



# Installing Fault Tolerance

---

## Introduction

The basic principle of fault tolerance is to keep your production schedule running continuously despite machine failures.

## Auto Mode

Auto mode is the default way of configuring fault tolerance. This mode allows the primary master to run in standby mode. If the primary master fails and the backup master assumes control, then the backup master assumes the active role. When the primary master that failed comes back online, it remains in standby mode. This type of fault tolerance does not care if the original primary master is actively controlling the production or if the configured backup master is in control. Regardless of the original configuration, each master is interchangeable and can operate in either an active or standby mode. See also, [“Operational Modes for Fault Tolerance”](#).

## Fixed Mode

When configured in Fixed mode, if the machine managing your production schedule fails, fault tolerance ensures that another machine is available to assume control over the production schedule. Scheduler’s fault tolerance ensures that a backup master can keep production going if the primary master should fail.



### Note

---

Fault tolerance does not protect against database failures. This is best left to your database administrator who can set up data mirroring based on the type of database being used.

---

Whenever the primary master is running while the backup master remains available to assume control, the system is in standby mode. If the primary master is unable to run, control of the production schedule passes to the backup master ensuring uninterrupted production. Whenever the backup master assumes control from the primary master, the system is the backup mode.

When the backup master assumes control, it continues the production schedule until control is manually switched back to the primary master. During the time the backup master controls the production schedule, fault tolerance is disabled. Fault tolerance is enabled again when the primary master resumes control. In the backup mode, fault tolerance is disabled because the backup master does not have a backup.

**Note**

---

Plan to spend approximately two hours for the installation and configuration of fault tolerance.

---

During a failover, the green light beside the fault monitor name (located in the first column of the Connections pane) turns red. This light indicates that fault tolerance is not operating.

The status lights warn users that without master redundancy, the network is vulnerable to failure. Returning the primary master to service and restoring your system to a normal fault tolerant status should be the highest priority. Use the switch back procedure to return the primary master to service. See [“Primary Master Switchback”](#).

In the Unix installation procedure after providing a directory location for the installation files, a screen asks if you wish to install a primary master or backup master. You should install the primary master first. Complete the primary master installation and then repeat the master installation on a different machine, selecting the backup master option for the second installation. For more information on installing the primary and backup masters for Unix, refer to [“Installing the Master for Unix”](#).

## Components of Fault Tolerance

Fault tolerance consists of the following main components:

- Client Manager – The Client Manager services requests from user initiated activities, such as through the Tidal Web Client.
- Primary Master – The primary master controls production scheduling during normal system operations.
- Backup Master – The backup master operates in standby mode until it takes over for the primary master. In case of a failover, the backup master becomes active and clients reconnect to the backup master.
- Fault Monitor – The fault monitor continuously monitors the status of the primary and backup masters. It initiates the transfer of scheduling control from the primary master to the backup master. The Tidal Web client provides an interface to the fault monitor service.


Both the primary master and the backup master are designed to communicate with a database. Responsibility for setting up and maintaining this database is left to your database administrator. *TES* does *not* provide fault tolerance for the database.

## Operational Modes for Fault Tolerance

### Normal (Sleep) Mode

Figure 5-1 shows normal operation, or the sleep mode. The backup master remains in the background until required though maintaining constant communication with both the primary master and the fault monitor.


**Figure 5-1** Normal Operation, Sleep Mode



### Backup Mode

Figure 5-2 shows fault tolerance operation when the primary master goes down (backup mode). The backup master becomes active, assuming control of the production schedule while the primary master is out of service. Both figures show only the main components of fault tolerance.


Figure 5-2 Primary Master Down, Backup Mode



## Network Configuration

For fault tolerance to operate properly, the physical network connections between the various components must be configured properly for reliable communication. Dedicated TCP/IP communication ports, configured during installation, are used to exchange messages between components and to verify whether the connections are up or down. Figure 5-3 shows the network connections and the communication ports.

Figure 5-3 Network Configuration



## System Requirements

The primary master, the backup master and the fault monitor must be installed on separate machines. There are different system requirements for the Windows and Unix platforms. See your Cisco Tidal Enterprise Scheduler User Guide for the system requirements.

