definition of the workstation {0}.
The workload broker server is currently unreachable.
```

Cause and solution:

This problem occurs if you removed a dynamic domain manager without following the procedure that describes how to uninstall a dynamic domain manager in the *IBM Workload Scheduler: Planning and Installation*.

To remove workstations connected to the dynamic domain manager, perform the following steps:

1. Verify that the dynamic domain manager was deleted, not just unavailable, otherwise when the dynamic domain manager restarts, you must wait until the workstations register again on the master domain manager before using them.
2. Delete the workstations using the following command:

```
composer del ws <workstation_name>;force
```

When using the composer add and replace commands, a Job Scheduler has synchronicity problems

The composer add and replace commands do not correctly validate the time zone used in the Job Scheduler definition at daylight savings; as a consequence, the following unexpected warning message is displayed:

```
AWSBIA148W WARNING: UNTIL time occurs before AT time for
<workstation>#<schedule>.
AWSBIA019E For <workstation>#<schedule> Errors 0, warnings 1.
AWSBIA106W The schedule definition has warnings.
AWSBIA015I Schedule <workstation>#<schedule> added.
```

The same might happen for the `deadline` keyword.

Cause and solution:

The problem is related to the C-Runtime Library `date` and `time` functions that fail to calculate the correct time during the first week of daylight savings time.

To ensure the accuracy of scheduling times, for the time argument of the `at`, `until`, or `deadline` scheduling keywords, specify a different value than that of the start time for the IBM Workload Scheduler production period defined in the global options file. These values must differ from one another by plus or minus one hour.

Exiting the composer or conman command line takes an abnormally long amount of time

Troubleshooting for a problem in composer or conman command line in a Windows environment

Cause and solution:

In a Windows environment, if there is a large number of files in the Conman-FFDC and Composer-FFDC folder, exit from composer or conman command line may take an abnormally long amount of time.

If you encounter this problem, you need to clean up the folders Conman-FFDC e Composer-FFDC to the trace files located in stdlist/JM

JnextPlan problems

The following problems could be encountered with JnextPlan:

- [JnextPlan fails to start on page 86](#)
- [JnextPlan fails with the database message "The transaction log for the database is full." on page 86](#)
- [JnextPlan fails with a Java out-of-memory error on page 87](#)
- [JnextPlan fails with the DB2 error like: nullDSRA0010E on page 87](#)
- [JnextPlan fails with message AWSJPL017E on page 87](#)
- [JnextPlan is slow on page 88](#)
- [A remote workstation does not initialize after JnextPlan on page 89](#)
- [A job remains in "exec" status after JnextPlan but is not running on page 89](#)
- [CreatePostReports.cmd, or Makeplan.cmd, or Updatestats.cmd, or rep8.cmd hang on Windows operating systems on page 91](#)
- [Possible performance impact on large production plans after JnextPlan on page 91](#)

JnextPlan fails to start

JnextPlan fails to start.

Cause and solution:

This error might be a symptom that your IBM Workload Scheduler network requires additional tuning because of a problem with the sizing of the `pobox` files. The default size of the `pobox` files is 10MB. You might want to increase the size according to the following criteria:

- The role (master domain manager, domain manager, or fault-tolerant agent) of the workstation in the network. Higher hierarchical roles need larger `pobox` files due to the larger number of events they must handle (since the total number of events that a workstation receives is proportional to the number of its connections). For a domain manager, also the number of sub domains under its control make a difference.
- The average number of jobs in the plan.
- The I/O speed of the workstation (IBM Workload Scheduler is IO- dependent).

JnextPlan fails with the database message "The transaction log for the database is full."

You receive a message from JnextPlan which includes the following database message (the example is from DB2®, but the Oracle message is very similar):

```
The transaction log for the database is full.
```

The JnextPlan message is probably the general database access error message AWSJDB801E.

Cause and solution:

The problem is probably caused by the number of Job Scheduler instances that JnextPlan needs to handle. The default database transaction log files cannot handle more than the transactions generated by a certain number of Job Scheduler instances. In the case of DB2® this number is 180 000; in the case of Oracle it depends on how you configured the database. If JnextPlan is generating this many instances, you need to change the log file creation parameters to ensure more log space is created. You might also need to increase the Java™ heap size on the application server. See "Scalability" in the *IBM® Workload Scheduler: Administration Guide* for a full description of how to perform these activities.

JnextPlan fails with a Java™ out-of-memory error

You receive the following message from JnextPlan:

```
AWSJCS011E An internal error has occurred.
The error is the following: "java.lang.OutOfMemoryError".
```

Cause and solution:

This error is a symptom that the processes running in WebSphere Application Server Liberty Base during JnextPlan phase need more Java virtual memory to run. To increase the Java virtual memory for the WebSphere Application Server Liberty Base, you must increase the default Java heap size values.

See the section about increasing application server heap size in *Administration Guide*

JnextPlan fails with the DB2® error like: nullDSRA0010E

JnextPlan has failed with the following messages:

```
AWSJPL705E An internal error has occurred. The planner is unable to create the preproduction plan.
```

```
AWSBIS348E An internal error has occurred. MakePlan failed while running: planman.
```

```
AWSBIS335E JnextPlan failed while running: tclsh84 The messages.log has an error like this:
```

```
AWSJDB801E An internal error has been found while accessing the database. The internal error message is:
>nullDSRA0010E: SQL State = 57011, Error Code = -912".
```

Cause and solution:

This indicates that the memory that DB2® allocates for its "lock list" is insufficient. To understand why the problem has occurred and resolve it, see the section in the *IBM® Workload Scheduler: Administration Guide* about monitoring the "lock list" value among the DB2® administrative tasks.

JnextPlan fails with message AWSJPL017E

You receive the following message from JnextPlan:

```
AWSJPL017E The production plan cannot be created because a
previous action on the production plan did not complete successfully.
See the message help for more details.
```

Cause and solution:

The problem might be caused by a JnextPlan being launched before the previous JnextPlan has run the SwitchPlan command.

The situation might not resolve itself. To resolve it yourself, perform the following steps:

1. Reset the plan by issuing the command `ResetPlan -scratch`
2. If the reset of the plan shows that the database is locked, run a `planman unlock` command.

On Windows operating systems JnextPlan fails with cscript error

On Windows operating systems, you receive the following message from JnextPlan or ResetPlan:

```
'cscript' is not recognized as an internal or external command,
operable program or batch file
```

Cause and solution:

JnextPlan or ResetPlan fails because the cscript utility that runs files with `.vbs` extension is not installed.

Install the cscript utility on your Windows operating system and rerun JnextPlan or ResetPlan.

JnextPlan is slow

You find that JnextPlan is unacceptably slow.

Cause and solution:

There are three possible causes for this problem:

Tracing too much

One possible cause is the tracing facility. It could be that it is providing too much trace information. There are three possible solutions:

- Reduce the number of processes that the tracing facility is monitoring. See [Quick reference: how to modify log and trace levels on page 21](#) for full details.
- Stop the tracing facility while JnextPlan is running. See [xcli command syntax on page 61](#) to find out how to activate or deactivate a program or segment in memory, or all programs and segments.

Application server tracing too much

Another possible cause is that the application server tracing is set to high. See [Log and trace files for the application server on page 40](#) for more details about the trace and how to reset it.

Database needs reorganizing

Another possible cause is that the database needs reorganizing. See "Reorganizing the database" in *IBM® Workload Scheduler: Administration Guide* for a description of how and why you reorganize the database, logically and physically.

A remote workstation does not initialize after JnextPlan

After running JnextPlan you notice that a remote workstation does not immediately initialize. The following message is seen:

```

+ AWSBCW037E Writer cannot initialize this workstation because mailman
+ is still active.
+ AWSBCW039E Writer encountered an error opening the Mailbox.msg file.
+ The total cpu time used is as follows: 0

```

Cause and solution:

If mailman is still running a process on the remote workstation, JnextPlan cannot download the Symphony file and initialize the next production period's activities. Instead, the domain manager issues a stop command to the workstation. The workstation reacts in the normal way to the stop command, completing those activities it must complete and stopping those activities it can stop.

After the interval determined in the `localopts` parameter `mm retrylink`, the domain manager tries again to initialize the workstation. When it finds that the stop command has been implemented, it starts to initialize the workstation, downloading the Symphony file and starting the workstation's activities.

A workstation does not link after JnextPlan

On UNIX and Windows operating systems, after running JnextPlan a workstation does not link to the domain manager.

Cause and solution:

The workstation does not recognize its own IP address and host name. To resolve this problem, add the workstation IP address and hostname to the `hosts` file located in `C:\Windows\System32\Drivers\etc\hosts` (Windows) or `/etc/hosts` (UNIX) file in the following format:

```
IPaddress hostname hostname.domain
```

For example, `9.168.60.9 nc060009 nc060009.romelab.it.ibm.com`

A job remains in "exec" status after JnextPlan but is not running

After running JnextPlan you notice that a job has remained in "exec" status, but is not being processed.

Cause and solution:

This error scenario is possible if a job completes its processing at a fault-tolerant agent just before JnextPlan is run. The detail of the circumstances in which the error occurs is as follows:

1. A job completes processing
2. The fault-tolerant agent marks the job as "succ" in its current Symphony file
3. The fault-tolerant agent prepares and sends a job status changed event (JS) and a job termination event (JT), informing the master domain manager of the successful end of job
4. At this point JnextPlan is started on the master domain manager
5. JnextPlan starts by unlinking its workstations, including the one that has just sent the JS and JT events. The message is thus not received, and waits in a message queue at an intermediate node in the network.
6. JnextPlan carries the job forward into the next Symphony file, and marks it as "exec", because the last information it had received from the workstation was the Launch Job Event (BL).
7. JnextPlan relinks the workstation
8. The fault-tolerant agent receives the new Symphony file and checks for jobs in the "exec" status.
9. It then correlates these jobs with running processes but does not make a match, so does not update the job status
10. The master domain manager receives the Completed Job Event that was waiting in the network and marks the carried forward job as "succ" and so does not send any further messages in respect of the job
11. Next time JnextPlan is run, the job will be treated as completed and will not figure in any further Symphony files, so the situation will be resolved. However, in the meantime, any dependent jobs will not have been run. If you are running JnextPlan with an extended frequency (for example once per month), this might be a serious problem.

There are two possible solutions:

Leave JnextPlan to resolve the problem

If there are no jobs dependent on this one, leave the situation to be resolved by the next JnextPlan.

Change the job status locally to "succ"

Change the job status as follows:

1. Check the job's stdlist file on the fault-tolerant agent to confirm that it did complete successfully.
2. Issue the following command on the fault-tolerant agent:

```
conman "confirm <job>;succ"
```

To prevent the reoccurrence of this problem, take the following steps:

1. Edit the JnextPlan script
2. Locate the following command:

```
conman "stop @!@;wait ;noask"
```

3. Replace this command with individual stop commands for each workstation (`conman "stop <workstation> ;wait ;noask"`) starting with the farthest distant nodes in the workstation and following with their parents, and so on, ending up with the master domain manager last. Thus, in a workstation at any level, a message

placed in its forwarding queue either by its own job monitoring processes or by a communication from a lower level should have time to be forwarded at least to the level above before the workstation itself is closed down.

4. Save the modified JnextPlan.

CreatePostReports.cmd, or Makeplan.cmd, or Updatestats.cmd, or rep8.cmd hang on Windows operating systems

CreatePostReports.cmd, or Makeplan.cmd, or Updatestats.cmd, or rep8.cmd hang on Windows operating systems

Cause and solution:

On Windows operating systems it might happen that, when running `CreatePostReports.cmd`, `Or Makeplan.cmd`, `Or Updatestats.cmd`, `Or rep8.cmd` jobs, the Tool Command Language interpreter hangs not returning an answer to the caller and the jobs do not complete. These jobs use by default the Tool Command Language interpreter but can be configured to not use it. To avoid that `CreatePostReports.cmd`, `Or Makeplan.cmd`, `Or Updatestats.cmd`, `Or rep8.cmd` jobs use the Tool Command Language interpreter, open each of them with a text editor, find the following line:

```
set USETCL=yes
```

and change it either commenting it:

```
REM set USETCL=yes
```

or setting the variable to `no`:

```
set USETCL=no
```

In this way, when `CreatePostReports.cmd`, `Or Makeplan.cmd` `Or Updatestats.cmd`, `Or rep8.cmd` jobs run, the Tool Command Language interpreter is not invoked.

Possible performance impact on large production plans after JnextPlan

This problem might occur on environments with production plans running around hundred thousand jobs.

Symptom:

After the removal of the production and pre-production plans using the `ResetPlan -scratch` command and the creation of a new plan using the `JnextPlan` command, CPU utilization on the database server is high, around 95% or 100%, for as long as the application server logs the following message:

```
AWSJPU009W The following event "event type"
for "scheduling object" has been skipped.
```

During this time, scheduling information about jobs and job streams displayed by the Dynamic Workload Console might not be updated for several minutes, and the actual start time for jobs and job streams might be delayed.

Cause:

The commands for removing a plan (`ResetPlan -scratch`) and creating a new one (`JnextPlan`) do not remove events related to the deleted plan from the message queues. Right after the new plan is created, the application server processes all the events related to the previous plan just deleted. For large production plans (in the range of hundred thousand jobs), there

might be several thousand events to process, and this processing could take several minutes, some times more than one hour, to complete.

Solution:

Wait for the application server to process all the events related to the deleted plan. The processing is completed when message `AWSJPU009W` is not logged anymore in application server logs. The duration of this processing is proportional to the number of jobs in the production plan.

If your production plan contains several thousand jobs, consider creating a new plan defining a future start time (using the `JnextPlan` command with the **-from** argument) to allow enough time for both the completion of the `JnextPlan` command and the processing of all the events related to the previous plan. This ensures the information displayed by the Dynamic Workload Console and the start time of jobs and job streams are not affected.

Conman problems

The following problems could be encountered when running conman:

- [On Windows, the message AWSDEQ024E is received on page 92](#)
- [Duplicate ad-hoc prompt number on page 94](#)
- [Submitting job streams with a wildcard loses dependencies on page 95](#)
- [Exiting the composer or conman command line takes an abnormally long amount of time on page 85](#)
- [Job log not displayed on page 96](#)
- [Cannot retrieve objects defined in folders on a workstation set to ignore on page 96](#)

On Windows™, the message AWSDEQ024E is received

When attempting to log in to conman on a Windows™ operating system, the following error is received:

```

+ AWSDEQ024E Error owner is not of type user in TOKENUTILS.C;1178
+

```

Cause and solution:

This problem can have a variety of causes related to users and permissions. On the server, perform the following checks:

<<TWS_user>> password

Make sure that the password that you supplied for the <<TWS_user>> user is correct, that the account is not locked out, and that the password has not expired.

Tokensrv service

Ensure that the Tivoli® Token Service (tokensrv) is started by the IBM Workload Scheduler administrative user (not the local system account). This must be verified in the properties of that service in the Services panel; see *IBM® Workload Scheduler: Administration Guide* for details of how to access that panel and view the details of the user that "owns" the service.

If the password to this user has changed on the workstation, check also that the password has been changed in the entry on the Services panel.

File ownerships

Check that the following ownerships are correct:

- All `.exe` and `.dll` files in the `<TWA_home>\TWS\bin` directory are owned by the `<<TWS_user>>`
- All `.cmd` files are owned by "Administrator"

If necessary, alter the ownership of these files as follows:

1. Stop any active IBM Workload Scheduler processes.
2. Change to the `<TWA_home>\TWS` directory.
3. Issue the following commands:

```
setown -u <TWS_user> .\bin\*.exe
setown -u <TWS_user> .\bin\*.dll c:\win32app\maestro
setown -u administrator .\bin\*.cmd
```

4. Issue a StartUp command on the affected server.
5. On the IBM Workload Scheduler master domain manager, launch conman.
6. Once conman is started, issue the following command sequence: **link @!@;noask**
7. Keep issuing the `sc` command to ensure that all the servers relink. A server is considered linked if the State shows "LTI JW"

Advanced user rights

Make sure that the `<<TWS_user>>` has the correct advanced user rights, as documented in the *IBM® Workload Scheduler: Planning and Installation Guide*. These are as follows:

- Act as part of the operating system
- Adjust memory quotas for a process
- Log on as a batch job
- Log on as a service
- Log on locally
- Replace a process level token
- Impersonate a client after authentication right

Resolving the problem by reinstalling

If none of the above suggestions resolve the problem, you might need to reinstall IBM Workload Scheduler. However, it might happen that the uninstallation fails to completely remove all of the Registry keys from the previous installation. In this case, remove the registry keys following the procedure in the *IBM® Workload Scheduler: Planning and Installation Guide*. Then make a fresh installation from the product DVD, subsequently reapplying the most recent fix pack, if there is any.

Duplicate ad-hoc prompt number

You issue a job or Job Scheduler that is dependent on an ad-hoc prompt, but conman cannot submit the job because the prompt number is duplicated.

Cause and solution:

On the master domain manager, prompts are created in the plan using a unique prompt number. This number is maintained in the file of the master domain manager. JnextPlan initially sets the prompt number to "1", and then increments it for each prompt that is to be included in the plan.

If you want to submit a job or Job Scheduler using an ad-hoc prompt on another IBM Workload Scheduler agent during the currency of a plan, the local conman looks in its own `runmsgno` file in its own `<TWA_home>/TWS/mozart/` directory, and uses the number it finds there. The value in the local file does not necessarily reflect the current value used in the `Symphony` file. For example, when the file is first created on an agent the run number is created as the highest run number used in the `Symphony` file at that time, plus `1000`. It is then incremented every time conman needs to assign a number to a prompt. Despite this interval of `1000`, it is still possible for duplicates to occur.

To resolve the problem, edit the file and change the number. An example of the file contents is as follows:

```
0      1236
```

The format is as follows:

- The 10-digit last Symphony run number, right-justified, blank filled. This should not be edited.
- A single blank.
- The 10-digit last prompt number, right-justified, blank filled.

For example:

```
123456789012345678901
0      98
```

When modifying the last prompt number, remember that the least significant digit must always be in character position 21. This means that if the current number is "98" and you want to modify it to display "2098" then you must replace two spaces with the "20", and not just insert the two characters. For example:

```
123456789012345678901
0      2098
```

Save the file and rerun the submit. No error should be given by conman.

During Jnextplan fault-tolerant agents cannot be linked

When you run the conman command stop, the command might take time to stop all the IBM Workload Scheduler processes on local fault-tolerant agents. If, in the meantime, the Symphony file was downloaded, it cannot be received by the IBM Workload Scheduler agent because some processes are still running and the following message is displayed:

```
AWSBCW037E Writer cannot initialize this workstation because
WRITER:+ mailman is still active.
```

Solution:

Insert `sleep 60` after `conman stop` in the Jnextplan script.

Submitting job streams with a wildcard loses dependencies

You issue a submit of interdependent job streams using a wildcard. In certain circumstances you lose the dependencies in an anomalous way.

Cause and solution:

To understand the cause, follow this example, in which the job streams are represented by A, B, C, and their instances are represented by 1, 2:

1. You have the following job streams and jobs in the `Symphony` file:

```
A1
B1 (A1,C1)
C1
```

where `B1` depends on `A1` and `C1`.

2. You submit all the jobs, using:

```
sbs @
```

The planner creates the following Job Scheduler instances:

```
A2
B2 (A2,C1)
C2
```

`B2` now depends on `A2` and `C1`. This is correct, because at the moment of submitting the `B2` Job Scheduler `C2` did not exist, so the highest instance available was `C1`.

3. The planner then asks you to confirm that you want to submit the instances:

```
Do you want to submit A2?
Do you want to submit B2?
Do you want to submit C2?
```

4. Assume that you do not want to submit the job streams `A2` and `C2`, yet, so you reply "No" to the first and last questions. In these circumstances you lose the dependency on `A2`, but not on `C1`. This behavior is correct and logical but could be seen by some as anomalous.

To correct the situation, stop the agent on the workstation where the Job Scheduler is running and cancel the Job Scheduler. Then determine the correct sequence of actions to achieve the objective you want, and submit the appropriate jobs.

Exiting the composer or conman command line takes an abnormally long amount of time

Troubleshooting for a problem in composer or conman command line in a Windows environment

Cause and solution:

In a Windows environment, if there is a large number of files in the Conman-FFDC and Composer-FFDC folder, exit from composer or conman command line may take an abnormally long amount of time.

If you encounter this problem, you need to clean up the folders Conman-FFDC e Composer-FFDC to the trace files located in stdlist/JM

Job log not displayed

The job log does not display when submitting conman sj;stdlist and a wildcard is used in place of the workstation name if a previous operation canceled the job stream related to that job. For example, consider the following scenario:

1. Schedule a job stream named "jobstreamA" to run on workstation "workstationA". The job stream contains a job named "jobA".
2. Cancel the job stream:

```
%cancel workstationA#jobsteramA
```

3. Display the job stream specifying a wild card, @, in place of the actual workstation name:

```
%sj @#jobstreamA.jobA;stdlist
```

The job log is not displayed.

Cause and solution:

In a scenario where a job stream has been canceled and a request is made to display the job log related to a job defined in that job stream, and a wild card is used in place of the workstation name, the job log cannot be displayed. To work around this problem, avoid using a wild card in place of the workstation name when requesting to display the job log and use instead the CPU name. For example, in the scenario described, the following is the correct way to successfully display the job log:

```
%sj workstationA#jobstreamA.jobA;stdlist
```

Cannot retrieve objects defined in folders on a workstation set to ignore

You search for scheduling objects, such as jobs or resources, defined on a workstation which is defined on a folder and set to ignore, but no objects are retrieved in the search.

Cause and solution:

This is a known problem related to the folders feature. To work around this problem, search for the workstation with the filter set to /@/@, or specify the workstation unique identifier, as described in the following example:

```
conman ss 0ASEA234AAE432#/FOLDER_1/JS1
```

where 0ASEA234AAE432 is the ID of the workstation.

You can retrieve the workstation unique identifier using the following composer command:

```
composer li ws=<workstation_name>;showid
```

Fault-tolerant agent problems

The following problems could be encountered with fault-tolerant agents.

- [A job fails in heavy workload conditions on page 97](#)
- [Batchman, and other processes fail on a fault-tolerant agent with the message AWSDEC002E on page 97](#)
- [Fault-tolerant agents unlink from mailman on a domain manager on page 98](#)

A job fails in heavy workload conditions

A job fails on a fault-tolerant agent where a large number of jobs are running concurrently and one of the following messages is logged:

- "TOS error: No space left on device."
- "TOS error: Interrupted system call."

Cause and solution:

This problem could indicate that one or more of the CCLog properties has been inadvertently set back to the default values applied in a prior version (which used to occasionally impact performance).

See [IBM Workload Scheduler logging and tracing using CCLog on page 26](#) and check that the `TWSCCLog.properties` file contains the indicated default values for the properties `twshnd.logfile.className` and `twslloggers.className`.

If the correct default values are being used, contact IBM® Software Support to address this problem.

Batchman, and other processes fail on a fault-tolerant agent with the message AWSDEC002E

The batchman process fails together with all other processes that are running on the fault-tolerant agent, typically mailman and jobman (and JOBMON on Windows™ 2000). The following errors are recorded in the stdlist log of the fault-tolerant agent:

```
+ ++++++
+ AWSBCV012E Mailman cannot read a message in a message file.
+ The following gives more details of the error:
+ AWSDEC002E An internal error has occurred. The following UNIX
+ system error occurred on an events file: "9" at line = 2212
+ ++++++
```

Cause and solution:

The cause is a corruption of the file `Mailbox.msg`, probably because the file is not large enough for the number of messages that needed to be written to it.

Consider if it seems likely that the problem is caused by the file overflowing:

- If you are sure that this is the cause, you can delete the corrupted message file.

All events lost: Following this procedure means that all events in the corrupted message file are lost.

Perform the following steps:

1. Use the `evtsize` command to increase the `Mailbox.msg` file. Ensure that the file system has sufficient space to accommodate the larger file.
 2. Delete the corrupt message file.
 3. Restart IBM Workload Scheduler by issuing the `conman start` command on the fault-tolerant agent.
- If you do not think that this is the answer, or are not sure, contact IBM® Software Support for assistance.

Fault-tolerant agents unlink from mailman on a domain manager

A message is received in the maestro log on the domain manager from mailman for each of the fault-tolerant agents to which it is connected. The messages are as follows:

```
MAILMAN:06:15/ + ++++++
MAILMAN:06:15/ + WARNING: No incoming from <<workstation>>
                  - disconnecting. [2073.25]
MAILMAN:06:15/ + ++++++
```

These messages usually occur in the 30 - 60 minutes immediately following JnextPlan.

Cause and solution:

This problem is normally caused by a false timeout in one of the mailman processes on the domain manager. During the initialization period immediately following JnextPlan, the `*.msg` files on the domain manager might become filled with a backlog of messages coming from fault-tolerant agents. While mailman is processing the messages for one fault-tolerant agent, messages from other fault-tolerant agents are kept waiting until the configured time interval for communications from a fault-tolerant agent is exceeded, at which point mailman unlinks them.

To correct the problem, increase the value of the `mm response` and `mm unlink` variables in the configuration file `~maestro/localopts`. These values must be increased together in small increments (60-300 seconds) until the timeouts no longer occur.

Symphony file on the master domain manager not updated with fault-tolerant agent job status

Jobs are running on a fault-tolerant agent and the local `Symphony` file has been updated, but due to a line failure with the master domain manager and the removal of `<TWS_home>\TWS*.msg` files, the `Symphony` file on the master domain manager has not been updated.

Solution:

Look at the latest `Symphony` file that was processed on the fault-tolerant agent, by using the `conman listsym` command. When used from a fault-tolerant agent command line, this command shows the latest `Symphony` file, saved as `MSymOldBackup`.

Dynamic agent problems

The following problems could be encountered with dynamic agent.

- [The dynamic agent cannot contact the server on page 99](#)
- [Error message AWKDBE009E is received on page 99](#)

The dynamic agent cannot contact the server

The dynamic agent cannot communicate with the server.

About this task

The dynamic agent cannot contact the IBM Workload Scheduler master domain manager or dynamic domain manager.

Cause and solution:

This problem might indicate that the list of URLs for connecting to the master domain manager or dynamic domain manager stored on the dynamic agent is incorrect. Perform the following steps:

1. Stop the dynamic agent
2. Delete the `BackupResourceAdvisorUrls` property from the `JobManager.ini` file
3. Edit the `ResourceAdvisorUrl` property in the `JobManager.ini` file and set the URL of the master domain manager or dynamic domain manager.
4. Start the dynamic agent.

Error message AWKDBE009E is received

Submission of an MSSQL job or of a Database job on an MSSQL database fails.

About this task

When you try to submit an MSSQL job or a Database job running on an MSSQL database, an error message similar to the following is returned, despite the required JDBC driver being installed in the correct directory:

```
AWKDBE009E Unable to create the connection - "
java.lang.UnsupportedOperationException:
Java Runtime Environment (JRE) version 1.6 is not supported by this driver.
Use the sqljdbc4.jar class library, which provides support for JDBC 4.0."
```

Cause and solution:

Verify that only the required `sqljdbc4.jar` driver is present in the JDBC driver directory. If unsupported JDBC drivers are also present in this directory, the dynamic agent might load them and cause the error message.

To solve the problem, perform the following steps:

1. Remove the unsupported JDBC drivers.
2. Stop the dynamic agent with command **ShutDownLwa**.
3. Restart the dynamic agent with command **StartUpLwa**.

For more information, see the section about configuring to schedule job types with advanced options in *IBM Workload Scheduler: Administration Guide*.

Error message AWSITA104E is received

System resources scan fails.

About this task

From the **dynamic workload broker > Tracking > Computers** entry of the Dynamic Workload Console, if you see the *status* of the agent *online*, but the *availability* is *unavailable*, you can see in the `JobManager_trace.log` file the following error:

```
AWSITA104E Unable to perform the system resources scan.  
The error is "Unable to parse hardware scan output".  
AWSITA105E Unable to notify scan results to the server  
because of a resources scanner error.
```

Cause and solution:

The problem is that the product is unable to perform correctly the resources scan because the hostname of the machine is not recognized.

To solve the problem, perform the following steps:

on UNIX operating systems

- verify that the hostname is listed in the `/etc/hosts` file.
- verify that the `ping hostname` command is performed successfully.

on Windows operating systems

- verify that the `ping hostname` command is performed successfully.

Event condition on dynamic agent does not generate any action

After the event rules are deployed on the dynamic agent, the dynamic agent is stopped and restarted automatically. The restarting process takes place only after some minutes, therefore an event condition issued shortly after the deployment of the event rule does not generate any action.

Job manager encounters a core dump

Job manager encounter a core dump when processing a high number of concurrent jobs on a UNIX workstation.

Cause and solution:

If you encounter this problem, you should adjust the supported number of concurrent jobs as needed, based on the requirements of your environment. To perform this operation, use the **ulimit** native UNIX command.

Problems on Windows™

You could encounter the following problems running IBM Workload Scheduler on Windows™.

- [Interactive jobs are not interactive using Terminal Services on page 101](#)
- [The IBM Workload Scheduler services fail to start after a restart of the workstation on page 101](#)
- [The IBM Workload Scheduler for user service \(batchup\) fails to start on page 101](#)
- [An error relating to impersonation level is received on page 103](#)

Interactive jobs are not interactive using Terminal Services

You want to run a job at a Windows™ fault-tolerant agent, launching the job remotely from another workstation. You want to use Windows™ Terminal Services to launch the job on the fault-tolerant agent, either with the Dynamic Workload Console or from the command line. You set the "is interactive" flag to supply some run time data to the job, and indicate the application program that is to be run (for example, `notepad.exe`). However, when the job starts running, although everything seems correct, the application program window does not open on the Terminal Services screen. An investigation at the fault-tolerant agent shows that the application program is running on the fault-tolerant agent, but Terminal Services is not showing you the window.

Cause and solution:

The problem is a limitation of Terminal Services, and there is no known workaround. All "interactive jobs" must be run by a user at the fault-tolerant agent, and cannot be run remotely, using Terminal Services. Jobs that do not require user interaction are not impacted, and can be run from Terminal Services without any problems.

The IBM Workload Scheduler services fail to start after a restart of the workstation

On Windows™, both the Tivoli® Token service and the IBM Workload Scheduler for user service (batchup) fail to start after a restart of the workstation on which they are running.

Cause and solution:

The user under which these services start might have changed password.

If you believe this to be the case, follow the procedure described in *IBM® Workload Scheduler: Administration Guide*.

The IBM Workload Scheduler for user service (batchup) fails to start

The IBM Workload Scheduler for `<<TWS_user>>` service (sometimes also called *batchup*) does not start when the other IBM Workload Scheduler processes (for example, mailman and batchman) start on workstations running Windows™ 2000 and 2003 Server. This problem occurs on a fault-tolerant agent, either after a `conman start` command or after a domain manager switch. The Tivoli® Token service and netman services are unaffected.

This problem does not impact scheduling, but can result in misleading status data.

Cause and solution:

The problem is probably caused either because the <<TWS_user>> has changed password, or because the name of the service does not match that expected by IBM Workload Scheduler. This could be because a change in the configuration of the workstation has impacted the name of the service.

To resolve the problem temporarily, start the service manually using the Windows™ Services panel (under **Administrative Tools**). The service starts and runs correctly. However, the problem could reoccur unless you correct the root cause.

To resolve the problem permanently, follow these steps:

1. If the <<TWS_user>> has changed password, ensure that the service has been changed to reflect the new password, as described in *IBM® Workload Scheduler: Administration Guide*.
2. Look at the Windows™ Event Viewer to see if the information there explains why the service did not start. Resolve any problem that you find.
3. If the failure of the service to start is referred to the following reason, there is a mismatch between the name of the installed service and the name of the service that the mailman process calls when it starts:

