TIDAL

Tidal Workload Automation
Getting Started Guide

Version 6.3.3

First Published: January 2018

tidalautomation.com

Contents

Preface

The document describes the basics of getting started with *Tidal Workload Automation* 6.3.3 after you've installed it including starting and stopping the components and using the basic features. If you are a new user, see these documents for useful information:

- *Tidal Workload Automation Installation and Configuration*–Provides complete installation and configuration instructions.

- *Tidal Workload Automation Compatibility Guide*–Provides platform, browser, and other important version support information.

- *Tidal Workload Automation User Guide*–Provides usage details.

Organization

This guide includes the following chapters:

Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, submitting a service request, and gathering additional information, see What's New in Tidal Product Documentation at:

http://docs.tidalautomation.com/rss

Subscribe to What's New in Tidal Product Documentation, which lists all new and revised Tidal technical documentation, as an RSS feed and deliver content directly to your desktop using a reader application. The RSS feeds are a free service.

Related Documentation

For a list of all Tidal Workload Automation guides, see the *Tidal Workload Automation Documentation Overview* of your release on tidalautomation.com at:

http://docs.tidalautomation.com/

Note: We sometimes update the documentation after original publication. Therefore, you should also review the documentation on tidalautomation.com for any updates.

Document Change History

The table below provides the revision history for the TWA Release Notes.

Table 1-1

Version Number	Issue Date	Reason for Change
6.2.1 (SP2)	May 2015	New Guide for CWA to help customers learn the essentials of working with CWA.
6.3	July 2016	Renamed guide from "Tidal Enterprise Scheduler Essential Knowledge Guide" to "Cisco Workload Automation Getting Started Guide". Rebranded Cisco Tidal Enterprise Scheduler (TES) to Cisco Workload Automation (CWA).
6.3.3	January 2018	Rebranded "Cisco Workload Automation (CWA)" to "Tidal Workload Automation (TWA)".

1
Tidal Workload Automation Overview

The purpose of this guide is to help you get started using Tidal Workload Automation (TWA) after you've installed it using the *Tidal Workload Automation Installation and Configuration.* This guide provides information about getting help if you have problems, how to help yourself, the basics of starting/stopping components, and some of the basic tasks you perform when you start using TWA to schedule and monitor jobs.

Note: Before you start to work through this document, use the *Tidal Workload Automation Installation and Configuration* to install TWA.

This chapter gives you a quick overview of the TWA components:

- ■ About Tidal Workload Automation
- ■ Tidal Workload Automation Components

About Tidal Workload Automation

Tidal Workload Automation is an automation platform for cross-application and cross-platform enterprise workloads, batch job scheduling, and data and application integration. Tidal Workload Automation can easily configure and run scheduled batch workloads and event-based business processes, integrate the commercial and custom applications these processes use, and determine which tasks to run, as well as where and when to run them, without the need to manage scripts or customize existing tools. Additionally, Tidal Workload Automation provides a single view and point of control over business processes and the jobs they comprise.

Based on a highly scalable multi-tier Java architecture, Tidal Workload Automation can scale to deliver the most demanding SLAs, because it is capable of handling hundreds of concurrent users, managing thousands of connections, and running hundreds of thousands of jobs a day. Tidal Workload Automation offers a distributed management architecture that works across many popular OS platforms and integrates with major enterprise applications and technologies.

Tidal Workload Automation can also manage complex application integrations that connect through web services and enterprise service bus (ESB) protocols. Tidal Workload Automation can be implemented quickly, allowing users to maximize time to value, build momentum across their organizations, and quickly simplify their entire workload processing environment.

Tidal Workload Automation Components

As described in the *Tidal Workload Automation Installation and Configuration*, a basic TWA system is composed of a number of required, optional, and 3rd-party components as shown here:

Figure 1   Basic TWA Components
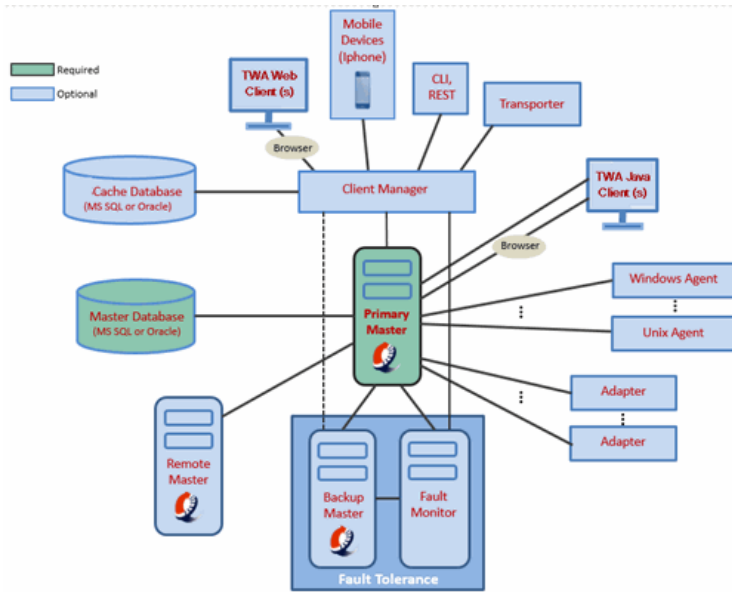
For the purposes of this document which is focused on giving you the basic information you need to get started and be successful with TWA, the components you will be working with are the TWA Web or Java Client, the Primary Master, the Client Manager, and the Agents and Adapters. A brief description of these components is described below. See the *Tidal Workload Automation Installation and Configuration* for descriptions of all components.

Master (or Primary Master)

The Master is the primary TWA component that conducts all scheduling tasks. You can have one or more Masters in your network. The Master can be installed on either the Windows platform or the Unix platform. The basic functionality of TWA remains the same regardless of the platform of the Master.

Each Master computer must supply a unique port number to which Client Managers connect. This port number ensures that communication between the Client Manager and Master is clear.

Client Manager

The Client Manager allows the Master to achieve higher performance and scalability. The purpose of the Client Manager is to service requests from user initiated activities, such as through the TWA Web Client, Tidal Workload Automation Transporter and from other external sources that utilize the Command Line Interface (CLI) or published TWA Web services. The Client Manager allows the TWA Master to focus more capacity on core scheduling needs related to job execution and job compilations, while the Client Manager addresses demands from activities such as users viewing/configuring scheduling data and output. The Client Manager constantly syncs the information from the Master database into its own TESCache database that it then uses to provide all TWA Web Client users with current information. Multiple Client Managers connected to the same Master can be deployed to address additional performance needs.

The Client Manager is required if you want to use the TWA Web Client, the Transporter, the Command Line Interface (CLI), REST, or mobile applications. The Client Manager is not required if you are only using the Java Client.

TWA Web Client

The TWA Web Client is the main user interface for managing TWA jobs, scheduling, connections, configuration, and so on. The TWA Web Client connects to the Client Manager using a browser.

You can use both the TWA Web Client and the TWA Java Client with the same Master.

TWA Java Client

The TWA Java Client is a standalone application that provides the same user interface as the TWA Web Client to manage TWA jobs, scheduling, connections, configuration, and so on. However, the Java Client connects directly to the TWA Master and not through the Client Manager. You can run the Java Client as an application, or you can launch it using a browser.

The TWA Java Client is typically more responsive than the TWA Web Client because it is connected directly to the Master. However, this increases the RAM used by the Master process as the Master is now doing the work instead offloading it to the Client Manager.