## User Account Requirements

The account used to install fault tolerance must meet the following requirements:

- The account must be a domain user account, not a local machine account.
- The account must be part of the local Administrators group.
- The account must have the advanced local user rights Logon as a service and Act as part of the operating system set by the system administrator on each server
- If you are installing on Unix, the masters and fault monitor must be installed under either root or a user created by root.



---

**If these user requirements are not met, fault tolerance will not install properly and your system will be unprotected in the case of a failure.**

---

## Prerequisites for Installation

The following points list the conditions that must be met before fault tolerance is installed. Failure to meet these preliminary conditions will result in unnecessary delay and may cause the installation process to fail.

- There must be at least three machines for a fault tolerance setup.
- All three machines must be in the same domain.
- The primary master must be installed, licensed and operational.
- Both the primary and backup masters must meet the minimum system requirements for TES.
- The primary master and the backup master must have the same operating system (including patches), database connectivity and hardware configuration. Their software and configuration should mirror the each other.
- The same version of JVM must be installed on the primary master, backup master and fault monitor machines.
- The backup master's clock should be as closely synchronized as possible with the primary master's clock.

## Installing Fault Tolerance for Windows

Installing fault tolerance on a network means adding another master to shadow an existing master. The first master becomes the primary master in TES while the second master is referred to as the backup master. This backup master, like the primary master, is controlled from the Service Manager. The procedure to install this backup master is similar yet different from the procedure detailed in the *Installing the Master* chapter.

A network component that monitors the operation of the two masters called the fault monitor is installed on a third machine. A Fault Monitor window is added to the **Navigator** pane in the Tidal Web client.

This chapter describes how to:

- install (or upgrade) the backup master on a Windows machine
- install the fault monitor on a Windows machine

## Prerequisites for Installation

The following items should be completed and ready prior to starting fault tolerance installation.

- Create a backup of your database. As a matter of general operating policy, it is recommended that a database be backed up at least once daily.
- Ensure that the primary and backup master clocks run no more than 15 seconds apart.
- The primary and backup master machines should mirror each other in hardware and software configurations.
- The primary and backup master machines must be able to ping each other, the fault monitor and the database.

The fault monitor machine must be able to ping the Client Manager, primary and backup masters.

Client Manager can reside on either the master machine or on a separate machine. In this case, there will be a fourth machine.

- Use three separate PCs, one as a primary master, one as a backup master and one as a fault monitor.

## Installation Check List

To ensure a successful and smooth installation of fault tolerance, prior to the start of installation collect the information about the machine components needed during installation. Use the following check list to collect the necessary information:

### Primary Master:

Computer Name \_\_\_\_\_

Host Name \_\_\_\_\_

### Backup Master:

Computer Name \_\_\_\_\_

Host Name \_\_\_\_\_

### Fault Monitor:

Computer Name \_\_\_\_\_

Host Name \_\_\_\_\_

Record the domain user account.

Domain Account \_\_\_\_\_

Record the port numbers to be used by the primary master, backup master and fault monitor.

Fault monitor to master \_\_\_\_\_

Master to master \_\_\_\_\_

Fault monitor to client\_\_\_\_\_

Scheduler provides default port numbers of **6703** for fault monitor to master, **6704** for master to master, and **6705** for fault monitor to client.

## Installing Components for Fault Tolerance

Before installing, get license files for your TES fault tolerance components from the Licensing Administrator for Scheduler. The fault tolerance setup consists of steps which must be performed in order. The procedures for each step are covered in this chapter.

A Scheduler primary master, Client Manager and agent(s) must be already installed, licensed and operational for a successful fault tolerance installation. The individual components can be installed on different machines, but they must all be in the same domain as your fault tolerance setup. You will need to refer to the information collected on the Installation checklist. The first master installed becomes your primary master.

