Tivoli® Federated Identity Manager
Version 6.2.1

Troubleshooting Guide

IBM®

GC27-2715-00

Tivoli® Federated Identity Manager
Version 6.2.1

Troubleshooting Guide

IBM®

GC27-2715-00

Note
Before using this information and the product it supports, read the information in "Notices" on page 61.

This edition applies to version 6, release 2, modification 1 of IBM Tivoli Federated Identity Manager (product number 5724-L73) and to all subsequent releases and modifications until otherwise indicated in new editions.

Contents

About this publication

IBM® Tivoli® Federated Identity Manager Version 6.2.1 implements solutions for federated single sign-on, Web services security management, and provisioning that are based on open standards. IBM Tivoli Federated Identity Manager extends the authentication and authorization solutions provided by IBM Tivoli Access Manager to simplify the integration of multiple existing Web solutions.

This guide describes how to troubleshoot IBM Tivoli Federated Identity Manager.

Intended audience

The target audience for this book includes network security architects, system administrators, network administrators, and system integrators. Readers of this book should have working knowledge of networking security issues, encryption technology, keys, and certificates. Readers should also be familiar with the implementation of authentication and authorization policies in a distributed environment.

This book describes an implementation of a Web services solution that supports multiple Web services standards. Readers should have knowledge of specific Web services standards, as obtained from the documentation produced by the standards body for each respective standard.

Readers should be familiar with the development and deployment of applications for use in a Web services environment. This includes experience with deploying applications into an IBM WebSphere® Application Server environment.

Publications

Read the descriptions of the IBM Tivoli Federated Identity Manager library, the prerequisite publications, and the related publications to determine which publications you might find helpful. The section also describes how to access Tivoli publications online and how to order Tivoli publications.

IBM Tivoli Federated Identity Manager library

The publications in the IBM Tivoli Federated Identity Manager library are:

- *IBM Tivoli Federated Identity Manager Quick Start Guide*

  Provides instructions for getting started with IBM Tivoli Federated Identity Manager.
- *IBM Tivoli Federated Identity Manager Installation Guide*

  Provides instructions for installing IBM Tivoli Federated Identity Manager.
- *IBM Tivoli Federated Identity Manager Configuration Guide*

  Provides instructions for configuring IBM Tivoli Federated Identity Manager.
- *IBM Tivoli Federated Identity Manager for z/OS Program Directory*

  Provides instructions for installing IBM Tivoli Federated Identity Manager on z/OS®.
- *IBM Tivoli Federated Identity Manager Administration Guide*

Provides instructions for completing administration tasks that are required for all deployments.

- *IBM Tivoli Federated Identity Manager Web Services Security Management Guide*

  Provides instructions for completing configuration tasks for Web services security management.
- *IBM Tivoli Federated Identity Manager Auditing Guide*

  Provides instructions for auditing IBM Tivoli Federated Identity Manager events.
- *IBM Tivoli Federated Identity Manager Error Message Reference*

  Provides explanations of the IBM Tivoli Federated Identity Manager error messages.
- *IBM Tivoli Federated Identity Manager Troubleshooting Guide*

  Provides troubleshooting information and instructions for problem solving.

You can obtain the publications from the IBM Tivoli Federated Identity Manager Information Center:

http://publib.boulder.ibm.com/infocenter/tivihelp/v2r1/index.jsp?toc=/com.ibm.tivoli.fim.doc_6.2.1/toc.xml

Prerequisite publications

To use the information in this book effectively, you should have some knowledge about related software products, which you can obtain from the following sources:

- IBM Tivoli Access Manager for e-business Information Center:

  http://publib.boulder.ibm.com/infocenter/tivihelp/v2r1/index.jsp?toc=/com.ibm.itame.doc/toc.xml
- IBM WebSphere Application Server Version 6.1 Information Center:

  http://publib.boulder.ibm.com/infocenter/wasinfo/v6r1/index.jsp

  You can obtain PDF versions of the IBM WebSphere Application Server documentation at:

  http://www.ibm.com/software/webservers/appserv/was/library/

Related publications

You can obtain related publications from the IBM Web sites:

- The IBM Tivoli Federated Identity Manager Business Gateway Information Center at
- *Enterprise Security Architecture Using IBM Tivoli Security Solutions*. This book is available in PDF (Portable Document Format) at http://www.redbooks.ibm.com/redbooks/pdfs/sg246014.pdf or in HTML (Hypertext Markup Language) at http://www.redbooks.ibm.com/redbooks/SG246014/
- *Federated Identity Management and Web Services Security with IBM Tivoli Security Solutions* (SG24-6394-01). This book is available in PDF at http://www.redbooks.ibm.com/redbooks/pdfs/sg246394.pdf or in HTML at http://www.redbooks.ibm.com/redbooks/SG246394/
- The Tivoli Software Library provides a variety of Tivoli publications such as white papers, datasheets, demonstrations, redbooks, and announcement letters. The Tivoli Software Library is available on the Web at: http://publib.boulder.ibm.com/tividd/td/tdprodlist.html

- The *Tivoli Software Glossary* includes definitions for many of the technical terms related to Tivoli software. The *Tivoli Software Glossary* is available at http://publib.boulder.ibm.com/tividd/glossary/tivoliglossarymst.htm

Accessing terminology online

The IBM Terminology Web site consolidates the terminology from IBM product libraries in one convenient location. You can access the Terminology Web site at http://www.ibm.com/software/globalization/terminology

Accessing publications online

IBM posts publications for this and all other Tivoli products, as they become available and whenever they are updated, to the Tivoli Information Center Web site at http://publib.boulder.ibm.com/infocenter/tivihelp/v3r1/index.jsp.

Note: If you print PDF documents on other than letter-sized paper, set the option in the **File → Print** window that allows Adobe® Reader to print letter-sized pages on your local paper.

Ordering publications

You can order many Tivoli publications online at http://www.ibm.com/e-business/linkweb/publications/servlet/pbi.wss.

You can also order by telephone by calling one of these numbers:

- In the United States: 800-879-2755
- In Canada: 800-426-4968

In other countries, contact your software account representative to order Tivoli publications. To locate the telephone number of your local representative, perform the following steps:

1. Go to http://www.elink.ibmlink.ibm.com/publications/servlet/pbi.wss.
2. Select your country from the list and click **Go**.
3. Click **About this site** in the main panel to see an information page that includes the telephone number of your local representative.

Accessibility

Accessibility features help a user who has a physical disability, such as restricted mobility or limited vision, to use software products successfully. With this product, you can use assistive technologies to hear and navigate the interface. You also can use the keyboard instead of the mouse to operate all features of the graphical user interface.

For additional information, see the "Accessibility" topic in the information center at http://publib.boulder.ibm.com/infocenter/tivihelp/v2r1/index.jsp?toc=/com.ibm.tivoli.fim.doc_6.2.1/toc.xml.

Tivoli technical training

For Tivoli technical training information, refer to the following IBM Tivoli Education Web site at http://www.ibm.com/software/tivoli/education.

Support information

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

Online
Go to the IBM Software Support site at http://www.ibm.com/software/support/probsub.html and follow the instructions.

IBM Support Assistant
The IBM Support Assistant (ISA) is a free local software serviceability workbench that helps you resolve questions and problems with IBM software products. The ISA provides quick access to support-related information and serviceability tools for problem determination. To install the ISA software, see the *IBM Tivoli Federated Identity Manager Installation Guide*. Also see: http://www.ibm.com/software/support/isa.

Troubleshooting Guide
For more information about resolving problems, see the *IBM Tivoli Federated Identity Manager Troubleshooting Guide*.

Conventions used in this book

This reference uses several conventions for special terms and actions and for operating system-dependent commands and paths.