You can use both the TWA Java Client and the TWA Web Client with the same Master. Note that an environment exclusively using the TWA Java Client will not be able to utilize features that require the Client Manager like the Transporter, Command Line Interface, and the WebService API.

Agents

An agent is a separate installation of TWA that runs jobs on behalf of the Master. Agents help you to automate the execution of jobs that you know need to be performed on a regular basis, Offloading jobs to agents frees the Master for intensive scheduling tasks such as production compiles. Agents exist for various platforms. The Windows Agent and the Unix Agent are the two most commonly used agents.

Each agent can connect to a Master by specifying the Master-to-agent communication port and the Master-to-agent file transfer port numbers.

Adapters

TWA provides adapters for many software products to enable connectivity to and access by TWA. The Master has an Adapter Host it uses to manage the adapters and is the interface that the adapters use to connect to the Master. Adapters are provided with the base TWA installation but must be licensed. In some cases, adapters need to be installed or configured.

2

Starting and Stopping TWA Components

This chapter contains the procedures for starting and stopping TWA components:

- Starting and Stopping TWA
- Starting and Stopping the Master
- Starting and Stopping the Client Manager
- Starting and Stopping an Agent
- Running the TWA Web Client
- Running the TWA Java Client
- Controlling the Fault Monitor

Starting and Stopping TWA

Several TWA components must be up and running and in a particular order before you can use TWA. Generally, the components must be started and stopping in this order.

To start your entire TWA system:

1. If using Fault Tolerance, start the Fault Monitor.

2. Start the TWA Master. If using Fault Tolerance, start the Master you wish to run in Active mode (commonly called the Primary Master).

3. If using Fault Tolerance, start the Master you want to run in Standby mode (commonly called the Backup Master.

4. Start the Client Manager.

5. Start the TWA Web or Java Client.

To stop your entire system:

1. If using Fault Tolerance, stop the Master running in Standby mode (commonly called the Backup Master).

2. Stop the TWA Master. If using Fault Tolerance, stop the Primary Master/Active Master.

3. If using Fault Tolerance, stop the Fault Monitor.

Starting and Stopping the Master

The TWA Master must be started before you can use the Tidal Workload Automation.

Note: See Troubleshooting TWA when the Master or Client Manager Will Not Start, page 6 if you have trouble starting the Master.

Starting and Stopping the Windows Master

To start or stop the TWA Master:

1. From the Windows Start menu on the master machine, select **All Programs > Tidal Workload Automation > Scheduler > Service Control Manager** to display the TWA Service Manager.

2. From the Service list, select **Scheduler Master**. The master status displays at the bottom of the dialog box.

3. Click **Start** to start the Master if it is not running.

   –or–

   Click **Stop** to stop the Master.

Starting and Stopping the UNIX Master

To start and stop the master on UNIX:

1. Open a comand prompt window.

2. Enter one of these commands:

   – Start the Master with this command:

   tesm start

   – Stop the Master with this command:

   tesm stop

   – Get the Master status with this command:

   tesm status

Starting and Stopping the Client Manager

Note: See Troubleshooting TWA when the Master or Client Manager Will Not Start, page 6 if you have trouble starting the Client Manager.

Starting and Stopping the Windows Client Manager

To start Client Manager on Windows:

1. From the Windows Start menu on the master machine, choose **All Programs > Tidal Workload Automation > Client Manager > TWA Service Manager** to display the TWA **Service Manager**.

2. From the Service list, choose **Client Manager**. The Client Manager status displays at the bottom of the dialog box.

3. Click **Start** to start the Client Manager.

To stop Client Manager on Windows:

1. From the Windows Start menu on the master machine, choose **All Programs > Tidal Workload Automation > Client Manager > TWA Service Manager** to display the TWA **Service Manager**.

2. From the Service list, select **Client Manager**. The Client Manager status displays at the bottom of the dialog box.

3. Click **Stop** to stop the Client Manager.

Starting and Stopping the Unix Client Manager

To start Client Manager on Unix:

1. Open a command prompt window.

2. Enter:

   ./cm start

   Note: ./ may not be required on some systems. Consult your system administrator to determine how the commands should be used.

3. Press **Enter**.

To stop Client Manager on Unix:

1. Open a command prompt window.

2. Enter:

   ./cm stop

3. Press **Enter**.

Starting and Stopping an Agent

Starting and Stopping an Agent on Windows

To start or stop an agent on Windows:

1. From the Windows Start menu, choose **All Programs > Tidal Workload Automation > Service Control Manager** to display the TWA Service Manager.

2. From the Services list, choose the name of the agent.

3. Click **Start** to start the agent.

   -or-

   Click **Stop** to stop the agent.

Starting and Stopping an Agent on Unix

To start or stop an agent on Unix:

1. Open a command prompt window.

2. Enter:

   ./tagent <agent name> start

   -or-

   .**/tagent <agent name> stop**

Note: You should stop all Unix agents before rebooting the Unix system. It is recommended to add the agent stop command to a Unix system shutdown script to be used when restarting a Unix system.

Note: When issuing the tagent start command, verify that you are logged on as the user intended to run the agent.

Running the TWA Web Client

The TWA Web Client connects to the Master via the Client Manager using a browser. The TWA Web Client is the main user interface that you use to manage your TWA system.

Note: Alternatively, you can use the TWA Java Client which provides the same user interface, but connects directly to the Master as an application or using a browser. See Running the TWA Java Client, page 14.

To run the TWA Web Client:
1. Open a TWA -supported web browser and enter the following URL:

    http://*master_hostname*:8080/client

    where *master_hostname* is the host name of the TWA Master.

Running the TWA Java Client

You can run the Tidal Workload Automation Java client as an application on your system, as well as via a web browser. The Java Client has the same user interface and features as the TWA Web Client, but it connects directly to the Master.

Running the Java Client as a System Application

The following prerequisites must be met to run the Java Client as a system application:

- The Java Client Host machine must be in DNS/NIS+ domain.

- The Java Client Host machine must be allowed to connect to port 6215 of Scheduler's host.

- The Master's master.props must have valid LDAP/AD configuration.

To run the Java Client as an application on your system:
1. Launch the Java Client that you have installed. The Login screen displays.

2. Enter the following details:

    - **Server**– The scheduler's hostname

    - **User**– AD/LDAP user name

    - **Password**–AD/LDAP password

3. Click **Connect**.

    The Java Client application window displays.

Note: The logs and help folders are created in your *temp* folder. You can view them by clicking **View > Client Logs**.

Note: Startup scripts of the Java client can be optionally modified to add jvm arguments for optimal performance.

Running the Java Client Via a Web Browser

The following prerequisites must be met to run the Java Client via a web browser:

- By default, the Master runs a web server at port 8080. The Java Client host must be allowed to access a configured port on the Master's host machine.

- On Windows, only Internet Explorer 64-bit (c:\Program Files\Internet Explorer\iexplorer.exe) is capable of running 64-bit Java 8. Only 64-bit Java 8 will support 8GB memory requirements.

- For all operating systems and browsers, you must enable Java content in the Java Control Panel.

  Note: Confirm that browser's security settings allow running Java applets.

To run the Java Client via a web browser:

1. Open a TWA -supported web browser and enter the following URL:

   http://_master_hostname_:8080/tesclient

   where _master_hostname_ is the host name of the TWA Master.

2. Click **Launch Enterprise Scheduler**.

3. Click **Run** to allow execution of the Java Client.

   The Java Client is launched.

If the version of Java Client does not match what has been installed on the master, remove all temporary Java files using options available in the **General** tab of the Java Control Panel.