## Installing the Backup Master

To install the backup master:

- 
- Step 1** Load the Tidal Enterprise Scheduler installation DVD-ROM into the DVD-ROM drive of the machine where the backup master is being installed. The Tidal Scheduler panel displays.
  - Step 2** On the Scheduler screen, click the **Backup Master** link and select the Run this program from its current location option in the File Download dialog box. The Welcome panel displays.
  - Step 3** Click **Next**. The Installation Type panel displays.
  - Step 4** Select **BackupMaster**, then click **Next**.
  - Step 5** On the Destination Folder panel, select the directory where the TES files will reside.
    - Click the **Change** button to search for a directory.
    - or-
    - Accept the default location *C:\Program Files\TIDAL*.
  - Step 6** Click **Next**. The Database Type panel displays.
  - Step 7** Select the type of database being used and click **Next**. The Database Server panel displays.
  - Step 8** Enter the name of the database server used by the primary master.
  - Step 9** Click **Next**. The Ready to Install the Program panel displays.
  - Step 10** Click **Install**. The Installshield Wizard Complete panel displays.
  - Step 11** Click **Finish** to close the wizard.
-

# Installing the Fault Monitor

**Warning**

---

**The fault monitor must be installed on a separate machine from the primary and backup masters.**

---

To install the fault monitor:

- 
- Step 1** Click the **Fault Monitor** link on the TIDAL Scheduler installation screen for Internet Explorer.
- or-
- If using the Netscape browser, copy the fault monitor files to a *temp* directory as directed in the TES Installation Steps section.
- The Welcome panel displays.
- Step 2** Click **Next**. The Installation Type panel displays.
- Step 3** Select **FaultMonitor**, then click **Next**. The **Destination Folder** panel displays.
- Step 4** Select the directory where the TES files will reside:
- Click the **Change** button to search for a directory.
- or-
- Accept the default location *C:\Program Files\*.
- Step 5** Click **Next**. The Enter requested data panel displays.
- Step 6** Enter the following:
- FM Port –The port number of the fault monitor.
  - Client Port – The port number of the Client Manager.
- Step 7** Click **Next**. The Ready to Install the Program panel displays.
- Step 8** Click **Install**. The Installing panel displays the progress of your fault monitor installation in the form of a progress bar.
- The Setup Completed panel displays.
- Step 9** Click **Finish** to complete fault monitor installation and return to the Scheduler installation dialog box.
- 

## Controlling the Fault Monitor

You can monitor the fault monitor from the Tidal Web 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 Tidal 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

To start the fault monitor, use the following command:

- 
- Step 1** From the Windows Start menu, and choose **Programs > Tidal Software > Tidal Service Manager** to display the Tidal Services Manager.
  - Step 2** From the Service list, choose **SchedulerFaultMon**.
  - Step 3** Click **Start**.
- 

## Stopping the Fault Monitor

To stop the fault monitor, use the following command:

- 
- Step 1** From the Windows Start menu, and choose **Programs > Tidal Software > Tidal Service Manager** to display the Tidal Services Manager.
  - Step 2** From the Service list, choose **SchedulerFaultMon**.
  - Step 3** Click **Stop**.
- 

## Checking the Fault Monitor Status

To check the operation status of the fault monitor, use the following command:

- 
- Step 1** On the Fault Monitor machine, click the Windows Start button and choose **Programs > Tidal Software > Tidal Service Manager** to display the Tidal Service Manager.
  - Step 2** From the Service list, select **Fault Monitor**. At the bottom of the Tidal Service Manager, the status of the selected service displays.
- 

# Installing Fault Tolerance for Unix

This section describes how to:

- install the fault monitor on a Unix machine
- verify that files were successfully installed
- start, stop and check the status of the fault monitor from the command line

While there is a fault monitor console for the Windows platform that displays activity messages about fault monitor components, there is no such fault monitor console for the Unix platform. The activity messages from the fault monitor can be displayed and controlled from the Fault Monitor pane in the Tidal Web client. For more information, refer to the [“Fault Monitor Interface”](#).

## Prerequisites for Installation

The following items should be completed and ready prior to starting fault tolerance installation.

- Create a backup of your database. As a matter of general operating policy, it is recommended that a database be backed up at least once daily.
- Ensure that the primary and backup master clocks run no more than 15 seconds apart.
- The primary and backup master machines must be able to ping each other, the fault monitor and the database.

The fault monitor machine must be able to ping the Client Manager, primary and backup masters.

Client Manager can reside on either the master machine or on a separate machine. In this case, there will be a fourth machine.

- Use three separate Unix machines, one as a primary master, one as a backup master and one as a fault monitor.