```
System error code 1060:
    The specified service does not exist as an installed service
```


The normal reason for this is that the user ID of the <<TWS_user>> has changed. The <<TWS_user>> cannot normally be changed by you, so this implies some change that has been imposed externally. A typical example of this is if you have promoted the workstation from *member server* to *domain controller*. When this happens, the local <<TWS_user>> is converted automatically to a domain user, which means that the domain name is prefixed to the user ID, as follows: <domain_name>\<TWS_user>.

The problem occurs because of the way IBM Workload Scheduler installs the service. If the workstation is *not* a domain controller the installation names the service: `tw_s_maestro_<TWS_user>`. If the workstation *is* a domain controller the installation names the service: `tw_s_maestro_<domain_name>_<TWS_user>`.

When batchman starts up it discovers that the <<TWS_user>> is a domain user. Batchman tries to use the domain user service name to start the batchup service. The action fails because the service on the workstation has the local user service name.

To resolve this problem you must change the name of this service, and to do this you are recommended to uninstall the IBM Workload Scheduler instance and install it again.

An alternative, but deprecated, method is to change the name of the service in the Windows™ registry.

 **Attention:** Making changes to the Windows™ Registry can make the operating system unusable. You are strongly advised to back up the Registry before you start.

If you decide to use this method you must edit the following keys:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\tw_s_maestro_<<TWS_user>>
HKEY_LOCAL_MACHINE\SYSTEM\ControlSet001\Services\tw_s_maestro_<<TWS_user>>
```

```
HKEY_LOCAL_MACHINE\SYSTEM\ControlSet002\Services\tws_maestro_<<TWS_user>>
```

and change them as follows:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\tws_maestro_<domain_name>_<<TWS_user>>
HKEY_LOCAL_MACHINE\SYSTEM\ControlSet001\Services\tws_maestro_<domain_name>_<<TWS_user>>
HKEY_LOCAL_MACHINE\SYSTEM\ControlSet002\Services\tws_maestro_<domain_name>_<<TWS_user>>
```

If you have changed the name of the service in the registry, you must ensure that the logon is correct. Open the **Log On** tab of the service in the Windows™ Services panel and change the account name, if necessary, to <domain_name>\<<TWS_user>>. You must also enter the password and confirm it.

An error relating to impersonation level is received

On Windows™, an error is received when you try to use any of the IBM Workload Scheduler commands (for example, conman, composer, datecalc). The error message is similar to the following example:

```
AWSDEQ008E Error opening thread token ../../src/libs/tokenutils.c:1380
message = Either a required impersonation level was not provided, or the
provided impersonation level is invalid
```

Cause and solution:

This issue occurs when the user account that is used to run the IBM Workload Scheduler command line does not have the user right: "Impersonate a client after authentication". This is a security setting that applies to a subset of Windows versions. For more information, see the related Microsoft documentation. The upgrade does not grant this right to existing users.

To resolve this problem, grant the user right "Impersonate a client after authentication" to all users that need to run IBM Workload Scheduler commands on the workstation. To do this, follow these steps:

1. Select **Start** → **Programs** → **Administrative Tools** → **Local Security Policy**
2. Expand **Local Policies**, and then click **User Rights Assignment**.
3. In the right pane, double-click **Impersonate a client after authentication**.
4. In the Local Security Policy Setting dialog box, click **Add**.
5. In the Select Users or Group dialog box, click the user account that you want to add, click **Add**, and then click **OK**.
6. Click **OK**.

Corrupted characters appear in the command shell when executing cli commands

On Windows™, corrupted characters appear in the command shell when you try to execute any of the IBM Workload Scheduler command line interface commands (for example, conman, composer, wappman). The error message is similar to the following example:

Cause and solution:

This issue is caused by the default font used in the command prompt window (Raster fonts). To change the font, do the following:

- Click the icon in the upper left corner of the command prompt window
- Select Properties, then click the Font tab
- The default font is Raster fonts. Change this to Lucida Console and click OK

Extended agent problems

The following problem could be encountered with extended agents:

The extended agent for MVS returns an error

The extended agent for MVS returns an unsatisfied link error.

Cause and solution:

Required libraries are missing. To find an updated list of required libraries, see [IBM Workload Scheduler Detailed System Requirements](#).

Planner problems

The following problems could be encountered with the planner:

There is a mismatch between Job Scheduler instances in the *Symphony* file and the preproduction plan

You notice that there are Job Scheduler instances in the *Symphony* file that are not in the preproduction plan.

Cause and solution:

Job streams are automatically deleted from the preproduction plan when they are completed. However, it is possible to set the "carryStates" global option (using *optman*) so that job streams with jobs in the *SUCC* status are carried forward. In this case such job streams are carried forward to the new *Symphony* file when the plan is extended, but are deleted from the preproduction plan if the job streams have been successfully completed. This is not an error. These job streams can remain in the current plan (*Symphony* file) and can even be run again.

To resolve the situation for a given plan, use *conman* or the Dynamic Workload Console to delete the Job Scheduler instances from the plan.

To prevent the problem reoccurring, consider why the "carryStates" global option is set so that job streams with jobs in the *SUCC* status are carried forward. If it has been set in error, or is no longer required, change the settings of the option (using *optman*) so that this no longer happens.

Planman deploy error when deploying a plug-in

When using the planman deploy command to deploy a plug-in, the deploy fails with the following error:

```
AWSJCS011E An internal error has occurred. The error is the following:
"ACTEX0019E The following errors
from the Java compiler cannot be parsed:
error: error reading <file_name>; Error opening zip file
<file_name>
```

Cause and solution:

The .jar file identified in the message is corrupt. Check and correct the format of the file before retrying the deploy.

An insufficient space error occurs while deploying rules

When using the planman deploy command with the -scratch option to deploy all non-draft rules, the following error occurs:

```
AWSJCS011E An internal error has occurred. The error is the following:
"ACTEX0023E The Active Correlation Technology compiler cannot
communicate with the external Java compiler.
java.io.IOException: Not enough space".
```

Cause and solution:

This error occurs when there is insufficient swap space (virtual memory) to perform the operation.

Create more swap space or wait until there are fewer active processes before retrying the operation.

UpdateStats fails if it runs more than two hours (message AWSJCO084E given)

When running the UpdateStats command in a large plan, if the job run time exceeds two hours, the job fails with messages that include the following:

```
AWSJCO084E The user "UNAUTHENTICATED" is not authorized to work with the
"planner" process.
```

Cause and solution:

This error occurs because the large number of jobs in the plan has caused the job run time to exceed two hours, which is the default timeout for the user credentials of the WebSphere® Application Server.

To increase the timeout so that the UpdateStats command has more time to run, perform the following steps:

1. Browse to the following path:

On Windows systems

```
<TWA_home>\usr\servers\engineServer\configDropins\defaults
```

On UNIX systems

```
<TWA_home>/usr/servers/engineServer/configDropins/defaults
```

2. Edit the expiration attribute for LTPA as necessary. The attribute is expressed in minutes and the default value is 24 hours:

```
<ltpa keysPassword="{xor}0zo5PiozKw==" keysFileName="{server.config.dir}/resources/security/ltpa.keys"
expiration="1440"/>
```

3. Copy the file to the following path:

On Windows systems

```
<TWA_home>\usr\servers\engineServer\overrides\
```

On UNIX systems

```
<TWA_DATA_DIR>/usr/servers/engineServer/overrides/
```

The planman showinfo command displays inconsistent times

The plan time displayed by the `planman showinfo` command might be incongruent with the time set in the operating system of the workstation. For example, the time zone set for the workstation is GMT+2 but `planman showinfo` displays plan times according to the GMT+1 time zone.

Cause and solution:

This situation arises when the WebSphere Application Server Liberty Base Java™ virtual machine does not recognize the time zone set on the operating system.

As a workaround for this problem, set the time zone defined in the `server.xml` file equal to the time zone defined for the workstation in the IBM Workload Scheduler database. Proceed as follows:

1. Stop WebSphere Application Server Liberty Base.
2. Browse to the following path:

On Windows operating systems

```
<TWA_home>\usr\servers\engineServer\configDropins\overrides
```

On UNIX operating systems

```
<TWA_DATA_DIR>/usr/servers/engineServer/configDropins/overrides
```

3. In the `jvm.options` file, add the following line:

```
-Duser.timezone=time_zone
```

4. Restart WebSphere Application Server Liberty Base.

Job stream duration might be calculated incorrectly as well as other time-related calculations

The duration of a submitted job stream might be calculated incorrectly as well as other time-related calculations.

Cause and solution:

This situation arises when the time set on the workstations where the master and the engine are installed are not aligned.

As a workaround for this problem, align the time of all the workstations belonging to the IBM Workload Scheduler network, even if they are in different time zones.

A bound z/OS shadow job is carried forward indefinitely

A z/OS shadow job, defined for the distributed environment, is successfully bound to a remote z/OS® job, but the z/OS shadow job never completes and is carried forward indefinitely.

Cause and solution:

A Refresh Current Plan operation was performed on the remote IBM Z Workload Scheduler instance. Because this operation scratches the current plan, the remote z/OS® job instance binding was removed.

To prevent the z/OS shadow job from being indefinitely carried forward, manually cancel the z/OS shadow job instance in the distributed engine plan.

For instructions on how to do this, see the *cancel job* topic in the *Managing objects in the plan - conman* chapter of the *IBM Workload Scheduler: User's Guide and Reference*.

Problems with DB2®

The following problems could be encountered with DB2®:

- [Timeout occurs with DB2 on page 107](#)
- [JnextPlan fails with the DB2 message "The transaction log for the database is full." on page 108](#)
- [UpdateStats fails if it runs more than two hours \(message AWSJCO084E given\) on page 105](#)
- [DB2 might lock while making schedule changes on page 108](#)

Timeout occurs with DB2®

You are trying to edit an object, but after a delay an error is issued by DB2® referring to a timeout:

```
AWSJDB803E
An internal deadlock or timeout error has occurred while processing a
database transaction. The internal error message is:
"The current transaction has been rolled back because of a deadlock or timeout.
Reason code "68".
```

Cause and solution:

In this case the object you are trying to access is locked by another user, or by you in another session, but the lock has not been detected by the application. So the application waits to get access until it is interrupted by the DB2® timeout.

By default, both DB2® and WebSphere Application Server Liberty Base have the same length timeout, but as the WebSphere Application Server Liberty Base action starts before the DB2® action, it is normally the WebSphere Application Server Liberty Base timeout that is logged:

```
AWSJC0005E WebSphere Application Server Liberty Base has given the following error:
CORBA NO_RESPONSE 0x4942fb01 Maybe; nested exception is:
org.omg.CORBA.NO_RESPONSE:
```

```
Request 1685 timed out vmcid:
IBM minor code: B01 completed: Maybe.
```

To resolve the problem, check if the object in question is locked. If it is, take the appropriate action to unlock it, working with the user who locked it. If it is not locked retry the operation. If the problem persists contact IBM® Software Support for assistance.

JnextPlan fails with the DB2® message "The transaction log for the database is full."

You receive a message from JnextPlan that includes the following DB2® message:

```
The transaction log for the database is full.
```

The JnextPlan message is probably the general database access error message AWSJDB801E.

Cause and solution:

This scenario is described in [JnextPlan fails with the database message "The transaction log for the database is full." on page 86](#).

DB2® might lock while making schedule changes

Multiple concurrent changes (modify, delete or create) to job streams or domains might cause a logical deadlock between one or more database transactions. This is a remote but possible problem you might encounter.

This deadlock might take place even if the objects being worked on are different (for example, different job streams).

The problem affects database elements (rows or tables), not IBM Workload Scheduler objects, so it is unrelated with the `Locked By` property of IBM Workload Scheduler objects.

The same problem might arise when making concurrent changes for plan generation.

When the deadlock occurs, DB2® rolls back one of the deadlocking threads and the following error is logged in the `messages.log` of WebSphere Application Server Liberty Base:

```
AWSJDB803E An internal deadlock or timeout error has occurred
while processing a database transaction. The internal error
message is: "The current transaction has been rolled back
because of a deadlock or timeout. Reason code "2"."
```

In general, this type of error is timing-dependent, and the transactions involved must overlap in very specific conditions to generate a deadlock. However it might easily occur during plan generation (either forecast, trial, or current), when the plan includes many objects and DB2® must automatically escalate locks from row to table level, as the number of locked objects exceeds the current maximum limit.

You can mitigate the error by increasing the maximum number of locks that DB2® can hold. Refer to the DB2® Information Center to learn more about the DB2® lock escalation mechanism and to find how to increase the maximum number of concurrent locks.

In the above scenarios, if an interactive user session is rolled back, the user gets an error message but is allowed to repeat the task. Instead, if a script session is rolled back (for example, a script that generates a forecast plan or updates a job stream definition), the script ends in failure.

Problems with Oracle

The following problems could be encountered with Oracle:

- [JnextPlan fails with the database message "The transaction log for the database is full." on page 109](#)
- [You cannot do Oracle maintenance on UNIX after installation on page 109](#)
- [Dynamic workload broker fails to start after switching DB2 to Oracle on page 109](#)

JnextPlan fails with the database message "The transaction log for the database is full."

You receive a message from JnextPlan that includes the following database message:

```
The transaction log for the database is full.
```

The JnextPlan message is probably the general database access error message AWSJDB801E.

Cause and solution:

This scenario is described in [JnextPlan fails with the database message "The transaction log for the database is full." on page 86](#).

You cannot do Oracle maintenance on UNIX™ after installation

You have installed IBM Workload Scheduler, creating the installation directory with the default root user permission. When you switch to the Oracle administration user and try and use the Oracle tools, you encounter access problems.

Cause and solution:

The problem could be that the Oracle administration user does not have "read" permission for the entire path of the IBM Workload Scheduler installation directory. For example, if you have created the IBM Workload Scheduler installation directory as `/opt/myProducts/TWS`, the Oracle administration user must have "read" permission for `/opt` and `/myProducts`, as well as `/TWS`.

Give the Oracle administration user read permission for the full path of the IBM Workload Scheduler installation directory.

Dynamic workload broker fails to start after switching DB2® to Oracle

After migrating the IBM Workload Scheduler database vendor from DB2® to Oracle, the Dynamic Workload Broker fails to start. Following the procedure to switch the database vendor from DB2® to Oracle results in missing Dynamic Workload Broker tables in Oracle. Also, the configuration file `DAOCommon.properties` is not updated with the Oracle specific information. The `DAOCommon.properties` file is located in the following path:

On Windows operating systems

```
<TWA_home>\broker\config
```

On UNIX operating systems

```
<TWA_DATA_DIR>/broker/config
```

Cause and solution:

To correct the set of database tables, follow the procedure *Planning and Installation Guide* about upgrading the database schema. Also, manually update the `DAOCommon.properties` so that the `rdbmsName=oracle`, and not `db2`, and ensure that all three of the values for the setting `*Schema=` are set to the Oracle schema name for the IBM Workload Scheduler database.

Problems with MSSQL

The following problem could be encountered with MSSQL:

- [An error is returned when deleting records with cascade option in MSSQL on page 110](#)

An error is returned when deleting records with cascade option in MSSQL

When performing a record deletion with cascade option in MSSQL, you receive a message similar to the following:

```
Maximum stored procedure, function, trigger, or view nesting level exceeded (limit 32).
```

Cause and solution:

Deleting with cascade option a row which contains more than 30 references is not supported in MSSQL.

Problems with Informix

The following problem could be encountered with Informix:

- [Solving deadlocks when using composer with an Informix database on page 110](#)

Solving deadlocks when using composer with an Informix database

You encounter a database deadlock when two concurrent transactions cannot make progress because each one waits for the other to release a lock. If you encounter an Informix deadlock while using composer commands, perform the following steps:

1. Stop WebSphere Application Server Liberty Base, as described in the topic about starting and stopping the application server in *Administration Guide*.
2. Browse to the `TWSConfig.properties` file, located in the following paths:

On Windows operating systems

```
<TWA_home>\usr\servers\engineServer\resources\properties
```

On UNIX operating systems

```
<TWA_DATA_DIR>/usr/servers/engineServer/resources/properties
```

3. Add the following property to the file:

```
com.ibm.tws.dao.rdbms.ids.lockTimeout=0
```

4. Start WebSphere Application Server Liberty Base, as described in the topic about starting and stopping the application server in *Administration Guide*.

```
Maximum stored procedure, function, trigger, or view nesting level exceeded (limit 32).
```

Application server problems

The following problems might occur:

- [Timeout occurs with the application server on page 111](#)
- [The application server does not start after changes to the SSL keystore password on page 111](#)
- [Master Domain Manager may stop when trying to retrieve a big job log on page 112](#)
- [AWK TSA050E error issued during submission on page 112](#)

The application server does not start after changes to the SSL keystore password

You change the password to the SSL keystore on WebSphere Application Server Liberty Base, or you change the security settings using WebSphere Application Server Liberty Base the `ssl_config.xml` or `authentication_config.xml` configuration files. The application server does not start. The following message is found in the application server trace file `trace.log` (the message is shown here on three lines to make it more readable):

```
JSAS0011E: [SSLConfiguration.validateSSLConfig] Java. exception
Exception = java.io.IOException:
Keystore was tampered with, or password was incorrect
```

Cause and solution:

The certificate has not been reloaded or regenerated. Any change to the keystore password on the server or connector requires the SSL certificate to be reloaded or regenerated to work correctly.

Reload or regenerate the certificate and restart the application server.

To regenerate the certificate issue this command:

```
openssl genrsa -des3 -passout pass:<your_password> -out client.key 2048
```

If you do not want to supply the password openly in the command, omit it, and you will be prompted for it.

Timeout occurs with the application server

You are trying to edit an object, but after a delay an error is issued by WebSphere Application Server Liberty Base referring to a timeout:

```
AWSJC0005E WebSphere Application Server Liberty Base has given the following error:
CORBA NO_RESPONSE 0x4942fb01 Maybe; nested exception is:
  org.omg.CORBA.NO_RESPONSE:
    Request 1685 timed out vmcid:
    IBM minor code: B01 completed: Maybe.
```

Cause and solution:

In this case the object you are trying to access is locked from outside IBM Workload Scheduler, maybe by the database administrator or an automatic database function. So the application waits to get access until it is interrupted by the application server timeout.

DB2®

By default, both DB2® and WebSphere Application Server Liberty Base have the same length timeout, but as the WebSphere Application Server Liberty Base action starts before the DB2® action, it is normally the WebSphere Application Server Liberty Base timeout that is logged.

If one or both of the timeouts have been modified from the default values, and the DB2® timeout is now shorter, the following message is issued:

```
AWSJDB803E
An internal deadlock or timeout error has occurred while processing a
database transaction. The internal error message is:
"The current transaction has been rolled back because of a
deadlock or timeout.
Reason code "68".
```

Oracle

There is no corresponding timeout on Oracle, so the Dynamic Workload Console hangs.

To resolve the problem, get the database administrator to check if the object in question is locked outside IBM Workload Scheduler. If it is, take the appropriate action to unlock it, if necessary asking the database administrator to force unlock the object.

If the object is not locked outside IBM Workload Scheduler, retry the operation. If the problem persists contact IBM® Software Support for assistance.

Master Domain Manager may stop when trying to retrieve a big job log

When trying to get a big job log (greater than 100 Mb) from the Dynamic Workload Console, the Master Domain Manager Liberty server could unexpectedly stop working due to a known issue affecting OpenJDK v8, which is related to a temporary issue with the /tmp filesystem.

After a short time, the WebSphere Application Server Liberty Base automatically restarts and, if the temporary issue with the /tmp filesystem is resolved, the operation can be performed again.

AWKTSA050E error issued during submission

The following error is written to the master domain manager `messages.log` file after submitting your workload to run:

```
[12/1/20 13:33:05:868 CET] 000000a2 TWSAgent E AWKTSA050E A problem with the JCL content,
prevent Dynamic Workload Bridge from submitting the job "null"
(alias"DYNAGT7#S_STD_EXTRA_1.S_STD_J_24.SCHEDID-0AAAAAAAAAAA353V2.JNUM-499584099".
The following error was returned:"java.lang.NullPointerException
```

Solution:

This error indicates that a restart of WebSphere Application Server Liberty Base is required. Restart WebSphere® Liberty.

Event management problems

This section describes problems that might occur with processing of events. The topics are as follows:

- [Troubleshooting an event rule that does not trigger the required action on page 113](#)
- [After replying to a prompt, the triggered action is not performed on page 122](#)
- [Actions involving the automatic sending of an email fail on page 122](#)
- [An event is lost on page 123](#)
- [Expected actions not triggered following an event on page 123](#)
- [Event rules not deployed after switching event processor on page 124](#)
- [Event LogMessageWritten is not triggered on page 125](#)
- [Deploy \(D\) flag not set after ResetPlan command used on page 125](#)
- [Missing or empty event monitoring configuration file on page 125](#)
- [Events not processed in correct order on page 126](#)
- [The stopeventprocessor or switcheventprocessor commands do not work on page 126](#)
- [Event rules not deployed with large numbers of rules on page 127](#)
- [Problem prevention with disk usage, process status, and mailbox usage on page 127](#)
- [On AIX operating systems the SSM agent crashes if you have a very large number of files to be managed using event-driven workload automation on page 127](#)
- [File creation and deletion actions not triggered on page 128](#)

Troubleshooting an event rule that does not trigger the required action

You have created an event rule, but the required action is not triggered when the event condition is encountered.

Cause and solution:

The cause and subsequent solution might be any of a number of things. Use the following checklist and procedures to determine what has happened and resolve the problem. The checklist uses a test event which has the following characteristics:

```
<eventRule name="TEST1" ruleType="filter" isDraft="no">
  <description>A Rule that checks the sequence of events</description>
  <eventCondition name="fileCreated1" eventProvider="FileMonitor"
    eventType="FileCreated">
    <scope>
      C:\TEMP\FILE5.TXT ON CPU_MASTER
    </scope>
    <filteringPredicate>
      <attributeFilter name="FileName" operator="eq">
        <value>c:\temp\file5.txt</value>
      </attributeFilter>
      <attributeFilter name="Workstation" operator="eq">
        <value>CPU_MASTER</value>
      </attributeFilter>
      <attributeFilter name="SampleInterval" operator="eq">
```

```

<value>60</value>
</attributeFilter>
</filteringPredicate>
</eventCondition>
<action actionProvider="TWSAction" actionType="sbj" responseType="onDetection">
<scope>
  SBJ CPU_MASTER#JOB1 INTO CPU_MASTER#JOBS
</scope>
<parameter name="JobUseUniqueAlias">
  <value>true</value>
</parameter>
<parameter name="JobDefinitionWorkstationName">
  <value>CPU_MASTER</value>
</parameter>
<parameter name="JobDefinitionName">
  <value>JOB1</value>
</parameter>
</action>
</eventRule>

```

The checklist is as follows:

Step 1: Is event management enabled?

Check if the event management feature is enabled (at installation it is enabled by default):

1. Run the following command:

```
optman ls
```

and look for the following entry:

```
enEventDrivenWorkloadAutomation / ed = YES
```

If the value is "YES", go to [Step 2: Is the workstation enabled for event processing? on page 114](#).

2. **Action:** If the property is set to *NO*, run the command:

```
optman chg ed=YES
```

3. To effect the change, run:

```
JnextPlan -for 0000
```

Check that the event rule is now being processed correctly. If not, go to [Step 2: Is the workstation enabled for event processing? on page 114](#).

Step 2: Is the workstation enabled for event processing?

Check that the workstation is enabled for event processing. By default the master domain manager and backup master domain manager are enabled for event processing, but the default value might have been changed.

Perform the following steps:

1. View the `localopts` file on the master domain manager with a text editor or viewer, and check for the following entry:

```
can be event processor = yes
```

If the value is "yes", go to [Step 3: Is the event processor installed, up and running, and correctly configured? on page 115](#).

2. **Action:** If the value is "no", set it to "yes". Save the `localopts` file and stop and start IBM Workload Scheduler. Check that the event rule is now being processed correctly. If not, go to [Step 3: Is the event processor installed, up and running, and correctly configured? on page 115](#).

Step 3: Is the event processor installed, up and running, and correctly configured?

1. Start `conman`.
2. Issue the `showcpus` command:

```
%sc @!@
```

You will see the following output:

CPUID	RUN	NODE	LIMIT	FENCE	DATE	TIME	STATE	METHOD	DOMAIN
CPU_MASTER	11	*WNT	MASTER	0	0	09/03/19 09:51	I JW MDEA		MASTERDM
FTA1	11	WNT	FTA	0	0		LT		MASTERDM

3. Check the STATE field for the presence of an M, a D, and an E (uppercase). In the example, the STATE field has a value of `I JW MDEA`, and the `MDE` is highlighted. If *all* are present, the event processor is installed, up and running, and correctly configured; go to [Step 9: Is the SSM agent running \(for rules with FileMonitor plug-in-related events only?\) on page 119](#).
4. **Actions:** If one or more of M, D, and E are not present, perform one or more of the following actions until they are all present:

The STATE field has neither an uppercase E nor a lowercase e

If there is neither an uppercase E nor a lowercase e, the event processor is not installed.

The event processor is installed by default on the master domain manager and backup master domain manager. If you are working on either, then the installation did not complete correctly. Collect the log files in the following paths:

Windows operating systems

```
<TWA_home>\TWS\stdlist\logs
```

UNIX operating systems

```
TWA_DATA_DIR/stdlist/logs/
```

and contact customer support for assistance.

The STATE field has a lowercase e

If the STATE field has a lower case e, the event processor is installed but not running. Start the event processor using the conman startevtproc command, or the Dynamic Workload Console. If you use conman, for example, you will see the following output:

```
%startevtproc
AWSJCL528I The event processor has been started successfully.
```

The STATE field has no M

If the STATE field has no M, monman is not running. Start monman using the conman startmon command. You will see the following output:

```
%startmon
AWSBHU470I A startmon command was issued for CPU_MASTER.
```

The STATE field has no D

If the STATE field has no D, the current monitoring package configuration is not deployed. Go to [Step 5: Has the rule been added to the monitoring configuration on the workstation? on page 116.](#)

5. Rerun the showcpus command.
6. When the M, D, and E are all present, check that the event rule is now being processed correctly. If not, go to [Step 9: Is the SSM agent running \(for rules with FileMonitor or plug-in-related events only?\) on page 119.](#)

Step 4: Is the workstation definition present in the plan?

1. Start conman.
2. Issue the showcpus command:

```
%sc @!e
```

If the workstation definition is not included in the plan, add it by generating the production plan with the JnextPlan - for 00 command. If the workstation definition is included, go to [Step 5: Has the rule been added to the monitoring configuration on the workstation? on page 116.](#)

Step 5: Has the rule been added to the monitoring configuration on the workstation?

1. Check if the rule is present in the workstation monitoring configuration by running the conman showcpus command with the ;getmon argument:

```
%sc ;getmon
Monitoring configuration for CPU_MASTER:
*****
*** Package date : 2019/09/03 07:48 GMT ***
*****
```

```
TEST1::FileMonitor#FileCreated:C:\TEMP\FILE5.TXT ON CPU_MASTER;
TEST1::TWSObjectsMonitor#JobSubmit:* # * . TEST*;
```

If the rule is present, go to [Step 7: Has the new monitoring configuration been deployed to the workstation? on page 117](#).

2. **Action:** If the configuration does not contain the expected rule, go to [Step 6: Is the rule active on page 117](#).

Step 6: Is the rule active

If the configuration does not contain the expected rule, check if it is active.

1. Check the rule status, using the composer list command or the Dynamic Workload Console. For example, if you use composer you will see the following output:

```
-list er=@

Event Rule Name          Type      Draft  Status          Updated On  Locked By
-----
TEST1                    filter    N      active          09/03/2019  -
```

If the rule is in *active* status go to [Step 7: Has the new monitoring configuration been deployed to the workstation? on page 117](#).

2. **Action:** If the rule is in *error* status, activate the IBM Workload Scheduler trace, collect the log files located in

Windows operating systems

```
<TWA_home>\TWS\stdlist\logs
```

UNIX operating systems

```
TWA_DATA_DIR/stdlist/logs/
```

and contact customer support for assistance.

Step 7: Has the new monitoring configuration been deployed to the workstation?

If the rule is active, check if the new monitoring configuration has been deployed to the workstation.

1. The deployment of a new monitoring configuration can be checked in either of these ways:
 - Check if the configuration is present in the following paths:

On Windows systems

```
<TWA_home>/TWS/monconf
```

On UNIX systems

```
<TWA_DATA_DIR>/monconf
```

- Check in the `SystemOut` file located in

On Windows systems

<TWA_home>\stdlist\appserver\engineServer\logs

On UNIX systems

<TWA_DATA_DIR>/stdlist/appserver/engineServer/logs

Look for the message:

```
[9/3/19 9:50:00:796 CEST] 00000020 sendEventReadyConfiguration(wsInPlanIds,
zipsToDeploy)
                                AWSDPM001I The workstation "CPU_MASTER" has been notified
about
                                a new available configuration.
```

If the message is present for the workstation in question after the time when the rule was made available for deployment, then the new configuration has been deployed.

If the configuration has been deployed, go to [Step 8: Has the deploy of the new monitoring configuration worked correctly? on page 118.](#)

- Action:** If the configuration has not been deployed, deploy it with the conman deploy command:

```
%deploy
AWSBHU470I A deployconf command was issued for MASTER_CPU.
```

Check that the event rule is now being processed correctly. If not, go to [Step 8: Has the deploy of the new monitoring configuration worked correctly? on page 118.](#)

Step 8: Has the deploy of the new monitoring configuration worked correctly?

If the new monitoring configuration has been deployed, check that the deployment was successful:

- Check in the <date>_TWSMERGE.log in the following paths:

Windows operating systems

<TWA_home>\TWS\stdlist\traces

UNIX operating systems

TWA_DATA_DIR/stdlist/traces/

and look for the most recent occurrence of these 2 messages:

```
09:51:57 03.09.2019|MONMAN:INFO:=== DEPLOY ==> CPU_MASTER has been notified
                                of the availability of the new monitoring configuration.
09:51:57 03.09.2019|MONMAN:INFO:=== DEPLOY ==> The zip file
d:\TWS\twsuser\monconf\deployconf.zip
                                has been successfully downloaded.
```

If you find these messages, referring to the workstation in question, and occurring after the time when the rule was deployed, then the rule has been successfully deployed to the workstation: go to [Step 9: Is the SSM agent running \(for rules with FileMonitor plug-in-related events only?\) on page 119.](#)

- Actions:** If you find messages that indicate an error, follow one of these actions:

Message indicates that the server could not be contacted or that the action has been resubmitted by monman

Message indicates a problem with decoding or extracting the compressed file

You find one of the following messages:

```
=== DEPLOY ===> ERROR decoding the zip file temporarily downloaded in
<TWA_home>\TWS\monconf

=== DEPLOY ===> ERROR unzipping the zip file <file_name>
```

Collect the log files and contact customer support for assistance.

Step 9: Is the SSM agent running (for rules with FileMonitor plug-in-related events only?)

1. If the rule has an event that uses the FileMonitor plug-in, check that the SSM Agent is running. Check in the log that when the conman startmon command was run (either when you ran it manually or when IBM Workload Scheduler started).
2. Search in the log for the following message:

```
11:13:56 03.09.2019|MONMAN:INFO:SSM Agent service successfully started
```

If it is present, or the rule does not use the FileMonitor plug-in, go to [Step 6: Is the rule active on page 117](#).

3. **Action:** If the SSM Agent message is not present, collect the log files. Log files are located in

Windows operating systems

```
<TWA_home>\TWS\stdlist\logs
```

```
<TWA_home>/ssm
```

UNIX operating systems

```
TWA_DATA_DIR/stdlist/logs/
```

```
<TWA_DATA_DIR>/EDWA/ssm
```

and contact customer support for assistance.

Step 10: Have the events been received?

You know the rule has been deployed, but now you need to know if the event or events have been received.

1. Check in the SystemOut of the server to see if the event has been received. The output is different, depending on the type of event:

FileMonitorPlugIn event

a. Following is the output of a FileMonitorPlugIn event:

```
[9/3/19 9:55:05:078 CEST] 00000035 EventProcessor A
com.ibm.tws.event.EventProcessorManager
    processEvent(IEvent)
    AWSEVP001I The following event has been
received:
    event type = "FILECREATED"; event provider =
"FileMonitor";
    event scope = "c:\temp\file5.txt on
CPU_MASTER".
    FILECREATED FileMonitor c:\temp\file5.txt on
CPU_MASTER
```

If the event has been received, go to [Step 11: Has the rule been performed? on page 121](#).

b. If the event has not been received check if it has been created by looking in the traps.log for the message that indicates that the event has been created:

```
.1.3.6.1.4.1.1977.47.1.1.4.25 OCTET STRING FileCreatedEvent event
```

c. **Action:** Whether the event has or has not been created, collect the information in

Windows operating systems

```
<TWA_home>/ssm
```

UNIX operating systems

```
<TWA_DATA_DIR>/EDWA/ssm
```

and contact customer support for assistance.

TWSObjectMonitorPlugIn event

a. Following is the output of a TWSObjectMonitorPlugIn event:

```
[9/3/19 12:28:38:843 CEST] 00000042 EventProcesso A
com.ibm.tws.event.EventProcessorManager
    processEvent(IEvent)
    AWSEVP001I The following event has been
received: event type = "JOBSSUBMIT";
    event provider = "'TWSObjectsMonitor'"; event
scope = "CPU_MASTER # JOBS .
(CPU_MASTER #) TEST". JOBSUBMIT
"TWSObjectsMonitor" CPU_MASTER # JOBS .
(CPU_MASTER #) TEST
```

b. **Action:** If the event has not been received, collect the log data and contact customer support for assistance.

c. If the TWSObjectMonitorPlugIn event has been received, check in the same log that the EIF event has been sent. Following is the output of an EIF event:

```
12:27:18 03.09.19|MONMAN:INFO:Sending EIF Event:
"JobSubmit;
TimeStamp="2019-09-03T12:26:00Z/";
```

```

EventProvider="TWSObjectsMonitor";
HostName="CPU_MASTER";
IPAddress="9.71.147.38";
PlanNumber="11";
Workstation="CPU_MASTER";
JobStreamWorkstation="CPU_MASTER";
JobStreamId="JOBS";
JobStreamName="JOBS";
JobStreamSchedTime="2019-09-03T12:26:00";
JobName="TEST";
Priority="10";
Monitored="false";
EstimatedDuration="0";
ActualDuration="0";
Status="Waiting";
InternalStatus="ADD";
Login="twsuser";END

```

d. If the EIF event has been sent, it might be cached in the following path:

On Windows operating systems

<TWA_home>/EIF

On UNIX operating systems

<TWA_DATA_DIR>/EIF

e. If the event is found there, check the communication with the agent and the server.

If no communication problem is present wait until the event is sent.

f. The event might also be cached in the machine where the event processor is located. Check this in the following paths:

On Windows operating systems

<TWA_home>TWS\stdlist\appserver\engineServer\temp\
TWS\EIFListener/eif.temp1

On UNIX operating systems

<TWA_DATA_DIR>/stdlist/appserver/engineServer/te
mp/TWS/EIFListener/eif.temp1

If the event is found there, check the communication with the agent and the server.

If no communication problem is present wait until the event is sent.

2. **Action:** If the problem persists, collect the log data and contact customer support for assistance.

Step 11: Has the rule been performed?

You now know that the event has been received, but that the action has apparently not been performed.

1. Check in the `SystemOut` of the server to see if the rules have been performed. Look for messages like these:

```

[9/3/19 9:55:05:578 CEST] 00000035 ActionHelper A
com.ibm.tws.event.plugin.action.ActionHelper

```

```

invokeAction(ActionContext,Map,EventRuleHeader)
AWSAHL004I The rule "TEST1" has been triggered. TEST1
[9/3/19 9:55:05:625 CEST] 00000036 ActionHelper A
com.ibm.tws.event.plugin.action.ActionHelper
AsynchAction::run()
AWSAHL002I The action "sbj" for the plug-in "TWSAction" has been
started.
sbj TWSAction
[9/3/19 9:55:06:296 CEST] 00000036 ActionHelper A
com.ibm.tws.event.plugin.action.ActionHelper
AsynchAction::run()
AWSAHL003I The action "sbj" for the plug-in "TWSAction" has
completed.
sbj TWSAction

```

If the rule has been triggered and the action completed, go to [Step 12: Is the problem in the visualization of the event? on page 122.](#)

2. **Action:** If the action has not been completed collect the log data and contact customer support for assistance.

Step 12: Is the problem in the visualization of the event?

Action: If the event has been received, but you cannot see it, there might be a problem with the console you are using to view the event. See [Troubleshooting Dynamic Workload Console problems on page 153.](#)

After replying to a prompt, the triggered action is not performed

An event rule of prompt status changed and recovery prompt is created and deployed, however, after a "yes" response to the prompt, the event rule action is not triggered. The following message is written to WebSphere Application Server Liberty Base traces:

```
AWSEVP008I The following event did not match any existing event condition
```

Cause and solution:

The event rule action is not triggered because no match is found for the event condition. The combination of conditions specified can never be met. To workaround this problem, you need to modify the event rule and ensure that when specifying an event rule of prompt status changed, if you want to specify the recovery prompt, then you must obligatorily set the prompt status to "Asked".

Actions involving the automatic sending of an email fail

An event rule is created, including as the required action the sending of an email. When the event occurs, the action fails with the following message:

```

AWSMSP104E The mail "<mailID>" has not been successfully
delivered to "<recipient>".
Reason: "Sending failed;
nested exception is:
?????class javax.mail.MessagingException: 553 5.5.4 <TWS>...
Domain name required for sender address TWS

```

Cause and solution:

The mail send action failed because the domain name of the SMTP server was not defined in the mail sender name global option: `mailSenderName (ms)`.

Use the `optman` command to specify the correct mail sender name including the domain. For example, if the mail sender name is `twS@alpha.ibm.com`, issue the following command:

```
optman chg ms=twS@alpha.ibm.com
```

An event is lost

You have sent a large number of events to the event processor. When you check the event queue you find that the most recent event or events are missing.

Cause and solution:

The event queue is not big enough. The event queue is circular, with events being added at the end and removed from the beginning. However, if there is no room to write an event at the end of the queue it is written at the beginning, overwriting the event at the beginning of the queue.

You cannot recover the event that has been overwritten, but you can increase the size of the queue to ensure the problem does not recur. Follow the instructions in "Managing the event queue" in *IBM® Workload Scheduler: Administration Guide*.

Expected actions not triggered following an event

In situations where a large number of events are generated and where an action or actions that are expected to be triggered on the master are not triggered, then most probably not all events arrived on the server. There are a number of steps you can perform to collect the necessary information customer support requires to assist you in solving the problem.

Collecting required information:

1. Set the keyword **LogEvents** = `YES` in the following files on the server and client. By default, the value assigned to this keyword is `NO`.

On the client

On Windows operating systems

```
<TWA_home>/EIF/monmaneif.templ
```

```
<TWA_home>/ssm/eif/tecad_snmp_eeif.templ
```

On UNIX operating systems

```
<TWA_DATA_DIR>/EIF/monmaneif.templ
```

```
<TWA_DATA_DIR>/EDWA/ssm/eif/tecad_snmp_eeif.templ
```

On the server

On Windows operating systems

```
<TWA_home>TWS\stdlist\appserver\engineServer\temp\TWS\EIFListen  
er/eif.templ
```

On UNIX operating systems

```
<TWA_DATA_DIR>/stdlist/appserver/engineServer/temp/TWS/EIFListen
er/eif.templ
```

2. On the client, stop and start monman by submitting the respective command.

Stop

```
conman stopmon
```

Start

```
conman startmon
```

3. On the server, stop and start the event processor submitting the respective command.

Stop

```
conman stopevtproc
```

Start

```
conman startevtproc
```

4. All generated events are logged in a file. By default, this file is stored in the path `/EIF`. Collect these logs and send them to customer support. Because this log file can become quite large very quickly, there are ways to filter the events that get logged in the file.

On the server, set the **LogEventsFilter** keyword using a regular expression to filter events written on the server. By default, the value of this keyword is `".*"` and therefore all events are logged. For example, if you are aware that problems originate from a specific client, then you can specify the CPU name of the client in the value of the **LogEventsFilter** keyword to limit the events logged to the file to this specific client as follows:

```
LogEventsFilter=.*<cpu_name>#MONMAN.*
```

Event rules not deployed after switching event processor

You have switched the event processor, but new or amended rules have not been deployed (the event states of the workstations that were affected by the new or amended rules do not show "D" indicating that the rules are not up-to-date, and the `getmon` command shows the old rules).

Cause and solution:

The probable cause is that you made some changes to the rules before running the `switcheventprocessor` command, and these rules were not deployed (for whatever reason) before the switch.

To remediate the situation, run the command `conman deployconf <workstation_name>`, for each affected workstation, and the rule changes will be deployed.

To avoid that this problem reoccurs, run `planman` with the `deploy` action before running `switcheventprocessor`.

Event *LogMessageWritten* is not triggered

You are monitoring a log file for a specific log message, using the *LogMessageWritten* event. The message is written to the file but the event is not triggered.

Cause and solution:

The SSM agent monitors the log file. It sends an event when a new message is written to the log file that matches the string in the event rule. However, there is a limitation. It cannot detect the very latest message to be written to the file, but only messages prior to the latest. Thus, when message line "n" is written containing the string that the event rule is configured to search for, the agent does not detect that a message has been written, because the message is the last one in the file. When any other message line is written, if or not it contains the monitored string, the agent is now able to read the message line containing the string it is monitoring, and sends an event for it.

There is no workaround to resolve this problem. However, it should be noted that in a typical log file, messages are being written by one or other processes frequently, perhaps every few seconds, and the writing of a subsequent message line will trigger the event in question. If you have log files where few messages are written, you might want to attempt to write a dummy blank message after every "real" message, in order to ensure that the "real" message is never the last in the file for any length of time.

Deploy (D) flag not set after *ResetPlan* command used

The deploy (D) flag is not set on workstations after the *ResetPlan* command is used.

Cause and solution:

This is not a problem that affects the processing of events but just the visualization of the flag which indicates that the event configuration file has been received at the workstation.

No action is required, because the situation will be normalized the next time that the event processor sends an event configuration file to the workstation.

However, if you want to take a positive action to resolve the problem, perform the following steps:

1. Create a dummy event rule that applies only to the affected workstations.
2. Perform a planman deploy to send the configuration file.
3. Monitor the receipt of the file on the agent.
4. When it is received, delete the dummy rule at the event processor.

Missing or empty event monitoring configuration file

You have received a MONMAN trace message on a workstation, similar to this:

```
MONMAN:INFO:=== DEPLOY ==> ERROR reading the .zip file
                /home/f_edwa3/monconf/deployconf.zip.
                It is empty or does not exist".
```

Cause and solution:

The IBM Workload Scheduler agent on a workstation monitors for events using a configuration file. This file is created on the event processor, compressed, and sent to the agent. If a `switcheventprocessor` action is performed between the creation of the file on the old event processor and the receipt on the new event processor of the request for download from the agent, the file is not found on the new event processor, and this message is issued.

To resolve the problem, perform the following steps:

1. Create a dummy event rule that applies only to the affected workstation.
2. Perform a planman deploy to send the configuration file.
3. Monitor the receipt of the file on the agent.
4. When it is received, delete the dummy rule at the event processor.

Events not processed in correct order

You have specified an event rule with two or more events that must arrive in the correct order, using the *sequence* event grouping attribute. However, although the events occurred in the required sequence the rule is not triggered, because the events arrived at the event processor in an order different from their creation order.

Cause and solution:

Events are processed in the order they arrive, not the order they are created. If they arrive in order different from the creation order, you will not get the expected result.

For example, consider a rule which is triggered if event A defined on workstation AA occurs before event B which is defined on workstation BB. If workstation AA loses its network connection before event A occurs, and does not regain it until after event B has arrived at the event processor, the event rule will not be satisfied, even though the events might have occurred in the correct order.

The solution to this problem is that if you need to define a rule involving more than one event, use the *set* event grouping attribute, unless you can be certain that the events will arrive at the event processor in the order they occur.

The `stopeventprocessor` or `switcheventprocessor` commands do not work

You have run `stopeventprocessor` or `switcheventprocessor` but the command has failed. The log indicates a communication problem.

Cause and solution:

If you issue the `stopeventprocessor` command from a workstation other than that where the event processor is configured, the command uses the command-line client, so the user credentials for the command-line client must be set correctly.

Similarly, if you use `switcheventprocessor`, it also uses the command-line client, so the user credentials for the command-line client must be set correctly also in this case.

Event rules not deployed with large numbers of rules

You have run **planman deploy** (or the equivalent action from the Dynamic Workload Console), with a very large number of event rules, but the command has failed. The log indicates a memory error.

Cause and solution:

A large number of event rules requires a Java™ heap size for the application server larger than the default. In this context, a large number would be 10 000 or more. Doubling the default size should be sufficient.

Full details of how to do this are described in the *IBM® Workload Scheduler: Administration Guide* in the section on *Increase application server heap size* in the *Performance* chapter.

Problem prevention with disk usage, process status, and mailbox usage

You can use event-driven workload automation (EDWA) to monitor the health of the IBM Workload Scheduler environment and to start a predefined set of actions when one or more specific events take place. You can prevent problems in the IBM Workload Scheduler environment by monitoring the filling percentage of the mailboxes, the status of IBM Workload Scheduler processes, and the disk usage of the IBM Workload Scheduler file system.

Full details of how to do this are described in the *IBM® Workload Scheduler: Administration Guide*, as follows:

- section on *Monitoring the disk space used by IBM Workload Scheduler* in the *Data maintenance* chapter
- sections on *Monitoring the IBM Workload Scheduler message queues* and *Monitoring the IBM Workload Scheduler processes* in chapter *Network administration*

See also [Configuring trace properties when the agent is running on page 37](#).

On AIX operating systems the SSM agent crashes if you have a very large number of files to be managed using event-driven workload automation

On AIX operating systems the SSM agent crashes if you use the have a very large number of files to be managed using the FileCreated and FileDeleted event types of the FileMonitor events provided by even-driven workload automation (EDWA) feature.

Cause and solution:

This problem is due to a missing option setting in the EDWA configuration file. To solve this problem, add the following line in the files `<TWS_INST_DIR>/TWS/ssm/bin/preload_ssmagent_0.sh` and `<TWS_INST_DIR>/TWS/EDWA/ssm/bin/preload_ssmagent_0.sh`:

```
export LDR_CNTRL=MAXDATA=0x80000000
```

where `<TWS_INST_DIR>` is the IBM Workload Scheduler installation directory.

File creation and deletion actions not triggered

You are monitoring the creation or the deletion of a file on Windows using the FileCreated and FileDeleted events. The file is created or deleted, but the event action is not triggered.

Cause and solution:

The SSM agent monitors the file creation and deletion. An event is sent when a file that matches the string in the event rule is created or is deleted. However, on Windows platforms, if the file path contains forward slashes ("/"), the event action is not triggered. Replace forward slashes ("/") with backward slashes ("\") and redeploy the rule.

Managing concurrent accesses to the Symphony file

This section contains two sample scenarios describing how IBM Workload Scheduler manages possible concurrent accesses to the `Symphony` file when running **stageman**.

Scenario 1: Access to Symphony file locked by other IBM Workload Scheduler processes

If IBM Workload Scheduler processes are still active and accessing the `Symphony` file when **stageman** is run, the following message is displayed:

```
Unable to get exclusive access to Symphony.
Shutdown batchman and mailman.
```

To continue, stop IBM Workload Scheduler and rerun **stageman**. If **stageman** aborts for any reason, you must rerun both **planman** and **stageman**.

Scenario 2: Access to Symphony file locked by stageman

If you try to access the plan using the command-line interface while the `Symphony` is being switched, you get the following message:

```
Current Symphony file is old. Switching to new Symphony.
Schedule mm/dd/yyyy (nnnn) on cpu, Symphony switched.
```

StartApp Server problems

The StartApp Server command checks if WebSphere Application Server Liberty Base is running, if the WebSphere Application Server Liberty Base is not running startApp Server starts it.

In case of failure:

- Rerun the job.

MakePlan problems

MakePlan performs the following actions:

- Replans or extends the preproduction plan.
- Produces the Symnew file.
- Generates preproduction reports in the joblog

The following problems could be encountered when running MakePlan:

- [MakePlan fails to start on page 129](#)
- [Unable to establish communication with the server on host - AWSBEH023E on page 129](#)
- [The user "twuser" is not authorized to access the server on host - AWSBEH021E on page 130](#)
- [The database is already locked - AWSJPL018E on page 130](#)
- [An internal error has occurred - AWSJPL006E on page 130](#)
- [The production plan cannot be created - AWSJPL017E on page 130](#)
- [An internal error has occurred - AWSJPL704E on page 130](#)

MakePlan fails to start

If MakePlan fails to start:

- Global lock might be left to 'set'. Use planman unlock to 'reset' it
- Rerun the job to recover:
 - Preproduction plan is automatically reverified and updated.
 - Symnew is created again.

How to stop it:

- Stopping the job might not stop the processing that is still running on WebSphere Application Server Liberty Base or in the database.
- Force the Database statement to close if a database statement runs for too long and causes MakePlan to abend.
- Restart the WebSphere Application Server Liberty Base if processing is still running on WebSphere Application Server Liberty Base and MakePlan does not finish.



Note: Check if the database statistics are enabled. If not, it is recommended to schedule the runstatistics script stored in the dbtools IBM Workload Scheduler directory.

Unable to establish communication with the server on host - AWSBEH023E

If you receive the following error message from MakePlan stdlist:

```
AWSBEH023E Unable to establish communication with the server on host
"127.0.0.1" using port "31116"
```

Cause and solution: This error means that the application server is down and MakePlan cannot continue. If this happens, start the WebSphere Application Server Liberty Base and check the WebSphere Application Server Liberty Base logs to identify the reason why the WebSphere Application Server Liberty Base has stopped.

The user "twuser" is not authorized to access the server on host - AWSBEH021E

```
AWSBEH021E The user "twuser" is not authorized to access the server on host "127.0.0.1" using port "31116"
```

Cause and solution: This is an authorization error. Check your IBM Workload Scheduler user name and password in the useropts file.

The database is already locked - AWSJPL018E

```
AWSJPL018E The database is already locked
```

Cause and solution: The previous operation of MakePlan is stopped and the global lock is not reset. To recover the situation run planman unlock.

An internal error has occurred - AWSJPL006E

```
AWSJPL006E An internal error has occurred. A database object "xxxx" cannot be loaded from the database.
```

Cause and solution: Usually "xxxx" is an object like a workstation, job, or job stream. This error means that a connection with the database is broken. In this case, check the error in the messages.log and the ffdc directory as additional information related to the error is logged there.

The production plan cannot be created - AWSJPL017E

```
AWSJPL017E The production plan cannot be created because a previous action on the production plan did not complete successfully. Refer to the message help for more details.
```

Cause and solution: This error might mean that a previous operation on the preproduction plan is performed but ended with an error. Generally it is present when "ResetPlan - scratch" is performed but does not end successfully.

An internal error has occurred - AWSJPL704E

```
AWSJPL704E An internal error has occurred. The planner is unable to extend the preproduction plan.
```

Cause and solution: This error might mean that MakePlan cannot extend the preproduction plan. Different root causes are associated with this issue, typically always related to the database, for example, no space for the tablespace or full transaction logs. Check for more information in the messages.log directory or in the directories where ffdc (*first failure data capture*) utility saves the extracted information.

SwitchPlan problems

SwitchPlan performs the following actions:

- Stops all the workstations
- Runs Stageman to:
 - Merge the old Symphony file with SymNew
 - Archive the old Symphony file in the schedlog directory
- Runs the planman confirm command to update the database plan status information . For example, the plan end date and the current run number.
- Restarts the master to distribute the Symphony file and restart scheduling.

The following problems could be encountered when running SwitchPlan:

- [When SwitchPlan fails to start on page 131](#)
- [The previous Symphony file and Symnew file have the same run number - AWSBHV082E on page 131](#)

When SwitchPlan fails to start

If SwitchPlan fails to start:

1. planman confirm is not running. Perform the following actions:
 - a. Check the logs
 - b. Run "planman showinfo
 - c. Rerun SwitchPlan
2. planman confirm failed. Perform the following action:
 - Manually run planman confirm and conman confirm.
3. planman confirm was already run and the plan end date has been updated. Perform the following action:
 - Run conman start

If conman stop hangs, kill the conman command. This might impact the plan distribution because it stops the agents left running before distributing the new Symphony.

The previous Symphony file and Symnew file have the same run number - AWSBHV082E

If SwitchPlan stdlist shows the following messages:

```
STAGEMAN: AWSBHV082E: The previous Symphony file and Symnew file have
the same run number. They cannot be merged to form the new symphony file.
```

Cause and solution: There are several possible causes because the Symphony and Symnew run numbers have the same values, common causes for this are:

1. MakePlan did not extend the run number in the Symnew file.
2. SwitchPlan ran before MakePlan.
3. The Stageman process ran twice on the same Symnew file without resetting the plan or deleting the Symphony file.

```
AWSJCL054E: The command "CONFIRM" has failed.
```

```
AWSJPL016E: An internal error has occurred. A global option "confirm  
run member" cannot be set.
```

Cause and solution: These error messages are present when the last step of the SwitchPlan, that is planman confirm fails. Analyze the messages.log for more information and rerun planman confirm.

Create Post Reports

Create Post Reports has the following function:

- General post production reports in the job output.

In case of failure:

- Rerun the job if reports are needed.

Update Stats problems

Update Stats has the following functions:

- Runs logman to update job statistics and history.
- Extends the Preproduction plan if its length is shorter than minLen.

In case of failure:

- Rerun the job or manually run "logman <file> " on the latest schedlog file.
- If it does not run, the statistics and history will be partial. Preproduction plan is updated at the beginning of MakePlan .

How to stop it:

- Kill the job or logman process, the statistics and history will be partial until the job or logman is rerun.

Miscellaneous problems

The following problems might occur:

- An error message indicates that a database table, or an object in a table, is locked on page 133
- Command line programs (like composer) give the error "user is not authorized to access server" on page 134
- The rmstdlist command gives different results on different platforms on page 134
- Question marks are found in the stdlist on page 135
- Deleting stdlist or one of its files when processes are still running on page 135
- A job with a "rerun" recovery job remains in the "running" state on page 136
- Job statistics are not updated daily on page 136
- A Job Scheduler dependency is not added on page 136
- Incorrect time-related status displayed when time zone not enabled on page 136
- Completed job or job stream not found on page 137
- Variables not resolved after upgrade on page 137
- Default variable table not accessible after upgrade on page 137
- Local parameters not being resolved correctly on page 137
- Deleting leftover files after uninstallation is too slow on page 138
- Corrupted special characters in the job log from scripts running on Windows on page 139
- Failover Cluster Command Interface deprecated on page 139
- Job stream creation failure due to missing rights on calendars on page 139
- Suppressing connector error messages on page 139

An error message indicates that a database table, or an object in a table, is locked

An error message indicates that a function cannot be performed because a table, or an object in a table, is locked. However, the table or object does not appear to be locked by another IBM Workload Scheduler process.

Cause and solution:

The probable cause is that a user has locked the table by using the database command-line or GUI:

DB2®

Just opening the DB2® GUI is sufficient to lock the database tables, denying access to all IBM Workload Scheduler processes.

Oracle

If the Oracle command-line is opened without the auto-commit option, or the GUI is opened, Oracle locks all tables, denying access to all IBM Workload Scheduler processes.

To unlock the table close the command-line or GUI, as appropriate.



Note: IBM Workload Scheduler provides all of the database views and reports you need to manage the product. You are strongly recommended to not use the facilities of the database to perform any operations, including viewing, on the database tables.

Command line programs (like composer) give the error "user is not authorized to access server"

You launch CLI programs (like composer) but when you try and run a command, the following error is given:

```
user is not authorized to access server
```

Cause and solution:

This problem occurs when the user running the command has a null password. Composer, and many of the other IBM Workload Scheduler CLI programs cannot run if the password is null.

Change the password of the user and retry the operation.

The rmstdlist command gives different results on different platforms

The rmstdlist command on a given UNIX™ platform gives results that differ from when it is used on other platforms with the same parameters and scenario.

Cause and solution:

This is because on UNIX™ platforms the command uses the *-mtime* option of the find command, which is interpreted differently on different UNIX™ platforms.

To help you determine how the *-mtime* option of the find command is interpreted on your workstation, consider that the following command:

```
<TWA_home>/TWS/bin/stdlist/rmstdlist -p 6
```

gives the same results as these commands:

```
find <TWA_home>/TWS/stdlist/ -type d ! -name logs ! -name traces -mtime +6 -print
find <TWA_home>/TWS/stdlist/logs/ -type f -mtime +6 -print
find <TWA_home>/TWS/stdlist/traces/ -type f -mtime +6 -print
```

Look at your operating system documentation and determine how the option works.

The rmstdlist command fails on AIX® with an exit code of 126

The rmstdlist command on AIX® fails with an exit code of 126 and no other error message.

Cause and solution:

This could be because there are too many log files in the `stdlist` directory.

On AIX®, you should regularly remove standard list files every 10-20 days. See the usage instructions in the *IBM® Workload Scheduler: User's Guide and Reference* for full details.

Question marks are found in the stdlist

You discover messages in the log or trace files that contain question marks, as in the following example (the message has been split over several lines to make it more readable and the question marks are highlighted to make them more obvious):

```
10:20:02 03.02.2008|BATCHMAN:+ AWSBHT057W
Batchman has found a non-valid run number in the Symphony
file for the following record type: "Jt" and object:
"F235011S3_01#???[()],(0AAAAAAAAAAAAZD)].A_7_13 (#J18214)".
```

Cause and solution:

This problem occurs when the process that needs to write the log message cannot obtain the Job Scheduler name. For example, when a Job Scheduler is dependent on a Job Scheduler that is not in the current plan (Symphony file). The process writes "???" in place of the missing Job Scheduler name.

The message contains the Job Scheduler ID (in the above example it is the string in the second set of parentheses: (0AAAAAAAAAAAAZD)). Use the Job Scheduler ID to identify the instance of the Job Scheduler, and take any action suggested by the message that contained the question marks.

Deleting stdlist or one of its files when processes are still running

Erroneously deleted the `stdlist` directory or one of its files while processes are running.

You erroneously deleted the `stdlist` directory or one of its files while processes are running and you receive the following error when performing an operation:

```
Permission denied
Bad file descriptor
```

Cause and solution:

This problem occurs because the directories or files with root ownership are not re-created during the initialization phase.

According to your operating system, perform one of the following actions:

UNIX

- Create the directory or file that you deleted with **twuser** and **group** ownership.
- Modify the ownership of the directory or file created with **root** ownership to **twuser** and **group** ownership.

Windows

- Ensure that you are running the command line with the **Run as Administrator** privilege level.

A job with a "rerun" recovery job remains in the "running" state

You have run a job specifying a recovery job using the "rerun" recovery method. The original job fails, but when the recovery job starts the original job shows that the recovery action has been completed successfully, but remains in the "running" state.

Cause and solution:

This problem would occur if the recovery job was specified to run on a different workstation and domain from the original job. The original job is then unable to detect the state of the recovery job, so it cannot determine if the recovery job has finished or what state it finished in.

To resolve the problem for the specific job that is still in "running" state, you must manually stop the job.

To avoid the recurrence of the problem specify the "rerun" recovery action on the same workstation in the same domain.

Job statistics are not updated daily

Job statistics are not updated daily, as they were with versions prior to version 8.3.

Cause and solution:

Job statistics are updated by JnextPlan. If you are running JnextPlan less frequently than daily, the statistics are only updated when JnextPlan is run.

A Job Scheduler dependency is not added

A dependency is added to a Job Scheduler instance and the Job Scheduler is saved. When the list of dependencies is reopened, the new dependency is not present.

Cause and solution:

This occurs when a Job Scheduler instance already has the maximum number (40) of dependencies defined. Normally, an error message would alert you to the limit, but the message might not be displayed if there is a delay propagating the Symphony updates across the network or if your update coincided with updates by other users.

Incorrect time-related status displayed when time zone not enabled

You are using IBM Workload Scheduler in an environment where nodes are in different time zones, but the time zone feature is not enabled. The time-related status of a job (for example, "Late") is not reported correctly on workstations other than that where the job is being run.

Cause and solution:

Enable the time zone feature to resolve this problem. See *IBM® Workload Scheduler: User's Guide and Reference* to learn more about the time zone feature. See *IBM® Workload Scheduler: Administration Guide* for instructions on how to enable it in the global options.

Completed job or job stream not found

A job or job stream that uses an alias has completed but when you define a query or report to include it, the job or job stream is not included.

Cause and solution:

Jobs and job streams in *final* status are stored in the archive with their *original* names, not their aliases, so any search or reporting of completed jobs must ignore the aliases.

Variables not resolved after upgrade

After performing an upgrade, global variables are not resolved.

Cause and solution:

During the upgrade, all the security file statements relating to your global variables were copied by the installation wizard into a default variable table in the new security file. Global variables are disabled and can only be used through the variable tables. If you subsequently rebuilt the security file using the output from your *previous* dumpsec as input to the new makesec, you will have overwritten the security statements relating to your default variable table, so no user has access to the default variable table.

If you have a backup of your security file from prior to when you ran makesec, run dumpsec from that, and merge your old dumpsec output file with your new one, as described in the upgrade procedure in the *IBM Workload Scheduler: Planning and Installation*.

If you do not have a backup, create the default variable table security statement, following the instructions about configuring the security file in the *IBM Workload Scheduler: Administration Guide*.

Default variable table not accessible after upgrade

After upgrading, your default variable table is not accessible by any user.

Cause and solution:

This problem has exactly the same Cause and solution: as the preceding - see [Variables not resolved after upgrade on page 137](#).

Local parameters not being resolved correctly

You have scheduled a job or job stream that uses local parameters, but the parameters are not resolved correctly.

Cause and solution:

One reason for this could be that one or both of the files where the parameters are stored have been deleted or renamed.

Check that the following files can be found in the *TWA_home* /TWS directory:

```
parameters
parameters.KEY
```

These files are required by IBM Workload Scheduler to resolve local parameters, so they must not be deleted or renamed. Fix the problem as follows:

1. If the files have been renamed, rename them to the original names.
2. If the files have been deleted, re-create them, using the parms utility.
3. To make the changes effective, restart the application server, using the stopappserver and startappserver commands.

Inconsistent time and date in conman and planman output

If you notice inconsistent times and dates in jobs and job streams on an AIX® master domain manager, ensure that the system time zone is set correctly. For example, you might notice this problem in the job schedtime or start time, or in other properties related to date and time.

Cause and solution:

The problem might be due to an incorrect setting of the time zone. To set the correct time zone, perform the following steps on the AIX® master domain manager:

1. Start smit (System Management Interface Tool).
2. Select System Environments > Change / Show Date, Time, and Time Zone > Change Time Zone Using User Entered Values.
3. Set the relevant time zone. For example, to set the Central European Time (CET) time zone, enter the following values:

```
* Standard Time ID(only alphabets)           [CET]
* Standard Time Offset from CUT([+|-]HH:MM:SS) [-1]
  Day Light Savings Time ID(only alphabets)  [CEST]
```

4. Restart the system to make the change effective.

For information about how to set the time zone, see *IBM Workload Scheduler: Administration Guide*. For a description of how the time zone works, see *IBM Workload Scheduler: User's Guide and Reference*.

Deleting leftover files after uninstallation is too slow

Deleting leftover *Onnnn.hhmm* files *TWA_installation_directory\TWS\stdlist\yyyy.mm.dd* after uninstalling IBM Workload Scheduler is too slow.

Cause and solution:

This problem is caused by a known Microsoft™ issue on Windows™ operating systems. It occurs when you try to delete the *Onnnn.hhmm* files in *TWA_installation_directory\TWS\stdlist\yyyy.mm.dd* on the Windows™ system after having uninstalled the master domain manager.

To prevent the problem, remove the *Onnnn.hhmm* files permanently using the Shift-Canc keys instead of using the Delete key or sending the files to the Recycle Bin.

Corrupted special characters in the job log from scripts running on Windows™

When you run scripts on Windows™ systems, any special characters resulting from the commands in the script might not be displayed correctly in the job log. This is a display problem that does not affect the correct run of the job. No workaround is currently available for this problem.

Failover Cluster Command Interface deprecated

The `cluster.exe` command-line tool for Failover Clustering was deprecated for Windows Server 2012 platforms. The commands `Startup_clu.cmd`, `ShutDown_clu.cmd`, and `clusterupg` do not work.

Cause and solution:

This occurs because of the deprecation of the `cluster.exe` command-line tool for Failover Clustering on Windows Server 2012 platforms. To avoid this problem, you must reinstall the deprecated Failover Clustering feature `cluster.exe`.

Job stream creation failure due to missing rights on calendars

When creating a job stream, if no calendar is specified in the job stream definition, an error occurs and a message similar to the following is displayed:

```
AWSJC0026E User <user_name>" is not authorized to perform the action "<action_name>" on an object
"<object_name>" and key "<key_name>".
```

In this case, modify the security file providing the ACCESS USE on CALENDAR HOLIDAYS to the user. For more information about configuring security, see the *Configuring user authorization (Security file)* chapter in the *Administration Guide*.

Suppressing connector error messages

Optionally suppress a subset of the error messages generated by the connectors

About this task

You can optionally suppress a subset of the error messages generated by the connectors. You can suppress the following messages:

- AWSJC0026E
- AWSJC0050E
- AWSJC0191E
- AWSJC0123E
- AWSJC0052E
- AWSJC0055E
- AWSJC0084E
- AWSJC0053E
- AWSJC0124E
- AWSJC0131E
- AWSJC0132E

To suppress the messages, perform the following steps:

1. Stop WebSphere Application Server Liberty Base, as explained in the section about Application server - starting and stopping in *Administration Guide*.
2. Browse to the `TWSConfig.properties` file, which is located in:

On Windows operating systems

`<TWA_home>\usr\servers\engineServer\resources\properties`

On UNIX operating systems

`<TWA_DATA_DIR>/usr/servers/engineServer/resources/properties`

3. Open the file in a flat text editor and set the `com.ibm.tws.conn.exception.avoidDetailedMessages` property to `true`. Add the property if it is not present.
4. Start WebSphere Application Server Liberty Base, as explained in the section about Application server - starting and stopping in *Administration Guide*.

StartUp shows an error after upgrade

Problem:

After upgrading to version 8.6, StartUp script shows the following error:

```
TWS for UNIX/STARTUP 8.5.1
Licensed Materials - Property of IBM*
5698-WSH
(C) Copyright IBM Corp. 1998, 2012 All rights reserved.
* Trademark of International Business Machines
Program code level: 20120510
Killed
ld.so.1: /export/home/svtUser/TWS/trace/atctl: fatal:
libatrc.so: open failed: No such file or directory
Killed
AWSBHU507I A start command was issued for NC121016.
```

Cause and solution:

During the upgrade to version 8.6, the following configuration files:

- `tws_env.sh`
- `tws_env.csh`
- `jobmanrc`
- `TWSCCLog.properties`
- `StartUp`
- `MakePlan`
- `SwitchPlan`
- `SwitchPlan`
- `CreatePostReports`
- `UpdateStats`

- ResetPlan
- Sfinal

are not overwritten but the 8.6 version of above files are installed under `twc_home/config` directory, and therefore to remove the above error message you must merge manually the two version of the files modifying the file under `twc_home` directory.

Chapter 8. Troubleshooting dynamic workload scheduling

This section provides information that is useful in identifying and resolving problems with dynamic workload scheduling, including how to tune the job processing rate and how to solve common dynamic scheduling problems.

It includes the following sections:

- [How to tune the rate of job processing on page 142](#)
- [Troubleshooting common problems on page 149](#)
- [Remote command job fails on page 145](#)
- [Monitoring and canceling jobs on page 145](#)

See also the section about auditing in the *Administration Guide*.

How to tune the rate of job processing

The processing of jobs submitted for dynamic scheduling is handled by the two subcomponents of dynamic workload broker, job dispatcher and resource advisor, through a mechanism of queues and a cache memory. Job dispatcher uses a system of queues into which jobs are placed according to their processing status and thus transmitted to the resource advisor. Resource advisor uses a system of time slots during which it takes a number of jobs from the job dispatcher and allocates them to the resources that will run them.

The `JobDispatcherConfig.properties` and `ResourceAdvisorConfig.properties` configuration files are tuned to suit most environments. However, if your environment requires a high job throughput or if jobs are processed too slowly, you can add the parameters listed below to the specified configuration files and provide customized values. The configuration files are created for dynamic workload broker at installation time and are documented in *IBM® Workload Scheduler: Administration Guide*.

By default, the parameters listed below are not listed in the configuration files to prevent unwanted modifications. Only expert administrators should set these parameters.

After modifying these parameters, stop and restart dynamic workload broker, as explained in the section about `startbrokerapp` in the *IBM® Workload Scheduler: Administration Guide*.

JobDispatcherConfig.properties

MaxProcessingWorkers

Job dispatcher queues the submitted jobs according to their processing status. By default the following 3 queues are already specified:

```
Queue.actions.0 = cancel,  
                 cancelAllocation,  
                 completed,  
                 cancelOrphanAllocation  
Queue.actions.1 = execute,  
                 reallocateAllocation  
Queue.size.1    = 20  
Queue.actions.2 = submitted,
```

```
notification,
updateFailed
```

Each queue is determined by the keywords:

Queue.actions.queue_number

Specifies the jobs added in this queue based on their processing status. The *queue_number* identifies the queue and ranges from 0 to 9. You can specify a maximum of 10 queues. The following table shows the entire list of process statuses you can specify in the queues.

Table 5. Job processing status to queue jobs for dispatching

Job processing statuses:		
activated	cancel	cancelAllocation
cancelJobCommand	cancelOrphanAllocation	childActivated
childCompleted	childDeactivated	childStarted
completed	deleteJobCommand	execute
getJobLogCommand	getJobPropertiesCommand	holdJobCommand
notification	reallocateAllocation	reconnect
resumeJobCommand	submitJobCommand	submitted
updateFailed	-	-

Unspecified job processing statuses are automatically placed in queue 0.

Queue.size.queue_number

Specifies the number of threads available to the queue identified by *queue_number*. You can specify 1 to 100 threads for each queue you define. The default is the number specified for `MaxProcessingWorkers`.

`MaxProcessingWorkers` specifies the default number of concurrent threads available to each queue. Each job dispatcher queue uses `MaxProcessingWorkers` threads, unless otherwise specified in `Queue.size.queue_number`. The `MaxProcessingWorkers` default is 10. Of the three default queues shown above, only queue 1 has its size specified to 20 threads (or workers). Queues 0 and 2 use the default defined in `MaxProcessingWorkers` (10 threads).

For example, in a test scenario with 250K jobs submitted through the workload broker workstation, the job allocation queues are re-configured as follows:

```
# Override default settings
Queue.actions.0 = cancel,
                 cancelAllocation,
```

```

cancelOrphanAllocation
Queue.size.0 = 10
Queue.actions.1 = reallocateAllocation
Queue.size.1 = 10
Queue.actions.2 = updateFailed
Queue.size.2 = 10

# Relevant to jobs submitted from
# workload broker workstation, when successful
Queue.actions.3 = completed
Queue.size.3 = 50
Queue.actions.4 = execute
Queue.size.4 = 50
Queue.actions.5 = submitted
Queue.size.5 = 50
Queue.actions.6 = notification
Queue.size.6 = 50

# Default for every queue size
MaxProcessingWorkers = 10

```

Tune this parameter carefully to avoid impairing product performance.

HistoryDataChunk

Specifies the number of jobs to be processed at the same time when moving job data to the archive database. This is applicable only to a DB2® RDBMS. This parameter prevents an overload on the job dispatcher. The unit of measurement is jobs. The default value is 1000 jobs.

ResourceAdvisorConfig.properties

MaxAllocsPerTimeSlot

Specifies the number of requests for job allocation to be processed for each time slot. The default value is 100 requests per time slot. By default, each time slot lasts 15 seconds. Increasing this number causes the resource advisor to process a higher number of resource allocation requests per time slot with consequent processor time usage. This also allows the processing of a higher number of jobs per time slot. Decreasing this number causes the resource advisor to process a lower number of resource allocation requests per time slot resulting in a smoother processor usage and slower job submission processing. You can also modify the time slot duration using the `TimeSlotLength` parameter available in this file.

MaxAllocsInCache

Specifies the number of requests for job allocation submitted by job manager to the resource advisor and stored in its cache. This number should be substantially higher than the value specified in the `MaxAllocsPerTimeSlot` parameter. The default value is 5000 allocation requests. Increasing this number causes the resource advisor to process a potentially higher number of resource reservations per time slot with consequent processor time usage. This also allows the processing of a higher number of jobs. Decreasing this number causes the resource advisor to process a lower number of resource reservations per time slot resulting in lower processor usage

and slower job submission processing. For optimal performance, this value should be at least 10 times the value specified in the `MaxAllocsPerTimeSlot` parameter.

Remote command job fails

You define and run a remote command job that performs a task on a remote Windows system. If the remote command job goes into ABEND state, and the job log contains a message like: **AWKRCE012E Could not establish a connection to "nc112134.romelab.it.ibm.com" target machine.**, see the following cause and solution.

Cause and solution:

A necessary Windows service might be stopped. Start the Remote Registry Windows service on the remote system.



Note: A Remote Command job that runs on a Windows workstation that is configured to use the samba protocol version 2 or 3, without an active SSH server, fails.

Monitoring and canceling jobs

This section explains how you can monitor and cancel jobs using the Dynamic Workload Console or the **conman** command.

You can use the Dynamic Workload Console or the **conman** command line to monitor the status of submitted jobs, retrieve the job output, and cancel jobs if necessary, as you normally do in IBM Workload Scheduler. You can also use the Dynamic Workload Console to view their status since it provides more detail on jobs processed through IBM® Workload Scheduler.

Job statuses in Dynamic workload broker correspond to the following statuses in IBM Workload Scheduler:

Table 6. Status mapping between Dynamic workload broker and IBM Workload Scheduler

Dynamic workload broker job status	IBM Workload Scheduler job status
1. Run failed	1. ABEND
2. Unable to start	2. FAILED
3. Resource allocation failed	3. FAILED
4. Unknown	4. ABEND

Table 6. Status mapping between Dynamic workload broker and IBM Workload Scheduler (continued)

Dynamic workload broker job status	IBM Workload Scheduler job status
1. Submitted	1. INTRO
2. Submitted to Agent	2. WAIT
3. Resource Allocation Received	3. WAIT
4. Waiting for Reallocation	4. WAIT
5. Waiting for Resources	5. WAIT
1. Running	1. EXEC
1. Completed Successfully	1. SUCC
1. Canceled	1. ABEND
2. Cancel Pending	2. The status is updated when the job reaches the Canceled state in IBM® Workload Scheduler
3. Cancel Allocation	3. The status is updated when the job reaches the Canceled state in IBM® Workload Scheduler



Note: The + flag written beside the INTRO and EXEC statuses means that the job is managed by the local **batchman** process.

You can view the job output by using both the Dynamic Workload Console or the **conman** command-line.

Consider the following Job Submission Description Language (JSDL) example, called BROKER_COMMAND_JOB:

```
<?xml version="1.0" encoding="UTF-8"?>
  <jSDL:jobDefinition xmlns:jSDL="http://www.abc.com/xmlns/prod/scheduling/1.0/jSDL"
  xmlns:jSDLE="http://www.abc.com/xmlns/prod/scheduling/1.0/jSDLE"
    name="BROKER_COMMAND_JOB">
    <jSDL:variables>
      <jSDL:stringVariable name="command">dir<jSDL:stringVariable>
      <jSDL:stringVariable name="params">d</jSDL:stringVariable>
    </jSDL:variables>
    <jSDL:application name="executable">
      <jSDLE:executable interactive="false">
      <jSDLE:script>${command} ${params}</jSDLE:script>
      </jSDLE:executable>
    </jSDL:application>
  </jSDL:jobDefinition>
```

and the associated IBM Workload Scheduler job definition, called TWS_COMMAND_JOB:

```

TTANCREDBRK#TWS_COMMAND_JOB
SCRIPTNAME "BROKER_COMMAND_JOB -var command=ping,
           params=ttancred.romelab.it.abc.com"
STREAMLOGON tws86master
TASKTYPE BROKER
RECOVERY STOP

```

The following example displays the output of the previous job, submitted from IBM Workload Scheduler to the workload broker workstation.

```

%sj TTANCREDDWB#JOBS.TWS_COMMAND_JOB;std

=====
= JOB      : TTANCREDDWB#JOBS[(0000 10/28/12),(JOBS)].TWS_COMMAND_JOB
= USER    : tws86master
= TASK     : <?xml version="1.0" encoding="UTF-8"?>
<jSDL:jobDefinition xmlns:jSDL="http://www.abc.com/xmlns/prod/scheduling/1.0/jSDL"
xmlns:jSDLE="http://www.abc.com/xmlns/prod/scheduling/1.0/jSDLE"
                name="BROKER_COMMAND_JOB">
  <jSDL:variables>
    <jSDL:stringVariable name="command">dir<jSDL:stringVariable>
    <jSDL:stringVariable name="params">d</jSDL:stringVariable>
  </jSDL:variables>
  <jSDL:application name="executable">
    <jSDLE:executable interactive="false">
      <jSDLE:script>${command} ${params}</jSDLE:script>
    </jSDLE:executable>
  </jSDL:application>
</jSDL:jobDefinition>
= AGENT    : TTAGENT
= Job Number: 318858480
= Wed Oct 24 15:58:24 CEST 2012

=====
Pinging ttancred.romelab.it.abc.com [9.168.101.100] with 32 bytes of data:

Reply from 9.168.101.100: bytes=32 time<1ms TTL=128
Reply from 9.168.101.100: bytes=32 time<1ms TTL=128
Reply from 9.168.101.100: bytes=32 time<1ms TTL=128
Reply from 9.168.101.100: bytes=32 time<1ms TTL=128

Ping statistics for 9.168.101.100:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 0ms, Average = 0ms

=====
= Exit Status      : 0
= Elapsed Time (Minutes) : 1
= Job CPU usage (ms) : 1406
= Job Memory usage (kb) : 3940
= Wed Oct 24 15:58:27 CEST 2012
=====

```

The keywords in the job output are as follows:

JOB

Is the host name of the IBM Workload Scheduler agent to which the job has been submitted and the job name.

USER

Is the IBM Workload Scheduler user who submitted the job to the workload broker workstation. When a scheduled job is submitted, the user name is retrieved from the **STREAMLOGON** keyword specified in the job definition. When an ad-hoc job is submitted from conman and the logon is not specified, the user name corresponds to the user who submitted the job.

TASK

Is the full JSDL job submitted with all the variables substituted.

AGENT

Is the agent that run the job.

Job Number

Is the job number.

Exit Status

Is the status of the job when completed.

System Time

Is the time the kernel system spent for the job.

User Time

Is the time the system user spent for the job.



Note: The same job output format is used for JSDL template job types.

You can also kill the job after submitting it. Killing a job in IBM Workload Scheduler performs the same operations as issuing the **cancel** command in IBM® Workload Scheduler.

On Windows 2012 the user interfaces for the interactive jobs are not visible on dynamic agents

On dynamic agent installed on Windows 2012 operating system, if you are running an interactive job, no user interface required by the interactive job, is displayed.

Cause and solution:

Ensure that the Interactive Services Detection service is running on Windows 2012 workstation where the dynamic agent is installed, and rerun the interactive job.



Note: Before start the Interactive Services Detection service on your operating system, ensure to read carefully the Windows documentation for the Interactive Services Detection running.

MASTERAGENTS workstation not updated in optman

After moving the MASTERAGENTS workstation to a folder, the workstation name is not updated in the resubmitJobName optman property.

When you modify the MASTERAGENTS workstation by storing it in a folder using the composer rename command, the command completes successfully. However, the resubmitJobName optman property is not updated and you cannot update it using the optman chg command. When you try to update the property, a message similar to the following is displayed:

```
AWSJCS037E The value of the string "<folder_name>/MASTERAGENTS" in property "RJ" is not correct. The value must be alphanumeric.
```

Cause and solution:

The `pools.properties` file is not updated automatically.

To solve the problem, perform the following steps:

1. Stop the agent using the ShutDownLwa command. For more information, see the section about the ShutDownLwa in *User's Guide and Reference*.
2. Browse to the following paths:

On Windows operating systems

```
<TWS_home>\ITA\cpa\config
```

On UNIX operating systems

```
TWA_DATA_DIR/ITA/cpa/config
```

3. Edit the `pools.properties` file and update the name of the MASTERAGENTS workstation as necessary.
4. Start the agent using the StartUpLwa command. For more information, see the section about the StartUpLwa in *User's Guide and Reference*.

Troubleshooting common problems

The following problems could be encountered with dynamic workload broker:

- [Dynamic workload broker cannot run after the IBM Workload Scheduler database is stopped on page 150](#)
- [Getting an OutofMemory exception when submitting a job on page 150](#)
- [Getting an error exception when submitting a job on a fresh agent on page 151](#)

On AIX operating systems the concurrent submission of one hundred or more jobs on the same agent can result in a core dump or in a resource temporarily unavailable message

On AIX operating systems, the concurrent submission of one hundred or more jobs on the same agent can result in a memory dump or in a resource temporarily unavailable message

On AIX operating systems if you submit concurrently one hundred or more jobs on the same agent you can receive a core memory dump or the following message:

```
resource temporarily unavailable
```

Cause and solution:

This problem is due to insufficient memory and the process number per user allocated to run the jobs concurrently. To solve this problem, verify the value of the following configuration settings and change them as follows:

Ulimit settings

The submission of a significant number of Java jobs requires a large amount of memory. Change the value for data, stack, and memory limits according to the number of jobs you want to submit. The submission of a significant number of native jobs requires a high number of file descriptors and processes. Change the values for nofiles and processes according to the number of jobs you want to submit. The following example gives possible setting values to submit 100 jobs concurrently:

```
time(seconds)      unlimited
file(blocks)       2097151
data(kbytes)       131072
stack(kbytes)      32768
memory(kbytes)     32768
coredump(blocks)   2097151
nofiles(descriptors) 4000
threads(per process) unlimited
processes(per user) unlimited
```

Process number per user

To submit a high number of jobs concurrently you must have a high value for the **maxuproc** setting. Use the **lsattr -E -l sys0 -a maxuproc** command to verify the number of concurrent processes that a user can create. Use the **chdev -l sys0 -a maxuproc=<value>** command to change the value for the **maxuproc** setting. For example, to submit 100 jobs concurrently use the following command:

```
chdev -l sys0 -a maxuproc=500
```

Dynamic workload broker cannot run after the IBM Workload Scheduler database is stopped

Getting an OutofMemory exception when submitting a job

If you get the following message after you submit a job for dynamic scheduling:

```
The job with ID job ID failed to start.
The error is "unable to create new native thread".
```

you must tune a property of the scheduling agent.

The property is named **ExecutorsMinThreads** and is located in the `JobManager.ini` file on the agent (for the path, see [Where products and components are installed on page 13](#)). Its default value is 38 but if this error occurs, you must decrease it to reduce the number of threads created when the job is launched.

The `JobManager.ini` file is described in the *IBM® Workload Scheduler: Administration Guide*.

Getting an error exception when submitting a job on a fresh agent

If you register a dynamic agent after the generation of a plan and you submit a job or a jobstream on a workstation class using this agent, you get an error message.

This behaviour is normal considering that the agent is not in the plan yet. To avoid getting this error message, do not use agents registered after the plan has been generated.

Chapter 9. Troubleshooting when automatically adding dynamic agent workstations to the plan

This section provides information that is useful in identifying and resolving problems in the IBM Workload Scheduler environment when you enable the automatic adding of dynamic agent workstations to the plan.

It includes the following section:

- [The dynamic agent workstation automatically added to the plan is not initialized on page 152.](#)

The dynamic agent workstation automatically added to the plan is not initialized

The global option `enAddWorkstation` is set to "yes". On the master domain manager, if the `Intercom.msg` file is full and you are stopping processes during the dynamic agent workstation addition to the plan, the dynamic agent workstation is added to the plan, but it might not be initialized, linked, or started. On the master domain manager, after the master domain manager processes restart, you cannot see the `LTI J` flags in the `conman` command-line output, if you perform the following command:

```
conman sc dynamic_agent_workstation_name
```

Where `dynamic_agent_workstation_name` is the dynamic agent workstation name that you inserted during the installation process.

Cause and solution:

The cause is that the master domain manager mailman stopping process is unable to manage the IBM Workload Scheduler event that communicates the workstation addition to the plan, also after the master domain manager processes restart.

To solve the problem, restart the master domain manager processes by performing in order the following commands:

```
conman stop mdm_workstation_name  
conman start mdm_workstation_name
```

Where `mdm_workstation_name` is the master domain manager workstation name.

Chapter 10. Troubleshooting Dynamic Workload Console problems

Describes how to troubleshoot problems with the Dynamic Workload Console related to connections, performance, user access, reports, and others.

This section describes the problems which could occur while using the Dynamic Workload Console:

The problems are described in these groups:

- [Troubleshooting connection problems on page 153](#)
- [Troubleshooting performance problems on page 160](#)
- [Troubleshooting user access problems on page 161](#)
- [Troubleshooting problems with reports on page 162](#)
- [Troubleshooting problems with graphical views on page 169](#)
- [Troubleshooting problems with database on page 170](#)
- [Troubleshooting other problems on page 171](#)

Troubleshooting connection problems

The following problems could occur with the connection to the engine or the database:

- [The engine connection does not work on page 153](#)
- [Test connection takes several minutes before returning failure on page 155](#)
- [Engine version and connection status not displayed on page 155](#)
- [Failure in testing a connection or running reports on an engine using an Oracle database on page 156](#)
- [Connection problem with the engine when performing any operation on page 156](#)
- [Engine connection does not work when connecting to the z/OS connector \(versions 8.3.x and 8.5.x\) on page 156](#)
- [Engine connection does not work when connecting to the z/OS connector V8.3.x or a distributed IBM Workload Scheduler engine V8.3.x on page 158](#)
- [Engine connection settings are not checked for validity when establishing the connection on page 159](#)
- [LDAP account locked after one wrong authentication attempt on page 159](#)

The engine connection does not work

You define an engine connection, you verify that the values entered for the engine connection are correct, and then you click **Test Connection**. The test fails and a connection error message is returned.

Cause and solution:

Assuming that `system_A` is where you installed the Dynamic Workload Console, and `system_B` is where you installed IBM Workload Scheduler, follow these verification steps to investigate and fix the problem:

1. Verify that there is no firewall between the two systems as follows:
 - a. Make sure the two systems can ping each other. If you are trying to connect to a z/OS® engine you must check that the system where the Dynamic Workload Console is installed and the system where the IBM Workload Scheduler z/OS® connector is installed can ping each other.
 - b. Make sure you can telnet from `system_A` to `system_B` using the port number specified in the engine connection settings (for example, 31117 is the default port number for distributed engine).
 - c. Make sure you can telnet from `system_A` to `system_B` using the CSIV2 authentication port numbers specified during installation (for example, 31120 is the default server port number and 31121 is the default client port number).

If either of these two steps fails then there might be a firewall preventing the two systems from communicating.

2. Check if you can connect using the **composer** command line interface, or the Dynamic Workload Console to the IBM Workload Scheduler engine on `system_B` using the same credentials specified in the engine connection. If you cannot, then check if the user definition on `system_B` and the user authorization specified in the IBM Workload Scheduler security file are correct.
3. If you are using LDAP or another User Registry on the Dynamic Workload Console make sure that:
 - a. The connection to the user registry works.
 - b. The User Registry settings specified on the Integrated Solutions Console in the **Security** menu under **Secure administration, applications, and infrastructure** are correct.

For more information about how to configure the Dynamic Workload Console to use LDAP or about how to test the connection to a User Registry, refer to the chapter on configuring user security in the *IBM Workload Scheduler: Administration Guide*.

4. If you set up to use Single Sign-On between the Dynamic Workload Console and the IBM Workload Scheduler engine, make sure you correctly shared the LTPA_keys as described in the chapter on configuring SSL in the *Administration Guide*.



Note: Make sure that you correctly shared the LTPA_keys also if you get errors AWSUI0766E and AWSUI0833E. The problem occurs when the realm values are the same for more than one WebSphere Application Server (Dynamic Workload Console, IBM Workload Scheduler z/OS® connector, or IBM Workload Scheduler engine). These steps are usually described only when you configure the Single Sign On, but they are required also when you have the same realm. You have the same realm when you configure all WebSphere Application Server Liberty Base instances with the same LDAP user registry and when you install all WebSphere Application Server Liberty Base instances on the same machine.

If this checklist does not help you in identifying and fixing your problem then activate tracing on the Dynamic Workload Console (adding also the Java™ packages `com.ibm.ws.security.*=all:com.ibm.tws.*=all`), and on the IBM Workload Scheduler engine. See the procedure in [Quick reference: how to modify log and trace levels on page 21](#).

After modifying the traces as necessary, connect to the Dynamic Workload Console again, test the connection to the IBM Workload Scheduler engine, and then check the information stored in the following trace logs:

- On the Dynamic Workload Console, browse to the following path:

On Windows operating systems

`DWC_home\stdlist\appserver\dwcServer\logs`

On UNIX operating systems

`DWC_DATA_dir/stdlist/appserver/dwcServer/logs`

- On the IBM Workload Scheduler engine, browse to the following path:

On Windows operating systems

`TWA_home\TWS\stdlist\appserver\engineServer\logs`

On UNIX operating systems

`TWA_DATA_DIR/stdlist/appserver/engineServer/logs`

In these files you see the information about the error that occurred. If useful, compare the connection information stored in the traces with the information set for WebSphere Application Server Liberty Base security on both sides. Refer to the *Administration Guide* to list the information about the security properties.

Test connection takes several minutes before returning failure

You select an engine connection and click **Test Connection** to check that the communication is working. The test takes several minutes to complete and then returns a failure.

Cause and solution:

When the **Test Connection** is run, the result is returned only after the timeout expires. The timeout for running the **Test Connection** operation cannot be customized. The connection failed because of one of the following reasons:

- The system where the IBM Workload Scheduler engine is installed is not active.
- The IP address or the hostname of the system where the IBM Workload Scheduler engine is installed was not correctly specified (in other words, the host name specified by the `showHostProperties` command must be capable of being contacted by the Dynamic Workload Console and vice versa)
- A network firewall prevents the system where the Dynamic Workload Console is installed and the system where the IBM Workload Scheduler engine is installed from communicating.

Check which of these reasons causes the communication failure, fix the problem, and then retry.

Engine version and connection status not displayed

The table listing your engine connections, in the Manage Engines panel of the Dynamic Workload Console does not display the engine version or the icon identifying the connection status.

Cause and solution:

Probably it happened because you tried to establish an engine connection using wrong credentials, then you entered the correct credentials and tested the connection again without checking the **Save** option. The correct credentials have not been saved and even though the connection is successful, data relating engine version and connection status is not loaded in the table.

Refresh the panel to make data displayed.

Failure in testing a connection or running reports on an engine using an Oracle database

You test the connection to an engine by specifying the user credentials for an Oracle database, or you run a report on that engine connection. The operation fails and the following error message is displayed:

```
AWSUI0360E The JDBC URL is not configured on the selected engine,
so the reporting capabilities cannot be used.
      Contact the IBM Workload Scheduler administrator."
```

Cause and solution:

Make sure that the IBM Workload Scheduler administrator has updated the `TWSConfig.properties` file by adding the following key:

```
com.ibm.tws.webui.oracleJdbcURL
```

specifying the JDBC Oracle URL. For example:

```
com.ibm.tws.webui.oracleJdbcURL=jdbc:oracle:thin:@//9.132.235.7:1521/orcl
```

Rerun the operation after the `TWSConfig.properties` has been updated. For more information about showing and changing database security properties for IBM Workload Scheduler, see *Administration Guide*.

Connection problem with the engine when performing any operation

Whatever operation you try to run in the Dynamic Workload Console, you get an error message saying that there is a connection problem with the engine.

Cause and solution:

Do the following steps:

1. Exit the Dynamic Workload Console.
2. Restart the WebSphere Application Server Liberty Base.
3. Log in again to the Dynamic Workload Console.

Continue with your activities on Dynamic Workload Console.

Engine connection does not work when connecting to the z/OS® connector (versions 8.3.x and 8.5.x)

If one of the following errors occurs when running the test connection, follow the steps described in the cause and solution section:

1. AWSUI0766E Test connection to *myengine*: failed. AWSUI0833E The operation did not complete. There was a communication failure. The internal message is: AWSJZC093E The requested engine zserver is not defined.
2. AWSUI0766E Test connection to *myengine* : failed. AWSUI0833E The operation did not complete. There was a communication failure. The internal message is: A communication failure occurred while attempting to obtain an initial context with the provider URL: "corbaloc:iiop:ZOS_CONNECTOR_HOSTNAME:31127".
3. AWSUI0766E Test connection to *myengine* : failed. AWSUI0833E The operation did complete. There was a communication failure. The internal message is: EQQPH26E TME user ID missing in TME user to RACF userid mapping table: myuser@hostname1.test.com

Cause and solution:

The possible causes for the case above are:

1. The name of the server startup job on host side must be defined on the z/OS® connector before you perform the test connection from the TDWC.
2. The WebSphere Bootstrap port is incorrect. Make sure that any bootstrap address information in the URL is correct and that the target name server is running. A bootstrap address with no port specification defaults to port 2809. Possible causes other than an incorrect bootstrap address or unavailable name server include the network environment and workstation network configuration.
3. The RACF® user ID has not been defined in the mapping table on host side.

You can solve the problem as follows:

Environment description example

The environment is composed of a z/OS® connector installed on the `hostname1.test.com`, a TDWC installed on either the same or another system, and a z/OS® engine installed on the `hostname2.test.com(port 445)`.

Steps on the z/OS® connector side

Define a connection from the z/OS® connector to the host side by running the following script located in the directory `<ZCONN_INST_DIR>/wastools` and then restart WebSphere®:

```
> createZosEngine -name zserver -hostName hostname2.test.com/portNumber 445
> stopWas
> startWas
```

where `zserver` is a logical name and can be changed to any other name.

Check the Bootstrap port by running the script `showHostProperties.bat (sh)` located in the directory

```
<ZCONN_INST_DIR>/wastools.
```

Steps on the TDWC side

On the TDWC web interface, define an engine connection from TDWC to the z/OS® connector, as follows:

Engine name

Choose any name.

Engine Type

z/OS®.

Host Name

Either `hostname1.test.com` or `localhost` depending on if TDWC is installed on the same host of Z/CONN or not.

Port Number

The z/OS® connector Bootstrap port.

Remote Server Name

`zserver` (or the name you used in step 2 - `createZosEngine`).

User ID / Password

For example, the credentials you specified when installing z/OS® Connector (that is, the user that owns the z/OS® Connector instance). The user can be any user that is authenticated by the User Registry configured on the embedded WebSphere® installed with the products.



Note: Bootstrap Port Number in version 8.5.x depends on which product is installed first. If TDWC is installed first, the Bootstrap port is 22809 and subsequent products installed on top of TDWC inherit that. If z/OS® Connector is installed first, the Bootstrap port is 31217. If the z/OS® connector version is 8.3 FPx, the default Bootstrap port is 31127.

Steps on the z/OS® side

Make sure that user `myuser@hostname1.test.com` is defined in the RACF® user ID mapping table on host side (USERMAP parameter in the SERVOPTS initialization statement).

Engine connection does not work when connecting to the z/OS® connector V8.3.x or a distributed IBM Workload Scheduler engine V8.3.x

If one of the following errors occurs when running the test connection, follow the steps described in the cause and solution section:

1. `AWSUI0766E Test connection to myengine: failed. AWSUI0833E The operation did not complete.`

```
Reason: AWSJC0005E WebSphere Application Server gives the following error:
CORBA NO_PERMISSION 0x0 No; nested exception is:
org.omg.CORBA.NO_PERMISSION: Trace from server: 1198777258 at host
myhostname.com >>
org.omg.CORBA.NO_PERMISSION: java.rmi.AccessException: ; nested exception is:
com.ibm.websphere.csi.CSIAccessException:
SECJ0053E: Authorization failed for /UNAUTHENTICATED while invoking (Bean)
ejb/com/ibm/tws/zconn/engine/ZConnEngineHome
getEngineInfo(com.ibm.tws.conn.util.Context):
1 securityName: /UNAUTHENTICATED;accessID:
```

```
UNAUTHENTICATED is not granted any of the required roles:
TWSAdmin vmcid: 0x0 minor code: 0 completed: No . . .
```

2. AWSUI0778E There was an authentication failure: the user name or password is incorrect.

Cause and solution:

The symptoms above are caused because on the z/OS® connector, or on the distributed engine side, the script **webui.sh** (bat) must be run to enable communication with the TDWC. Under the `wastools` directory of the home directory of the installation directory, run these commands:

```
./webui.sh -operation enable -user wasuser
          -password waspwd -port soap_port
          -pwdLTPA anypassword -server server1

./stopWas.sh -user wasuser -password waspwd

./startWas.sh
```

where:

`user` and `password` are those specified at installation time.

`port` is the WebSphere® SOAP port (display it by running the command **showHostProperties.sh**).

`pwdLTPA` is any password used to export and encrypt the LTPA keys.

`server` is the WebSphere® server name. The default is `server1`.

Engine connection settings are not checked for validity when establishing the connection

You incorrectly defined an engine connection to a distributed engine specifying a value for **Remote Server Name**. The **Remote Server Name** is not a valid setting for a connection to a distributed engine.

The check runs when you save the engine connection definition or when you run a test connection to that engine, but no exception about the incorrect setting is returned.

Cause and solution:

Whenever the test connection is run, only the mandatory fields for that specific type of engine, distributed rather than z/OS®, are used to test the connection. Fields that are not mandatory, such as **Remote Server Name** for distributed engine connections are not taken into account.

LDAP account locked after one wrong authentication attempt

LDAP accounts might be blocked even after only one login attempt when connecting using the web user interface or Dynamic Workload Console through LDAP/AD authentication, if wrong credentials are provided because of internal LDAP/AD security policy. This happens because one login attempt with wrong credentials using the web user interface or Dynamic Workload Console, is transformed into several authentication requests to LDAP.

Cause and solution:

When a single LDAP hostname is mapped to multiple IP addresses in a network configuration, if an invalid password is entered during the login, WebSphere makes as many login attempts as the number of associated IP addresses plus 1. If the resulting number exceeds the maximum number of failed logins allowed by local LDAP/AD security policy, the LDAP account is blocked. In the log file `messages.log` the following error shows an authentication error because of wrong credentials:

```
ECJ0369E: Authentication failed when using LTPA. The exception is
javax.naming.AuthenticationException: [LDAP: error code 49 -
80090308: LdapErr: DSID-0C090334, comment:
AcceptSecurityContext error, data 52e, vece
```

The WebSphere APAR PK42672 addresses this problem in the following way:

Two new custom properties are available to prevent this issue; use the one suitable for your LDAP configuration:

1. If LDAP is configured using the `wsadmin` command to register backend LDAP server hostnames, in the administration console click **Security > User Registries > LDAP > Custom Properties** and set to `true` the property `com.ibm.websphere.security.ldap.retryBind`. If this property is set to `false`, the Application Server does not retry LDAP bind calls. The default value is `true`.
2. If LDAP is configured associating a hostname with multiple IP addresses using the network configuration, in the administration console click **Security > User Registries > LDAP > Custom Properties** and set to `false` the property `com.ibm.websphere.security.registry.ldap.singleLDAP`. If this property is set to `true`, the Application Server does not resolve an LDAP hostname to multiple IP addresses. The default value is `false`.

Troubleshooting performance problems

- [With a distributed engine the responsiveness decreases overtime on page 160](#)
- [Running production details reports might overload the distributed engine on page 160](#)

With a distributed engine the responsiveness decreases overtime

When working with a distributed engine the responsiveness decreases overtime.

Cause and solution:

The problem might be related to multiple production plan report request running on that IBM Workload Scheduler engine, since those operations are CPU consuming. Ensure to wait until the report completion before running again other requests of the same kind.

Running production details reports might overload the distributed engine

The WebSphere Application Server Liberty Base on the distributed engine where the production details reports run, is overloaded and the temporary directory is full.

Cause and solution:

The amount of memory used by the WebSphere Application Server Liberty Base to extract the data varies depending on the number of objects to be extracted. For example, to extract 70 000 objects required almost 1 GB of RAM. To change

the WebSphere Application Server Liberty Base heap size, see the section about increasing application server heap size in *Administration Guide*.



Note: As a general recommendation, use filters to avoid extracting huge production report files.

Troubleshooting user access problems

- [Wrong user logged in when using multiple accesses from the same system on page 161](#)
- [Unexpected user login request after having configured to use Single Sign-On on page 161](#)
- [Authentication problem when opening the Workload Designer on page 162](#)

Wrong user logged in when using multiple accesses from the same system

You try to access the Dynamic Workload Console as *user2* using Firefox or Internet Explorer 7, where a connection as *user1* is already active in the same browser. In the case of Firefox the problem occurs if *user1* is active in *any* other Firefox window or tab. In Internet Explorer 7 the problem only occurs if the other user is active in a different tab of the same browser instance. But in both cases the result is the same: the browser logs you in to the Dynamic Workload Console as *user1* instead of *user2*.

Cause and solution:

This is a browser limitation. If you have an active connection through Internet Explorer 7 to the Dynamic Workload Console, and you want to open another session on the same system, you need only to open a different browser window. If the active connection is on Firefox, however, you must use a different browser. For a list of supported browsers, see the Dynamic Workload Console [System Requirements Document](#) at [Dynamic Workload Console Detailed System Requirements](#).

Unexpected user login request after having configured to use Single Sign-On

It might happen that, after running successfully all the steps required to configure the Single Sign-On between the Dynamic Workload Console and a IBM Workload Scheduler engine, when you try to test the connection or run a task on that engine, you are unexpectedly prompted to enter your user credentials to connect. This behavior means that the Single Sign-On method is not working properly on that engine and you receive the following exception:

Cause and solution:

Locate the `authentication_config.xml` configuration files on both the Dynamic Workload Console and the master domain manager.

The file is located in the following path for the master domain manager:

On UNIX operating systems

```
TWA_DATA_DIR/usr/servers/engineServer/configDropins/overrides
```

On Windows operating systems

```
TWA_home\usr\servers\engineServer\configDropins\overrides
```

The file is located in the following path for the Dynamic Workload Console:

On UNIX operating systems

`DWC_DATA_dir/usr/servers/dwcServer/configDropins/overrides`

On Windows operating systems

`DWC_home\usr\servers\dwcServer\configDropins\overrides`

Make sure that the files have identical values assigned to the *name* field of the *primaryRealm* section. This setting belongs to the WebSphere Application Server Liberty Base profile configuration.

For example, even though you ran all the required steps to configure the Single Sign-On, it might still not work if you set `name="myHost.myDomain:389"` on the Dynamic Workload Console and `name="myHost:389"` on the IBM Workload Scheduler engine. To make it work, you must enter: `name="myLDAPHost.myDomain:389"`.

Authentication problem when opening the Workload Designer

You try to open the Workload Designer but the action fails and you receive a message saying that the engine credentials are wrong, even though you entered the correct user name and password. This happens in a configuration with the Dynamic Workload Console connected to IBM Z Workload Scheduler connector on z/OS WebSphere Application Server Liberty Base.

The problem is due to the definition of the engine connection that contain the engine IP address instead of the hostname.

To solve the problem, open the Engine Connection Properties panel and replace the IP address with the hostname in the Connection Data section.

Troubleshooting problems with reports

- [The validate command running on a custom SQL query returns the error message AWSWUI0331E on page 162](#)
- [The output of report tasks is not displayed in a browser with a toolbar installed on page 163](#)
- [WSWUI0331E error when running reports on an Oracle database on page 163](#)
- [CSV report looks corrupted on Microsoft Excel not supporting UTF8 on page 163](#)
- [Insufficient space when running production details reports on page 164](#)
- [After IBM Workload Scheduler upgrades from version 8.3 to version 8.5 some fields in the output of reports show default values \(-1, 0, unknown, regular\) on page 164](#)
- [Report error: the specified run period exceeds the historical data time frame on page 165](#)

The validate command running on a custom SQL query returns the error message AWSWUI0331E

You are creating a Custom SQL report, and you run the **Validate** command to check your query. The validate fails and the following error message is returned:

```
AWSWUI0331E The SQL query could not be validated. The database internal message is:
[ibm][db2][jcc][10103][10941] Method executeQuery cannot be used for update.
```

Cause and solution:

The validate failure is caused by a syntax error in the query statement, for example, a typing error, such as:

```
sele Workstation_name,Job_name,Job_start_time from MDL.JOB_HISTORY_V
where Workstation_name like 'H%'
```

In this query, *sele* is written in place of *select*.

Verify the SQL query is correct and, optionally, try to run the same query from the DB2® command line to get additional details.

The output of report tasks is not displayed in a browser with a toolbar installed

You tested that the connection to the database set in the engine connection works properly but, after you run a report task, no window opens in your browser to display the task results. You have a third-party toolbar installed on your browser.

Cause and solution:

A third-party toolbar (such as Yahoo! or Google or similar) installed on top of the browser might conflict with the correct operation of the Dynamic Workload Console reporting feature. To make the reporting feature work correctly you must uninstall the toolbar and then rerun the report task.

WSWUI0331E error when running reports on an Oracle database

You try to run a report on an engine connection where an Oracle database has been referenced. The report task fails and the following error is displayed:

```
WSWUI0331E SQL validate failure.The database internal message is:ORA-00942:
table or view does not exist
```

If you try to run an SQL query statement in the Oracle database on the same table or view using the userid specified for the database connection in the engine connection properties, the query runs successfully.

Cause and solution:

On Oracle databases only, you must run these steps, as Oracle database administrator, to allow the database user specified in the engine connection properties to run reports from the Dynamic Workload Console:

1. Assign to the database user the "CREATE TABLE" Oracle System privilege.
2. Run the following script:

On Windows™

```
TWA_home\TWS\dbtools\oracle\script\dbgrant.bat
```

On UNIX™:

```
TWA_home/dbtools/oracle/script/dbgrant.sh
```

CSV report looks corrupted on Microsoft™ Excel not supporting UTF8

You run a report asking to save the result in a CSV file. When you open the CSV file using Microsoft™ Excel, the content of the file looks corrupted.

Cause and solution:

To bypass this problem, make sure that the version of Microsoft™ Excel you are using supports the UTF8 character set. If it does not, install a more recent version that supports UTF8. Then, follow these steps to correctly open CSV reports from Microsoft™ Excel:

1. Open Microsoft™ Excel.
2. In the **Data** menu entry, select **Import External Data** and then **Import Data**.
3. Select the CSV file saved and click **Open**.
4. In the field File Origin, select **UTF8**.

Insufficient space when running production details reports

When running production details reports the temporary directory on the IBM Workload Scheduler engine where the reports run, could be full.

Cause and solution:

You need to free some space in the temporary directory on the IBM Workload Scheduler engine before continuing to work on that engine.

After IBM Workload Scheduler upgrades from version 8.3 to version 8.5 some fields in the output of reports show default values (-1, 0, unknown, regular)

After migrating IBM Workload Scheduler from version 8.3 to version 8.5, the output on the Dynamic Workload Console of reports run on old migrated jobs show default values for the new fields introduced since version 8.3.

Cause and solution:

This is not a problem or a limitation but the result of migrating data from old tables to new tables containing newly created fields. After migration, it is necessary to assign a value to the new fields introduced since version 8.3 for job runs that occurred before migrating. The values assigned by default to these new fields are:

For job run statistic reports:

Table 7. Default settings for new job run statistic reports

Value	Field
0	Number of "Long Duration" job runs
0	Number of "Suppressed" job runs
0	Number of "Started Late" job runs
0	Number of "Ended late" job runs
0	Total Reruns
-1	Average CPU Time

Table 7. Default settings for new job run statistic reports (continued)

Value	Field
-1	Average Duration

For job run history reports:

Table 8. Default settings for new job run history reports

Value	Field
unknown	Workstation Name (Job Stream)
-1	Started Late (delay hh:mm)
-1	Ended Late (delay hh:mm)
-1	Estimated Duration (hh:mm)
No	Long Duration
Regular	Run Type
-1	Iteration Number
0	Return Code
0	Job Number
unknown	Login

Report error: the specified run period exceeds the historical data time frame

You define a report specifying a valid date range as execution period in the filter criteria. When you run the report, you receive the following warning message:

```
AWSUI2003The specified run period exceeds the historical data time frame. The database contains historical data from...
```

Cause and solution:

This problem occurs if different time zones are used when creating the report and when running it. To solve the problem, edit the report task, change the time zone making it equal to the time zone currently specified in the user preferences and run the report again.

Troubleshooting problems with browsers

- [Default tasks are not converted into the language set in the browser on page 166](#)
- ["Access Error" received when launching a task from the browser bookmark on page 166](#)
- [If you close the browser window, processing threads continue in the background on page 166](#)
- [Unresponsive script warning with Firefox browser on page 167](#)

- [Plan View panel seems to freeze with Internet Explorer version 7 on page 167](#)
- [Workload Designer does not show on foreground with Firefox browser on page 168](#)
- [Some panels in Dynamic Workload Console might not be displayed correctly in Internet Explorer, version 8 and 9 on page 168](#)
- [Web page error with Internet Explorer, version 9 on page 168](#)
- [Dynamic Workload Console problems with Internet Explorer developer tools on page 169](#)
- [Some Simplified Chinese characters are missing or corrupted when using Google Chrome or Apple Safari browser on page 169](#)

Default tasks are not converted into the language set in the browser

An existing user logs in to the Dynamic Workload Console using a browser where the language set is different from the language that was set in the browser the first time he logged in. In the Manage Tasks window, the default tasks are not translated into the new language.

Cause and solution:

The default tasks are created, using the current language set in the browser, when the new user logs into the Dynamic Workload Console for the first time. To have the default tasks translated into a different language, the WebSphere Application Server Liberty Base administrator must create a new Dynamic Workload Console user, and use that to login to the Dynamic Workload Console for the first time using a browser configured with the requested language. By doing this the default tasks are created using the requested language.

"Access Error" received when launching a task from the browser bookmark

A Dynamic Workload Console task has been saved in the list of bookmarks of the browser. You try to launch the task using the bookmark but you receive the following error message:

"User does not have access to view this page, use the browser back button to return to previous page."

Cause and solution:

You do not have the necessary role required to run the task. To run a task you must have a role that allows you to access the Dynamic Workload Console panels that are relevant to the type of task you need.

For more information about setting roles to work with the Dynamic Workload Console, see the Administration Guide, under the section about Configuring new users to access Dynamic Workload Console

If you close the browser window, processing threads continue in the background

You perform an action or make a selection and immediately close the browser window. You expect that processing terminated but the messages stored in the `messages.log` file show that processing continued in the background.

Cause and solution:

This is normal behavior for any WEB application, when the client browser is closed no notification is delivered to the server according to the HTTP protocol specifications. This is the reason why the last triggered thread continues to process even after the browser window was closed. You do not need to run any action, just allow the thread to end.

Unresponsive script warning with Firefox browser

When opening the Workload Designer with Firefox, the following warning message might appear:

```
Warning: Unresponsive script
A script on this page may be busy, or it may have stopped responding.
You can stop the script now, or you can continue to see if the script will complete.
```

Cause and solution:

This is caused by a Firefox timeout. If prompted with this warning message, select the **Continue** option.

This behavior of Firefox is ruled by its `dom.max_script_run_time` preference, which determines the timeout that the browser must wait for before issuing the warning. The default value is 10 seconds, and might be changed to another value according to your needs.

To change this value, perform the following steps:

1. Type `about:config` in the address field of the browser.
2. Scroll down to the preference, select it, change the value, and click **OK**.

Plan View panel seems to freeze with Internet Explorer version 7

When using Internet Explorer version 7, some actions performed in sequence might cause the **Plan View** browser window to freeze and stay frozen for about 5 minutes. After this timeframe the browser window resumes.

Cause and solution:

Action sequences that might cause this problem typically include opening multiple **Plan View** panels at the same time and refreshing the **Plan View** panels that were already open.

To avoid or limit this behavior add the Dynamic Workload Console website to the **Local intranet** security zone of Internet Explorer 7, with its default security level.

Blank page displayed (in High availability disaster recovery configuration)

You are in a high availability disaster recovery (HADR) configuration and the Dynamic Workload Console displays blank panels when trying to retrieve information from DB2.

Cause and solution:

When DB2 primary node stops, every Dynamic Workload Console request waits for a manual switch to a standby node.

If you have a HADR DB2 configuration related to your IBM Workload Scheduler engine, and you get an empty or blocked panel in the Dynamic Workload Console, verify that your primary node is up and running.

Workload Designer does not show on foreground with Firefox browser

With Firefox, if you open the Workload Designer from a graphical view (with the `Open Job definition` or the `Open Job stream definition` commands), and the Workload Designer window is already open, this window might not be moved to the foreground.

Solution:

To fix this problem, change the Firefox settings as follows:

1. On the Firefox action bar select **Tools**, then **Options**, then **Content**, and finally **Advanced**
2. Enable the `Raise or lower windows` option

Some panels in Dynamic Workload Console might not be displayed correctly in Internet Explorer, version 8 and 9

When using Internet Explorer version 8 or 9, some panels in Dynamic Workload Console, might not display as expected, for example

- The Graphical View or some Dashboard graphics, might not be displayed correctly.
- When duplicating a Monitor task, the entire Dashboard Application Services Hub navigation toolbar appears duplicated.

Cause and solution:

This problem can be due to incorrect settings in Internet Explorer.

To avoid or limit this behavior, use the following workarounds:

- Add the Dynamic Workload Console web site to the **Local intranet** security zone of Internet Explorer, with its default security level.
- Add the hostname of the Dynamic Workload Console to the web sites used in Compatibility View by Internet Explorer. To do it, from Internet Explorer toolbar, click **Tools > Compatibility View Settings** and add the Dynamic Workload Console hostname to the list.
- Turn off Internet Explorer Enhanced Security mode as described in the following documentation: *IBM Tivoli Security Information and Event Manager*, section about Disabling Enhanced Security Configuration.. In fact, Dashboard Application Services Hub does not support Internet Explorer with Enhanced Security mode active.

Web page error with Internet Explorer, version 9

When using Internet Explorer version 9, some panels in Dynamic Workload Console, might not be displayed correctly and the page issues the following web page error: "object Error".

Cause and solution:

To solve this problem, clear the browser cache.

Dynamic Workload Console problems with Internet Explorer developer tools

When using Internet Explorer version 8 or 9, the Dashboard might not display correctly when working with Internet Explorer developer tools open. Some sections of the Dashboard might remain loading without completing the update.

Cause and solution:

This problem is due to a conflict between how events are managed respectively by Internet Explorer developer tools and the Dashboard.

To avoid or limit this behavior, close the developer tools and try again the operation.

Some Simplified Chinese characters are missing or corrupted when using Google Chrome or Apple Safari browser

When you access the Self-Service Catalog or Self-Service Dashboard from a mobile device that uses the Google Chrome or Apple iPad Safari browser, if you use the GB18030 Simplified Chinese character set, the characters that you type might be missing or corrupted.

Cause:

Google Chrome and Apple iPad Safari do not fully support GB18030 Simplified Chinese.

Solution:

Ensure that you are using a browser that supports the GB18030 Simplified Chinese character set.

Troubleshooting problems with graphical views

- [Language-specific characters are not correctly displayed in graphical views on page 169](#)
- [Plan View limit: maximum five users using the same engine on page 169](#)
- [AWSITA122E or AWKRAA209E error while working with jobs in the Workload Designer on page 170](#)

Language-specific characters are not correctly displayed in graphical views

When working with the graphical views some language specific characters might not be displayed correctly.

Cause and solution:

This might occur because the necessary language files have not been installed on the computer on which the Dynamic Workload Console is running. To solve the problem, install the operating system language files on the system hosting the Dynamic Workload Console.

Plan View limit: maximum five users using the same engine

If you try to open the **Plan View** when five users are already concurrently using it, with the same engine, your request is rejected with the following error message: `AWSJCO136E No more than 5 users are allowed to perform this operation at the same time. The maximum number of concurrent requests has been reached: please try again later.`

Cause and solution:

The maximum number of users that can use the **Plan View** connected to the same engine is five.

If needed, you can modify this limit by editing the `com.ibm.tws.conn.plan.view.maxusers` property in the `TWSConfig.properties` file.

AWSITA122E or AWKRAA209E error while working with jobs in the Workload Designer

While you are working with a job in the Workload Designer, you might receive one of the following errors:

AWSITA122E - A Java exception occurred while calling the command

Cause and solution:

An unexpected error occurred while running a Java method.

In case it happens, check the following log files:

1. Check the `JobManager_message.log` file.
2. Check the latest `/opt/ibm/TWA/TWS/JavaExt/eclipse/configuration/*.log` file

AWKRAA209E - The job with advanced options with ID "application_type" was not found.

Cause and solution:

The job with advanced options cannot be found.

In case it happens, perform the following steps:

1. Ensure that the job plug-in is present in the `/opt/IBM/TWA/TWS/JavaExt/eclipse/plugins` directory.
2. Ensure that the job plug-in is listed in the `/opt/IBM/TWA/TWS/JavaExt/eclipse/configuration/config.ini` file
3. Check the latest `/opt/IBM/TWA/TWS/JavaExt/eclipse/configuration/*.lma` log file.

Troubleshooting problems with database

- [Import preferences fails on page 170](#)

Import preferences fails

When trying to import your settings repository from the XML file by specifying the **Cancel and recreate** option, the import operation fails and you receive the following message: *AWSUI0924E Preferences import operation failed: Unable to create database.* Moreover, the following error is logged in the `SystemErr.log` file: `DB2 SQL error: SQLCODE: -601, SQLSTATE: 42710, SQLERRMC: TDWC.TDWC_PREFERENCEABLE;TABLE.`

Cause and solution:

This might occur because the database user with administrative authority specified to import the settings does not have the privileges required to drop the existing Dynamic Workload Console tables that were created with Dynamic Workload Console V8.6.0.0.

To solve this problem, provide the specified user with `CONTROL` privilege on all Dynamic Workload Console tables.

For example,

```
db2 "GRANT CONTROL ON TABLE TDWC.TDWC_CONFIGURATIONPROPERTY TO USER myuser"
db2 "GRANT CONTROL ON TABLE TDWC.TDWC_CREDENTIAL TO USER myuser"
db2 "GRANT CONTROL ON TABLE TDWC.TDWC_ENGINECONNECTION TO USER myuser"
db2 "GRANT CONTROL ON TABLE TDWC.TDWC_MEQUERYTASK TO USER myuser"
db2 "GRANT CONTROL ON TABLE TDWC.TDWC_PREFERENCEABLE TO USER myuser"
db2 "GRANT CONTROL ON TABLE TDWC.TDWC_QUERYTASK TO USER myuser"
db2 "GRANT CONTROL ON TABLE TDWC.TDWC_REPORTTASK TO USER myuser"
```

Troubleshooting other problems

- [The deletion of a workstation fails with the "AWSJOM179E error on page 85](#)
- [Data not updated after running actions against monitor tasks results on page 172](#)
- ["Session has become invalid" message received on page 172](#)
- [Actions running against scheduling objects return empty tables on page 173](#)
- [Default tasks are not converted into the language set in the browser on page 166](#)
- ["Access Error" received when launching a task from the browser bookmark on page 166](#)
- [The validate command running on a custom SQL query returns the error message AWSWUI0331E on page 162](#)
- [If you close the browser window, processing threads continue in the background on page 166](#)
- [The list of Available Groups is empty in the Enter Task Information window on page 174](#)
- [Blank page displayed \(in High availability disaster recovery configuration\) on page 167](#)
- [Extraneous exception logged in SystemOut on page 175](#)
- [Filtering task results might not work as expected on page 176](#)
- [Sorting task results might not work as expected on page 178](#)
- [Monitoring job streams on multiple engines does not respect the scheduled time range on z/OS on page 178](#)
- [Java exception when performing a query on job streams in plan on page 179](#)
- [Import a property file bigger than 100MB on page 179](#)
- [Query results not displayed on page 180](#)

Jobs in READY status do not start

When monitoring a job, it appears to be ready to run, but it does not start. All of the dependencies have been satisfied and the start time has passed, but something is holding it back.

Cause and solution:

There are a few reasons why a job cannot start. The most common reasons are the following:

- The workstation limit is set to zero.
- The workstation is stopped.
- The workstation on which the job should run is not linked.
- The number of jobs running on the workstation is greater than the limit set on the workstation.

The Dynamic Workload Console provides the capability to detect the problem and provide the solution for these most common reasons.

From the Dynamic Workload Console, when monitoring the job status, you can request problem determination for a job in READY status that does not start. To determine why a job does not start:

1. From Monitor Workload, run a query to monitor jobs.
2. Select one or more jobs in READY status and then click **More Actions > Why a job does not start**.
3. A window displays an explanation and a solution for each selected job in READY status.

The deletion of a workstation fails with the "AWSJOM179E error

You want to delete a workstation either using Composer or the Dynamic Workload Console and the following error occurs:

```
AWSJOM179E An error occurred deleting definition of the workstation {0}.  
The workload broker server is currently unreachable.
```

Cause and solution:

This problem occurs if you removed a dynamic domain manager without following the procedure that describes how to uninstall a dynamic domain manager in the *IBM Workload Scheduler: Planning and Installation*.

To remove workstations connected to the dynamic domain manager, perform the following steps:

1. Verify that the dynamic domain manager was deleted, not just unavailable, otherwise when the dynamic domain manager restarts, you must wait until the workstations register again on the master domain manager before using them.
2. Delete the workstations using the following command:

```
composer del ws <workstation_name>;force
```

Data not updated after running actions against monitor tasks results

After you run an action on a list of objects returned from running a monitor task, the list is not updated.

Cause and solution:

The scheduling objects lists are not automatically updated after running actions. Click the **Refresh** button to update the list of objects.

"Session has become invalid" message received

You try to use the Dynamic Workload Console user interface, your working session closes, and you get the following warning:

```
Session has become invalid
Your session has become invalid. This is due to a session timeout, an administrator
has logged you out, or another user has invalidated your session by logging on with
the same User ID.
```

Cause and solution:

Check which reason among those listed in the warning has occurred, solve the issue, and then log in again to continue your working session.

If the session expired because either the HTTP session or the Lightweight Third Party Authentication (LTPA) session timeout was exceeded, you might decide to customize the timeout settings to values that are appropriate for your environment.

For instructions on how to do this, see the topic on session timeout settings in the *Performance* chapter of the *IBM® Workload Scheduler: Administration Guide*.

Actions running against scheduling objects return empty tables

After running a monitor task, you run an action against the scheduling objects listed in the result table, but you get, as a result of the action, an empty table or window, and no error message is displayed. This occurs regardless of which action you try to run against the listed scheduling objects.

Cause and solution:

Check if the connection with the IBM Workload Scheduler engine where you run the task failed by performing the following steps:

1. In the **Configuration** window select Scheduler Connections.
2. Select in the list the engine used to run the browse task and click **Test Connection**.



Note: The user ID you use to connect to the Dynamic Workload Console must belong to the **Administrator** groups to test the engine connection.

If the connection with the IBM Workload Scheduler engine is not active, ask the IBM Workload Scheduler administrator to restart the connection as described in the *IBM Workload Scheduler: User's Guide and Reference*, and then rerun the action.

If the connection with the IBM Workload Scheduler engine is active, then, on that engine, check that:

- The IBM Workload Scheduler user running the command to list scheduling objects is authorized to do so. For more information about how to set user authorization, see *IBM Workload Scheduler: User's Guide and Reference*.
- The global property *enListSecChk* is set to enable on the IBM Workload Scheduler master domain manager. For more information about how to set global properties, see *IBM Workload Scheduler: Planning and Installation*.

Then rerun the action.

Default tasks are not converted into the language set in the browser

An existing user logs in to the Dynamic Workload Console using a browser where the language set is different from the language that was set in the browser the first time he logged in. In the Manage Tasks window, the default tasks are not translated into the new language.

Cause and solution:

The default tasks are created, using the current language set in the browser, when the new user logs into the Dynamic Workload Console for the first time. To have the default tasks translated into a different language, the WebSphere Application Server Liberty Base administrator must create a new Dynamic Workload Console user, and use that to login to the Dynamic Workload Console for the first time using a browser configured with the requested language. By doing this the default tasks are created using the requested language.

"Access Error" received when launching a task from the browser bookmark

A Dynamic Workload Console task has been saved in the list of bookmarks of the browser. You try to launch the task using the bookmark but you receive the following error message:

"User does not have access to view this page, use the browser back button to return to previous page."

Cause and solution:

You do not have the necessary role required to run the task. To run a task you must have a role that allows you to access the Dynamic Workload Console panels that are relevant to the type of task you need.

For more information about setting roles to work with the Dynamic Workload Console, see the Administration Guide, under the section about Configuring new users to access Dynamic Workload Console

If you close the browser window, processing threads continue in the background

You perform an action or make a selection and immediately close the browser window. You expect that processing terminated but the messages stored in the `messages.log` file show that processing continued in the background.

Cause and solution:

This is normal behavior for any WEB application, when the client browser is closed no notification is delivered to the server according to the HTTP protocol specifications. This is the reason why the last triggered thread continues to process even after the browser window was closed. You do not need to run any action, just allow the thread to end.

The list of Available Groups is empty in the Enter Task Information window

You are creating a task, and you notice that in the Enter Task Information the list of Available Groups is empty. You are using an LDAP user registry.

Cause and solution:

Log into the Integrated Solutions Console as administrator and check the advanced LDAP configuration settings are correct as follows:

1. In the Navigation tree click **Security**.
2. Click **Secure administration, applications, and infrastructure**.
3. Check that the **Available realm definitions** field is set to **Standalone LDAP registry**.
4. Click **Configure**.
5. Click **Advanced Lightweight Directory Access Protocol (LDAP) user registry settings** under **Additional Properties**.
6. Verify that the settings for groups and users are correct for your configuration.

For more information about how to set these values, refer to: http://publib.boulder.ibm.com/infocenter/wasinfo/v6r0/topic/com.ibm.websphere.express.doc/info/exp/ae/usec_advldap.html

Blank page displayed (in High availability disaster recovery configuration)

You are in a high availability disaster recovery (HADR) configuration and the Dynamic Workload Console displays blank panels when trying to retrieve information from DB2.

Cause and solution:

When DB2 primary node stops, every Dynamic Workload Console request waits for a manual switch to a standby node.

If you have a HADR DB2 configuration related to your IBM Workload Scheduler engine, and you get an empty or blocked panel in the Dynamic Workload Console, verify that your primary node is up and running.

Exceptions might not be displayed in Language-specific in the Dynamic Workload Console

When working with the Dynamic Workload Console, some exception might not be displayed correctly in the specific Language.

Cause and solution:

This might occur because if the master domain manager is installed in English language, all the exception that the master domain manager returns, are in English. To solve the problem, you need to change the language of the machine where the engine is installed and restart master domain manager.

Extraneous exception logged in SystemOut

While working with the Dynamic Workload Console, the following exception message might be reported in the `messages.log` file:

```
ConnException E   AWSJC0005E WebSphere Application Server Liberty Base
has given the following error:
CORBA NO_PERMISSION 0x0 No;
nested exception is:
  org.omg.CORBA.NO_PERMISSION:
>> SERVER (id=4773e3aa, host=axrsgpar0612.metlife.com) TRACE START:
```

```
>> org.omg.CORBA.NO_PERMISSION: java.rmi.AccessException: ;
nested exception is:
  com.ibm.websphere.csi.CSIAccessException: SECJ0053E:
Authorization failed for
/UNAUTHENTICATED while invoking (Bean)ejb/com/ibm/tws/conn/engine
/ConnEngineHome getEngineInfo(com.ibm.tws.conn.util.Context):1 securityName:
/UNAUTHENTICATED;accessID: UNAUTHENTICATED is not granted any
of the required roles:
TWSAdmin vmcid: 0x0 minor code: 0 completed: No
```

Cause and solution:

This exception message might be logged if you are using an engine for which the credentials are not stored. You can ignore the exception message. It is not an indication that the product is not functioning correctly.

Filtering task results might not work as expected

When you use the quick filtering feature to filter the list of results shown in Dynamic Workload Console tables, with engines version 9.1 or later, you must consider the following limitations:

Filtering on dates and duration

- You cannot filter by timezone and offset.
- You cannot filter using text strings in date columns.
- Even though the table of results shows dates in the following format: `mm/dd/yyyy`, the leading "0" is not considered when filtering. For example, `6/8` when filtering is considered as if it was `06/08`.
- Even though the table of results shows duration time in the following format: `hh:mm`, the leading "0" is not considered when filtering. For example, `06:08` when filtering is considered as if it was `6:08`. In durations like `00:01` the `00` part cannot be matched with the search string `00`, because all non significant digits are discarded. The proper way to search for that is with the string `0`.

Filtering not supported

Quick filtering feature is not supported on the following columns:

- Information
- Node type

Filtering on job types

To filter for the following job types you must use the specified text:

Shadow Distributed

In the filter field, enter: `distributedShadowJob`

Shadow z/OS

In the filter field, enter: `zShadowJob`

Remote Command

In the filter field, enter: `remotecommand`

Database

In the filter field, enter: `database`

Executable

In the filter field, enter: `executable`

File Transfer

In the filter field, enter: `filetransfer`

IBM i

In the filter field, enter: `ibmi`

J2EE

In the filter field, enter: `j2ee`

Java

In the filter field, enter: `java`

z/OS

In the filter field, enter: `zosl`

MS SQL

In the filter field, enter: `mssqljob`

Provisioning

In the filter field, enter: `provisioning`

Web Services

In the filter field, enter: `ws`

Access Method

In the filter field, enter: `xajob`

OSLC Automation

In the filter field, enter: `oslcautomation`

Cause and solution:

This is due to a mismatch between how data is stored in the database and how it is shown from the Dynamic Workload Console.

You must set your user preferences so that dates are shown in short format (`6/27/08 5:59`) and you must use the specified strings to filter for job types.

Sorting task results might not work as expected

When you use the sorting feature to sort the list of results shown in Dynamic Workload Console tables, with engines version 9.1 or later, you must consider the following limitations:

- Sorting is not supported on Information column.
- When sorting on Node type column of Monitor workstations tasks, the sorting might result not correct.

Cause:

This is due to a mismatch between how data is stored in the database and how it is shown from the Dynamic Workload Console.

Monitoring job streams on multiple engines does not respect the scheduled time range on z/OS

If you created a task to monitor job streams on multiple engines with Dynamic Workload Console V9.1 or earlier, the scheduled time range is not respected on z/OS engines.

Solution:

This problem was solved on Dynamic Workload Console V9.2. Create the task again by using this version of the console.

Dynamic Workload Console 9.x login or graphical view pages do not display

If you access the Dynamic Workload Console from a supported browser, the following error message is displayed:

```
java.io.IOException: Too many open files
```

Another symptom related to the same problem is the graphical view opening a blank page.

Cause:

The number of open files for the WebSphere Application Server Liberty Base user has reached its limit. You can check the maximum number of open files with the command:

```
ulimit -n
```

You can list the open files and sockets with the command:

```
lsof -p <process id>
```

Solution:

Increase the limit for the number of open files for the WebSphere Application Server Liberty Base user, with the command:

```
ulimit -n <number_of_open_files>
```

It is recommended to set the limit to 8000.

For additional details, see the WebSphere Application Server Liberty Base technote: [Too Many Open Files error message](#).

Java exception when performing a query on job streams in plan

If the plan data replication feature (also known as mirroring) is disabled on the master domain manager, there is a time interval that makes a monitoring query fail and the following exception is displayed by the Dynamic Workload Console:

```
java.lang.UnsupportedOperationException
```

at

```
com.ibm.tws.dao.plan.SymphonyJobStreamDAO.queryJobStreams(SymphonyJobStreamDAO.java:1090)
    at com.ibm.tws.conn.plan.PlanImpl.queryPlanObjectPage(PlanImpl.java:2602)
```

Cause and solution:

This problem occurs if the plan data replication feature (also known as mirroring) is disabled on the master domain manager.

The workaround is to log-out and log-in from the Dynamic Workload Console.

Import a property file bigger than 100MB

You would need to upload a property file bigger than 100MB which is the limit set by default.

Solution:

Edit the settings that are specified in the `TdwcGlobalSettings.xml.template`.

By default, the customizable file is copied into the following path after you install the Dynamic Workload Console:

On Windows operating systems:

```
DWC_home\usr\servers\dwcServer\registry\TdwcGlobalSettings.xml.template
```

On UNIX and Linux operating systems:

```
DWC_home/usr/servers/dwcServer/configDropins/templates/TdwcGlobalSettings.xml.template
```

In the section User Registry, set the limit of the file size you want to import in the property as follow:

```
<property name="importSettingsMaxFileSize" value="102400"></property>
```



Note: the value is expressed in KB



Note: For security purposes, it is recommended to revert the file size back to the default value after the import has been performed.

See the topic about managing user settings in the *Dynamic Workload Console User's Guide*.

Query results not displayed

If you are running a query or displaying a table and switch to another page before the query or table has fully loaded, the results might not display correctly.

Cause and solution:

This happens because the page did not fully load the results before you switched to another page.

To resolve this problem, perform the following steps:

1. Right click on the page content.
2. Select the **Reload frame** option.
3. Wait until the results are fully loaded.

Chapter 11. Troubleshooting workload service assurance

Gives you troubleshooting information about workload service assurance by explaining how it works and how it exchanges information between the modules. In addition, it provides solutions to common problems.

This chapter provides information that is useful in identifying and resolving problems with the Workload Service Assurance feature. It includes the following sections:

- [Components involved in workload service assurance on page 181](#)
- [Exchange of information on page 182](#)
- [Common problems with workload service assurance on page 182](#)

Components involved in workload service assurance

Workload service assurance uses the following components to plan, monitor, and if necessary, intervene in the processing of jobs that are part of a critical network:

Planner

The planner component is triggered by the **JnextPlan** command. It includes a series of actions that result in the creation of the `Symphony` file on the master domain manager.

When workload service assurance is enabled, the planner calculates job streams and job networks, taking into consideration all "follows" dependencies in the new plan.

The planner then identifies all the jobs and job streams that are part of a critical network. These are jobs that are direct or indirect predecessors of a critical job. For each job, a critical start time is created and added to the `Symphony` file. It represents the latest time at which the job can start without putting the critical job deadline at risk. The plan information is then replicated in the database.

The `Symphony` file is subsequently distributed to all agents.

Plan monitor

The plan monitor component is introduced with the workload service assurance feature. It runs in the WebSphere Application Server Liberty Base on the master domain manager and is responsible for keeping track of the job streams and job network and for updating it when changes to the plan occur either because of the normal running of jobs or because of manual operations.

The plan monitor holds the information that is required to monitor the progress of the jobs involved in a critical network, for example critical start, planned start, estimated start, and risk level. It changes these values in response to changes in the plan, identified by the batchman process running on the master domain manager and communicated to the plan monitor using the `server.msg` file.

The information maintained by the plan monitor can be viewed on the Dynamic Workload Console in specialized views for critical jobs, allowing you easily to identify real and potential problems.

Agent processes (batchman and jobman)

Jobs in the critical network that are approaching the critical start time and have not started are promoted. The time at which the job is considered to be approaching its critical start time is determined by the global options setting `promotionOffset`.

The **batchman** process monitors the critical start time to determine if promotion is required and if so to schedule it at the highest job priority available in IBM Workload Scheduler. The **batchman** process also communicates with the **jobman** process, which is responsible for promoting the job at operating system level so that it receives more system resources when it starts. The operating system promotion is controlled by the local options settings `jm promoted nice` (UNIX™) and `jm promoted priority` (Windows™).

Exchange of information

Initially, the critical start time for jobs in the critical network is calculated by the planner and then recalculated, as required, by the plan monitor. Both of these components run on the master domain manager.

The critical start time is used by agents to determine when to promote a job. It is initially sent to the agent when the new `Symphony` file for the plan is distributed. Subsequent changes to critical start times are sent by the plan manager to agents using an IBM Workload Scheduler message. The agents update the local copy of the `Symphony` file.

The most common situations in which the plan monitor updates critical start times are:

- The Workload Designer functions on the Dynamic Workload Console or the `conman` command are used to modify jobs in the critical network. For example, predecessor jobs are added or cancelled.
- When **JnextPlan** is run to create the plan extension that includes the critical job, jobs in the original plan might be predecessors of the critical job and so be part of the critical network. In this case, critical start times are calculated by the plan monitor and sent in messages to the agents. This information is updated in both the local `Symphony` files and in the database where the plan data is replicated.

Common problems with workload service assurance

The following problems might occur when you are using IBM Workload Scheduler with workload service assurance enabled:

- [Critical start times not aligned on page 182](#)
- [Critical start times inconsistent on page 183](#)
- [Critical network timings change unexpectedly on page 183](#)
- [A high risk critical job has an empty hot list on page 184](#)

Critical start times not aligned

The values for critical start times in a critical network obtained from the appropriate `conman` commands on an agent are different from those displayed on the Tivoli® Dynamic Workload Console.

Cause and solution:

Changes that affect the critical start times have been made to the plan since the Symphony file was sent to the agent. The changes are calculated on the master domain manager and sent to agents in messages. It is probable that the message has not reached the affected agent.

Check that the agent is active and linked to the master domain manager, either directly or by other domain managers.

Critical start times inconsistent

The values for critical start time in the chain of jobs in the critical network appears to be inconsistent. There are predecessor jobs that have critical start dates that are later than their successors.

Cause and solution:

This inconsistency occurs when critical start times are recalculated after some of the jobs in the critical network have completed. To optimize the calculation, new critical start times are only recalculated and updated for jobs that have not yet completed. The completed jobs retain the original critical start time. If a completed job is subsequently selected to be rerun, its critical start date will be recalculated.

Critical network timings change unexpectedly

Timings for jobs in the critical network change even though there have been no user actions related to the timing of jobs.

Cause and solution:

Changes can be made to timings because of a plan extension or because of the submission of jobs or job streams.

A critical job is consistently late

A job that is defined as critical is consistently late despite promotion mechanisms being applied to it and its predecessors.

Cause and solution:

Using the successful predecessors task, compare the planned start, the actual start, and the critical start of all the predecessors of the late job. Check if any of them have time values that are too close together or have a planned start time that is later than the critical start time.

In such a case, you can:

- Consider changing the timings of these jobs. For example, postpone the deadline if possible, or if the deadline must be maintained anticipate the start of some of the jobs.
- Consider redesigning your job streams to optimize the paths that are causing delays.
- Increase the value of the `promotionOffset` global option, so that jobs are promoted earlier.
- On the workstations where jobs are tending to be late, increase the `jm promoted nice` (UNIX™) and `jm promoted priority` (Windows™) local options, so that promoted jobs receive more system resources.

A high risk critical job has an empty hot list

A job that is defined as critical is shown to be at high risk, but its hot list is empty.

Cause and solution:

This normally only occurs if you have designed a critical job or a critical predecessor with a conflict which means it will always be late, for example a start restriction after the critical job deadline. The hot list is empty if either the job or job stream that is causing the problem doesn't have its follows dependencies resolved, or the job stream that is causing the problem is empty.

The only solution is to examine the critical path in detail and determine where the problem lies. The steps to resolving this problem are the same as those documented in [A critical job is consistently late on page 183](#).

Chapter 12. Troubleshooting the fault-tolerant switch manager

Provides troubleshooting information about the fault-tolerant switch manager in terms of the event counter, the Ftbox, and link problems. It also provides solutions to some common problems with the backup domain manager.

This section describes how to address the potential problems related to the use of the fault-tolerant switch manager.

It is divided into the following sections:

- [Event counter on page 185](#)
- [Ftbox on page 186](#)
- [Troubleshooting link problems on page 186](#)
- [Common problems with the backup domain manager on page 190](#)

Event counter

The messages displayed in the log file concerning the event counter table are of three types:

- Messages that confirm the successful event counter initialization. No action is needed.
- Messages that the event counter reports related to problems not related to it. For example, they could reveal that the workstation received a message out of sequence. If action is required it does not impact the event counter.
- Messages that indicate that the event counter has failed. User action is needed to restore the counter.

This section concerns itself with this third type of messages.

Two processes can display this kind of error message:

Writer

When an error message of this type is received from writer, the event counters stops. All messages received from the workstation which asked netman to activate writer, and from all its children, are ignored. This can lead to two situations:

- The workstation talking to writer is switched to a new manager. In this case the new manager asks for a counter table and receive a corrupt counter table. The replay protocol proceeds following the default behavior.
- Before the switchmgr operation can be performed, writer fails and is automatically restarted. In this case the counter mechanism partially repairs itself. New messages received by the process are stored in the counter, but the messages received by the writer from the moment the error message was displayed up to the point at which writer restarted are not tracked. The situation for a given workstation might be considered as reset only when the new instance of writer receives a message from it.

The situation is recovered after the next scheduled JnextPlan. If you need to recover more urgently, run JnextPlan -for 0000 to refresh the Symphony file.

Mailman

When an error message of this type is received from mailman, the event counters stops. Mailman sets the IDs of all messages to 0. This means that there is a risk of duplication, because without the event counter, mailman is unable to properly sequence and process messages.

When the switchmgr is performed, and the new domain manager commences the replay protocol mechanism, for each message in the ftbox it looks at the position of the target workstation with respect to its own position in the tree:

- If the position of the target workstation in the workstation tree is higher than the new domain manager's (the workstation is either the domain manager or a full-status member of the parent domain of the domain where the switchmgr operation took place), the message is sent.
- If the position of the target workstation in the workstation tree is lower than the new domain manager's (the workstation either belongs to the domain where the switchmgr operation took place and it is not the new domain manager or is the domain manager or a full-status member of one of the child domains), the message is *not* sent.

The situation is recovered after JnextPlan.

Ftbox

If, on a full-status agent, you receive an error message concerning the ftbox, it means that the fault-tolerant backup domain manager feature is not working properly on that agent. Do not make this agent the new domain manager.

To restore the correct functionality of the feature on the instance, solve the problem as described in the error message, and restart the agent.

Troubleshooting link problems

When troubleshooting a link problem, the analysis is started from the master domain manager. The loss of the "F" flag at an agent indicates that some link had a problem. The absence of a secondary link can be located by matching the "W" flags found on the full-status fault-tolerant agent on the other side.

Consider the network shown in [Figure 1: ACCT_FS has not linked on page 187](#), where the workstation ACCT_FS, which is a full-status fault-tolerant agent, is not linked:

You might become aware of a network problem in a number of ways, but if you believe that a workstation is not linked, follow this procedure to troubleshoot the fault:

1. Use the command `conman sc @!@` on the master domain manager, and you can see that there is a problem with `ACCT_FS`, as shown in the example command output in [Figure 2: Example output for conman sc @!@ run on the master domain manager on page 188](#):

Figure 2. Example output for conman sc @!@ run on the master domain manager

```

$ conman sc @!@
Installed for user 'eagle'.
Locale LANG set to "C"
Schedule (Exp) 01/25/11 (#365) on EAGLE. Batchman LIVES. Limit: 20, Fence: 0,
Audit Level: 1
sc @!@
CPUID      RUN   NODE      LIMIT FENCE   DATE   TIME   STATE  METHOD  DOMAIN
EAGLE      365  *UNIX MASTER   20    0    01/25/11 05:59  I J
FS4MDM     365  UNIX  FTA     10    0    01/25/11 06:57 FTI JW  MASTERDM
ACCT_DM    365  UNIX  MANAGER  10    0    01/25/11 05:42 LTI JW  DM4ACCT
ACCT011    365  WNT   FTA     10    0    01/25/11 06:49 L I J  DM4ACCT
ACCT012    365  WNT   FTA     10    0    01/25/11 06:50 L I J  DM4ACCT
ACCT013    365  UNIX  FTA     10    0    01/25/11 05:32 L I J  DM4ACCT
ACCT_FS    363  UNIX  FTA     10    0
VDC_DM     365  UNIX  MANAGER  10    0    01/25/11 06:40 L I J  DM4VDC
FS4VDC     365  UNIX  FTA     10    0    01/25/11 06:55 F I J  DM4VDC
GRIDFTA    365  OTHR  FTA     10    0    01/25/11 06:49 F I J  DM4VDC
GRIDXA     365  OTHR  X-AGENT  10    0    01/25/11 06:49 L I J  gridage+ DM4VDC
LLFTA      365  OTHR  FTA     10    0    01/25/11 07:49 F I J  DM4VDC
LLXA       365  OTHR  X-AGENT  10    0    01/25/11 07:49 L I J  llagent  DM4VDC
$
    
```

2. From the `ACCT_DM` workstation run `conman sc`. In this case you see that all the writer processes are running, except for `ACCT_FS`. These are the primary links, shown by the solid lines in [Figure 1: ACCT_FS has not linked on page 187](#). The output of the command in this example is as shown in [Figure 3: Example output for conman sc run on the domain manager on page 189](#):

Figure 3. Example output for conman sc run on the domain manager

```

$ conman sc
TWS for UNIX (SOLARIS)/CONMAN 8.6 (1.36.2.21)
Licensed Materials Property of IBM
5698-WKB
(C) Copyright IBM Corp 1998,2011
US Government User Restricted Rights
Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM
Corp.
Installed for user 'dm010'.
Locale LANG set to "C"
Schedule (Exp) 01/25/11 (#365) on ACCT_DM. Batchman LIVES. Limit: 10, Fence: 0
, Audit Level: 1
sc
CPUID      RUN   NODE      LIMIT FENCE   DATE   TIME   STATE  METHOD  DOMAIN
EAGLE      365   UNIX MASTER  20    0    01/25/11 05:59 LTI JW   MASTERDM
ACCT_DM    365   *UNIX MANAGER  10    0    01/25/11 05:42 I J     DM4ACCT
ACCT011    365   WNT FTA     10    0    01/25/11 06:49 LTI JW   DM4ACCT
ACCT012    365   WNT FTA     10    0    01/25/11 06:50 LTI JW   DM4ACCT
ACCT013    365   UNIX FTA     10    0    01/25/11 05:32 LTI JW   DM4ACCT
ACCT_FS    363   UNIX FTA     10    0                                DM4ACCT
VDC_DM     365   UNIX MANAGER  10    0    01/25/11 06:40 LTI JW   DM4VDC
$

```

3. From the ACCT_FS workstation run conman sc. In this case you see that there are no writer processes running. These are the secondary links, shown with the dashed lines in [Figure 1: ACCT_FS has not linked on page 187](#). The output of the command in this example is as shown in [Figure 4: Example output for conman sc run on the unlinked workstation on page 189](#):

Figure 4. Example output for conman sc run on the unlinked workstation

```

$ conman sc
Installed for user 'dm82'.
Locale LANG set to "C"
Schedule (Exp) 01/24/11 (#364) on ACCT_FS. Batchman LIVES. Limit: 10, Fence: 0
, Audit Level: 1
sc @!@
CPUID      RUN   NODE      LIMIT FENCE   DATE   TIME   STATE  METHOD  DOMAIN
EAGLE      363   UNIX MASTER  20    0                                MASTERDM
FS4MDM     363   UNIX FTA     10    0                                MASTERDM
ACCT_DM    363   UNIX MANAGER  10    0                                DM4ACCT
ACCT011    363   WNT FTA     10    0                                DM4ACCT
ACCT012    363   WNT FTA     10    0                                DM4ACCT
ACCT013    363   UNIX FTA     10    0                                DM4ACCT
ACCT_FS    363   *UNIX FTA     10    0                                DM4ACCT
VDC_DM     363   UNIX MANAGER  10    0                                DM4VDC
FS4VDC     363   UNIX FTA     10    0                                DM4VDC
GRIDFTA    363   OTHR FTA     10    0                                DM4VDC
GRIDXA     363   OTHR X-AGENT  10    0                                gridage+ DM4VDC
$

```

4. If a network problem is preventing ACCT_FS from linking, resolve the problem.

5. Wait for ACCT_FS to link.
6. From the ACCT_FS workstation, run `conman sc @!@`. If the workstation has started to link, you can see that a writer process is running on many of the workstations indicated in [Figure 1: ACCT_FS has not linked on page 187](#). Their secondary links have now been made to ACCT_FS. The workstations that have linked have an "F" instead of their previous setting. This view also shows that the master domain manager has started a writer process running on ACCT_FS. The output of the command in this example is as shown in [Figure 5: Example output for conman sc @!@ run on the unlinked workstation on page 190](#):

Figure 5. Example output for `conman sc @!@` run on the unlinked workstation

```

$ conman sc @!@
Installed for user 'dm82'.
Locale LANG set to "C"
Schedule (Exp) 01/24/11 (#364) on ACCT_FS. Batchman LIVES. Limit: 10, Fence: 0
, Audit Level: 1
sc @!@
CPUID      RUN   NODE      LIMIT FENCE   DATE   TIME   STATE  METHOD  DOMAIN
EAGLE      371   UNIX MASTER  20    0    01/25/11 10:16 F I JW   MASTERDM
FS4MDM     370   UNIX FTA    10    0
ACCT_DM    371   UNIX MANAGER 10    0    01/25/11 10:03 LTI JW   DM4ACCT
ACCT011    369   WNT FTA    10    0
ACCT012    371   WNT FTA    10    0    01/25/11 11:03 F I JW   DM4ACCT
ACCT013    371   UNIX FTA    10    0    01/25/11 09:54 F I JW   DM4ACCT
ACCT_FS    371   *UNIX FTA    10    0    01/25/11 11:08 F I J    DM4ACCT
VDC_DM     371   UNIX MANAGER 10    0    01/25/11 10:52 F I JW   DM4VDC
FS4VDC     371   UNIX FTA    10    0    01/25/11 11:07 F I J    DM4VDC
GRIDFTA    371   OTHR FTA    10    0    01/25/11 11:01 F I J    DM4VDC
GRIDXA     371   OTHR X-AGENT 10    0    01/25/11 11:01 L I J    gridage+ DM4VDC
LLFTA     371   OTHR FTA    10    0    01/25/11 12:02 F I J    DM4VDC
LLXA      371   OTHR X-AGENT 10    0    01/25/11 12:02 L I J    llagent  DM4VDC
$

```

7. Another way of checking which writer processes are running on ACCT_FS is to run the command: `ps -ef | grep writer` (use Task Manager on Windows™). The output of the ps command in this example is as shown in [Figure 6: Example output for ps -ef | grep writer run on the unlinked workstation on page 190](#):

Figure 6. Example output for `ps -ef | grep writer` run on the unlinked workstation

```

$ ps -ef | grep writer
dm82 1363 616 0 06:43:11 ? 0:01 /usr/local/Tivoli/dm82/bin/write -- 2001 EAGLE MAILMAN UNIX 8.6 9
dm82 1317 616 0 06:42:21 ? 0:01 /usr/local/Tivoli/dm82/bin/write -- 2001 ACCT_DM MAILMAN UNIX 8.6 9
dm82 1337 616 0 06:42:25 ? 0:01 /usr/local/Tivoli/dm82/bin/write -- 2001 ACCT013 MAILMAN UNIX 8.6 9
dm82 1338 616 0 06:42:27 ? 0:01 /usr/local/Tivoli/dm82/bin/write -- 2001 VDC_DM MAILMAN UNIX 8.6 9
dm82 1364 616 0 06:51:48 ? 0:01 /usr/local/Tivoli/dm82/bin/write -- 2001 ACCT012 MAILMAN WNT 8.6 9
dm82 1336 616 0 06:42:24 ? 0:00 /usr/local/Tivoli/dm82/bin/write -- 2001 ACCT011 MAILMAN WNT 8.6 9
$

```

8. To determine if a workstation is fully linked, use the **Monitor Workstations** list in the Dynamic Workload Console.

Common problems with the backup domain manager

The following problems could be encountered with the fault-tolerant backup domain manager (note that a backup domain manager is an agent with the *full status* attribute set):

- [The Symphony file on the backup domain manager is corrupted. on page 191](#)
- [Processes seem not to have been killed on previous UNIX domain manager after running switchmgr on page 191](#)
- [In a scenario involving more than one switchmgr command, agent cannot relink on page 191](#)

The Symphony file on the backup domain manager is corrupted.

When switching to the backup domain manager from the master domain manager, the Symphony file on the backup domain manager might become corrupted.

Cause and solution:

The "thiscpu" variable in the `localopts` file does not match the workstation name. Change the variable to match the workstation name and the problem no longer occurs.

Processes seem not to have been killed on previous UNIX™ domain manager after running switchmgr

You want to use the switch manager facility. You first stop all IBM Workload Scheduler processes on the domain manager and then you run `switchmgr`, which completes successfully. However, after running `%sc @!@`, the J flag state is given for the domain manager where you stopped the processes.

Cause and solution:

When a shutdown command is sent to a workstation, some unexpected output might be shown by the status of the processes shown by `conman`, as follows:

- The J flag relative to the shut workstation remains active (no message indicating that jobman is not running can be transmitted because mailman is also not running).
- `Conman` output on the shutdown workstation is not up-to-date (the Symphony file is not updated on the shutdown workstation).
- The shutdown workstation seems linked from its parent and child workstations (no unlink operation is run by the writers on the workstation that is shutting down).
- Both F or L flags might be displayed, depending on the messages processed by mailman before unlinking and stopping.

The correct link situation is restored as soon as a new link attempt is made to the workstation, either manually, or automatically (after 10 minutes).

The shutdown command must be sent only in critical situations (where a workstation is shutting down, for example).

To avoid these problems, precede the shutdown command with an `unlink @!@` or `stop` command.

In a scenario involving more than one switchmgr command, agent cannot relink

You have been using the `switchmgr` command to switch to backup master domain manager, and then back to the master domain manager, but an agent might not have relinked to the original master domain manager.

Cause and solution:

The complex interaction of variables, environments, network conditions, and linking and relinking events can sometimes prevent an agent from relinking correctly.

No events or messages are lost, you can repeat the use of `switchmgr`, if necessary, and the performance of the network is not normally impacted because one agent is out of communication.

If only one agent is involved the easiest solution is to manually relink it.

However, to avoid having to identify and specifically relink the non-linked agent or agents, you can, in any case, issue the following command, which automatically relinks all agents without needing to specifically identify the unlinked ones:

```
JnextPlan -for 0000
```

Chapter 13. Synchronizing the database with the Symphony file

If you suspect the plan data loaded in the database is not up-to-date, you can run `planman resync` to update the database with the latest information in the `Symphony` file.



Note: If the message box file, `mirrorbox.msg`, responsible for synchronizing the database with the `Symphony` file becomes full, for example, the database is unavailable for a long period of time, then a `planman resync` is automatically issued so that the plan is fully reloaded in the database.

The procedure requires you to run the following step on the master domain manager. If you run the command on the backup master domain manager when it is not acting as the master domain manager, the plan data is not replicated in the database.

On the master domain manager, issue the following command:

```
planman resync
```

All of the plan data currently stored in the `Symphony` file is replicated in database tables in the database.

For the complete command-line syntax for the `planman` command, refer to the *IBM Workload Scheduler: User's Guide and Reference*.

To see the database views containing information about the objects in the plan, see the views beginning with "`PLAN_`" in *IBM Workload Scheduler: Database Views*.

Chapter 14. Corrupt Symphony file recovery

Explains the symptoms of Symphony file corruption and links you to tasks that can recover the file on the master domain manager, fault-tolerant agent, or lower domain manager.

Symphony file corruption is a rare event, and a potential corruption must be verified before taking action. Following are common symptoms of the file corruption:

- A specific message informing you that the Symphony file is corrupt.
- A shutdown of various processes (especially batchman) with error messages referring to problems with the Symphony file in the stdlist.

The normal reason for the corruption of the Symphony file is a full file system. This can be avoided by regular monitoring of the file system where IBM Workload Scheduler is installed.

The procedure is different, depending on the location of the corrupt Symphony file.

Recovery procedure on a master domain manager

If a Symphony file is corrupt on a master domain manager, there are several ways in which you can regenerate it.

The Symphony file can be regenerated in the following ways:

1. Using the backup master domain manager
2. Using the logman and ResetPlan commands
3. Using the latest archived plan

Recovering using the backup master domain manager

If the `Symphony` file is corrupt on a master domain manager, it can be regenerated using the backup master domain manager.

The regeneration of the `Symphony` file causes some minor loss of data. The following procedure indicates what is lost.

The prerequisite for the procedure is to have a backup master domain manager already available. A backup master domain manager is a fault-tolerant agent in the master domain with its `fullstatus` attribute set to `yes`.



Note: If you have not already created a backup master domain manager, the `Symphony` file cannot be recovered and the processing it contains is lost.

The procedure requires you to take the following steps on either the master domain manager or the backup master domain manager:



Note: The steps must be followed in strict order; each step description below is prefaced by the identification of the workstation on which it must be performed.

1. On the backup master domain manager, do the following:
 - a. Issue the `switchmgr` command.
 - b. Verify that the backup master domain manager is acting as the master domain manager.
2. From the new master domain manager set the job "limit" on the old master domain manager to "0", using `conman` or the Dynamic Workload Console.

This prevents jobs from launching.

3. On the original master domain manager do the following:
 - a. Shut down all IBM Workload Scheduler processes
 - b. Rename the `Sinfonia` file and the corrupt `Symphony` file (any names will do).
4. On the current master domain manager (previous backup master domain manager) do the following:
 - a. Verify that it is linked to all agents *except* the old master domain manager.
 - b. Shut down all IBM Workload Scheduler processes (unlink from all agents).
 - c. Rename `Sinfonia` as `Sinfonia.orig`
 - d. Copy `Symphony` to `Sinfonia`

You now have identical `Symphony` and `Sinfonia` files.

5. On the original master domain manager do the following:
 - a. Issue a `StartUp` from the operating system's command line, to start the `netman` process.
 - b. Verify that the process remains active.
6. On the current master domain manager (previous backup master domain manager) do the following:
 - a. Issue a `StartUp` from the operating system's command line, to start the `netman` process.
 - b. Issue a `conman start`, or use the Dynamic Workload Console to start the current master domain manager.
 - c. Issue a link to the original master domain manager.

This action sends the `Symphony` file to the original master domain manager.

7. On the original master domain manager do the following:
 - a. Verify that the `Symphony` file is present and is the correct size (same as on the current master domain manager (previous backup master domain manager))
 - b. Verify that all IBM Workload Scheduler processes are active.
8. On the current master domain manager (previous backup master domain manager) verify that the original master domain manager is linked.
9. On the original master domain manager do the following:
 - a. Set the job "limit" on the old master domain manager to the previous level, using `conman` or the Dynamic Workload Console.

Jobs can commence launching.

- b. Verify that the original master domain manager has the current job status for all agents.
- c. Issue the `switchmgr` command to switch control back to the original master domain manager.

Following this procedure some information is lost, in particular, any events that were suspended on the master domain manager when you started the recovery procedure.

If this procedure cannot be performed, try using the procedure that uses the `logman` and `ResetPlan` commands: [Recover using the logman and ResetPlan commands on page 196](#).

Recover using the logman and ResetPlan commands

The following procedures can also be used to recover a corrupt Symphony file on the master domain manager.

These procedures do not recover as much data as [Recovering using the backup master domain manager on page 194](#), but they might be useful if that procedure cannot be performed.

The procedure that makes use of `ResetPlan` might result in a more complete recovery, but it is more demanding in time since it scratches both the production and the preproduction plans. The preproduction plan will be created again based on the modeling information stored in the database when you later generate a new production plan. This means that the new production plan will contain all job stream instances scheduled to run in the time frame covered by the plan regardless of whether or not they were already in COMPLETE state when the plan was scratched.

You should first run the recovery procedure that makes use of `logman`. If you do not obtain satisfactory results, run the other one.

Neither procedure requires the use of a backup master domain manager.

Recovering the Symphony file using the logman command

Describes how to recover from a corrupt Symphony file using the `logman` command.

Perform these steps on the master domain manager:

1. Set the job "limit" to "0" on all the workstations by using `conman`, the Dynamic Workload Console. If you are using `conman`, run the following command:

```
conman "limit cpu=@!@;0;noask"
```

This command prevents jobs from launching.

2. Shut down all IBM Workload Scheduler processes on the master domain manager.
3. Run `logman -prod` to update the preproduction plan with the information for the job streams that are in COMPLETE state.
4. Run `planman showinfo` and check for the first incomplete job stream instance.
5. Run `ResetPlan`.
6. Run `JnextPlan`, setting the `-from` parameter to the start time of the first incomplete job stream instance in the preproduction plan (acquired from the output of `planman showinfo`) and the `-to` parameter to the end date of your plan (or to the following day). Only incomplete job stream instances will be included in the new Symphony file. If the

instance of the first incomplete job stream is very old the new plan creation can take a long time. The incomplete jobs and job streams that are created again with the JnextPlan -from parameter are those present in the database when the command is run.

7. Check the created plan and verify that all the jobs and job streams in the plan have the correct status.
8. Ensure that you want to run all the instances in the plan, deleting those that you do not want to run.
9. All the submitted job streams are not carried forward. Resubmit them.
10. Reset the job "limit" to its previous value. The Symphony file is distributed and production starts again.



Note:

- - The status of the jobs and the job streams after you run the recovery procedure is reset to HOLD or READY.
 - Some IBM® Workload Scheduler events that were triggered before applying the recovery procedure, might be triggered again after the recovery procedure has completed. This limitation concerns those events that are not managed through a message queue, for example, UNTIL, DEADLINE, and MAXDUR.
 - Jobs in USERJOBS Job Stream are not subject to resource controlling. As a result, affected resources should be adjusted and attended manually.
 - Prompts in the recovered plan might have a prompt number different from the prompt number in the original plan. To prevent mismatches, prompt reply events are not recovered.

Recovering with the use of the ResetPlan command

Perform these steps on the master domain manager:

1. Set the job "limit" to "0", using conman or the Dynamic Workload Console.

This prevents jobs from launching.


2. Shut down all IBM Workload Scheduler processes on the master domain manager.
3. Run ResetPlan -scratch.
4. Run JnextPlan, setting the -from and -to parameters to cover the period for which there are still outstanding jobs.
5. Check the created plan and ensure that you want to run all the instances it contains, deleting those that you do not want to run.
6. Reset the job "limit" to the previous value. The Symphony file is distributed and production recommences.

Recovering the plan from the latest archived plan

You can recover a corrupted plan on a master domain manager using the latest archived plan, however; this is possible only if you have performed some configuration steps prior to the occurrence of the file corruption.

About this task

The following procedure recovers a corrupted plan using ResetPlan. The plan is recovered using the latest archived plan and any events logged throughout the day are written to a new event message file. The last archived Symphony file is copied into the current `Symphony` file and then JnextPlan is run to apply the events from the `evtlog.msg` file.

 **Restriction:** Before you can perform the recovery procedure, you must have completed some configuration steps prior to the file corruption occurrence.


 **Note:**

- - Some IBM® Workload Scheduler events that were triggered before applying the recovery procedure, might be triggered again after the recovery procedure has completed. This limitation concerns those events that are not managed through a message queue, for example, UNTIL, DEADLINE, and MAXDUR.
 - Jobs in USERJOBS Job Stream are not subject to resource controlling. As a result, affected resources should be adjusted and attended manually.
 - Prompts in the recovered plan might have a prompt number different from the prompt number in the original plan. To prevent mismatches, prompt reply events are not recovered.

1. Complete the following configuration steps so that you can use the recovery procedure in the future if it becomes necessary:

- a. In the `localopts` file, add the following attribute and value: **bm log events = ON**.
- b. Optionally, customize the path where IBM Workload Scheduler creates the `evtlog.msg` event file by setting the **bm log events path** property in the `localopts` file. If you do not modify this setting, the `evtlog.msg` event file is created in the following default location: `<TWA_INST_DIR>/TWS`.
- c. Stop and start all IBM Workload Scheduler processes or run JnextPlan to create the `evtlog.msg` file.
- d. If necessary, you can configure the maximum size of both the `evtlog.msg` and `Intercom.msg` event files as follows:

```
evtsize -c evtlog.msg 500000000
evtsize -c Intercom.msg 550000000
```

 **Note:** The default size of these event files is 10 MB. When the maximum size is reached, events are no longer logged to these files and the recovery procedure is unable to recover them and any that follow. Moreover, the following BATCHMAN warning is logged to `<TWA_INST_DIR>/TWS/stdlist/traces/YYYYMMDD_TWSMERGE.log`:

```
13:11:51 18.10.2012|BATCHMAN:+ WARNING:Error writing in evtlog:
      AWSDEC003I End of file on events file.
13:11:51 18.10.2012|BATCHMAN:*
13:11:51 18.10.2012|BATCHMAN:* AWSBHT160E The EvtLog message file is full,
      events will not be logged until a new Symphony is produced. Recovery with
```



```
event reapply is no more possible until that time.
13:11:51 18.10.2012|BATCHMAN:*
```

If you encounter this problem increase the size of the `evtlog.msg` and `Intercom.msg` event files.

Consider that for 80,000 jobs and a Symphony file of size 40 MB, the `evtlog.msg` file is approximately 70 MB in size.



Important: The `Intercom.msg` maximum size should always be set to a value greater than the maximum size of `evtlog.msg`

In the `<TWA_INST_DIR>/TWS/stdlist/traces/YYYYMMDD_TWSMERGE.log` trace file, the BATCHMAN process logs an informational line containing the expected size of the `evtlog.msg` queue. For example:

```
19:02:06 14.10.2012|BATCHMAN:INFO:0.25 MB of events to log during this
batchman run
```

If `Intercom.msg` reaches the maximum size during the recovery procedure, batchman stops.

- e. If the file system where `evtlog.msg` resides runs out of space, a BATCHMAN warning is logged to `<TWA_INST_DIR>/TWS/stdlist/traces/YYYYMMDD_TWSMERGE.log` as follows:

```
13:10:36 16.10.2012|BATCHMAN:+ WARNING>Error writing in evtlog:
AWSDEC002E An internal error has occurred.
The following UNIX system error occurred on an events file:
"No space left on device" at line = 3517.
```

2. Complete the recovery procedure:

- Ensure the IBM Workload Scheduler processes are stopped. Run `conman stop` to stop them.
- Copy the information retrieved by running the `planman showinfo` command.
- Run `ResetPlan`. The corrupted Symphony file is archived in the `schedlog` folder.
- Copy the second last Symphony file archived in the `schedlog` folder, and not the most recent one which is the corrupted file. For example, on UNIX, submit the following command:

```
cp -p /opt/ibm/TWA/TWS/schedlog/MYYYYMMDDhhmm /opt/ibm/TWA/TWS/Symphony
```

- e. Run `JnextPlan` as follows using the information retrieved when you ran `planman showinfo`:

```
JnextPlan -from MM/DD/YYYY hhmm TZ Timezone -for hhhmm
```

where,

-from

Production plan start time of last extension.

-for

Production plan time extension.

Results

When running this procedure, consider that the run number, that is, the total number of times the plan was generated, is automatically increased by one.

Job stream instances that have already completed successfully at the time this procedure is run are not included in the recovered plan.

After you have performed the recovery procedure, the workstation limit is set to 0 and the `evtlog.msg` queue is cleared with each successive run of JnextPlan.

Recovery procedure on a fault-tolerant agent or lower domain manager

If the `Symphony` file is corrupt on a lower level domain manager, or on a fault-tolerant agent, it can be replaced.

Complete removal and replacement of the `Symphony` file causes some loss of data. The following procedure minimizes that loss and indicates what is lost.

The procedure involves two agents, the agent where the `Symphony` file is corrupt and its domain manager.



Note: Where the agent is a top level domain manager (below the master), or a fault-tolerant agent in the master domain, the manager is the master domain manager.

The procedure is as follows:

1. On the domain manager, unlink the agent which is having the `Symphony` file problem.
2. On the agent do the following:
 - a. Stop the agent if it has not yet failed. You do not need to shut it down.
 - b. Delete the `Symphony` and the `Sinfonia` files from the agent workstation. Alternatively you can move them to a different location on the agent workstation, or rename them.
3. On the domain manager do the following:
 - a. Back up the `Sinfonia` file if you want to be able to restore the original situation after completion. This is not an obligatory step, and no problems have been reported from not performing it.
 - b. Ensure that no agent is linking with the domain manager, optionally stopping the domain manager agent.
 - c. Copy the `Symphony` file on the domain manager to the `Sinfonia` file, replacing the existing version.
 - d. Restart the domain manager agent if necessary.
 - e. Link the agent and wait for the `Symphony` file to copy from the domain manager to the agent. The agent automatically starts.
 - f. Optionally restore the `Sinfonia` file from the backup you took in [step 3.a on page 200](#). This restores the original situation, but with the agent now having an uncorrupted `Symphony` file. This is not an obligatory step, and no problems have been reported from not performing it.

Following this procedure some information is lost, in particular, the contents of the `Mailbox.msg` message and the `tomaster.msg` message queues. If state information about a job was contained in those queues, such that the `Symphony`

file on the domain manager was not updated by the time the `Sinfonia` file is replaced (step 3.c on page 200), that job is rerun. To avoid that event, add these steps to the procedure immediately before step 3.a on page 200:

1. Make a list of jobs that ran recently on the agent.
2. At the domain manager, change their states to either SUCC or ABEND, or even cancel them on the domain manager.



Note: if you set the states of jobs to SUCC, or cancel them, any successor jobs would be triggered to start. Ensure that this is the acceptable before performing this action.

This way these jobs are not rerun.

Recovery procedure on a fault-tolerant agent with the use of the `resetFTA` command

If the `Symphony` file is corrupt on a fault-tolerant agent, you can use the `resetFTA` command to automate the recovery procedure.

Complete removal and replacement of the `Symphony` file causes some loss of data, for example events on job status, or the contents of the `Mailbox.msg` message and the `tomaster.msg` message queues. If state information about a job was contained in those queues, that job is rerun. The following procedure minimizes that loss and indicates what is lost. It is recommended that you apply this procedure with caution.

The procedure renames the `Symphony`, `Sinfonia`, `*.msg` files on the fault-tolerant agent where the `Symphony` corruption occurred and generates an updated `Sinfonia` file, which is sent to the fault-tolerant agent. You can therefore resume operations quickly on the affected fault-tolerant agent, minimize loss of job and job stream information, and reduce recovery time.

The procedure involves two agents, the fault-tolerant agent where the `Symphony` file is corrupt and its domain manager.

You can start the command from any IBM Workload Scheduler workstation, with the exception of the fault-tolerant agent where the corruption occurred. Connection to the target fault-tolerant agent and to its domain manager is established using the `netman` port number. The default port number is 31111.

When you start the `resetFTA` command, the following operations are performed in the specified order:

on the fault-tolerant agent

- The following files are renamed:
 - `Appserverbox.msg`
 - `clbox.msg`
 - `Courier.msg`
 - `Intercom.msg`
 - `Mailbox.msg`
 - `Monbox.msg`
 - `Moncmd.msg`

- Symphony
- Sinfonia

The operations are performed asynchronously, to ensure that all target files have been renamed before starting the procedure on the domain manager.

on the domain manager

1. A backup of the `Sinfonia` file is created.
2. The `Symphony` file is copied to the `Sinfonia` file.
3. The target fault-tolerant agent is linked.
4. The updated `Sinfonia` file is sent to the target fault-tolerant agent.

The syntax of the command is as follows:

Syntax

resetFTA `cpu`

Arguments

cpu

Is the fault-tolerant agent to be reset.

This command is not available in the Dynamic Workload Console.

For more information, see the section about the `resetfta` command in *IBM Workload Scheduler: User's Guide and Reference*.

Appendix A. Support information

If you have a problem with your IBM® software, you want to resolve it quickly. This section describes the following options for obtaining support for IBM® software products:

- [Searching knowledge bases on page 203](#)
- [Obtaining fixes on page 204](#)
- [Receiving support updates on page 205](#)

Searching knowledge bases

You can search the available knowledge bases to determine if your problem was already encountered and is already documented.

Search online product documentation in IBM Documentation

IBM® provides extensive documentation that you can search and query for conceptual information, instructions for completing tasks, and reference information.

The online product documentation can be found in IBM Documentation at: <http://www-01.ibm.com/support/knowledgecenter/SSGSPN/welcome>.

Search the Internet

If you cannot find an answer to your question in the information center, search the Internet for the latest, most complete information that might help you resolve your problem.


To search multiple Internet resources for your product, use the **Web search** topic in your information center. In the navigation frame, click **Troubleshooting and support** → **Searching knowledge bases** and select **Web search**. From this topic, you can search a variety of resources, including the following:

- IBM® technical notes (Technotes)
- IBM® downloads
- IBM® Redbooks®
- IBM® developerWorks®
- Forums and newsgroups
- Google


Search the IBM® support website

The IBM® software support website has many publications available online, one or more of which might provide the information you require:

1. Go to the IBM® Software Support website (<http://www.ibm.com/software/support>).
2. Select **Tivoli** under the **Select a brand and/or product** heading.

3. Select **IBM Workload Scheduler** under **Select a product**, and click the "Go" icon: . The IBM Workload Scheduler support page is displayed.
4. In the **IBM Workload Scheduler support** pane click **Documentation**, and the documentation page is displayed.
5. Either search for information you require, or choose from the list of different types of product support publications in the **Additional Documentation support links** pane:
 - Information center
 - Manuals
 - IBM® Redbooks®
 - White papers

If you click **Information center**, the IBM Workload Scheduler Information Center page opens, otherwise a search for the selected documentation type is performed, and the results displayed.



6. Use the on-screen navigation to look through the displayed list for the document you require, or use the options in the **Search within results for** section to narrow the search criteria. You can add **Additional search terms** or select a specific **Document type**. You can also change the sort order of the results (**Sort results by**). Then click the search icon to start the search: .

To access some of the publications you need to register (indicated by a key icon beside the publication title). To register, select the publication you want to look at, and when asked to sign in follow the links to register yourself. There is also a FAQ available on the advantages of registering.

Obtaining fixes

About this task

A product fix might be available to resolve your problem. To determine what fixes are available for your IBM® software product, follow these steps:

1. Go to the IBM® Software Support website (<http://www.ibm.com/software/support>).
2. Select **Tivoli** under the **Select a brand and/or product** heading.
3. Select **IBM Workload Scheduler** under **Select a product** and click the "Go" icon: . The IBM Workload Scheduler support page is displayed.
4. In the **IBM Workload Scheduler support** pane click **Download**, and the download page is displayed.
5. Either choose one of the displayed most-popular downloads, or click **View all download items**. A search for the downloads is performed, and the results displayed.
6. Use the on-screen navigation to look through the displayed list for the download you require, or use the options in the **Search within results for** section to narrow the search criteria. You can add **Additional search terms**, or select a specific **Download type**, **Platform/Operating system**, and **Versions**. Then click the search icon to start the search: .
7. Click the name of a fix to read the description of the fix and to optionally download the fix.

For more information about the types of fixes that are available, see the *IBM® Software Support Handbook* at <http://www14.software.ibm.com/webapp/set2/sas/f/handbook/home.html>.

Receiving support updates

About this task

To receive email notifications about fixes and other software support news, follow these steps:

1. Go to the IBM® Software Support website at <http://www.ibm.com/software/support>.
2. Click **My notifications** under the **Stay informed** heading in the upper-right corner of the page.
3. If you have already registered for **My support**, sign in and skip to the next step. If you have not registered, click **register now**. Complete the registration form using your email address as your IBM® ID and click **Submit**.
4. Follow the instructions on the page for subscribing to the information you require, at the frequency you require, for the products you require.

If you experience problems with the **My notifications** feature, you can obtain help in one of the following ways:

Online

Send an email message to erchelp@ca.ibm.com, describing your problem.

By phone

Call 1-800-IBM-4You (1-888 426 4409).

Appendix B. Date and time format reference - strftime

IBM Workload Scheduler uses the *strftime* standard method for defining the presentation of the date and time in log files generated by CCLog. There is a parameter in the properties file of CCLog, where you define the format (see [IBM Workload Scheduler logging and tracing using CCLog on page 26](#)).

This parameter uses one or more of the following variables, each of which is introduced by a "%" sign, separated, if required, by spaces or other character separators.

For example, to define a date and time stamp that would produce the following (12-hour time, followed by the date) "7:30:49 a.m. - November 7, 2008", you would use the following definition:

```
%L:%M:%S %P - %B %e, %G
```

The full details of the parameters you can use are as follows:

Table 9. strftime date and time format parameters

Parameter	Description	Example
%a	The abbreviated weekday name according to the current locale.	Wed
%A	The full weekday name according to the current locale.	Wednesday
%b	The abbreviated month name according to the current locale.	Jan
%B	The full month name according to the current locale.	January
%c	The preferred date and time representation for the current locale.	
%C	The century number (year/100) as a 2-digit integer.	19
%d	The day of the month as a decimal number (range 01 to 31).	07
%D	Equivalent to %m/%d/%y. (This is the USA date format. In many countries %d/%m/%y is the standard date format. Thus, in an international context, both of these formats are ambiguous and must be avoided.)	12/25/04
%e	Like %d, the day of the month as a decimal number, but a leading zero is replaced by a space.	7
%G	The ISO 8601 year with century as a decimal number. The 4-digit year corresponding to the ISO week number (see %V). This has the same format and value as %y, except that if the ISO week number belongs to the previous or next year, that year is used instead.	2008
%g	Like %G, but without century, i.e., with a 2-digit year (00-99).	04
%h	Equivalent to %b.	Jan
%H	The hour as a decimal number using a 24-hour clock (range 00 to 23).	22
%I	The hour as a decimal number using a 12-hour clock (range 01 to 12).	07
%j	The day of the year as a decimal number (range 001 to 366).	008

Table 9. strftime date and time format parameters (continued)

Parameter	Description	Example
%k	The hour (24-hour clock) as a decimal number (range 0 to 23); single digits are preceded by a blank. (See also %H.)	7
%l	The hour (12-hour clock) as a decimal number (range 1 to 12); single digits are preceded by a blank. (See also %I.)	7
%m	The month as a decimal number (range 01 to 12).	04
%M	The minute as a decimal number (range 00 to 59).	58
%n	A newline character.	
%p	Either `AM' or `PM' according to the given time value, or the corresponding strings for the current locale. Noon is treated as `pm' and midnight as `am'.	AM
%P	Like %p but in lowercase: `am' or `pm' or a corresponding string for the current locale.	am
%r	The time in a.m. or p.m. notation. In the POSIX locale this is equivalent to `%l:%M:%S %p'.	07:58:40 am
%R	The time in 24-hour notation (%H:%M). For a version including the seconds, see %T below.	07:58
%s	The number of seconds since the Epoch, i.e., since 1970-01-01 00:00:00 UTC.	1099928130
%S	The second as a decimal number (range 00 to 61). the upper level of the range 61 rather than 59 to allow for the occasional leap second and even more occasional double leap second.	07
%t	A tab character.	
%T	The time in 24-hour notation (%H:%M:%S).	17:58:40
%u	The day of the week as a decimal, range 1 to 7, Monday being 1. See also %w.	3
%U	The week number of the current year as a decimal number, range 00 to 53, starting with the first Sunday as the first day of week 01. See also %V and %W.	26
%V	The ISO 8601:1988 week number of the current year as a decimal number, range 01 to 53, where week 1 is the first week that has at least 4 days in the current year, and with Monday as the first day of the week. See also %U and %W.	26
%w	The day of the week as a decimal, range 0 to 6, Sunday being 0. See also %u.	5
%W	The week number of the current year as a decimal number, range 00 to 53, starting with the first Monday as the first day of week 01.	34
%x	The preferred date representation for the current locale without the time.	
%X	The preferred time representation for the current locale without the date.	
%y	The year as a decimal number without a century (range 00 to 99).	04
%Y	The year as a decimal number including the century.	2008

Table 9. strftime date and time format parameters (continued)

Parameter	Description	Example
%z	The time-zone as hour offset from GMT. Required to emit RFC822-conformant dates (using "%a, %d %b %Y %H:%M:%S %z").	-2
%Z	The time zone or name or abbreviation.	GMT
%%	A literal '%' character.	%

Notices

This document provides information about copyright, trademarks, terms and conditions for product documentation.

© Copyright IBM Corporation 1993, 2016 / © Copyright HCL Technologies Limited 2016, 2024

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

© (your company name) (year).

Portions of this code are derived from IBM Corp. Sample Programs.

© Copyright IBM Corp. 2024

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM® or other companies. A current list of IBM® trademarks is available on the web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Adobe™, the Adobe™ logo, PostScript™, and the PostScript™ logo are either registered trademarks or trademarks of Adobe™ Systems Incorporated in the United States, and/or other countries.

IT Infrastructure Library™ is a Registered Trade Mark of AXELOS Limited.

Linear Tape-Open™, LTO™, the LTO™ Logo, Ultrium™, and the Ultrium™ logo are trademarks of HP, IBM® Corp. and Quantum in the U.S. and other countries.

Intel™, Intel™ logo, Intel Inside™, Intel Inside™ logo, Intel Centrino™, Intel Centrino™ logo, Celeron™, Intel Xeon™, Intel SpeedStep™, Itanium™, and Pentium™ are trademarks or registered trademarks of Intel™ Corporation or its subsidiaries in the United States and other countries.

Linux™ is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft™, Windows™, Windows NT™, and the Windows™ logo are trademarks of Microsoft™ Corporation in the United States, other countries, or both.



Java™ and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Cell Broadband Engine™ is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

ITIL™ is a Registered Trade Mark of AXELOS Limited.

UNIX™ is a registered trademark of The Open Group in the United States and other countries.

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Index

Special Characters

- @ (atsign) key setup incorrectly on UNIX 84
- <TWS_user>
- unable to login to conman 92

A

- about this guide xi
- access permission problem for Oracle administration user 109
- access problems for user on DWC 161
- access to Symphony locked by stageman 128
- accesses, multiple, from TDWC, wrong user logged in 161
- accessibility xi
- action on TDWC, list not updated after running 172
- actions return empty tables in TDWC 173
- add, command, validating time zone incorrectly 85
- administration user, Oracle, access permission problem 109
- advanced user rights (incorrect), causing login failure to conman 93
- agent
 - log and trace files 35
 - traces 33
- agent log and trace files
 - twstrace syntax 37
- agent traces
 - modifying 37
 - viewing settings 37
- agents
 - down 82
 - not linking after first JnextPlan on HP-UX 79
 - not linking after repeated switchmgr 191
 - not linking to master domain manager 79
- AIX
 - rmstdlist fails with an exit code of 126 134
- An internal error has occurred - AWSJPL006E 130, 130, 130
- APARs
 - IY50132 85
 - IY50136 26
 - IY60841 88
- application server
 - creating core dump 52
 - does not start after keystore password change 111
 - hanging, creating core dump 52
 - log and trace files 40
 - times out 111
 - trace settings 40
 - troubleshooting 111
- at keyword, validating time zone incorrectly 85
- authentication
 - wrong attempt 159
- authentication problem with UpdateStats 105
- Autotrace
 - stopping while running JnextPlan 88

- available groups list is empty in enter task information window, using LDAP with TDWC 174
- AWKRAA209E 170
- AWKRCE012E
 - connection failure 78
- AWSBCV012E received 97
- AWSBCW037E received 89
- AWSBCW039E received 89
- AWSBIA015I received 85
- AWSBIA019E received 85
- AWSBIA106W received 85
- AWSBIA148W received 85
- AWSDEB003I
 - Writing socket Resource temporarily unavailable 77
- AWSDEC002E received 97
- AWSDEQ008E received 103
- AWSDEQ024E received 92
- AWSSECM003E message received 125
- AWSSEDW001I received 78
- AWSSEDW020E received 78
- AWSITA122E 170
- AWSITA245E
 - agents down 82
- AWSJCO005E
 - CORBA NO_PERMISSION 175
- AWSJCO084E message issued 105
- AWSJCS011E message using planman deploy not enough space 105
- zip file error 105
- AWSJPL017E received 87
- AWSMSP104E message, failed mail send 122
- AWSUI0924E
 - problems with import operation 170
- AWSWUI0331E error returned from custom SQL query with validate command on TDWC 162

B

- background threads continue if browser window closed 166, 174
- backup domain manager
 - agents not linking after repeated switchmgr 191
 - common problems 190
 - Symphony file becomes corrupted 191
 - troubleshooting 185
- batchman
 - fails on a fault-tolerant agent 97
 - in workload service assurance 182
- batchup service fails to start 101
- behind firewall, attribute in fault-tolerant agents 78
- blank page in DWC
 - High availability disaster recovery 167, 175
- bound
 - z/OS shadow job is carried forward indefinitely 107
- browser window closing leaves background threads running 166, 174
- built-in troubleshooting features 19

C

- can be event processor, used to check workstation event enablement 114
- canceling TWS jobs
 - kill command 145
- carry forward z/OS bound shadow job never completes 107
- ccg_basiclogger, CCLog parameter value 30
- ccg_filehandler, CCLog parameter value 30
- ccg_multiproc_filehandler, CCLog parameter value 30
- ccg_pdlogger, CCLog parameter value 30
- CCLog
 - causing jobs to fail on fault-tolerant agent 97
 - date and time format 206
 - description 26
 - parameters 26, 31
 - performance 31
 - switching 27
- character corruption 139
- cli
 - command line
 - Windows, problems with 103
- CLI
 - for composer
 - gives server access error 84
 - programs (like composer) do not run 134
- Cloud & Smarter Infrastructure technical training xi
- cluster.exe 139
- clusterupg 139
- collected data
 - data capture utility 49
- command
 - startappserver 128
- command line (Windows), problems executing cli commands 103
- commands
 - xcli 60
- commands and scripts
 - add, validating time zone incorrectly 85
 - cpuname 78
 - deldep 74
 - evtsize, to enlarge Mailbox.msg file 98
 - release 74
 - replace, validating time zone incorrectly 85
 - rmstdlist, fails on AIX with an exit code of 126 134
 - rmstdlist, gives different results 134
 - shutdown 191
 - start, not working with firewall 78
 - stop, not working with firewall 78
 - submit job 74
 - submit schedule 74
- completed jobs or job streams not found 137
- composer
 - CLI gives server access error 84
 - display cpu=@ fails on UNIX 84
 - gives a dependency error with interdependent object definitions 83
 - gives the AWSJOM179E error when deleting a workstation 85, 172
 - troubleshooting 83

- Composer deletion of a workstation fails with the AWSJOM179E error 85, 172
- configuration
 - data capture utility 44
- configuration file, event monitoring, empty or missing 125
- conman
 - login fails on Windows 92
 - troubleshooting 92
- conman command
 - monitoring TWS jobs 145
 - viewing job output 145
- conman sj
 - job log not displaying 96
- connection from DWC
 - not working 153
- connection from TDWC
 - fails if Oracle database in use 156
 - fails when performing any operation 156
 - settings not checked 159
 - test, takes several minutes before failing 155
 - troubleshooting 153
- connectivity. troubleshooting 186
- core dump
 - Job manager 100
- core dump of application server, creating 52
- corrupt CSV report generated from TDWC as seen in MS Excel 163
- corrupt Symphony file recovery 194
- corrupt Symphony recovery 196
 - automated procedure 201
 - command-line command 201
 - corrupt Symphony recovery on FTA 201
 - on
 - fault-tolerant agent 201
 - resetFTA command 201
- corrupted
 - Symphony file 197
- corrupted characters in the command shell 103
- corrupted characters(Windows) 103
- cpuname, command 78
- Create Post Reports 132
 - Create Post Reports 132
 - Create Post Reports 132
- credentials error with
 - IBM Z Workload Scheduler connector z/OS
 - WebSphere
 - Application Server 162
- critical job
 - high risk, has an empty hot list 184
 - is consistently late 183
- critical network timings changing unexpectedly 183
- critical start times
 - inconsistent 183
 - not aligned 182
- cscript error 88
- CSV report generated from TDWC is corrupted in MS Excel 163
- custom SQL query returns the error message AWSWUI0331E with validate command on TDWC 162
- customization
 - CCLog 27

D

- data capture in event of problems 44
- data capture tool
 - used for ffdc 52
- data capture utility 44
 - command syntax 46
 - data collection 49
 - parameters 46
 - prerequisites 45
 - syntax 46
 - tasks 48
 - when to run 44
- database
 - table locked 133
- database jobs
 - supported JDBC drivers 99
 - troubleshooting 99
- database query returns the error message AWSWUI0331E with validate command on TDWC 162
- database transaction log full on Oracle - JnextPlan
- fails
 - 109
- date and time format, CCLog
 - parameter 26
 - reference 206
- date errors in jobs
 - time zone incorrect setting 138
- date inconsistency
 - AIX
 - master domain manager 138
- date inconsistency in job streams
 - time zone incorrect setting 138
- db2 170
- DB2
 - deadlock 108
 - error causing
 - JnextPlan
 - to fail 87
 - full transaction log causing
 - JnextPlan
 - to fail 108
 - table locked 133
 - timeout 108
 - times out 107
 - transaction log full 86
 - troubleshooting 107
- DB2 error message 170
- DB2 to Oracle
 - problems 109
- deadline keyword, validating time zone incorrectly 85
- default tasks not converted into language set in browser, in TDWC 166, 174
- default values in report fields in TDWC after upgrade 164
- deldep, command 74
- delete
 - workstation fails with the AWSJOM179E error 85, 172
- deleted
 - stdlist erroneously 135
- deleting MSSQL records
 - with cascade option 110
- dependencies
 - giving error with interdependent object definitions 83
 - lost when submitting
 - job streams
 - with wildcards 95
 - of
 - Job Scheduler
 - instance not updated 136
- deploy (D) flag not set after ResetPlan
- command used 125
- deploying event rule 100
- developer tools
 - freeze panels 169
- diagnostic tools 21
- directories
 - pobox, storing messages 74
- disk filling up
 - EDWA 127
- disk usage problems
 - EDWA 127
- display cpu=@ fails, on UNIX 84
- distributed engine
 - responsiveness of TDWC decreasing with 160
 - when running production details reports, might overload TDWC 160
- domain manager
 - agents not linking after repeated switchmgr 191
 - cannot link to
 - fault-tolerant agent 78
 - mailman unlinking from
 - fault-tolerant agents 98
 - not connecting to
 - fault-tolerant agent
 - dynamic domain manager 77
 - not connecting to
 - fault-tolerant agent
 - using SSL 76
 - not shut down on UNIX after
 - switchmgr 191
 - recovering corrupt Symphony file 194
 - running as standalone 73
 - standalone running 73
 - start and stop, commands not working 78
 - Symphony file on backup becomes corrupted 191
 - UNIX, system processes not killed after
 - switchmgr 191
 - workstation not linking after
 - JnextPlan 89
- domain name
 - not included for mail sender 122
- duplicate user ID invalidating session on TDWC 172
- duration might be calculated incorrectly 106
- DWC connection status blank 155
- DWC engine version empty 155
- dynamic agent 148
 - event condition does not generate any action 100
 - job status continually in running state 81
 - not found from console 81

- not running submitted job 81
- server connection 99
- troubleshooting 99, 99
- dynamic workload broker
 - cached jobs
 - increasing 142
 - concurrent threads on server
 - configuring 142
 - job archiving
 - configuring 142
 - job throughput
 - increasing 142
- Dynamic Workload Console 167, 175
 - 9.x login or graphical view pages do not display 178
 - access error launching task from bookmark 166, 174
 - accessibility xi
 - actions return empty tables 173
 - available groups list is empty in enter task information window, using LDAP 174
 - CSV report corrupted in MS Excel 163
 - db2 170
 - default tasks not converted into language set in browser 166, 174
 - engine connection
 - fails if Oracle database in use 156
 - fails when performing any operation 156
 - not working 153
 - settings not checked 159
 - test, taking several minutes before failing 155
 - troubleshooting 153
 - insufficient space when running production details reports 164
 - internet explorer problem 168, 168
 - list not updated after running action 172
 - other problems 171
 - performance problems 160
 - plan view 169
 - problems with browsers 165
 - problems with reports 162
 - processing threads continue in background if browser window closed 166, 174
 - production details reports, running, might overload distributed engine 160
 - report fields show default values after upgrade 164
 - reports not displayed when third party toolbar in use 163
 - responsiveness decreasing with distributed engine 160
 - run period exceeds the historical data time frame 165
 - scheduled time range not respected on z/OS engines 178
 - session has become invalid message received 172
 - Simplified Chinese characters missing or corrupted when using Google Chrome or Apple Safari 169
 - SQL query returns the error message AWSWUI0331E with validate command 162
 - troubleshooting 13, 153
 - unexpected login request when using single sign-on 161
 - unresponsive script warning with Firefox browser when opening Workload Designer 167
 - user access problems 161
 - Workload Designer

- 170
- wrong user logged in when making multiple accesses 161
- WSWUI0331E error when running reports on an Oracle database 163
- dynamic workload scheduling
 - log files 34
 - trace files 34

E

- education xi
- EIF event, checking it has been sent 120
- email send action fails
 - for event rule 122
- empty or missing event monitoring configuration file 125
- empty panel in DWC
 - HADR 167, 175
- empty tables returned in TDWC from actions 173
- enEventDrivenWorkloadAutomation, used to check event management enablement 114
- engine connection from DWC
 - not working 153
- engine connection from TDWC
 - fails if Oracle database in use 156
 - fails when performing any operation 156
 - settings not checked 159
 - test, takes several minutes before failing 155
 - troubleshooting 153
- enter task information window, has available groups list empty, using LDAP with TDWC 174
- erroneously deleted
 - stdlist 135
- error AWSJOM179E Composer deletion of a workstation fails 85, 172
- error given with interdependent object definitions 83
- error
 - java.lang.UnsupportedOperationException 179
- error launching tasks from browser 166, 174
- error messages MakePlan 129, 129, 130, 130, 130
- error messages SwitchPlan 131, 131, 131
- error opening IPC, error message 78
- error opening zip file
 - in planman deploy 105
- error using add task to bookmark, in TDWC 166, 174
- event
 - lost 123, 123
 - no match to event condition 122
- event condition on dynamic agent does not generate any action 100, 100
- event counter
 - troubleshooting 185
- event management
 - check if enabled 114
 - checking
 - EIF event has been sent 120
 - FileMonitorPlugIn event has been received 119
 - monconf directory 117
 - that SSM Agent is running 119
 - TWSObjectMonitorPlugIn event has been received 120
 - deploy (D) flag not set after ResetPlan command used 125
 - LogMessageWritten not triggered 125
 - monman deploy messages 118

- not processed in correct order 126
- showcpus state values 115
- troubleshooting 113
- using getmon 116
- event monitoring configuration file, empty or missing 125
- event processor
 - commands not working 126
 - not deploying rules after switching 124
- event rule deployment 100
- event rules
 - action not triggered 122, 128
 - do not trigger 113
 - email send action fails 122
 - many, causing planman deploy to fail 127
 - not deployed after switching event processor 124
- evtsize, command to enlarge Mailbox.msg file 98
- Excel showing corrupt CSV report generated from TDWC 163
- exclusive access to Symphony, not possible with stageman 128
- extended agent for MVS does not work 104
- extended agent, troubleshooting 104
- extraneous exception 175

F

- F flag state given for domain manager on UNIX after switchmgr 191
- Failover Cluster Command Interface 139
- fault-tolerant
 - agent, cannot be linked 94
- fault-tolerant agent
 - cannot link to domain manager 78
 - job status not updated
 - master domain manager 98
 - jobs failing in heavy workload conditions 97
 - not connecting to domain manager using SSL 76
 - not linking to master domain manager 79
 - not obeying start and stop commands 78
 - recovering corrupt Symphony file 194
 - running as standalone 73
 - troubleshooting 97
 - unlinking from mailman on domain manager 98
- fault-tolerant agents
 - cannot be linked 94
- ffdc
 - SOE first failure data capture
- file create
 - action not triggered 128
- file delete
 - action not triggered 128
- FileMonitorPlugIn event, checking it has been received 119
- files
 - CPU, full 91
 - localopts, thiscpu option not set correctly 191

- Mailbox.msg corrupt 97
- pobox, full 86
- Sinfonia
 - in recovery of corrupt Symphony file 194
 - to delete after SSL mode change 77
- Symphony
 - becomes corrupted on backup domain manager 191
 - to delete after SSL mode change 77
- TWSCCLog.properties 26
- filling percentage of the mailboxes
 - EDWA 127
- final status, jobs or job streams in, not found 137
- Firefox browser giving unresponsive script warning when using the TDWC Workload Designer 167
- firewall, between domain managers 78
- first failure data capture 52
- fix packs
 - keeping up-to-date 20
 - obtaining 204
- fixes 204
- fomatters.basicFmt.dateTimeFormat, CCLog parameter 30
- fomatters.basicFmt.separator, CCLog parameter 30
- forced logout invalidating session on TDWC 172
- freeze
 - panels 169
- ftbox, troubleshooting 186
- full mailboxes
 - EDWA 127

G

- getmon, used to check workstation monitoring configuration 116
- getting a new socket, error message 78
- graphical view problems 169
- groups available list is empty in enter task information window, using LDAP with TDWC 174

H

- hang
 - Windows 91
- hang of application server, creating core dump 52
- high risk critical job has an empty hot list 184
- host name
 - not recognized 89
- hot list, empty, for high risk critical job 184
- HP-UX
 - agents not linking after first JnextPlan 79

I

- IBM Workload Automation
 - overview 13
- IBM Workload Scheduler agent
 - trace files 33
 - traces 33
- IBM Workload Scheduler agent traces
 - from
 - Dynamic Workload Console 33
- IBM Workload Scheduler
 - for z/OS
 - troubleshooting 13

- IBM Workload Scheduler
 - service for
 - <TWS_user>
 - fails to start 101, 101
- IBM Workload Scheduler: Troubleshooting 104
- impersonation level errors (Windows) 103
- Import a property file bigger than 100MB 179
- import settings 170
- ImportSettingsMaxFileSize 179
- In-Flight Trace
 - facility
 - 54
- inconsistent times in planman showinfo 106
- increase job processing 142
- increase processed jobs 142
- information centers
 - at IBM support website, searching for problem resolution 203
- Informix
 - troubleshooting 110
- Informix deadlocks
 - solving 110
 - with composer 110
- initialization problems 73
- installation
 - log files 31
- interactive job 148
- interactive job is not visible 148
- interactive jobs not interactive using Terminal Services 101
- internet explorer
 - developer tools 169
 - graphic panels not displayed 168
 - object error 168
- Internet, searching for problem resolution 203
- invalid session message received on TDWC 172
- IP address
 - not recognized 89
- IY50132, APAR 85
- IY50136, APAR 26
- IY60841, APAR 88

J

- J flag state given for
 - domain manager
 - on UNIX after switchmgr 191
- Java compiler error
 - using planman deploy 105
- Java exception
 - not enough space
 - using planman deploy 105
- Java exception when performing a query on job streams in plan 179
- Java out of memory when running JnextPlan
 - 87
- Jnextplan
 - fault-tolerant agents
 - cannot be linked 94
- JnextPlan
 - CreatePostReports.cmd 91
 - fails
 - AWSJPL017E 87
 - because database log is full 86
 - cscript 88
 - Java out of memory 87
 - to start 86
 - with DB2 error: nullDSRA0010E: SQL State = 57011, Error Code = -912 87

- fails because database transaction log full 109
- fails because DB2 transaction log full 108
- job remains in "exec" status after 89
- large production plans 91
- Makeplan.cmd
 - rep8.cmd 91
 - Updatestats.cmd 91
- not initializing remote workstation 89
- slow 88
- troubleshooting 86
- workstation not linking 89
- job
 - bound z/OS shadow, is carried forward indefinitely 107
 - remains in "exec" status 89
- job canceling
 - TWS kill command 145
- job failure
 - remote command 145
- job log
 - not displayed 96
- Job manager
 - core dump 100
- job number increase 142
- job output character corruption 139
- job rate increase 142
- Job Scheduler
 - instance
 - dependency not updated 136
 - predecessor not updated 136
- Job Scheduler
 - instance mismatch between Symphony and preproduction plan 104
- job shows as not running 81
- job shows as running 81
- job status
 - problems with, on
 - fault-tolerant agent 98
- job status mapping 145
 - TWS job status 145
- Job stream duration might be calculated incorrectly 106
- job streams
 - completed, not found 137
- job types with advanced options
 - database jobs error 99
 - MSSQL jobs error 99
- jobman and JOBMAN
 - fails on a
 - fault-tolerant agent 97
 - in
 - workload service assurance 182
- jobmon and JOBMON
 - fails on a
 - fault-tolerant agent 97
- jobs
 - completed, not found 137
 - failing on
 - fault-tolerant agent in heavy workload conditions 97
 - interactive, not interactive using Terminal Services 101
 - limit problem 171

- not starting 171
- statistics are not updated daily 136
- with a "rerun" recovery job remains in the "running" state 136

jobs ready

- do not start 171

K

- keystore password changed, WebSphere Application Server does not start 111
- kill command
 - job canceling 145
- knowledge bases, searching for problem resolution 203

L

- L flag state given for domain manager on UNIX after switchmgr 191
- language
 - of log messages 24
- language not being set for default tasks in TDWC 166, 174
- late, consistently, critical job 183
- late, job status, incorrectly reported when time zones not enabled 136
- LDAP
 - account lock 159
 - LDAP, using when available groups list is empty in enter task information window (TDWC) 174
- limit
 - jobs 171
- Limited fault-tolerant agents on IBM i
 - troubleshooting 13
- link problems, troubleshooting 186
- linking
 - agent not found 81
 - no resources available 78, 81
 - problems 74
 - problems with, in dynamic environment 81
 - problems with, on fault-tolerant agent 79
- links
 - cannot be made
 - after SSL mode change 77
 - between fault-tolerant agent and domain manager 78
- links, agents not making after repeated switchmgr 191
- list not updated after running action on TDWC 172
- local parameters not resolving correctly 137
- localopts
 - merge stdlists 26
 - nm port 78
 - SSL port setting 77
 - thiscpu option not set correctly 191
- locked, database table 133
- locklist problem causing JnextPlan 87
- log and trace files
 - agent 35
- log file
 - content 31

- location 31
- log files
 - database, full 86
 - for application server 40
 - location 26
 - question marks found in 135
 - Self-Service Dashboard 34
 - Self-Service Monitoring 34
 - separate from trace files 24
- logging
 - dynamic workload scheduling 34
 - engine log file switching 27
 - file locations 26
 - modify logging level (quick reference) 21
 - overview 21
- login request, unexpected, when using single sign-on 161
- login to conman fails on Windows 92
- LogMessageWritten event not triggered 125
- logout (forced) invalidating session on TDWC 172
- low disk space EDWA 127

M

- Mailbox.msg
 - file, corrupt 97
- mailman
 - fails on a fault-tolerant agent 97
 - initialization phase
 - backup domain manager 77
 - message from, stops event counter 186
 - messages
 - when SSL connection not made 76
 - no incoming message from 98
- mailSenderName option
 - not defined 122
- makeplan 128, 128
- MakePlan problems 129, 129, 130, 130, 130
- master domain manager
 - recovering corrupt Symphony file 194
- Master Domain Manager may stop when trying to retrieve a big job log 112, 112
- MASTERAGENTS workstation
 - moving to folder 149
- memory
 - problem, Java, when running JnextPlan 87
- messages
 - concerning ftbox on full-status agent 186
 - from mailman, stopping event counter 186
 - from writer, stopping event counter 185
 - log, described 24
 - not being tracked 185
 - trace, described 24
 - xcli 68
- migrate database
 - DB2 to Oracle, errors 109
- mirroring disabled
 - query error 179
- mismatch of Job Scheduler instances between Symphony and preproduction plan 104
- missing or empty event monitoring configuration file 125

- modifying
 - agent traces 37
- monconf directory, checking for monitoring configuration availability 117
- monitoring TWS jobs
 - conman command 145
- monman deploy messages 118
- MS Excel showing corrupt CSV report generated from TDWC 163
- MSSQL
 - troubleshooting 110
- MSSQL jobs
 - supported JDBC drivers 99
 - troubleshooting 99
- multiple accesses from TDWC, wrong user logged in 161
- mustgather tool 44

N

- netman
 - two instances listening on the same port 78
- network
 - common problems 76
 - link problems 74
 - problems, common 76
 - recovery 73
 - troubleshooting 73
- network timings, critical, changing unexpectedly 183
- nm port, localopts parameter 78
- no results from query 180
- nullDSRA0010E error causing JnextPlan to fail 87

O

- objects in folders, cannot retrieve 96
- online product documentation
 - searching for problem resolution 203
- Onnnn.hhmm files
 - deleting 138
- opening Workload Designer from graphical view with Firefox 168
- Oracle
 - transaction log full 86
 - troubleshooting 109
- Oracle database giving WSWUI0331E error when running reports in TDWC 163
- order of events not respected 126
- organization parameter, in CCLog 26

P

- panels not displayed with internet explorer 168
- parameters, local, not resolving correctly 137
- parms, not resolving local parameters correctly 137
- performance
 - CCLog 31
 - logging 31
 - troubleshooting for TDWC 160
- performance - troubleshooting 72
- permissions problem for Oracle administration user 109
- plan monitor
 - in workload service assurance 181
- planman
 - deploy, failing with many rules 127
- planman deploy
 - fails with Java compiler error 105

- insufficient space error 105
- planman showinfo displays inconsistent times 106
- planner
 - in workload service assurance 181
 - troubleshooting 104
- plug-in deploy
 - fails with Java compiler error 105
- pobox
 - directory, storing messages 74
 - file, full 86
 - high CPU usage 91
- post-uninstallation clean up 138
- predecessor to Job Scheduler instance
 - not updated 136
- preproduction plan has different Job Scheduler instances than Symphony file 104
- problems MakePlan 129, 129, 130, 130, 130
- problems SwitchPlan 131, 131, 131
- problems, other, on TDWC 171
- processing threads continue in background if browser window closed 166, 174
- product
 - parameter, in CCLog 26
- production details reports run in TDWC, insufficient space to complete 164
- production details reports, running from TDWC, might overload distributed engine 160
- prompts, duplicate numbers 94

Q

- query does not complete 180
- question marks found in the stdlist 135

R

- record deletion in MSSQL fails 110
- recover
 - corrupt Symphony 197
- recovering
 - corrupt Symphony file 194
 - network failures 73
- recovering a corrupt Symphony file 196
- release, command 74
- remote command 145
- remote command job
 - connection failure 78
- remote workstation not initializing after JnextPlan 89
- replace, command, validating time zone incorrectly 85
- replay protocol, after switchmgr 185
- report fields show default values in TDWC after upgrade 164
- report problems, on DWC 165
- report problems, on TDWC 162
- reports getting WSWUI0331E error when running on an Oracle database in TDWC 163
- reports not displayed in TDWC when third party toolbar in use 163
- reports, not including completed jobs or job streams 137
- rerun recovery job, original job remains in the "running" state 136
- ResetPlan command

- not setting deploy (D) flag 125
- responsiveness of TDWC decreasing with distributed engine 160
- resubmitJobName optman property update failure 149
- rights problem for Oracle administration user 109
- rmstdlist command
 - fails on AIX with an exit code of 126 134
 - gives different results 134
- rules (event)
 - do not trigger 113
- rules deploy
 - insufficient space error 105
- run time
 - log files 31
- runmsgno, reset 94
- running but not visible on dynamic agent 148
- running state, original job remains in, with a "rerun" recovery job 136

S

- scratch option
 - in planman deploy insufficient space 105
- Self-Service Catalog or Self-Service Dashboard from mobile device, characters missing or corrupted 169
- Self-Service Dashboard
 - log files 34
- Self-Service Monitoring
 - log files 34
- separator parameter, in CCLog 26
- service pack (Windows), problems after upgrading with 103
- services (Windows)
 - fail to start 101, 101
 - Tivoli Token Service, causing login failure to conman 92
- session has become invalid message received on TDWC 172
- setting trace levels for application server 40
- settings
 - agent traces 37
- shadow bound z/OS job is carried forward indefinitely 107
- showinfo (planman) displays 106
- Shutdown_clu.cmd 139
- shutdown, command 191
- Simplified Chinese character set, not fully supported by Google Chrome and Apple Safari 169
- Sinfonia, file
 - in recovery of corrupt Symphony file 194
 - to delete after SSL mode change 77
- single sign-on, unexpected login request received 161
- Software Support
 - receiving weekly updates 205
- space insufficient when running production details reports in TDWC 164
- Special characters
 - corruption 139
- SQL query returns the error message AWSWUI0331E with validate command on TDWC 162
- SSL
 - no connection between fault-tolerant agent and its domain manager

- 76
 - port setting in localopts 77
 - workstation cannot link after changing mode 77
- SSM Agent, checking for event processing 119
- stageman, unable to get exclusive access to Symphony 128
- standalone mode for fault-tolerant agents and domain managers 73
- standard list file 26
- start times, critical
 - inconsistent 183
 - not aligned 182
- start-of-plan-period
 - problems 73
- start, command not working with firewall 78
- startappserver
 - command 128
- Startup_clu.cmd 139
- statistics are not updated daily 136
- status of TWS processes
 - EDWA 127
- stdlist
 - erroneously deleted 135
 - restricting access to 24
- stdlist, question marks found in 135
- stop, command
 - not working with firewall 78
- stopeventprocessor, not working 126
- strftime (date and time format) 206
- submit
 - job streams with wildcards loses dependencies 95
 - submit job, command 74
 - submit schedule, command 74
 - support website, searching to find software problem resolution 203
- swap space problem
 - using planman deploy 105
- switcheventprocessor, not working 126
- switching logs in CCLog 27
- switchmgr
 - used repeatedly 191
- switchmgr command, UNIX system processes not being killed after 191
- SwitchPlan problems 131, 131, 131
- Symphony
 - Symphony download timeout 77
- Symphony corruption 196
- symphony file
 - master domain manager 98
- Symphony file
 - becomes corrupted on backup domain manager 191
 - corrupt 194
 - corrupted 197
 - different Job Scheduler instances than preproduction plan 104
 - managing concurrent access to 128
 - recovery 194
 - to delete after SSL mode change 77
 - troubleshooting 194
- Symphony recovery 196
- system resources scan
 - notify scan 100
 - troubleshooting 100
- systemout exception 175

T

- table does not render correctly 180
- table, database, locked 133
- task information entry window, has available groups list empty, using LDAP with TDWC 174
- TDWC test connection failure 156, 158
- TDWC Workload Designer does not show on foreground with Firefox browser 168
- technical training xi
- Terminal Services, interactive jobs not interactive when using 101
- test connection to engine from TDWC takes several minutes before failing 155
- The database is already locked - AWSJPL018E 130, 130, 130
- the specified run period exceeds the historical data time frame 165
- thiscpu option not set correctly in localopts file 191
- threads continue in background if browser window closed 166, 174
- time and date format, CCLog
 - parameter 26
 - reference 206
- time errors in jobs
 - time zone incorrect setting 138
- time inconsistency
 - AIX
 - master domain manager 138
- time inconsistency in job streams
 - time zone incorrect setting 138
- time zone 106
 - not enabled, causing time-related status problems 136
 - not recognized by WebSphere Application Server 106
 - not validated correctly by composer 85
- time-related calculations 106
- time-related status
 - incorrect when time zone not enabled 136
- timeout of session on TDWC 172
- timeout on application server 111
- timeout on DB2 107
- times inconsistent in planman showinfo 106
- timings, network, critical, changing unexpectedly 183
- Tivoli Token Service
 - causing login failure to conman 92
 - fails to start 101
- Tivoli Workload Dynamic Broker
 - troubleshooting 13
- too many concurrent jobs 100
- toolbar, third party, stopping display of reports in TDWC 163
- tools
 - CCLog 26
- tools, for troubleshooting 21
- TOS errors, on fault-tolerant agent 97
- trace and log files
 - agent 35
- trace and log files agent
 - agent twstrace syntax 37
- trace file
 - activation 32
- trace files
 - for application server 40
 - IBM Workload Scheduler agent

- 33, 33
- question marks found in 135
- separate from log files 24
- trace information
 - gathering 44
- trace levels
 - application server
 - setting 40
- tracing
 - dynamic workload scheduling 34
 - modify logging level (quick reference) 21
 - overview 21
- tracing facility 54
- training
 - technical xi
- transaction log for database is full 86
- transaction log for the database is full message received from DB2, causing JnextPlan
 - to fail 108
- troubleshooting 128
 - application server 111
 - built-in features 19
 - common problems 83
 - composer 83
 - concurrent accesses to the Symphony file 128
 - conman 92
 - database jobs 99
 - DB2 107
 - DWC
 - database 170
 - graphical view 169
 - problems with browsers 165
 - user access problems 161
 - dynamic agent
 - 99, 99
 - event management 113
 - extended agents 104
 - fault-tolerant agents 97
 - fault-tolerant switch manager 185
 - finding information in other manuals 13
 - IBM Workload Dynamic Broker 13
 - Informix 110
 - JnextPlan
 - 86
 - Limited fault-tolerant agents on IBM i
 - 13
 - miscellaneous problems 132
 - MSSQL 110
 - MSSQL jobs 99
 - networks 73
 - Oracle
 - 109
 - performance 72
 - planner 104
 - Symphony file corruptions 194
 - system resources scan 100
 - TDWC 153
 - engine connections 153
 - other problems 171
 - performance problems 160
 - problems with reports 162
 - tools 21
 - TWS for z/OS 13
 - Windows 101
 - workload service assurance 181
- TWS job status

- job status mapping 145
- TWS kill command
 - job canceling 145
- tws.loggers.className, CCLog parameter 30
- tws.loggers.msgLogger.level, CCLog parameter 28
- tws.loggers.organization, CCLog parameter 30
- tws.loggers.product, CCLog parameter 31
- tws.loggers.trc<component>.level, CCLog parameter 29
- TWSCCLog.properties
 - customization 28
- TWSCCLog.properties, file 26
- twsHnd.logFile.className, CCLog parameter 30
- TWSObjectMonitorPlugIn event, checking it has been received 120
- twstrace syntax
 - agent log and trace files 37

U

- UNIX
 - display cpu=@ fails 84
 - rmstdlist, fails on AIX with an exit code of 126 134
 - rmstdlist, gives different results 134
 - system processes not killed on ex domain manager after switchmgr 191
- unlinking
 - fault-tolerant agents from mailman on domain manager 98
- unresponsive script warning with Firefox browser when using the TDWC Workload Designer 167
- unsatisfied link error when using extended agent for MVS 104
- until keyword, validating time zone incorrectly 85
- Update Stats 132
 - Update Stats 132
 - Update Stats 132
- UpdateStats, fails if longer than two hours 105
- upgrade
 - Windows, problems after 103
 - your whole environment 20
- upgrade, making report fields show default values in TDWC 164
- user access problems, on DWC 161
- user, wrong, logged in when making multiple accesses from TDWC 161
- users
 - <TWS_user>
 - unable to login to conman 92
 - not authorized to access server, error given by CLI programs 134
 - rights
 - causing login failure to conman 93
 - Windows, problems with 103

V

- validate command returns the error message AWSWUI0331E from TDWC database query 162
- validation error given with interdependent object definitions 83
- variable tables
 - default not accessible 137
- variables

- not resolved after upgrade 137, 140
- viewing
 - agent traces 37
- viewing job output
 - conman command 145
- virtual memory problem
 - using planman deploy 105

W

- wa_pull_info
 - new version of old tws_inst_pull_info tool 45
- web page error with internet explorer 168
- Windows
 - conman login fails 92
 - Terminal Services, interactive jobs not interactive when using 101
 - troubleshooting 101
 - upgrading, problems after 103
 - user rights, problems with 103
- workload
 - fault-tolerant agent causing jobs to fail 97
- workload designer does not open 162
- workload service assurance
 - critical job
 - is consistently late 183
 - critical network timings changing unexpectedly 183
 - critical start times
 - inconsistent 183
 - not aligned 182
 - high risk critical job has an empty hot list 184
 - troubleshooting 181
 - use of batchman 182
 - use of jobman 182
 - use of plan monitor 181
 - use of planner 181
- workstation ID
 - retrieving 96
- workstation not linking after JnextPlan 89
- workstation set to ignore, cannot find objects 96
- workstations
 - not linking after JnextPlan 89
 - not shut down on UNIX after switchmgr 191
 - remote, not initializing after JnextPlan 89
- writer
 - message from, stops event counter 185
 - messages
 - when SSL connection not made 76
- Writing socket messages
 - Resource temporarily unavailable 77
- wrong duration 106
- wrong schedtime in jobs
 - time zone incorrect setting 138
- wrong start time in jobs
 - time zone incorrect setting 138
- WSWUI0331E error when running reports on an Oracle database in TDWC 163

X

- xcli 54
 - messages 68
- xcli command 60
- xtrace.ini
 - description 56
 - modify 57
 - syntax 57

Z

- z/OS bound shadow job is carried forward indefinitely 107