

IBM Tivoli Workload Scheduler 8.2 – Fault Tolerant Switch Manager



Date:	April 28, 2004 – Final
Component:	IBM Tivoli Workload Scheduler Fault Tolerance Switch Version 8.2
General Description:	IBM Tivoli Workload Scheduler Fault Tolerance Switch Version 8.2
Dependencies:	Tivoli Workload Scheduler 8.2.

This document provides the following information:

- Introduction
- Feature
- Installation
- Configuration
- Use of the feature
- Troubleshooting
- Performance

Introduction

The aim of this document is to describe the new Fault Tolerant Switch Manager functionality.

This new functionality replaces the traditional Fault Tolerant Switch Manager with an enhanced one. In this document, the general characteristics of the feature are described, as well as use information, validity conditions, limits, and configuration.

A flag is provided to activate the functionality for the whole Tivoli Workload Scheduler network. When the flag is not enabled, operations are the same as in previous versions of the product.

This functionality, when active, prevents the loss of messages when a domain manager stops working.

The correct situation will be restored performing a switch manager operation on a backup domain manager.

The full functionality of the feature requires a few conditions which will be explained in this document.

Feature

The goal of this new functionality is to improve fault tolerance in Tivoli Workload Scheduler networks. This has been achieved meeting the conditions of "no loss of messages" as follows:

- Ensuring that there is no single point of failure for the events. This means that since an event is sent from a workstation, until it has reached its final target FTA, that event is not stored on a single workstation only.
- Ensuring that, even if a domain manager stops working, any backup domain manager in the same domain can replace the damaged domain manager and restore the network functions.

Traditional Fault Tolerance Mechanism

The traditional Tivoli Workload Scheduler Backup domain manager architecture is based on the full-status FTA mechanism, where the primary domain manager forwards all the Tivoli Workload Scheduler events that it receives from the network (or that it generates internally) to the full-status FTAs in its domain. Theoretically, as a full-status FTA receives exactly the same events as are sent to the primary domain manager, it should have the most updated information and virtually be able to take over the primary domain manager functionality in case of failure. The interactions between the primary domain manager and the full-status FTA acting as the backup domain managers can be summarized as:

- The incoming events (for example a job successful termination event) that are received from the network on the primary domain manager are buffered into a file named mailbox.
- The Tivoli Workload Scheduler mailman process on the primary domain manager scans the mailbox file and processes the events by sending them to the backup domain manager and to all the other target FTAs in the Tivoli Workload Scheduler network. It then removes the event from the mailbox file. If mailman is not able to route the event to a particular FTA or to the backup domain manager, the event is buffered into pobox files on the primary domain manager.

As a consequence of this traditional behavior, when the primary domain manager fails, all the events in the mailbox that are waiting to be processed and sent to the backup domain manager or to other FTAs/domain managers in the hierarchy, remain unprocessed and are not sent to any other FTA until the primary domain manager restarts. If the Tivoli Workload Scheduler primary domain manager has a permanent failure, all the events in its mailbox file and poboxes are permanently lost.

Enhanced Fault Tolerance Mechanism

In the current Tivoli Workload Scheduler architecture, the entry and exit points of a Tivoli Workload Scheduler domain are the domain managers. Each event entering or leaving a particular domain must flow through the corresponding domain manager, in this way creating a single point of failure in the domain manager itself.

Figure 1 shows the current architecture, where DMA and DMB are two distinct domain managers A and B.