## Installation Check List

To ensure a successful and smooth installation of fault tolerance, prior to the start of installation collect the information about the machine components needed during installation. Use the following check list to collect the necessary information:

### Primary Master:

Computer Name \_\_\_\_\_

Host Name \_\_\_\_\_

### Backup Master:

Computer Name \_\_\_\_\_

Host Name \_\_\_\_\_

### Fault Monitor:

Computer Name \_\_\_\_\_

Host Name \_\_\_\_\_

- Record the user account.

User Account \_\_\_\_\_

- Record the port numbers to be used by the primary master, backup master and fault monitor.

Fault monitor to master \_\_\_\_\_

Master to master \_\_\_\_\_

Fault monitor to CM \_\_\_\_\_

TES provides default port numbers of **6703** for fault monitor to master, **6704** for master to master, and **6705** for fault monitor to Client Manager.

## Installing Components for Fault Tolerance

Before installing, get license files for your TES fault tolerance components from the Licensing Administrator for TES. The fault tolerance setup consists of steps which must be performed in order. The procedures for each step are covered in detail in this chapter.

A TES primary master, Client Manager and agent(s) must be installed, licensed and operational for a successful fault tolerance installation. The individual components can be installed on different machines, but they must all be in the same domain as your fault tolerance setup. The first master installed will be your primary master.

## Installing the Backup Master

Instructions for installing the backup master are the same instructions provided in this guide for installing the master (primary) for Unix. The hardware and software requirements for a backup master are the same as the requirements for a primary master. During the installation procedure a screen is displayed to designate whether the installation is for a primary or backup master. Selecting the **Backup** option, ensures that a backup master is installed. Complete the described procedure to install and verify successful installation of the backup master.

## Installation Prerequisites for the Fault Monitor

See your *Cisco Tidal Enterprise Scheduler User Guide* for the requirements that must be met for successful installation of the Unix fault monitor.

## Installing the Fault Monitor



Warning

---

**The fault monitor must be installed on a separate machine from the primary and backup master machines. Only one fault monitor can be installed on a machine.**

---

To install the fault monitor:

- 
- Step 1** If you are copying the installation files from the network, FTP the *install.bin* file to the directory you created.
- If you are using the DVD-ROM, locate the *install.bin* file for your operating system on the DVD-ROM and copy it to a directory you created. The file can be found on the DVD-ROM at `<DVDROMDRIVE>\UnixFaultMon\<operating system>\install.bin`
- Step 2** Change the permissions on the *install.bin* file in the directory to make the file executable:  
**chmod 755 install.bin**
- Step 3** After copying the file to the directory, begin the installation program by entering:  
**sh ./install.bin**
- The Introduction panel displays.
- Step 4** After reading the introductory text that explains how to cancel the installation or modify an previous entry on a previous screen, click **Next**. The Choose Install Folder panel displays.

- Step 5** Select the directory where the TES files will reside:
- Click **Choose** to search for a directory. If you change your mind after selecting a different destination, select **Restore Default Folder** to revert back to the default installation location.
  - or-
  - Accept the default location: */opt/unixsa*
- Step 6** Click **Next**. The Port Numbers panel displays.
- Step 7** Enter the port number of the Fault Monitor and the Client Manager, then click **Next**. The Pre-Installation Summary panel displays the destination location selected for the fault monitor files.
- If the location is not where you intended, click **Previous** until you return to the Choose Install Folder panel and correct the installation location.
- Step 8** Click **Install** to begin the installation of files. The Installing UnixFM panel displays.
- The Install Complete panel displays when the installation process is completed.
- Step 9** Click **Done** to exit the installation program.
- 

## Verifying Successful Installation of the Fault Monitor

You should verify that the installation program installed all of the necessary files.

Go to the *bin* directory location where you installed the fault monitor files and list the contents of the directory with the following command:

```
ls -l
```

You must have two files called *tesfm* and *tmkdea* before the fault monitor can operate correctly.