Typeface conventions

This publication uses the following typeface conventions:

Bold

- Lowercase commands and mixed case commands that are otherwise difficult to distinguish from surrounding text
- Interface controls (check boxes, push buttons, radio buttons, spin buttons, fields, folders, icons, list boxes, items inside list boxes, multicolumn lists, containers, menu choices, menu names, tabs, property sheets), labels (such as **Tip:**, and **Operating system considerations**:)
- Keywords and parameters in text

Italic

- Citations (examples: titles of publications, diskettes, and CDs
- Words defined in text (example: a nonswitched line is called a *point-to-point line*)
- Emphasis of words and letters (words as words example: "Use the word *that* to introduce a restrictive clause."; letters as letters example: "The LUN address must start with the letter *L*.")
- New terms in text (except in a definition list): a *view* is a frame in a workspace that contains data.
- Variables and values you must provide: ... where *myname* represents....

`Monospace`

- Examples and code examples
- File names, programming keywords, and other elements that are difficult to distinguish from surrounding text
- Message text and prompts addressed to the user
- Text that the user must type

- Values for arguments or command options

Operating system-dependent variables and paths

This publication uses the UNIX® convention for specifying environment variables and for directory notation.

When using the Windows® command line, replace **$***variable* with **%** *variable***%** for environment variables and replace each forward slash (**/**) with a backslash (\) in directory paths. The names of environment variables are not always the same in the Windows and UNIX environments. For example, %TEMP% in Windows environments is equivalent to $TMPDIR in UNIX environments.

Note: If you are using the bash shell on a Windows system, you can use the UNIX conventions.

Chapter 1. Troubleshooting and support

The troubleshooting process, in general, requires that you isolate and identify a problem, then seek a resolution. For Tivoli Federated Identity Manager you can use a troubleshooting checklist to help you. If the checklist does not lead you to a resolution, you can collect additional diagnostic data and analyze it yourself or submit the data to IBM Software Support for analysis.

Troubleshooting topics for Tivoli Federated Identity Manager are organized according to the sequence of these steps:

1. Learn more about a symptom or feature.

   Before you can successfully troubleshoot a symptom, or a problem with a specific product feature, you need a basic understanding of that symptom or feature.
2. Follow the troubleshooting checklist for the appropriate feature or symptom.

   The troubleshooting checklist offers a series of questions to guide you through the process of isolating and identifying a problem. If the problem is known to IBM, the checklist guides you to a published fix, solution, or workaround.

   If the troubleshooting checklist has not led you to a resolution, continue to the next step.
3. Collect diagnostic data.

   This information explains how to gather the necessary information that you, or IBM Software Support, must have in order to determine the source of a problem.
4. Analyze diagnostic data.

   This information explains how to analyze the diagnostic data that you collected.

Chapter 2. Learning more about problem symptoms

The first step in the troubleshooting process is to learn more about the problem symptoms, or about the affected product feature.

The following topics can help you to acquire the conceptual information that you need to effectively troubleshoot problems with IBM Tivoli Federated Identity Manager:

About troubleshooting

Troubleshooting is a systematic approach to solving a problem. The goal is to determine why something does not work as expected and how to resolve the problem.

The first step in the troubleshooting process is to describe the problem completely. Without a problem description, neither you nor IBM can know where to start to find the cause of the problem. This step includes asking yourself basic questions, such as:

- What are the symptoms of the problem?
- Where does the problem occur?
- When does the problem occur?
- Under which conditions does the problem occur?
- Can the problem be reproduced?

The answers to these questions typically lead to a good description of the problem, and that is the best way to start down the path of problem resolution.

What are the symptoms of the problem?

When starting to describe a problem, the most obvious question is "What is the problem?" This might seem like a straightforward question; however, you can break it down into several more-focused questions that create a more descriptive picture of the problem. These questions can include:

- Who, or what, is reporting the problem?
- What are the error codes and messages?
- How does the system fail? For example, is it a loop, hang, crash, performance degradation, or incorrect result?
- What is the business impact of the problem?

Where does the problem occur?

Determining where the problem originates is not always easy, but it is one of the most important steps in resolving a problem. Many layers of technology can exist between the reporting and failing components. Networks, disks, and drivers are only a few components to be considered when you are investigating problems. The following questions can help you to focus on where the problem occurs in order to isolate the problem layer.