Controlling the Fault Monitor

You can monitor the Fault Monitor from the TWA Web Client or Java Client. If you have installed fault tolerance, then a Fault Monitor tab displays inside the _Master Status_ folder under the _Operations_ folder in the Navigator pane of the TWA Web Client.

Note: To see the Fault Monitor option, you must be properly licensed for fault tolerance and your security policy must include access to the Fault Monitor option.

The Fault Monitor can also be accessed from the command line of the machine it is installed on

Starting the Fault Monitor on Windows

To start the Fault Monitor on Windows:

1. From the Windows Start menu, and choose **All Programs > Tidal Workload Automation > Service Control Manager** to display the TWA Services Manager.

2. From the Service list, choose **SchedulerFaultMon**.

3. Click **Start**.

Stopping the Fault Monitor on Windows

To stop the Fault Monitor on Windows:

1. From the Windows Start menu, and choose **All Programs > Tidal Workload Automation > Service Control Manager** to display the TWA Services Manager.

2. From the Service list, choose **SchedulerFaultMon**.

3. Click **Stop**.

Starting the Fault Monitor on Unix

To start the Fault Monitor on Unix:

1. Open a command prompt window.

2. Enter:

   tesfm start

Stopping the Fault Monitor on Unix

To stop the Fault Monitor on Unix:

1. Open a command prompt window.

2. Enter:

   tesfm stop

3

Initial Configuration Recommendations

This chapter describes some initial configuration tasks that are recommended for most TWA systems when getting started with TWA:

- History Retention, page 1

- Performance Tuning, page 2

History Retention

It is important to decide as soon as possible what history retention levels you want to maintain for your TWA system because:

- History retention values play a LARGE part in TWA performance.

- Purging history files is a very manual process.

So, it is recommended to only keep the amount of history required for live operation. Often users set up exports of certain fields in the database to be offloaded to an archival database in case audits need to be stored in and accessed in the future.

Note: Warning: If you are upgrading an existing system and already have a large amount of history you wish to purge, please open a case with TAC to review strategy for reducing history. There are minor mistakes that can bring your system down. There are methods for avoiding purging too much history at once.

To configure history retention:

1. Open the TWA Web or Java Client.

2. From the Activites menu, choose **System Configuration**.

3. On the Master tab, set these parameters:

   Operator Alert Retention—The number of days to keep the alert history. Affects the Operations-> Alerts panel.

   Trigger History Retention—The number of days to keep event trigger history. Affects the Event definitions, History tab.

   Automatic Daily History Cleanup—The option that checks whether or not to purge history daily. If it is unchecked, you will keep history forever, but the information in the database will still be made with an expiry date in mind, it will just never be removed.

Warning: Since all information has an expiry date, if the **Automatic Daily History Cleanup** is checked, the first purge will wipe out a MASSIVE amount of data all at once, most likely clogging the transaction log and bringing down the system one way or another. The exception to this rule is msglog, which does not have an expiry date, so any changes to the "error" or "audit" values in "operations->logs" should be done incrementally, or massive amounts of msglog table history will be purged all at once, based on the value input during the previous date, filling up the transaction log and bringing your system completely down.

4. On the Logging tab, set these parameters:

Audits–The number of days to keep audits. Affects the audit log of Job Activity jobs, or the Operations->Logs. This corresponds to the msglog table in the database.

Errors–The number of days to keep the messages that show up as the "error" type, in Operations->logs. Also the msglog table.

5. On the Defaults tab, set this parameter:

Job History Retention–The number of days to keep the job history retention for each of the job instances in the job activity. Every job instance that runs is assigned a date on which it is to be purged based off of it's job definition's Options tab -> History Retention value. So if a job ran when that number was 100, it won't expire until 100 days from now, even if the job's definition is edited, since the job instance information has already been written. Mass edits on the backend would compromise the stability of the database, so to avoid that, edits to the job retention values in the job definition do not affect job instances that have already run.

Performance Tuning

See the *Tidal Workload Automation Performance and Tuning Guide* for details about performance and tuning. This guide describes:

- How to optimize TWA performance.

- How to monitor TWA scheduling, connections, the cache, and performance.

- How to tune TWA.

- Best practices for Transporter performance.

4

Logging and Troubleshooting

This chapter describes how to set TWA log levels and provides troubleshooting information:

- TWA Logging

- Troubleshooting TWA when the Master or Client Manager Will Not Start

TWA Logging

TWA supports logging of all actions by all TWA components. This section provides the location of the log files, and describes how to control the logging levels for each:

- Log Files and Location

- Controlling the Logging Levels

Log Files and Location

Here is the location of the log files for TWA components:

Master Logs
<master install dir>\logs

Client Manager process logs:
<CM install dir>\logs\

Client Manager sync logs (for information between the CM and it's cache DB):
<CM install dir>\plugins\tes-6.0\logs

Adapter Host logs:
<master install dir>\logs

Adapter logs:
<master install dir>\services\<GUID of adapter name, found in the adapter button in the client>\log\

Agent logs:
<agent install dir>\<agent instance name>\logs\ES_<name of master>_<date optional>.log

Controlling the Logging Levels

Logging is controlled in different ways depending on the TWA component as described in these sections:

- [Setting the Logging Level for the Master](#)
- [Setting the Logging Level for the Client Manager](#)
- [Setting the Logging Level for Agents](#)
- [Setting the Logging Level for Adapters](#)
- [Setting the Logging Level for Transporter](#)
- [Setting the Logging Level for Fault Monitor](#)

Setting the Logging Level for the Master

To set the logging level for the Master:

1. Open the TWA Web or Java Client.

2. Select **Activities > System Configuration**.

3. Select the **Logging** tab.

4. Choose the logging level from the drop down menus for each of the objects listed below. Logging levels include: None, Severe, Warning, Info, Low Debug, Medium Debug, or High Debug.

   - Schedule Log
   - Client Manager Log
   - Agent Manager Log
   - Compiler Log
   - Job Manager Log
   - Event Manager Log
   - Queue Manager Log
   - Database Log
   - Communication Log

Setting the Logging Level for the Client Manager

To set the logging level for the Client Manager:

1. Add the following to the bottom of the ClientManager/config/tesXXX.dsp:

   #AND clientmgr.props files

   ```
   DspLog=FINEST
   CacheLog=FINER
   RequestLog=FINE
   RpcLog=WARNING
   ClientLog=SEVERE
   MaxLogFiles=1000 #allocates a gig
   ```

2. Restart the Client Manager.

Setting the Logging Level for Agents

To set the logging level for Windows agents:

1. Open the TWA Service Manager on the agent machine.

2. Select the agent in the drop down list.

3. Click on the dots [...].

4. In the Service tab, change the path to include the DEBUG option to read "DEBUG=HIGH".

5. Click OK to save changes.

6. Restart the agent.

To set the logging level for Unix agents:

1. Go to the agent machine and go to the directory the agent is installed in.

2. Go to "Bin/" and edit "tagent.ini".

3. Under "[config]" put "ovb=tidaldebug".

4. Restart the agent.

Setting the Logging Level for Adapters

To set the logging level for adapters:

1. Open the TWA Web or Java Client.

2. In the Navigator, click **Connections**.

3. Locate and double-click the Master connection (Type=Master) to open its editor.

4. For Host Environment Log, choose the logging level.

Setting the Logging Level for Transporter

To set the logging level for Transporter:

1. Edit transporter.props in the Transporter's config directory:

```
XPORTER_DEBUG=YES

TransporterLog=FINEST
TransporterUILog=FINEST
TransporterJobLog=FINEST
TransporterDataLog=FINEST.
```

Setting the Logging Level for Fault Monitor

To set the logging level for Fault Monitor:

1. Edit Primary Master and Backup Master master.props file.

```
FaultToleranceLog=INFO
```

2. Edit Fault Monitor master.props file:

```
FaultMonitorLog=INFO
FaultToleranceLog=INFO
```

Troubleshooting TWA when the Master or Client Manager Will Not Start

If your Master or Client Manager will not start and no logs are created in the logs directories, your system probably has a problem loading the Java Virtual Machine (JVM). Java uses a JVM to execute its processes inside, so if the system cannot launch the JVM, then the TWA code cannot run.