## Controlling the Fault Monitor

You can monitor the fault monitor from the Tidal Web 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 Tidal 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

To start the fault monitor, use the following command:

```
tesfm start
```

## Stopping the Fault Monitor

To stop the fault monitor, use the following command:

```
tesfm stop
```

## Checking the Fault Monitor Status

To check the operation status of the fault monitor, use the following command:

```
tesfm status
```

## Modifying the Fault Monitor Configuration

You can change the properties of the fault monitor that were set during the installation. Circumstances may force you to change the configuration of the fault monitor as it was originally installed or you may need to change the logging levels of various components for diagnostic purposes.

The properties of the fault monitor are managed in a file called *master.props* that resides in the *config* directory on the fault monitor machine.

The *master.props* file on the fault monitor looks like the following example:

```
FMMasterPort=6703
```

```
FMClientPort=6705
```

Be careful when changing the properties of the fault monitor, incorrect entries to the *master.props* file may prevent the proper operation of the Unix fault monitor.



### Note

If you change the Fault Monitor Client Port number in the Connection Definition dialog box in the Tidal Web client, you must manually change the FMClientPort number in the fault monitor *master.props* file also.

The properties options that are managed in the *master.props* file are listed below:

| Property                 | Default Value | What it Controls   |
|--------------------------|---------------|--|
| FMMasterPort             | 6703          | Number of the port used by the master to connect to the fault monitor. The default number is 6703.   |
| FMClientPort             | 6705          | Number of the port used by the Client Manager to connect to the fault monitor. The default number is 6705. This port number must match the port number in the fault monitor's Connection Definition dialog box in the Tidal Web client. If you change the port number in one place, you must manually change the port number in the other place. |
| CMDMasterPort (Optional) | 6600          | Number of the port that the command line program uses to connect to the Unix master machine. (This property is only used to modify the port if it is being used by another application.)   |

## Fixing a Port Number Conflict

A port number conflict may occasionally occur in the fault monitor. Certain port numbers are used by default in the fault monitor. If another application is using the same port numbers then the fault monitor will not work and you must change the port numbers. Some port numbers can be changed from the Connection Definition dialog box for that component but others must be manually changed on the fault monitor machine. This port conflict may occur with either the port being used by the Client Manager to connect to the fault monitor (**6705**) or with the port used by the command line program to connect to the Unix master machine (**6600**).

To fix a port number conflict:

- 
- Step 1** On the fault monitor machine, locate the *config* directory.
- Step 2** Open the *master.props* file to see the various properties that control the port numbers used by the fault monitor.
- FMClientPort is for the port used by the Client Manager to connect to the fault monitor.
  - CMDMasterPort is for the port used by the command line program.
- Step 3** Change the port number to a port number not in use by any other application.

**Note**

Be sure that the port numbers in the *master.props* file match the port numbers in the component's Connection Definition dialog box.

---

## Using Fault Tolerance

Fault tolerance in TES is configured on the Fault Tolerance tab in the System Configuration dialog box of the Tidal Web client. Messages about fault tolerance are displayed in both the Fault Monitor console and the Tidal Web client Fault Monitor pane. The following sections explain how to configure and verify fault tolerance.

## Licensing Fault Tolerance

The Fault Tolerance function cannot be used unless it is properly licensed. You must shut down fault tolerance to load the license file.

**Note**

Fault tolerance cannot be turned off if the backup master is active. The primary master must be in control to turn off fault tolerance.

---

Obtain the license code from the licensing manager at Cisco. Registering the license for Fault Tolerance is performed from the Tidal Web client.

To load a production license for Fault Tolerance, you need the proper license file.

To license Fault Tolerance with a Full license:

- 
- Step 1** Stop the master:
- For Windows:
- a. From the Windows Start menu, choose **Programs > TIDAL Software > Scheduler > Master > Service Control Manager** to display the Tidal Services Manager.
  - b. Verify that the master is displayed in the Service list and click on the **Stop** button to stop the master.
- For Unix:
- Enter **tesm stop**.
- Step 2** Rename your Full license file to *master.lic*.

- Step 3** Place the file in the *C:\Program File\TIDAL\Scheduler\Master\config* directory.
- Step 4** For Windows, restart the master by clicking **Start** in the Service Control Manager. For Unix, restart the master by entering **tesm start**.
- The master will read and apply the license when it starts.
- 

## Failover Configuration

From the Activities menu, choose **Configure Scheduler** to display the System Configuration dialog box.