- Is the problem specific to one platform or operating system, or is it common across multiple platforms or operating systems?
- Is the current environment and configuration supported?

Remember that, even though one layer might report the problem, this does not mean that the problem originates in that layer. Part of identifying where a problem originates is understanding the environment in which it exists. Take some time to completely describe the problem environment, including the operating system, its version, all corresponding software and versions, and hardware information. Confirm that you are running within an environment that is a supported configuration; many problems can be traced back to incompatible levels of software that are not intended to run together or have not been fully tested together.

When does the problem occur?

Develop a detailed timeline of events leading up to a failure, especially for those cases that are one-time occurrences. You can most easily do this by working backward: Start at the time an error was reported (as precisely as possible, even down to the millisecond), and work backward through the available logs and information. Typically, you need to look only as far as the first suspicious event that you find in a diagnostic log; however, this is not always easy to do and takes practice. Knowing when to stop looking is especially difficult when multiple layers of technology are involved, and when each has its own diagnostic information.

To develop a detailed timeline of events, try to answer these questions:

- Does the problem happen only at a certain time of day or night?
- How often does the problem happen?
- What sequence of events leads up to the time that the problem is reported?
- Does the problem happen after an environment change, such as upgrading or installing software or hardware?

Responding to questions like this can help to provide you with a frame of reference in which to investigate the problem.

Under which conditions does the problem occur?

Knowing what other systems and applications are running at the time that a problem occurs is an important part of troubleshooting. These and other questions about your environment can help you to identify the root cause of the problem:

- Does the problem always occur when the same task is being performed?
- Does a certain sequence of events need to occur for the problem to surface?
- Do any other applications fail at the same time?

Can the problem be reproduced?

From a troubleshooting standpoint, the ideal problem is one that can be reproduced. Typically with problems that can be reproduced, you have a larger set of tools or procedures at your disposal to help you investigate. Consequently, problems that you can reproduce are often easier to debug and solve. However, problems that you can reproduce can have a disadvantage: If the problem is of significant business impact, you do not want it to recur! If possible, re-create the problem in a test or development environment, which typically offers you more flexibility and control during your investigation.

- Can the problem be re-created on a test machine?
- Are multiple users or applications encountering the same type of problem?
- Can the problem be re-created by running a single command, a set of commands, or a particular application, or a stand-alone application?

About connectivity problems

Connectivity problems typically involve multiple systems, including software, hardware, and communications. The best way to troubleshoot connectivity problems is through a process of elimination.

First, collect relevant data and determine what you know, what data you have not yet collected, and what paths you can eliminate. At a minimum, answer the following questions.

- Are the communication paths operational?
- Has the initial connection been successful?
- Is the problem intermittent or persistent?
- Have changes been made to the communication network that would invalidate the previous directory entries?
- Where is the communication breakdown encountered? For example, was the breakdown between the client and a server?
- Is the problem encountered only within a specific application?
- What can you determine by the content of the message and the tokens that are returned in the message?
- Are other systems able to perform similar tasks successfully? If this is a remote task, is it successful when performed locally?

Next, try to isolate the problem by answering the questions in the Chapter 3, “Troubleshooting checklist for Tivoli Federated Identity Manager,” on page 11.

About Tivoli Federated Identity Manager

The first step to troubleshooting a problem is to learn about the affected feature of the software.

You can learn more about Tivoli Federated Identity Manager from the following sources:

- The IBM Tivoli Federated Identity Manager information center:
  http://publib.boulder.ibm.com/infocenter/tivihelp/v2r1/index.jsp?toc=/com.ibm.tivoli.fim.doc_6.2.1/toc.xml

- IBM Tivoli Education for Tivoli Federated Identity Manager. See the Tivoli education catalog at http://www-306.ibm.com/software/tivoli/education/edu_prd.html
- *Federated Identity Management and Web Services Security with IBM Tivoli Security Solutions* (SG24-6394-01). This book is available in PDF at http://www.redbooks.ibm.com/redbooks/pdfs/sg246394.pdf or in HTML at http://www.redbooks.ibm.com/redbooks/SG246394/

About fixes and updates

If you encounter a problem with IBM Tivoli Federated Identity Manager software, first check the list of recommended updates to confirm that your software is at the latest maintenance level. Next, check the list of problems fixed to see if IBM has already published an individual fix to resolve your problem.