To troubleshoot and resolve a Java problem

1. Make sure Java is installed and in the PATH. Run this on a command line:

   ```
   java -version
   ```

2. Find where your java lives:

   Unix:

   ```
   which java
   ```

   Windows:

   ```
   where.exe java (only works on win 2003 and later)
   ```

3. Make sure that your master.props or clientmgr.props files match this location.

   Unix:

   No further Unix troubleshoot should be required. Open a TAC case if you still have no success running the master or Client Manager.

   Windows:

   Make sure regedit matches this location in the following keys:

   a. Launch regedit.exe

   b. HKEY_LOCAL_MACHINE\SOFTWARE\JavaSoft\Java Runtime Environment\1.8 support

   The JavaHome and the RuntimeLib should have their locations verified for accuracy. Make sure that "client" and "server" in the RuntimeLib string are correct. This is often the issue.

If you are trying to install the Java client, edit the tesclient.bat file and verify the paths inside. These can be edited if they are incorrect as discovered in steps 1 and 2.

5

Getting Started with the TWA Client

This chapter describes the how to get started managing your jobs with the TWA Web or Java Client:

- Introducing the TWA Client Interface

- Adding Connections

- Defining Calendars

- Adding Jobs

- Monitoring Jobs

Introducing the TWA Client Interface

There is one primary client interface you use to monitor and work with TWA. You can run the client interface in two different ways:

- **TWA Web Client**–Run using your browser, this Client connects to the Client Manager which then connects to the TWA Master.

- **TWA Java Client**–Can run as a standalone application on your workstation or via a browser. The Java Client connects directly to the TWA Master and does not require a Client Manager.

The two interfaces are essentially identical in their features and operation. The TWA Client interface is shown below:

Host machine name

Main Menus

Operations

Definitions

Administration

Navigator

Selected object details/grid
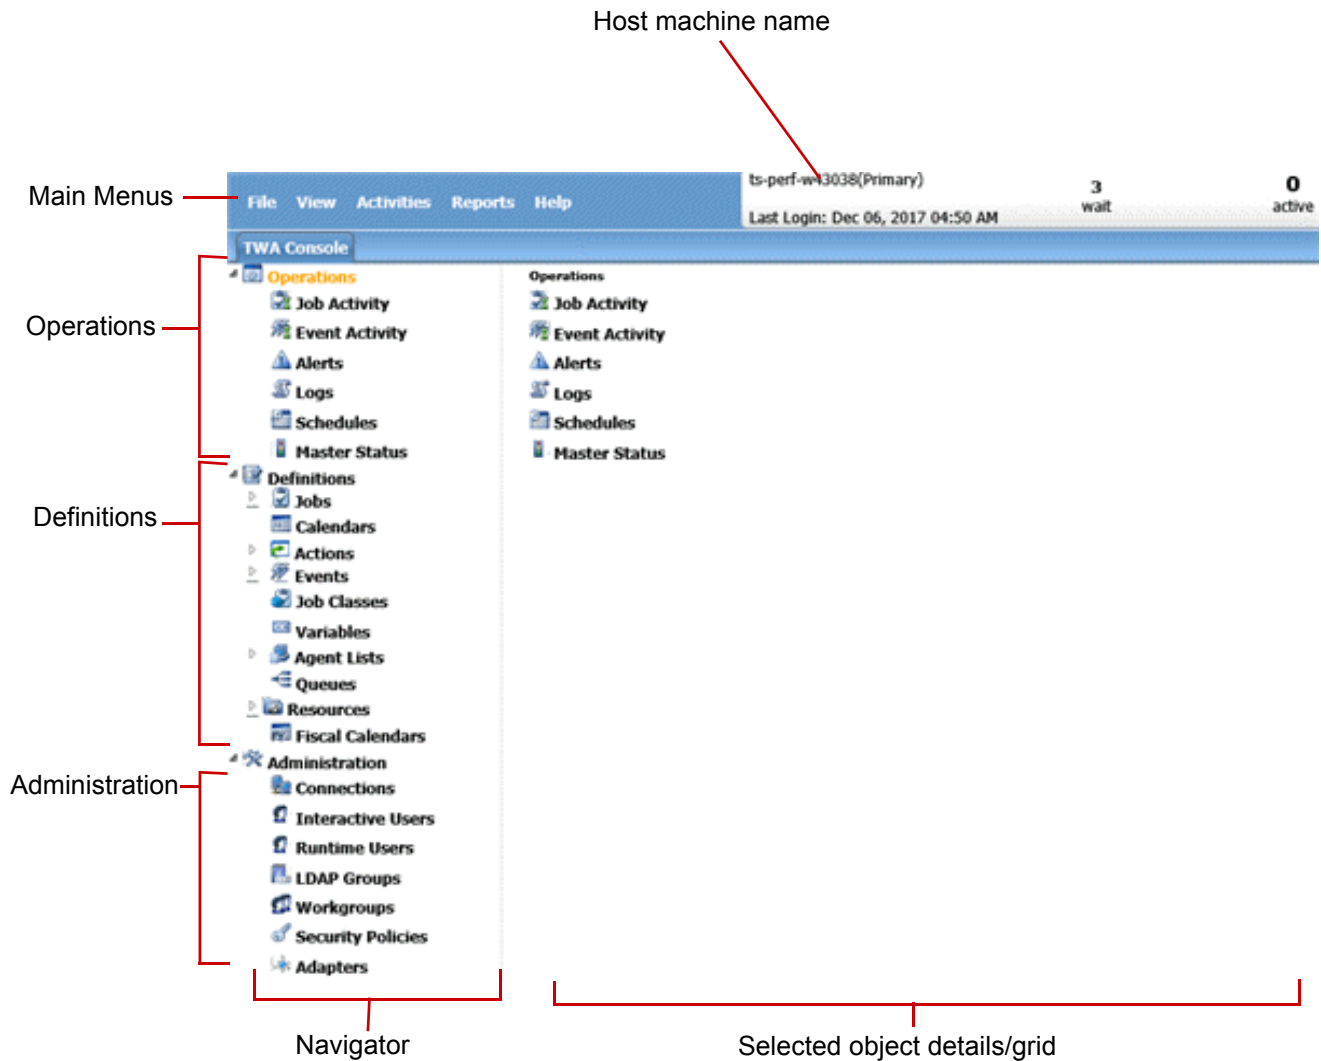
Adding Connections

You use the **Connections** link in the **Navigator** pane to add and manage connections. These connection types can be added:

Note: Agents and adapters (except Email and SSH) require separate licenses, so the options in this list might vary at your installation. See the specific agent or adapter guide for more information about licensing, configuration, and usage of each agent/adapter in the menu.