## Enabling Fault Tolerance

To enable fault tolerance:

- 
- Step 1** Before enabling fault tolerance, stop the backup master.
- From the Windows Start menu, choose **Programs > TIDAL Software > Scheduler > Master > Service Control Manager** to display the Tidal Services Manager.
  - In the Service list, select **Backup Master** if it is not selected.
  - Click the **Stop** button to stop the backup master.
- Step 2** In the Tidal Web client, choose **Configure Scheduler** from the Activities main menu to display the System Configuration dialog box.
- Step 3** Select the Fault Tolerance tab.
- Step 4** Click the Enable Failover option to add the check mark and enable fault tolerance operation.
- Step 5** To complete the Enable Failover process, verify that the Fault Monitor and the Backup Master are started.
- 

## Fault Tolerance Tab Options

The Fault Tolerance tab of the System Configuration dialog box contains the following options:

- Failover Enable – Enables fault tolerance. If this option is not selected, then no action (failover) is taken if the master fails. If this option is selected, control of production switches over to the designated backup master if a failure occurs on the primary master. The default configuration is disabled. Selecting this check box, displays the following options to enter the information required to configure fault tolerance.
- Machine Name (backup master only) – The name of the machine where the backup master resides.
- Backup-To-Master Port (backup master only) – The port number used for communication between backup master and primary master.
- Machine Name (fault monitor only) – The name of the machine where the fault monitor resides.
- Fault Monitor Master Port (6703 is the default) – The port number used by the fault monitor to communicate with masters.

- Fault Monitor Client Port (6705 is the default) – The port number used by the fault monitor to communicate with the Client Manager.

## Starting Fault Tolerance

To start fault tolerance:

- 
- Step 1** Verify that Failover is enabled. See [“Enabling Fault Tolerance”](#).
  - Step 2** Close all Clients.
  - Step 3** Stop the Client Manager, Masters, Primary and Backup, if running, via the Service Control Manager (Windows) or the command line (Unix).
  - Step 4** Verify that the Fault Monitor is running.
  - Step 5** Start the Primary Master and verify that it is running via the Service Control Manager (Windows) or the command line (Unix).
  - Step 6** Start the Backup Master and verify it is running via the Service Control Manager (Windows) or the command line (Unix).
  - Step 7** Start the Client Manager via the Service Control Manager (Windows) or the command line (Unix).
  - Step 8** Log into the application via the Tidal Web Client and choose **Operations > Master Status**.
  - Step 9** Select the Fault Monitor tab and validate Poll Activity. You should see *Primary OK and Backup OK. [Standby Mode]*.

More options are displayed to add the information required to configure fault tolerance. Refer to [“Fault Tolerance Tab Options”](#) for more information on the options used to configure fault tolerance.

---

## Verifying Fault Tolerance Operation

To verify fault tolerance operation:

- 
- Step 1** Launch the Tidal Web client and from the Navigator pane, select **Operations > Fault Monitor** to display the Fault Monitor pane.
  - Step 2** Check the activity messages displayed in the Fault Monitor pane to verify that all components of fault tolerance are operating correctly.
- 

## Setting Failover Time

By default, failover takes three minutes. You can adjust this time period for more or less primary master recovery time.

To set failover time:

- 
- Step 1** Locate the *master.props* file in the **config** directory where you installed the fault monitor files on the fault monitor machine.



- Step 2** Open the *master.props* file in a text editor.
- Step 3** On a separate line in the file, enter:  
**ToleranceTime=<number of minutes>**  
 where the brackets are replaced with the number of minutes to pass without contact with the primary master before failover to the backup master.
- Step 4** Stop the fault monitor.
- Step 5** Start the fault monitor to enable the new parameter.

## Modifying Fault Tolerance Parameters

You can change the properties of the fault monitor that were set during the installation. Circumstances may force you to change the configuration of the fault monitor as it was originally installed or you may need to change the logging levels of various components for diagnostic purposes.

The properties of the fault monitor are managed in a file called *master.props* that resides in the *config* directory on the fault monitor machine.