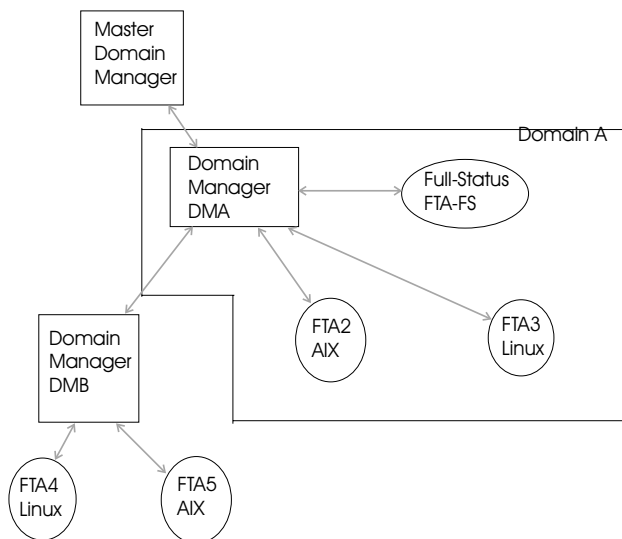


Figure 1. Current architecture

The enhanced fault tolerance mechanism is based on a "Multiple Inbound Connections" approach.

Inside a domain, the events are no longer routed from the primary domain manager, but arrive directly from the originating FTAs. When an FTA sends an event to a primary domain manager, it also sends the same event to all the full-status FTAs in that domain. If it is unable to deliver the event to any of them, then the event is buffered in the corresponding pobox file on the FTA.

Figure 2 shows the "Multiple Inbound Connections" architecture

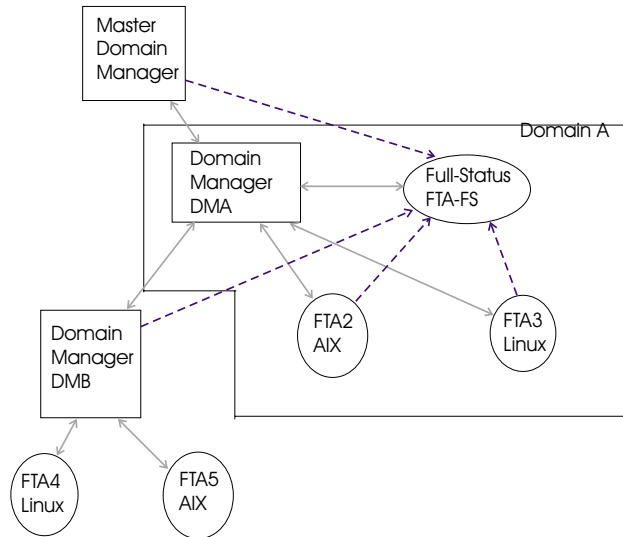


Figure 2. "Multiple Inbound Connections" Architecture

The plain arrows represent the connections that are created with the current Tivoli Workload Scheduler architecture. The dashed arrows represent the additional inbound connections that are created (to the full-status FTA) in a domain with the new "Multiple Inbound Connections" architecture.

This enhancement to the Tivoli Workload Scheduler architecture adds to the canonical entry and exit point to and from a Tivoli Workload Scheduler domain (the primary domain manager), further entry and exit points in the full-status FTAs of that domain, thus eliminating the single point of failure from the primary domain manager.

When the fault-tolerant switch is active, the **link** and **unlink** commands, issued from the primary domain manager, act both on the primary and on the secondary connections. For details, see "Use of the Feature" on page 5.

When a full-status FTA receives an event, it processes it and does not route it further in the Tivoli Workload Scheduler network, but buffers it locally in a cyclical queue, called **ftbox** and which acts as a recovery queue.

The "Multiple Inbound Connections" architecture ensures that all events that are received and processed by the primary domain manager are also received and processed by the full-status FTA (or will be received or processed later if the events are still in some FTA pobox). If the primary domain manager fails, a user can use the **switch-manager** command to switch the domain manager functionality from the primary-domain manager to a selected full-status FTA.

When the **switch-manager** command is received, all the FTAs in that domain disconnect from the primary domain manager and connect to the full-status FTA. During the link establishment phase, the new manager re-synchronizes with each connecting workstation by re-sending and re-generating the last events that were buffered on the **ftboxes**, thus ensuring that none of the events that were still in the primary domain manager message boxes are lost.

So, the full-status FTAs are always updated with the latest status information and all the unprocessed or partially processed events are stored in at least two machines (the original FTA if it was not able to deliver it, and the domain manager or the full-status FTA) ready to be re-sent and re-processed, thus eliminating the single point of failure of the primary-domain manager backup-domain manager communication.