To add a connection to a new adapter:

1. Open the TWA Web or Java Client.

2. From the **Navigator** pane, choose **Administration > Connections**.

3. Above the **Connections** table, click the **Add** button or choose **Add Connection** from a context menu.

4. Select the connection type: one of the TWA Agents, the Remote Master, or one of the Adapters.

5. In the Connection Definition dialog, provide the General, Connection, and Description information on each tab,

   The tabs and options displayed on each tab vary depending on the connection type. See the *Tidal Workload Automation User Guide* for more information.

6. Click **OK**.

   Here's an example of the Connections panel with a number of defined connections:

Note: This is only a partial display of the available columns. You can change what columns are displayed using the Preferences option for each pane.

Defining Calendars

A job is assigned a calendar which specifies when that job will run. TWA provides many default calendars which you can use as is or edit for your purposes. You can also define new calendars for your TWA system. Here's a partial display of some of the defined calendars in TWA:

This section gives a brief overview of how to define and edit calendars. See the *Tidal Workload Automation User Guide* for details about the calendar interface and how to use it to control your production environment.

All tasks are performed using the TWA Web Client or the Java Client:

- Adding a Calendar

- Editing Calendars

Adding a Calendar

To add a calendar:

1. Open the TWA Web or Java Client.

2. From the **Navigator** pane, select **Definitions > Calendars** to display the **Calendars** pane.

3. Click the **Add** button on the TWA toolbar or right-click and select **Add Calendar** from the context menu.

4. Select the type of calendar you want to add, then follow the instructions for editing a calendar in Editing Calendars, page 12. You can add these calendar types:

   - **Group Calendar**

   - **List Calendar**

   - **Daily Calendar**

   - **Weekly Calendar**

   - **Monthly Calendar with Days**

– **Monthly Calendar with Weeks**

– **Subset Calendar**

You can also create fiscal calendars. See the "Controlling Production" chapter in the *Tidal Workload Automation User Guide* for information about each calendar type and other calendar details.

Editing Calendars

Existing calendars can be customized for your use as described in this section.

Note: Whenever you edit a calendar that has already been assigned to a job, or is being used by another calendar (for example, as a subset) you will be presented with a warning message informing you that the calendar is in use. When you modify a calendar assigned to a job, the **Effective Date** dialog displays providing options for when the changes to the already compiled production schedule should take effect. You can select a date during the current production schedule for the calendar changes to take effect or let the changes occur during a manual or automatic compilation.

To edit a calendar:

1. Open the TWA Web or Java Client.

2. From the **Navigator** pane, select **Definitions > Calendars** to display the **Calendars** pane.

3. Double-click the calendar you want to edit or select the calendar and click the **Edit** button or right-click the calendar and select **Edit Calendar** from the context menu.

4. Edit the name of the calendar in the **Calendar Name** field (up to 60 characters). The name must be unique.

5. In the **Starting On** field, select or type the date that the calendar starts.

   A seven-day period begins on the date you specify. For example, if you select an interval of every two weeks, starting on Wednesday of this week, and include Monday through Friday, you will get Wednesday through Friday of this week, and Monday and Tuesday of next week, which completes the seven day period. The seven-day period is then repeated every other Wednesday. The calendar has no ending date.

6. To make the calendar public, select the **Public** option. Public calendars can be used by all TWA users, within the restrictions of their Security Policy.

7. On the **Forecast** tab, select dates for a calendar:

   a. Use the calendar control bar to select the month for which to add or delete dates.

   b. The outer buttons move the calendar in one year increments.

   c. The inner buttons move the calendar in one month increments.

   d. Double click the dates in the calendar section to select and deselect the dates for your calendar. The date cell text turns red, and the date appears in the **Selected Dates** section.

   Note: Calendar options vary by type; some of the following steps only apply to some calendar types.

8. To remove dates from a calendar:

   a. Click the **Forecast** tab.

   b. Select the date to remove from the **Selected Dates** section.

   c. Click **Delete**. The date moves back to the **Calendar** section and the date cell text turns black.

   Note: This option is not available for subset calendars.

9. Based on the calendar type, edit the calendar options on the tabs as described below:

Calendar Group

   a. Click the **Details** tab.

   b. Optionally, click **Add** to add a calendar to the group:

   – Select an existing calendar.

   – In the **Adjust by... days** field, you select the number of days, if any, by which to shift the list calendar forward (positive numbers) or backward (negative numbers).

   – Select **Include Dates** or **Exclude Dates**.

   – Click **OK**.

   c. Optionally, select a calendar and click **Edit** to edit a calendar group definition.

   d. Optionally, select a calendar and click **Delete** to remove the calendar from the group.

   e. Click **OK**.

List Calendars

   a. Click the **Details** tab.

   b. Navigate the calendar and double-click dates to add or remove them.

Daily Calendars

   a. Click the **Details** tab.

   b. In the **Frequency** drop-down menu, select the frequency of the interval (in days).

   c. In the **Include the following days** section, select the days to include in the seven day period.

Weekly Calendars

   a. Click the **Details** tab.

   b. In the **Frequency** drop-down menu, select the frequency of the interval (in weeks).

   c. In the **Include the following days** section, select the days to include in the seven day period.

Monthly with Days Calendars

   a. Click the **Details** tab.

   b. In the **Months** section, select the months when the job should run.

   c. In the **Days in Month** section, select the dates in the month(s) when the job should run.

Monthly with Weeks Calendars

   a. Click the **Details** tab.

   b. In the **Months** section, select the months when the job should run.

   c. In the **Days in Week** section, select the days of the week when the job should run.

   d. In the **Occurrence** section, select the occurrence for the day(s) selected. For example, **2nd** means the second time the day(s) selected in the **Days of Week** section occurs in the month.

   e. In the **Adjust by** field, select the number of days by which to offset all the dates forecasted. For example, three days after the last Friday of each month. If the last Friday is the 31st, the job will run on the 3rd day of the next month.

  f. Click the **Forecast** tab. The results of the computation from the specification made on the **Details** tab are shown.

Subset Calendars

  a. Click the **Details** tab.

  b. From the **Occurrence** drop-down menu, select the instance of the dates of the calendar that you want to use. For example, if you want to use all the days this calendar includes, select **Every**.

  c. From the **Calendar** field drop-down menu, select the calendar for which this calendar will be a subset. The drop-down menu lists all available calendars.

  d. From the **Time Frame** field, select the exact time period for your calendar to run. For example, you can select the **13th week** or the month of **March**.

  e. If you want to base your list calendar on a fiscal calendar, you must already have a fiscal calendar defined. If you have any fiscal calendars defined, they will be listed and available from the **Based on Fiscal** field drop-down menu.

  f. In the **Adjust by... days** field, you select the number of days, if any, by which to shift the list calendar forward (positive numbers) or backward (negative numbers).

10. Click the **Conditions** tab, and set any conditions you want to place on your calendar.

11. To type a description for the calendar, click the **Description** tab. You can add or edit a description for the calendar (up to 255 characters).

Adding Jobs

This section describes some of the basic tasks you need to perform to add jobs and job groups. Here's an example of what defined jobs and job groups look like in the TWA Web Client. Note that many of the job definition columns are not visible here; you need to scroll to see all fields.

All tasks are performed using the TWA Web Client or the Java Client:

- Adding a Job or Job Group

- Adding Jobs Automatically to the Production Schedule

- Adding Jobs Automatically to the Production Schedule

■ Adding a File Dependency

■ Adding a Variable Dependency

See the *Tidal Workload Automation User Guide* for information about the panels referenced in each task.

Adding a Job or Job Group

You can add a job or group rule and have it added to the production schedule simultaneously. Adding a job to the production schedule is optional. When jobs are added to groups, many properties can be inherited from the parent group.