The *master.props* file on the fault monitor looks like the following example:

**FMMasterPort=6703**

**FMClientPort=6705**

Be careful when changing the properties of the fault monitor, incorrect entries to the *master.props* file may prevent the proper operation of the fault monitor.



### Note

If you change the Fault Monitor Client Port number in the Connection Definition dialog box in the Tidal Web client, you must manually change the FMClientPort number in the fault monitor *master.props* file also.

The rest of the fault tolerance parameter options that are managed in the *master.props* file are listed below:

| Property          | Default Value | What it Controls   |
|-------------------|---------------|--|
| FaultMonitorLog   | INFO          | Sets the level of detail for recording messages about the fault monitor to the Fault Monitor log.  |
| FaultToleranceLog | INFO          | Sets the level of detail for recording messages about fault tolerance components to the Fault Monitor log.   |
| FMMasterPort      | 6703          | Number of the port used by the master to connect to the fault monitor. The default number is <b>6703</b> .   |
| FMClientPort      | 6705          | Number of the port used by the Client Manager to connect to the fault monitor. The default number is <b>6705</b> . This port number must match the port number in the fault monitor's Connection Definition dialog box in the Tidal Web client. If you change the port number in one place, you must manually change the port number in the other place. |

| Property   | Default Value | What it Controls   |
|--|---------------|--|
| ToleranceTime  | 3             | Number of minutes the fault monitor will go without communication with the primary master before having the backup master assume control.  |
| CMDMasterPort (Optional)   | 6600          | Number of the port that the command line program uses to connect to the Unix master machine. (This property is only used to modify the port if it is being used by another application.) |
| <b>Tuning for DSP to FM message traffic (all DSP connections).</b> |               |  |
| MinSessionPoolSize   | 2             | —  |
| MaxSessionPoolSize   | 5             | —  |
| MaxConcurrentMessages  | 5             | —  |

## Fixing a Port Number Conflict

A port number conflict may occasionally occur in the fault monitor. Certain port numbers are used by default in the fault monitor. If another application is using the same port numbers then the fault monitor will not work and you must change the port numbers. Some port numbers can be changed from the Connection Definition dialog box for that component but others must be manually changed on the fault monitor machine. This port conflict may occur with either the port being used by the Client Manager to connect to the fault monitor (**6705**) or with the port used by the command line program to connect to the Unix master machine (**6600**).

To fix a port number conflict:

- 
- Step 1** On the fault monitor machine, locate `config > master.props`.
  - Step 2** Use a text editor to open the file to see the various properties that control the port numbers used by the fault monitor.
    - FMClientPort is for the port used by the Client Manager to connect to the fault monitor.
    - CMDMasterPort is for the port used by the command line program.
  - Step 3** Start the Client Manager.
  - Step 4** Change the port number to a port number not in use by any other application.
  - Step 5** Stop the fault monitor.
  - Step 6** Start the fault monitor to enable the new parameter.



### Note

Be sure that the port numbers in the `master.props` file match the port numbers in the component's Connection Definition dialog box.

---

# Fault Monitor Interface

The fault monitor can also be displayed from the Tidal Web client console pane. To display messages from the fault monitor in the client, click the Fault Monitor tab within the *Master Status* folder. The Fault Monitor pane displays any messages from the fault monitor.

**Note**

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

## Fault Monitor Pane Context Menu

Various functions for the fault monitor can be accessed from the context menu of the Fault Monitor pane. To display the menu, right-click anywhere in the Navigator pane or the Fault Monitor pane.

The Fault Monitor pane context menu.

- Refresh – Updates the information displayed in the Fault Monitor tab.
- Print – Prints the messages displayed in the Fault Monitor tab.
- Print Selected – Prints the selected messages displayed in the Fault Monitor tab.
- Stop All – Stops the operation of the primary master, the backup master, and the fault monitor. When you select this option, a Confirm dialog box displays. Click **Yes** to continue and **No** to abort.
- Stop Fault Monitor – Stops the fault monitor. When you select this option, a Confirm dialog box is displayed. Click **Yes** to continue and **No** to abort. If the fault monitor is not running, failover is not possible.
- Stop Backup and FaultMon – Stops the operation of the backup master and the fault monitor. When you select this option, a Confirm dialog box displays. Click **Yes** to continue and **No** to abort. When the backup master and fault monitor are stopped, the primary master continues without being fault tolerant.