These lists are located at the Tivoli Support Web site for Tivoli Federated Identity Manager http://www.ibm.com/software/sysmgmt/products/support/IBMTivoliFederatedIdentityManager.html

Individual fixes are published as often as necessary to resolve defects in IBM Tivoli Federated Identity Manager. In addition, two kinds of cumulative collections of fixes, called fix packs and refresh packs, are published periodically for IBM Tivoli Federated Identity Manager, in order to bring users up to the latest maintenance level. You should install these update packages as early as possible in order to prevent problems.

To receive weekly notification of fixes and updates, subscribe to My Support e-mail updates. For more information, refer to "Receiving fix notifications" on page 37.

The following table describes the characteristics of each maintenance delivery vehicle.

Table 1. Maintenance types

Name	Characteristics
Fix	• A single fix that is published between updates to resolve a specific problem. • After you install a fix, test any functions that the fixed component has an impact on.
Fix pack	• A cumulative fix package that contains all fixes that have been published since the previous fix pack or refresh pack; a fix pack might also contain new fixes. • Fix packs increment the modification level of the product and are named accordingly, for example, 5.0.1 • A fix pack can update specific components, or it can update the entire product image. • During fix pack installation, all previously applied fixes are automatically uninstalled. • After you install a fix pack, you should regression-test all critical functions. • The most recent two fix packs are available for download (for example, 5.0.2 and 5.0.1). Earlier fix packs are not available.

Table 1. Maintenance types (continued)

Name	Characteristics
Refresh pack	• A cumulative fix package that contains all fixes that have been published since the previous fix pack or refresh pack, as well as new fixes. • A refresh pack typically contains new function, in addition to fixes, and it updates the entire product image. • Refresh packs increment the modification level of the product and are named accordingly, for example, 5.0.1. • During refresh pack installation, all previously applied fixes are automatically uninstalled. • After you install a refresh pack, you should regression-test all critical functions.

About messages

When you receive a message from IBM Tivoli Federated Identity Manager, you can often resolve the problem by reading the entire message text and the recovery actions that are associated with the message.

You can find the full text of messages, their explanations, and the recommended recovery actions by searching for the message identifier in the *IBM Tivoli Federated Identity Manager Error Message Reference*.

About performance problems and hangs

Performance problems arise in many different situations. A hang is one type of performance problem in which users wait for a response for an indefinite period of time. Troubleshooting techniques for hangs are similar to the techniques you would use for other performance problems.

Here are some examples of situations in which performance problems become evident:

- Query performance is slower than expected.
- The workload or a batch job is not completing as soon as expected, or a reduction in the transaction rate or throughput occurs.
- The overall system slows down.
- A bottleneck is suspected in one of the system resources such as CPU, I/O, or memory.
- Query or other workload processing is consuming more resource than is expected or available.
- One system is performing better than another.
- A query, application, or system hangs.

Hangs can be particularly difficult to troubleshoot because the symptoms often seem to match the symptoms of other problems. For example, if a user is waiting for a long time for a response from a query, that user might think that the system is hung. In many cases, the query might be extremely complex, and the system might also be heavily used at the time, so the system is not actually hung, but it is just very slow to respond. Also, during a severe system shutdown, a significant buildup of activity can result in most or all commands appearing to hang.

In addition to characterizing the problem correctly in terms of what the symptoms are (slowness, too much resource used, and so on), and where the symptoms are observed (in a query, application, system resource, and so on), you need several other pieces of information to put the problem in context.

Answer the following questions to quickly determine the best place to start looking for the cause of the performance problem.