To add a job or job group:

1. Open the TWA Web or Java Client.

2. From the **Navigator** pane, select **Definitions > Jobs** to display the **Jobs** pane.

3. Click the location where you want to insert your job as follows:

   – To add a job or job group under a job group, right-click the job group under which you want to add your job or job group.

   – To add a job or job group at the root level of the hierarchy, right-click a root-level job or on a white section of the pane.

4. Select **Job** or **Job Group** from the context menu or click the **Add**, then choose **Job** or **Job Group** on the TWA toolbar.

   Note: Depending on what Agents and Adapters you have licensed and configured, you can choose to create a job of a specific adapter type.

5. Depending on your selection, the **Job Definition** or the **Job Group Definition** dialog appears. For more information on the definition dialogs, see the "Jobs Pane Interface" section in the *Tidal Workload Automation User Guide* for details.

Adding Jobs Automatically to the Production Schedule

Jobs with associated calendars are automatically added to the production schedule through the TWA automatic compilation process. Jobs with calendars are only added to the schedule for days that are in the set of days defined for the calendar. No intervention is necessary, but you can customize its operation to tailor compilation to your needs. For more information, the "Defining Jobs" section in the *Tidal Workload Automation User Guide* for details. Another way of automatically adding jobs to the production schedule is through new job actions. For more information about new job actions, see the "Actions and Alerts" section in the *Tidal Workload Automation User Guide* for details.

Adding Jobs Manually

To add a job manually:

1. From the **Navigator** pane, select **Definitions > Jobs** to display the **Jobs** pane.

2. Right-click the job or job group to add to the production schedule and select **Insert Into Schedule** from the context menu.

   Or, select the job or job group to add to the production schedule, and from the **Activities** main menu select **Insert Into Schedule**.

The **Insert Job Into Schedule** dialog displays. The **Job/Group** field should contain the name of the job or group you right-clicked, and the **Date** field should contain the current date. If the job or group you selected has a time window, this will be displayed in the **From** and **Until** fields. Any parameters that were set in the job's or group's definition will be listed in the **Parameters** field. If the job or group has dependencies, you may want to select the **Override job's dependencies** option so that your job or group will enter the schedule without checking for its dependencies.

3. Click **Yes** at the confirmation dialog.

   The job is added to today's production schedule regardless of its calendar dates (if any). If the job is defined to repeat, only one instance of the job will be added to the schedule. Only jobs with the **Unscheduled Allowed** option selected (definition dialog, **Options** tab) can be added in this manner.

Adding a Job or Job Group Dependency

Jobs can depend on the status of other jobs and job groups. For example, you can set **Job B** to run when **Job A** completes normally using a job dependency for **Job B**.

To add a job or job group dependency:

1. From the **Navigator** pane, select **Definitions > Jobs** to display the **Jobs** pane.

2. Double-click the job or job group that you want to edit.

3. Click the **Dependencies** tab.

4. Click **Add** and select **Add Job Dependency** from the drop-down menu.

5. In the **Job/Group** field, click the **Browse** button to open the **Search** dialog to search for the job that this job will depend on. You can also view a drop-down list of jobs by clicking the down-arrow button. If you used the **Browse** button to narrow your search for jobs, the drop-down list will be that job subset.

6. In the **Operator** field, select whether the file dependency **Equals** or **Does not equal** the following status to satisfy the dependency. For example, you can set the job dependency to **Equal** the **Completed Normally** status.

7. In the **Status** field, select the status to use to satisfy the dependency. You can choose between **Active**, **Completed Abnormally**, **Completed Normally**, **Error Occurred**, **Externally Defined** and **Completed.**

   Note: A job group becomes active when any of its associated jobs become active. If all jobs in a job group depend on the job group becoming active, no jobs will launch.

8. If the job repeats during the day, select which instance of the job will trigger the dependency from the **Occurrence** drop-down menu.

   a. Select **First Occurrence** if you want the first instance of the preceding job to match the status criterion.

   b. Select **Last Occurrence** if you want the last instance of the preceding job to match the status criterion.

   c. Select **Match Occurrence** when both jobs repeat during the day, and the dependency should match instance numbers.

   There are two ways to apply the **First/Last/Match** dependency logic: by day instance or group instance.

9. To specify an instance offset, you can do so in the **Offset** field. This field only applies to **First Occurrence** and **Last Occurrence**. When applied to **First Occurrence**, specifies which instance after the first to use in satisfying the dependency. When applied to **Last Occurrence**, specifies which instance from the last.

10. If you want to specify an instance of a job that occurred a certain number of days in the past, go to the **Date Offset** field, and type the number of days in the past for the required job dependency. For example, if **Job A** repeats daily, but you want your job to be dependent on **Job A's** instance 3 days ago, specify **3** in this field.

11. Select **Ignore this dependency if Job not in schedule** if the dependency only applies when the job is part of the production schedule.

12. Click **OK** to add the job dependency.

Note: If your job has more than one dependency (file, job, variable or time), all dependencies must be satisfied for the job to run. It is possible for a dependency's state to change from satisfied to unsatisfied. If this occurs, the job will only run when all dependencies have been satisfied at the same time.

Adding a File Dependency

You can have a job that depends on the status of a file on any system in your network. For example, **Job A** can be defined to run only when file *data.txt* exists in the **c:\data** directory.

To add a file dependency to a job:

1. From the **Navigator** pane, select **Definitions > Jobs** to display the **Jobs** pane.

2. Double-click the job or job group that you want to edit to display the **Job** or **Job Group Definition** dialog.

3. Click the **Dependencies** tab.

4. Click **Add** and select **Add File Dependency** from the dependency context menu to display the **File Dependency Definition** dialog.

5. Type the path and filename for the required file or click the **Browse** button to select a file on the local TWA Web client.

   –or–

   Click the **Variables** button to add **System** or **Job** variables as a file name.

6. Select the agent where the file needs to exist.

7. Select the dependency type for the file from the following options:

   – **Exists** – The file exists at the path and on the agent specified.

   – **Does Not Exist** – The file no longer exists, or is not found at the path or on the agent specified.

   – **Has Changed In DD:HH:MM** – The file dependency is satisfied when the file changed within the specified time in days, hours, and minutes after the job entered the production schedule. For example, if the job entered the schedule at 1:00 PM, the period specified is 6 hours, and the file changed after 7:00 PM (or later), the dependency is met.

   – **Stable For DD:HH:MM** – The file dependency is satisfied when the file size has not changed for the specified time in days, hours, and minutes from the present time. For example, if the file's modified time is 1:00 PM, the period specified is 6 hours, and the job enters the schedule at 3:00 PM, the dependency is satisfied in 4 hours, i.e., 7:00 PM.

   – **Size >=** – The size of the file is greater than or equal to the specified amount in bytes.

   – **Size <=** – The size of the file is less than or equal to the specified amount in bytes.