Notice that there are no menu options to start the fault monitor or the primary master. These components are started from the Service Manager or if you are using the Unix version, the components are started from the command line of each machine hosting that component.

## Fault Tolerance Operation

Fault tolerance is only available if a backup machine is available to assume control. This means that if a primary master fails and the backup master assumes control during a failover, your system is no longer fault tolerant. If you are using the backup master because the primary master failed then there is no backup protection in case the backup master also fails. You must restore the failed master to operation to return to a fault tolerant mode. Only when both masters are operational, with one master running and the other master on standby, is your system fault tolerant.

**Note**

Fault tolerance cannot be turned off if the backup master is active. The primary master must be in control to turn off fault tolerance.

The default way of configuring fault tolerance though, allows the primary master to run in standby mode. If the primary master fails and the backup master assumes control, then the backup master assumes the active role. When the primary master that failed comes back online, it remains in standby mode. This type of fault tolerance does not care if the original primary master is actively controlling the production or if the configured backup master is in control. Regardless of the original configuration, each master is interchangeable and can operate in either an active or standby mode. Fault tolerance is configured to run in this manner by using the AUTO value for the FT\_OPERATION property in the *master.props* file.

**Note**

The default value for the FT\_OPERATION property in the *master.props* file is AUTO.

This duality of roles can be confusing to keep track of, but the messages displayed in the Fault Monitor pane note in which mode a master is operating. If a master is in control, it is considered active and if it is in standby mode, this is also noted. For example, a message in the Fault Monitor pane may read “*Backup OK [Active]*” to denote that the designated backup master is in control. A similar message concerning the backup master in standby mode would read “*Backup OK [Standby]*.”

Fault tolerance can operate in a different manner if needed. One of the masters can be designated to always be the primary master. Its primary master role is fixed. The machine that is configured as the primary master must be in control with the backup master on standby before the system is considered fault tolerant. The primary master cannot run in standby mode. In this configuration, the system can never be fault tolerant if the backup master is in control. Once a failover occurs, the primary master cannot be restarted until the backup master is stopped. Fault tolerance can be configured to run in this manner by using the FIXED value for the FT\_OPERATION property in the *master.props* file.

## Stopping Scheduler in Fault Tolerant Mode

If TES is running in fault tolerant mode, all of the Scheduler components can be conveniently stopped from the Fault Monitor pane in the Tidal Web client. Scheduler will automatically stop the components in the proper sequence.

To stop Scheduler in Fault Tolerant mode:

- 
- Step 1** Stop all fault tolerance components from the Navigator pane of the Tidal Web client by selecting **Operations > Fault Monitor** to display the Fault Monitor pane.
  - Step 2** Right-click the Fault Monitor pane and from the displayed context menu, choose **Stop All**.
- 

## Starting Scheduler in Fault Tolerant Mode

**Note**

It is a recommended practice to prevent any new jobs from being submitted during this procedure by setting the system queue to 0. Let the active jobs complete.

To start Scheduler in Fault Tolerant mode:

- 
- Step 1** On the fault monitor machine, verify that the fault monitor is running.
  - Step 2** Start the primary master.

- Step 3** Start the backup master.
  - Step 4** Launch the Tidal Web client and from the Fault Monitor pane, verify that both masters are running.
- 

## Primary Master Switchback

Primary master switchback is the process of switching scheduling duties from the backup master back to the primary master and restoring normal fault tolerance operation.

To switch back to the primary master on the Windows platform:

**Note**

It is a recommended practice to prevent any new jobs from being submitted during this procedure by setting the system queue to 0. Let the active jobs complete before beginning the switchback.

---

- Step 1** From the fault monitor machine, verify that the fault monitor is running.
- Step 2** If the primary master is not running, start it.
- Step 3** Stop the backup master.
- Step 4** The primary master will leave standby mode and assume control.
- Step 5** Start the backup master.
- Step 6** Launch the Tidal Web client and verify in the Fault Monitor pane that both masters are running.

Switchback is complete once the primary master is actively controlling the production schedule and the backup master is in Standby mode. Be sure to reset the system queue to its original setting.

---