Note: This approach applies both to top-down and bottom-up traffic. The idea of "inbound" does not depend on the direction of the traffic, but is domain-centric, and is repeated the same way for each domain where at least one full-status FTA resides.

Replay Protocol and Event Counter Tables

The above-mentioned enhancement is possible thanks to the concept of replay protocol, using event counter tables.

When the manager role is transferred from the current domain manager to a full-status FTA, the link procedure with a target agent is completed with two additional steps: a global message counter table is passed between the target agent and the new domain manager and the messages the workstation had to receive from the domain manager are sent to the target agent.

With this feature each message on the network is uniquely defined by two terms, *<event-generating workstation>* and *<ID number>*.

To enable the replay protocol, the new domain manager must know which of the messages it manages has to be sent, and which has already been sent by the previous domain manager. This piece of information can be retrieved from the event counter table.

Each workstation manages a counter. This is an array large enough to contain the number of workstations belonging to the network, and where the latest message number both generated locally, externally, or simply passing through the workstation, is recorded for each specific workstation. A message sent to more than one agent contains the same counter value.

The processes in charge of keeping this table updated are Writer (all) and Mailman (only the main mailman, not the server mailman).

After the switch manager operation, the new domain manager contacts the workstations in the network and, for each one, considers whether or not to start the replay protocol.

The replay protocol is coded by a series of message exchanges, but essentially consists of the following steps:

1. The new domain manager requests the protocol
2. If the protocol is recognized, the contacted workstation sends its own event counter table
3. Then the new domain manager compares the messages about to be sent to the workstation, with the values contained in the counter, and decides which messages to send.

When the replay protocol ends, the linking procedure proceeds as usual.

Installation

The installation of this feature is provided with Tivoli Workload Scheduler, Version 8.2, Fix Pack 4 .

To activate this feature, the fix pack must be installed on every agent of the network.

The fix pack can be installed in any order over the Tivoli Workload Scheduler network, but the feature can only be activated after every agent has been upgraded to Fix Pack 4 or later.

Configuration

To enable the feature, follow these steps:

1. Open the `<HOME>/mozart/globalopts` file on the Master
2. Add the line **enable switch fault tolerance = yes** at the bottom
3. Save and close
4. The feature will be active from the following **Jnextday** job.

In the End-to-End functionality, a new parameter has been introduced in the TOPOLOGY statement, `ENABLESWITCHFT (Y/N)`.

The `ftbox` is a cyclical message queue where each Full Status agent stores the messages it would send if it acted as a domain manager. The definition of "cyclical" means that when the queue fills up, the new messages will overwrite the old ones, starting back from the beginning.

It is possible to change the `ftbox` size as any other message queue, using the `evtsize` command.

Use of the Feature

The `showcpus` command standard format displays the linking status in a different way from before because the link is established in several directions.

The old L flag (for Linked) may be replaced by the new F flag (for Fully linked, showing all the open connections, both primary and secondary).

The semantic of the `link` command has changed, in the sense that if you request to link a workstation, other links will also be open. Namely, the linking/unlinking of a domain manager implies the opening or closing of the secondary connections towards the Full Status agents in the domain manager's domain.

If the main link is broken, the secondary links are terminated.

In case a secondary link is opened on an agent which is not initialized yet, the link will be terminated and writer will stay active for the timeout duration. This is a normal behavior.