8. Click **OK**.

Note: If your job has more than one dependency (file, job, variable or time), all dependencies must be satisfied for the job to run. It is possible for a dependency's state to change from satisfied to unsatisfied. If this occurs, the job will only run when all dependencies have been satisfied at the same time.
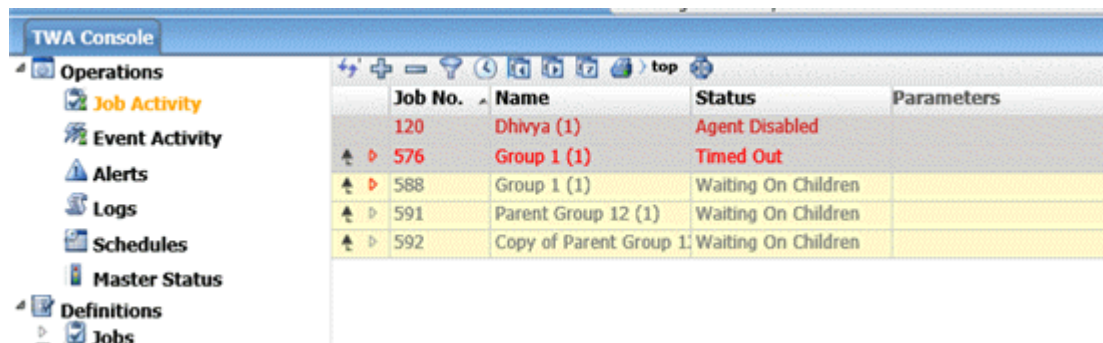
Adding a Variable Dependency

You can make a job depend on the value of a user-defined variable. For example, **Job A** should only run when the variable `RunVar` is equal to ten. For more information on variables, see the "Jobs" chapter in the *Tidal Workload Automation User Guide*.

To add a variable dependency to a job:

1. From the **Navigator** pane, select **Definitions > Jobs** to display the **Jobs** pane.

2. Double-click the job or job group that you want to edit.

3. Click the **Dependencies** tab from the **Job Definition** dialog.

4. Click **Add** and select **Add Variable Dependency** from the drop-down menu to display the **Variable Dependency Definition** dialog.

5. If you are creating an intermaster dependency, select from the **Master** drop-down list a master to manage the variable dependency. Leave the default field selection, if you are not creating an intermaster dependency.

6. In the **Variable Name** field, choose a variable from the drop-down menu that the job or job group will depend on.

7. In the **Operator** field, select from the drop-down list an operator to make the appropriate comparison to the value of the variable.

8. When text strings are used in comparison, "lower" letters of the alphabet are of greater value than "higher" letters. For example, Z > A. If the first letters of the string match, succeeding letters are used for comparison. For example, AZ > AA. The operation works similar to sorting strings in ascending order. Upper versus lower case is not considered (i.e., a=A, b=B, etc.).

9. In the **Variable Value** field, enter the value of the variable required for the dependency to be met. You can also select from a list of **system**, **job**, **user-defined** and **public** variables to which the variable should be compared. For example, suppose you are using a variable dependency to repeat a job a specific number of times, and this amount changes periodically. You can define how many times to repeat the job by changing its variable value instead of changing its job definition.

Monitoring Jobs

This section gives an overview of some of the typical tasks you perform to monitor jobs. Here is an example of the Job Activity panel:

All tasks are performed using the TWA Web Client or the Java Client:

- Adding or Removing Columns in the Job Activity Pane

- Adding a New Job or Group Occurrence to the Schedule

See the *Tidal Workload Automation User Guide* for information about the panels referenced in each task.

Adding or Removing Columns in the Job Activity Pane

To add or remove columns:

1. From the **Navigator** pane, select **Operations > Job Activity** to display the **Job Activity** pane.

2. Select **View>Preferences** from the main menu bar or right-click the **Navigator** pane and select **Preferences** from the context menu to display the **Preferences** dialog.

3. Click the **Columns** tab. By default, all columns are used.

   – To remove columns, clear the check box to the left of the column name.

   – To include columns, click the check box so that a black checkmark appears.

Adding a New Job or Group Occurrence to the Schedule

To add a new job or group occurrence:

1. From the **Navigator** pane, select **Operations > Job Activity** to display the **Job Activity** pane.

2. Right-click the job or job group occurrence to add to the production schedule and select **Insert Job Into Schedule** from the context menu or click the job or job group, and from the **Activities** main menu, select **Insert Job Into Schedule**.

3. Click **Yes** at the confirmation prompt.

Another occurrence of the selected job or job group is added to today's production schedule regardless of its calendar dates (if any). If the job is defined to repeat, only one occurrence of the job enters the schedule. Note that a job must have the Unscheduled Allowed option selected in its Job Definition to be added in this manner.

Adjusting Column Widths in the Job Activity Pane

To adjust a column's width:
1. Place the cursor on the right vertical border of the column heading. The mouse pointer turns into the separator icon.
2. Hold down the left mouse button and resize the width.

Deleting Jobs or Groups from the Production Schedule

To remove jobs that have not yet launched from the production schedule:
1. From the **Navigator** pane, select **Operations > Job Activity** to display the **Job Activity** pane.
2. Right-click the job or job group occurrence to remove, and from the context menu select **Remove Job(s) from Schedule**. In the resulting confirmation dialog, select **Yes** if you only want to delete the selected job occurrence. If you wish to delete all future scheduled occurrences of the job, select the **All** option.

   You can select multiple jobs or job groups simultaneously.

   – To select more than one job or job group that are adjacent to each other, hold down the Shift key and select the first and last jobs.

   – To select more than one job or job group that are not adjacent to each other, hold down the Control key while selecting the desired job occurrences.

3. Click **Yes** at the confirmation prompt.

Remove Job(s) from Schedule will remove any pre-launch job occurrence from the schedule, for any date. Only the job occurrence you select is removed. If a job group occurrence is selected, all of the group's child jobs are removed.

Filtering Jobs in the Job Activity Pane Based on Job Status

To filter a job:
1. From the **Navigator** pane, select **Operations > Job Activity** to display the **Job Activity** pane.
2. Click the **Filter** button on the TWA toolbar or right-click the Navigator pane and select **Filter** from the context menu to display the **Job Filter** dialog.

   The statuses that are presently displayed are the ones with a check mark.

3. Select the agent, owner and job statuses to view.

Moving to Different Dates and Times in the Job Activity Pane

You can view job occurrences for past, present or future days. Each job's retention history value in its job definition determines how far back you can see its job occurrences. The **Future Days** value in the **System Configuration** dialog determines how far you can see into the future.

Note: When you are viewing the current day in the Job Activity pane, you may also see jobs from previous dates if those jobs have not yet completed.

To move to different dates and times:

From the **Navigator** pane, select **Operations > Job Activity** to display the **Job Activity** pane.

- To move backward one day in the production schedule, right-click in the **Job Activity** pane and select **Prior Day** from the context menu.

- To move forward one day in the production schedule, right-click in the **Job Activity** pane and select **Next Day** from the context menu.

- To move to today's date and time, right-click the **Job Activity** pane and select **Goto Now** from the context menu.

- To move to a job's graphical start location, right-click the job's record and select **Goto Job Time** from the context menu to display the job's start time.

Pausing the Production Schedule

You can pause and then resume the production schedule at any time. You can also pause and resume job launching without having any affect on scheduling.

When you pause a production schedule, the master service is stopped and waiting jobs are prevented from running, even if their dependencies are met. Jobs that are already in **Active** status will continue to run on their respective agents; however, their updated status (for example, **Completed Normally**) will not be displayed in the **Job Activity** pane until the master service is restarted.

When you pause job launching, jobs can still be scheduled and calendars compiled, but jobs will not actually launch until job launching is resumed.

To pause the production schedule or job launching:

1. From the **Navigator** pane, select **Operations > Job Activity** to display the **Job Activity** pane.

2. From the **Activities** main menu, select **Pause Scheduler** or **Pause Job Launching**.

Rearranging Columns in the Job Activity Pane

To rearrange columns:

1. From the **Navigator** pane, select **Operations > Job Activity** to display the **Job Activity** pane.

2. From the **View** main menu, select **Preferences** or right-click the **Navigator** pane and select **Preferences** from the context menu to display the **Preferences** dialog.