1. When did the problem first occur?

   If the problem has been occurring for some time, you might be able to use historical data to find differences. This will allow you to focus on changes in system behavior and then focus on why these changes were introduced. You also need to consider whether any recent changes occurred, such as hardware or software upgrades, a new application rollout, additional users, and so on.
2. Is the performance issue constant or intermittent?

   If the poor performance is continual, check to see if the system has started to handle a larger workload or if a shared database resource has become a bottleneck. Other potential causes of performance degradation include increased user activity, multiple large applications, or removal of hardware devices. If performance is poor only for brief periods, begin by looking for common applications or utilities that run at these times. If users report that a group of applications are experiencing performance issues, you can begin your analysis by focusing on these applications.
3. Does the problem appear to be system-wide or isolated to Tivoli Federated Identity Manager or its components?

   System-wide performance problems suggest an issue outside of Tivoli Federated Identity Manager. Something at the operating system level might need to be addressed.
4. If the problem is isolated to one component, does one particular activity appear to be causing the problem?

   If one component seems to be causing the problem, you can evaluate whether users who are reporting that specific activity are experiencing a slowdown. You might be able to isolate the issue to one component and a specific activity.
5. Do you notice any common characteristics of the poor performance, or do the problems appear to be random?

   You should determine if any common functions are involved. If so, this suggests that these functions are a point of contention.

About traps, crashes, and abends

The terms *trap*, *crash*, and abnormal end (*abend*) are often used synonymously.

If Tivoli Federated Identity Manager cannot continue processing as the result of a trap, segmentation violation, or exception, it generates an error.

Most traps, crashes, and abends for Tivoli Federated Identity Manager result in an exception, which is included in the message log, and typically does not require a trace to be enabled in order for it to be reported. However, these errors can be recorded in a trace log, if you are instructed to enable trace logging by IBM Support personnel. If you open a problem report with IBM, you might need to provide the trace log for analysis.

Although Tivoli Federated Identity Manager can generate trace logs on demand, you should generate trace files only when IBM Software Support asks you to do so. Refer to "Trace logs" on page 43 for more information.

Chapter 3. Troubleshooting checklist for Tivoli Federated Identity Manager

The following questions can help you to identify the source of a problem that is occurring with Tivoli Federated Identity Manager.

1. Are your fixes and fix packs up to date?

   See "Obtaining fixes" on page 37 for more information.
2. Is the problem documented in the *Tivoli Federated Identity Manager Release Information*?

   See the release information in the Tivoli Federated Identity Manager information center: http://publib.boulder.ibm.com/infocenter/tiv2help/index.jsp
3. Does the IBM Knowledge Base contain additional information about the problem?

   See Chapter 6, "Searching knowledge bases," on page 39.
4. Are you receiving any error messages?

   See the *IBM Tivoli Federated Identity Manager Error Message Reference* for information about error messages.
5. Do the logs contain any messages about the problem?

   See "Message logs" on page 42 and "Trace logs" on page 43 for more information.
6. Does the problem occur while installing or uninstalling one of the following features?
   - Tivoli Federated Identity Manager or its components

     See the *IBM Tivoli Federated Identity Manager Installation Guide*.
   - Common Auditing and Reporting Service

     See the *IBM Tivoli Federated Identity Manager Auditing Guide*.
   - Tivoli Access Manager for e-business

     See the *IBM Tivoli Access Manager for e-business Problem Determination Guide*.
   - WebSphere Application Server

     See the installation troubleshooting topics in the IBM WebSphere Application Server information center at http://www.ibm.com/software/webservers/appserv/was/library/.
7. Does the problem occur when you are configuring Tivoli Federated Identity Manager?

   See "Tivoli Federated Identity Manager configuration issues" on page 16.
8. Does the problem occur when you are trying to run the administration console?

   See "Integrated Solutions Console issues" on page 24 and "WebSphere Application Server console issues" on page 26.
9. Does the problem involve a deployment failure?

   See "Error received about deployment operation" on page 32.
10. If you could not resolve the problem in the preceding steps, you will need to try to determine additional information about the location of the problem or conditions during which the problem occurs:
    - Did the problem occur during runtime processing?