For example, in a network like the one shown in Figure 3 on page 6

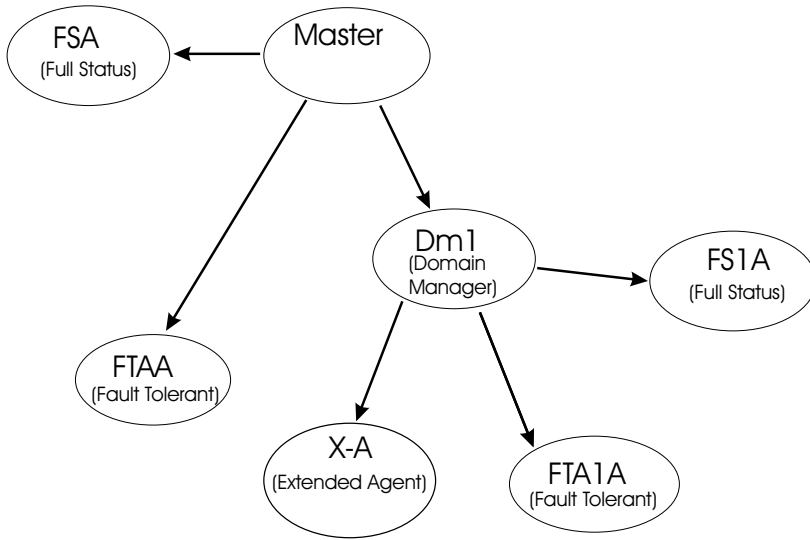


Figure 3. Example of network

the output of the **sc** command from the different workstations, in an ideal situation, is like the following:

Agent	Traditional	Enhanced	Agent	Traditional	Enhanced
sc command from Master			sc command from DM1		
Master		F	Master	LT W	FT W
FSA	LT W	FT W	FSA		
FTAA	LT W	FT W	FTAA		
DM1	LT W	FT W	DM1		F
FS1A	L	F	FS1A	LT W	FT W
FTA1A	L	F	FTA1A	LT W	FT W
X-A	L	L	X-A	L	L
sc command from FSA			sc command from FS1A		
Master	LT W	FT W	Master	L	F W
FSA		F	FSA		
FTAA	L	F W	FTAA		
DM1	L	F W	DM1	LT W	FT W
FS1A	L	F	FS1A		F
FTA1A	L	F	FTA1A	L	F W
X-A	L	L	X-A	L	L
sc command from FTAA			sc command from FTA1A		
Master	LT W	FT W	Master		
FSA			FSA		
FTAA		F	FTAA		
DM1			DM1	LT W	FT W
FS1A			FS1A		
FTA1A			FTA1A		F
X-A			X-A		

The Job Scheduling Console displays the information corresponding to the Full flag, without the possibility of filtering on it (as is the case with the Linking flag).

In a network where all the connections are correctly open, all the Fault Tolerant Agents should have an F flag.

The switch manager operation can be performed only if the feature is correctly working, and no error messages have been logged, by simply entering **switchmgr** from the command line.

To restore the correct functionality of the feature after receiving an error message, see “Troubleshooting” on page 11.

New demgr option

A new option, *demgr*, has been added to the conman command **start**, only for local workstations.

This option prevents the opening of external connections during the transition time between an agent starts as an old domain manager, and when the **switchmgr** command is run, depriving the agent of the domain manager function.

This option is run automatically, but as long as the old domain manager has not processed the **switchmgr** event (in case, for example, of delayed restart or restart after repairing a damaged agent), the *demgr* option **must** be used to start the old domain manager from the local command line.

Messages

New Messages

This is the list of the new messages introduced with the feature.

AWSBAT001I Event Counter successfully initializing:
<Message>

Explanation: Information message to report that the event counter is initializing.

AWSBAT002I Event Counter successfully initializing.
The workstation counter on the hard drive updated : <Message>

Explanation: While initializing its event counter, the process found that the file on the hard drive pointed to an old Symphony run number. The file has been updated. This message should be shown once for each jnextday.

AWSBAT003E Event Counter initialization failed:
<Error code>

Explanation: For the reason detailed in the Error message, the Event Counter initialization failed to initialize.

Operator Response: This error prevents the “Enable switch fault tolerance” feature from functioning as designed. Correct the problem reported, and restart the process. For more information, see the event counter troubleshooting section.

AWSBAT004E Due to the following problem, the event counter is no longer running:
<Error code>

Explanation: Event Counter processing encountered a problem. The counter has been deactivated and the

fault tolerance feature may not be able to run.