3. Click the **Columns** tab.

4. Select the column and click the up or down arrows to arrange the titles into the desired order.

Recreating the Production Schedule

To recreate the production schedule:

1. From the **Navigator** pane, select **Operations > Job Activity** to display the **Job Activity** pane.

2. From the **Activities** menu, select **Create Schedule** to display the **Create Schedule** dialog.

3. Select the date range for which to create the schedule.

4. If you want repeating jobs with no time window to start immediately, rather than after midnight, select the **For today's repeating jobs with no time window, start repeating ASAP** option.

5. TWA searches for all qualified jobs to add to the production schedule for the selected dates.

Note: Warning: All job occurrences that were added manually and all records of jobs that ran will be lost. Furthermore, the Create Schedule recreates the schedule at the time you issue the command. This means that jobs whose dependencies have not been met can qualify to run immediately. Use this command with caution.

Refreshing the Job Activity Pane

To refresh:

1. From the **Navigator** pane, select **Operations>Job Activity** to display the **Job Activity** pane.

2. From the **View** main menu, select **Refresh**.

Resuming the Production Schedule or Job Launching

To resume the production schedule or job launching:

1. From the **Navigator** pane, select **Operations>Job Activity** to display the **Job Activity** pane.

2. From the **Activities** main menu, select **Resume Scheduler** or **Resume Job Launching**.

Saving a Job's Output on the Master

You can view the output of a job after the job finishes running. In order to view the output of a job from the **Job Activity** pane, you must save the output for the job on the master (this is the default).

To save a job's output on the master:

1. From the **Navigator** pane, select **Definitions>Jobs** to display the **Jobs** pane.

2. Select the job for which to save output.

3. Click the **Edit** button to display the **Job Definition** dialog.

4. Click the **Options** tab.

5. In the **Save Output** option, select **Append** or **Replace**.

6. Click **OK**. The next time the job completes, you can view its output from the **Job Detail** dialog.

Selecting or Removing All Job Filters in the Job Activity Pane

To select or remove all job filters:

1. From the **Navigator** pane, select **Operations>Job Activity** to display the **Job Activity** pane.

2. Click the **Filter** button on the TWA toolbar or right-click the Navigator pane and select **Filter** from the context menu to display the **Job Filter** dialog.

3. Right-click in the **Statuses** section and select **Check All** or **Uncheck All** from the context menu.

Setting the Completion Status of a Job

Note: You should only set the completion status of jobs that have already finished running. Setting completion status of a job that is still active, may adversely affect job activity.

To set the completion status:

1. From the **Navigator** pane, select **Operations>Job Activity** to display the **Job Activity** pane.

2. Select the job on which to perform job control.

3. Right-click the job to open the **Job Activity** context menu.

4. From the **Job Control** submenu, select **Set** and from the **Set** submenu, select the desired job completion status.

   Note: When job completion statuses are changed in this way, an asterisk appears to the right of it, for example, Completed Normally*.

Sorting Jobs in the Job Activity Pane Using Column Headers

You can sort the job occurrences displayed in the **Job Activity** pane by two different sorting criteria. You can sort on a primary criteria and in turn, organize the primary sort by secondary criteria. The sorts are done by column header in either alphabetical or numerical order depending on the type of data contained in the column.

To sort a job using a column header:

1. From the **Navigator** pane, select **Operations>Job Activity** to display the **Job Activity** pane.

2. Click the column header to select it for a primary sort. The jobs sort alphabetically (in text data) or numerically (if numerical data). A single arrow displays beside the column header to denote the primary sort.

   If you click the header more than once, the sort order toggles back and forth from ascending to descending order.

3. To perform a secondary sort, hold down the **ALT** key while clicking the column header. The secondary sort displays a double arrow in the column header.

   For example, click the **Name** column heading. The items in the column are alphabetized from a to z (because the column contains alphabetical data). Clicking again reverses the order z to a. Click the **Time** column header while holding down the **ALT** key and the jobs remain in alphabetical order but are now organized also by numerical order (since this column contains numerical data). Now the jobs are arranged in alphabetical order and each numerical segment is organized from earliest time to latest time. Clicking again on the secondary sort column while holding down the **ALT** key reverses the secondary sort to latest time to earliest time.

   If you click the header more than once, the sort order toggles back and forth from ascending to descending order.

Sorting Jobs Using the Job Filter Dialog

To sort jobs using the Job Filter dialog:

1. From the **Navigator** pane, select **Operations>Job Activity** to display the **Job Activity** pane.

2. Click the **Filter** button on the TWA toolbar or right-click the **Navigator** pane and select **Filter** from the context menu to display the **Job Filter** dialog.

3. Checking or clearing the options for status or hours enables you to sort jobs using those criteria. You can also sort using other criteria listed including by job name, by agent, by queue, by command, etc.

   Note: The Job Filter dialog includes filter options that affect sorting. For example, if you want to sort by the status column, and some statuses have been filtered out, those statuses will not be used for sorting.

   Note: When a Status column is clicked for sorting within the Job Activity pane, the order is defined by the settings on the Job Status Order tab in the System Configuration dialog.

Stopping or Starting Jobs in the Job Activity Pane

To stop or start a job:

1. From the **Navigator** pane, select **Operations>Job Activity** to display the **Job Activity** pane.

2. Select the job on which to perform job control.

3. Right-click the job to display the **Job Activity** context menu.

4. From the **Job Control** submenu, select one of the following job control options:

    – Select **Override** to override a job's dependencies. The job enters a queue and becomes active. **Override** ignores all dependencies and launches the job immediately.

    – Select **Hold/Stop** to stop a job from running or to keep a job from launching. The job enters the **Held** or **Stopped** status. If the job cannot be held or stopped, you cannot select this menu item.

    Note: Windows jobs cannot be suspended, therefore this menu item is unavailable for active Windows jobs.

    – Select **Release/Resume** to release a held job, or a job waiting for operator release or to start a stopped job. The job enters the **Active** or **Waiting On Dependencies** status. If the job can't be released or resumed, you won't be able to select this menu item.

    – Select **Cancel/Abort** to cancel a job before running or to abort a job while running (**Canceled** jobs cannot be rerun). The job enters the **Canceled** or **Aborted** status. If the job can't be canceled or aborted, you won't be able to select this menu item.

    – Select **Cancel/Abort All** to cancel all occurrences of a job before running or to abort all occurrences of a job while running. All occurrences of the same job enter the **Canceled** or **Aborted** status.

    – Select **Rerun** to rerun a job. The job re-enters the schedule and is re-executed with the same job ID number. A new occurrence ID is not created. If the job can't be rerun, you won't be able to select this menu item.

Viewing a Job's Output

You can view the output of a job after the job finishes running. In order to view the output of a job from the **Job Activity** pane, you must first have the **Save Output** option selected in the job's or job group's definition.

Note: Scheduler's default is to discard job output.

To view a job's output:

1. From the **Navigator** pane, select **Operations>Job Activity** to display the **Job Activity** pane.

2. Double-click the completed job to display the **Job Detail** dialog.

3. Click the **Output** tab to view job output.

Viewing Other Dates in the Production Schedule

To view other dates:

1. From the **Navigator** pane, select **Operations>Job Activity** to display the **Job Activity** pane.

2. Click the **Select Day** button or right-click in the **Job Activity** pane and select the **Select Day** option to display a **Calendar** tab.

3. Click the date to go to. You can use the arrow controls in the title bar to move to different months.