  - Did it fail to connect?
  - Did it crash?
  - Did it have a performance problem such as slow response, or a "hang"?
  - Did it abend, trap, or throw a Java™ exception?
- Does the problem occur while you configure a specific function?
- Does the problem occur when you perform a specific task?

The answers to these questions might help you determine the location of the problem and assist you in locating additional information about the problem. For example, if the problem occurs during configuration of a specific function or performance of a specific task, you might find a solution in the documentation of that function or task.

If the checklist does not guide you to a resolution, you can collect additional diagnostic data. The additional data might be necessary to IBM Support personnel to help you continue troubleshooting the problem. Refer to Chapter 7, “Collecting data,” on page 41.

Chapter 4. Known problems and solutions

Common known problems and their solutions are described in the subsequent sections.

Installation issues

This topic describes issues associated with installing Tivoli Federated Identity Manager.

Federated Identity Manager installation fails on Solaris systems

Under rare circumstances, the Solaris volume management daemon (vold) might prevent a Tivoli Federated Identity Manager installation. After prompting for the installation parameters, the installation program cannot find the proper JVM.

Note: Mounting the Tivoli Federated Identity Manager ISO image using a loopback device does not exhibit this problem.

To fix this problem, you can either restart the volume management daemon (vold) or mount the CD-ROM manually.

To stop and restart vold:

1. Stop vold using the following command:
   /etc/init.d/volmgt stop
2. Start vold using the following command:
   /etc/init.d/volmgt start
3. Run the Tivoli Federated Identity Manager installation.

To manually mount the CD-ROM:

1. Stop vold using the following command:
   /etc/init.d/volmgt stop
2. Find your CD-ROM device using the following command:
   ls -la /dev/sr0
   This command will output the c#t#d#s2 device for the CD-ROM.
3. Mount the device using the following command:
   mount -F hsfs /dev/c#t#d#s2 *mount_point*
4. Run the Tivoli Federated Identity Manager installation.
5. Unmount the device after you install Tivoli Federated Identity Manager and restart vold using the following commands:
   umount *mount_point*
   /etc/init.d/volmgt start

Federated Identity Manager CLI commands do not work immediately after product installation

The Tivoli Federated Identity Manager command line interface (CLI) does not work immediately after the product is installed because the installation program does not stop and restart WebSphere Application Server. You must manually stop

and then restart WebSphere Application Server before you can begin using Tivoli Federated Identity Manager CLI commands.

Uninstallation issues

This topic describes issues associated with uninstalling Tivoli Federated Identity Manager.

Embedded WebSphere Application Server JAR files are deleted after a system restart

On the Windows platforms only, if you install Tivoli Federated Identity Manager and then uninstall and reinstall Tivoli Federated Identity Manager, restarting your system causes the embedded version of the WebSphere Application Server JAR files in the ewas\java\jre\lib directory, and other files not deleted during the uninstallation process, to be deleted after the system is restarted.

Note: This will happen only if:

- Tivoli Federated Identity Manager is installed in the same location as the files that are locked,

  -AND-

- The directory location is marked for deletion.

These JAR files should be deleted during uninstallation, but they could not be deleted because a process had locked the files. Therefore, these files are marked for deletion by the operating system at a later time, and a system restart causes the files to be deleted.

You must restart your system after a Tivoli Federated Identity Manager uninstallation is completed.

Federated Identity Manager fails to uninstall on AIX®

On AIX, if a failure occurs during the installation of Federated Identity Manager, sometimes the uninstaller executable is not placed in the Federated Identity Manger installation directory. If the uninstaller executable is not installed during product installation, follow the procedure described in this section to manually uninstall Federated Identity Manager on AIX. After product installation, the AIX fileset listing contains the following entry:

```
FIM620 6.2.0.0 C P ISMP installed entry (/opt/IBM/FIM)
```

Note: Product uninstallation can fail also if you do not use the same utility to install and uninstall Federated Identity Manager. For example, if you install the product using ISMP and then attempt to uninstall the product using `smitty` or `installp`, the following error can occur:

```
A lexical error occurred for 0e0febbe7616f2ea257999d5d7f41e9b 0.0.0.0
```

To remove Federated Identity Manager entries from the AIX fileset, you must use commands that administrate the Object Data Manager (ODM) database. ODM is the AIX software registry that is used to track software installations. Below are steps to remove Federated Identity Manager entries from ODM. These steps should be executed by an AIX administrator. If you need further assistance on removing ODM entries, contact IBM Software Support.