Operator Response: This error prevents the “Enable switch fault tolerance” feature from functioning as designed. Correct the problem reported, and restart the process. For more information, see the event counter troubleshooting section.

AWSBAT005E workstation insertion error, error code:
<System error>.

Explanation: Initialization of the event counter failed while trying to fill the event counter table with the workstation names stored into the Symphony file. The event counter is no longer running. The “Enable switch fault tolerance” feature may not be able to run.

Operator Response: This is an internal error. For more information, see the event counter troubleshooting section. If the problem persists, call IBM Support.

AWSBAT006I The event counter table was successfully filled. Initialization is complete.

Explanation: The event counter was successfully initialized and the table filled. The event counter is now working correctly.

Operator Response: No user action is needed.

AWSBAT008E Could not retrieve value from counter.

Explanation: A value could not be retrieved from the counter. Either the counter is damaged or the searched

workstation does not exist. The event counter will no longer work.

Operator Response: This is an internal error. For more information, see the event counter troubleshooting section in the product documentation. If the problem persists, call IBM Support.

**AWSBAT011W Event Counter received 'old'
id-number <ID number>; value updated
on counter.**

Explanation: The process had previously received a message coming from the same workstation but with a higher id number.

Operator Response: This is not an error. The message may be displayed while switch manager executes or in presence of some other Tivoli Workload Scheduler error. The Fault Tolerance feature is working.

AWSBAT012E Could not set value in counter.

Explanation: The process could not write the received id value in the event counter table.

Operator Response: This is an internal error. All the events routed through this workstation from the time this message is displayed until the time the process is restarted will be ignored by the "Enable switch fault tolerance" feature. See the event counter troubleshooting section in the product documentation. If the problem persists, call IBM Support.

**AWSBAT015E Run into<Error message> while trying
to save the event counter table.**

Explanation: An error occurred while trying to save the event counter table.

Operator Response: This error prevents the "Enable switch fault tolerance" feature from functioning as designed. Correct the problem reported, and restart the process. For more information, see the event counter troubleshooting section in the product documentation.

AWSBAT050I File instance created on the hard drive.

Explanation: During its initialization, the process found no saved instances of the event counter. A new instance has been created.

Operator Response: This is an information message. No action is needed

AWSBAT051I Counter table filled from file.

Explanation: During its initialization, the process copied the records for the event counter table from the version stored into EventCounter file.

Operator Response: This is an informational message. No action is needed.

**AWSBAT052I Symphony run-number of process
bigger than filed.**

Explanation: During its initialization, the process found that the file located on the hard drive referred to an old Symphony run number.

Operator Response: This is an informational message. No action is needed. This message is normally shown once for every jnextday.

**AWSBAT053I Symphony run-number of process
smaller than the saved one.**

Explanation: The process found that the file on the hard drive referred to a larger Symphony run number than the one stored in the process.

Operator Response: Either an occurrence of jnextday took place while the process was running, without the process being re-started, or someone manually changed the Symphony file or the EventCounter file. In the latter case, either replace the original version of the file or delete the EventCounter file. In the latter case, either reinstate the original version of the Symphony file or delete the EventCounter file before restarting the process. Events routed through this workstation since the last jnextday up to when the process is restarted will be ignored by the "Enable switch fault tolerance" feature. For more information, see the event counter troubleshooting section in the product documentation. If the problem persists, call IBM Support.

AWSBAT054I Failed to open existing counter file.

Explanation: A saved version of the event counter exists, but the process was not able to open the file.

Operator Response: Delete the event counter file and restart the process. Events routed through this workstation since the last jnextday up to when the process is restarted will be ignored by the "Enable switch fault tolerance" feature. For more information, see the event counter troubleshooting section in the product documentation. If the problem persists, call IBM Support.

**AWSBAT055I Failed to create a new event counter
file (and none are present).**

Explanation: A saved version of the event counter file does not exist, but the process is not able to create one.

Operator Response: If an event counter file does not exist, check that enough free disk space is available that the security permissions on the installation directory are correct. Restart the instance to reset the "Enable switch fault tolerance" feature. Events routed through this workstation since the last jnextday up to when the process is restarted will be ignored by the "Enable switch fault tolerance" feature.

If an event counter file does exist, call IBM support.

AWSBAT056E Operating System error number:
<Error code>.

Explanation: An operating system error occurred.

Operator Response: An operating system error occurred. Fix the error before trying to restore the "Enable switch fault tolerance" feature. For more information, see the event counter troubleshooting section in the product documentation.

AWSBAT057I The event counter file is corrupted.

Explanation: A saved version of EventCounter exists, but seems to be corrupted.

Operator Response: Delete EventCounter file and restart the process. Events routed through this workstation since the last jnextday up to when the process is restarted will be ignored by the "Enable switch fault tolerance" feature. For more information, see the event counter troubleshooting section in the product documentation.

AWSBAT058I Some fields hold incorrect values.

Explanation: Initialization of the event counter failed. The event counter is no longer working.

Operator Response: Restart the process. All the events routed through this workstation since the last jnextday up to when the process is restarted will be ignored by the "Enable switch fault tolerance" feature. For more information, see the event counter troubleshooting section in the product documentation.

AWSBAT059I Some fields have not been filled.

Explanation: Initialization of the event counter failed. The event counter is no longer working.

Operator Response: Restart the process. All the events routed through this workstation since the last jnextday up to when the process is restarted will be ignored by the "Enable switch fault tolerance" feature. For more information, see the event counter troubleshooting section in the product documentation.

AWSBCV125E Error writing to fbox for <Ftbox name> for <Agent>, Error: <Error code>.

Explanation: This error occurs when the program encounters a problem writing into an fbox queue. The possible fbox queues are ftup and ftdown. They are located inside the fbox directory. **Agent** is the local agent.

Operator Response: Verify that the fbox directory was correctly created and that it has access rights to the agent displaying the error. Verify that no program was requesting an exclusive access to that file at the time of the error.

AWSBCV126E Error: <Error code>. The event counter is out of service.

Explanation: Due to problems encountered while processing the event counter table, the event counter process was stopped.

Operator Response: Restart the process. For more information, see the event counter troubleshooting section in the product documentation.

AWSBCV127W Warning:<Warning>. X-agent message sent with id-number=0.

Explanation: An event with id-number 0 was sent from the workstation hosting an X-agent. This situation may inhibit the correct replay of events on this X-agent if a switch manager command is run later.

Operator Response: This is an internal error. Call IBM support for assistance.

AWSBCV128E Error while sending ACK message back to writer on <Agent>.

Explanation: During the replay protocol with the writer on Agent, mailman could not acknowledge the reception of a portion of the event counter table from the remote writer. As a consequence, the replay protocol was interrupted.

Operator Response: This is an internal error. Call IBM support for assistance.

AWSBCV129E Error sending NACK message back to writer on <Agent>.

Explanation: During the replay protocol with the writer on Agent, mailman could not communicate that a problem was encountered by mailman while transferring the event counter table from the remote writer. Because of this error, the replay protocol was interrupted.

Operator Response: Verify that on the other side writer does not have any problem loading the event counter table. Otherwise call IBM support for assistance.

AWSBCV130I Starting replay protocol with <Agent>.

Explanation: Mailman started the first phase of the replay protocol. The remote event counter table is being transferred.

AWSBCV131I Finished replay protocol with <Agent>.

Explanation: Mailman completed the last phase of the replay protocol, consisting in the transmission of the missing messages to the remote agent. A file named (agent name)_replayed was created in the fbox directory.

AWSBCV132E Could not create replay file for <Agent>.

Explanation: At the end of the replay protocol, mailman could not create an empty file with name (agent name)_replayed inside the ftbox directory. The replayed file indicates that the replay protocol was already run for that agent and prevents the protocol from restarting next time the link is established.

Operator Response: Verify that no problem due to incorrect access rights prevents mailman from writing into the ftbox directory. If you fix this problem, and can ascertain that the state of the remote agent has been synchronized by the mailman, you can then create the replayed file manually.

AWSBCV133W Could not truncate ftbox <Ftbox name>.

Explanation: During plan creation the ftdown message file is truncated (i.e the reader and writer pointers to the file are set to the beginning). For some reason, the process cannot perform this task.

Operator Response: Verify that the ftbox directory was correctly created and that it has the access rights on the agent displaying the error. Verify that no program was requesting an exclusive access to that file at the time of the error.

AWSBCV134W The replay table could not be sent. The link to <Agent> will proceed with a reduced replay protocol.

Explanation: This is the explanation message.

Operator Response: This is an internal error. Call IBM support for assistance.

AWSBCV135W Mailman is in isolated state and will not open any outbound connections.

Explanation: This agent is an old domain manager that still has to perform demotion from domain manager to full status FTA. During this phase, it is not safe for this agent to connect with other agents in the network, and the agent was restarted in isolated mode. This state will continue until the switchmanager command is processed.

This warning explains the error that generally precedes in the log about a failed link attempt.

Operator Response: Wait for the old manager to process the switchmanager command and give up its status of domain manager. If for urgent reasons the old manager must connect with the rest of the network, please stop and restart the agent. It will restart in non-isolated state.

AWSBCV137W An event loss was likely detected at start of replay protocol on event !1 from !2.

Explanation: This warning indicates a potential loss of a message during the replay protocol. After exchanging the counter table, the first message drawn from the fault tolerance queue (ftbox) is passing through the filter, indicating that maybe other previous messages were lost. This message might appear only after a switch manager.

Operator Response: If this message occurs after Jnextday, this is a normal condition. If it occurs during the day, it is likely to have a loss of messages, so a check on the state of the jobs run on the agent displaying the message or on other agents under it in the hierarchy is suggested. To avoid reproduction of this error, the administrator should tune the length allocated to the ftbox.

AWSBCW048E Decompress Symphony block with size: <I1>

Explanation: This is the explanation message

Operator Response: This is an internal error. Call IBM support for assistance.

AWSBCW049E Decompressed Symphony block size is: <I1>

Explanation: This is the explanation message

Operator Response: This is an internal error. Call IBM support for assistance.

AWSBCW052E Error while sending the event counter table header: <Error code>.

Explanation: During the exchange of the event counter table, it was not possible to transfer the header correctly. The error message more accurately describes the problem.

Operator Response: If this was due to a network problem, the connection will be attempted again according to configuration parameters.

AWSBCW053E Error sending event counter table values: <Error code>.

Explanation: While transmitting the event counter table, it was not possible to transfer the table values correctly. The error message more accurately describes the problem.

Operator Response: If this was due to a network problem, the connection will be attempted again according to configuration parameters.

AWSBCW054E Error while mailman sends confirmation for the event counter table:
<Error code>.

Explanation: During the exchange of the event counter table, writer did not receive the confirmation message it was expecting from mailman. The process was stopped.

Operator Response: If this was due to a network problem, the connection will be attempted again according to configuration parameters.

AWSBDY201E Unable to build directory <Directory name>. **Error =** <Error code>

Explanation: The ftbox directory could not be created.

Operator Response: Check if *twuser* is able to create the directory.

AWSBDY202E Unable to change access rights to directory <Directory name>. **Error =** <Error code>

Explanation: The access rights to the ftbox directory could not be changed.

Operator Response: Check if *twuser* has the authority to change the access rights to the ftbox directory.

AWSBDY203E Unable to change the owner of directory <Directory name>. **Error =** <Error code>

Explanation: The owner of the ftbox directory cannot be changed.

Operator Response: Check if *twuser* has the authority to change the ownership of the ftbox directory.

Troubleshooting

This section describes how to address the potential problems related to the use of the feature.

Event counter

The messages displayed in the log file concerning the event counter table can be divided into 3 sections:

1. "Everything right" messages. They concern the counter initialization. No action is needed.
2. "old-id" messages. They can spot other problems not concerning event counter problems. They reveal that the workstation received two messages switched in order.
3. "counter-breaking" messages. They are displayed when severe failure happens. The counter is not working anymore. User action is needed to restore the counter.

Two processes can display this kind of message:

1. **Writer**

From the moment the message is displayed, the counter is not working anymore. The id of the messages received from the writer (i.e. from the workstation which asked netman for writer's activation and from all its children) will not be "stored" anymore. This can lead to two situations:

- The workstation talking with the writer is switched to new manager. In this case the new manager will ask for a counter table and receive a "corrupt" counter table. The replay protocol will proceed following the default behavior.
- Before switchmanager, Writer goes down and is restarted. In this case the counter mechanism "heals" itself (at least partially). The id of the new messages received by the process will be stored into the counter, but the messages received by the Writer from the moment the message has been displayed up to the point in which it was restarted have not been tracked. If any message coming from workstation A has been received during this time gap, the situation may be considered re-set only when the new writer receives a message from workstation A. The situation will be recovered after *Jnextday*.

2. **Mailman**

From the moment the message is displayed, the counter is not working anymore. The id of the messages coming out from workstation A will all be set to 0. Messages coming out from that workstation will not be tracked anymore and therefore may be duplicated during replay protocol in any place in the network.

The situation will be recovered after Jnextday.

When the new master contacts a workstation, the **Replay protocol default behavior** is the following:

- If that workstation position into the workstation tree is **higher** than the new master's (the workstation is either the DM or a FS member of the father domain of the domain where the switch manager took place), **all** messages stored into ftbox, and that, according to the new tree structure, should be sent to that workstation, are actually sent there.
- If that workstation position into the workstation tree is **lower** than the new master's (the workstation either belongs to the domain where the switch manager took place and it is not the new domain manager or is the DM or a FS of one of the child domains), **no** message stored into ftbox is sent to that workstation.

Ftbox

If, on a Full Status agent, you receive any message concerning the ftbox, the feature is not working properly in the instance. **Do not** make this agent the new domain manager.

To restore the correct functionality of the feature on the instance, solve the problem as described in the message, and restart the agent.

Performance

This section describes the impact that the installation and use of this new feature has on the performance of the general architecture and the single system. It impacts:

- Connections
- Network traffic
- Disk space

Connections

To increment fault tolerance, the basic tree structure of the connections inside a Tivoli Workload Scheduler network has been changed so that no single connection interruption breaks the network. Obviously these are additional connections, and the number of connections opened by an agent is greater with this new feature.

From a point of view of the network security structure (firewall, SSL connections), all these new connections need to be opened and prepared in the same way as the connection to the domain manager.

These connections are in any case necessary for the backup domain manager to be able to act as a real domain manager after a switch manager command. The only difference brought by this feature is that the connections are always open.

Network Traffic

The network traffic is unchanged under normal conditions, and increased during the replay phase, according to the user's choice and only under special conditions.

Under normal conditions, the outbound reliability does not create any additional network traffic, because the messages are only stored for an eventual replay operation. The multiple inbound connections part does not generate additional traffic because the traffic that was previously copied by the domain manager to the FS member is now copied to the FS members directly by the FTAs.

During the replay phase, the connection protocol initiated by mailman on the backup domain manager includes a new phase for the replay of messages not sent by the failed domain manager. The impact of the message replay might be important, depending on the number of messages "trapped" inside the old domain manager.

Disk Space

There are two points in the network where disk space use increases following the activation of the additional fault tolerance.

These points are:

- On the single FTA. Here, in addition to the *tomaster.msg* queue, new queues are created for the other full-status FTAs. These queues should not be of any concern, because the impact on a single agent is small.
- On the full-status FTAs acting as backup domain managers. Here the new ftboxes for up and down traffic are created. These boxes contain all the traffic upward for *ftbox_up* and downward for *ftbox_down*, flowing through the corresponding domain manager.