1. Determine which lexical entries from the ISMP installation remain in the ODM using the `lslpp` command, for example:

   ```
   lslpp -l > /tmp/out.txt
   ```

2. Look in the output for entries that are similar to the following 32-character alphanumeric strings; these are the lexical numbers. Below are examples of some lexical numbers:

   ```
   11317ad90fc7a29df5c0bb762f59634e
                                0.0.0.0  COMMITTED  ISMP installed entry
   1188c43b206f32f0aefda112c87f1ea1
                                0.0.0.0  COMMITTED  ISMP installed entry
   1db8e322a4c716d975d865f055754736
                                0.0.0.0  COMMITTED  For the Configuration
   ...
   1e2bb97d8e16691143874a5410a42701
                                0.0.0.0  COMMITTED  ISMP installed entry
   27a07408c198e30e990f03da160c11f4
                                0.0.0.0  COMMITTED  ISMP installed entry
   2c1e1f7d0a0bd2da6975d31d9fc5e446
                                0.0.0.0  COMMITTED  ISMP installed entry
   335d2c0fd3c337c1b360a5e67df1a4ff
                                0.0.0.0  COMMITTED  ISMP installed entry
   3362ee349a200978dc4e1d262a00a304
                                0.0.0.0  COMMITTED  ISMP installed entry
   35e0d7afda41130690897b3e3e26fe1c
                                0.0.0.0  COMMITTED  ISMP installed entry
   36acdfe900b4761a4571c2ba54e78281
                                0.0.0.0  COMMITTED  ISMP installed entry
   36defc93f537a680a64504555c8a3ab7
                                0.0.0.0  COMMITTED  ISMP installed entry
   ```

3. Back up your ODM /usr directory, for example:

   ```
   #cd /
   #tar -cvf /tmp/odm.tar ./usr/lib/objrepos ./etc/objrepos
   ```

4. Remove the entries from the /usr subdirectory of the ODM, as follows, using the `odmget` and `odmdelete` commands. When specifying the `odmget` command, take note of the lpp_id# value that is returned because you specify that value for the inventory and history parameters in the `odmdelete` command. In some cases, no (0) objects are deleted, which is OK because the ISMP entries might not be in every part of the ODM.

   ```
   #export ODMDIR=/usr/lib/objrepos
   #odmget -q name=32_digit_lex_val lpp
   #odmdelete -q name=32_digit_lex_val -o lpp
   #odmdelete -q lpp_name=32_digit_lex_val -o product
   #odmdelete -q lpp_id=# -o inventory
   #odmdelete -q lpp_id=# -o history
   ```

5. Remove the entries from the /etc subdirectory of the ODM, as follows, using the `odmget` and `odmdelete` commands. When specifying the `odmget` command, take note of the lpp_id# value that is returned because you specify that value for the inventory and history parameters in the `odmdelete` command.

   ```
   #export ODMDIR=/etc/objrepos
   #odmget -q name=32_digit_lex_val lpp
   #odmdelete -q name=32_digit_lex_val -o lpp
   #odmdelete -q lpp_name=32_digit_lex_val -o product
   #odmdelete -q lpp_id=# -o inventory
   #odmdelete -q lpp_id=# -o history
   ```

6. Run the `lslpp` command again. The `lslpp` output should now be free of the ISMP entries. If any entries remain, ensure that you ran all of the `odmget` and `odmdelete` operations correctly. Note that you must run these commands for *every* entry that needs to be deleted.

Tivoli Federated Identity Manager configuration issues

The following issues and solutions are related to the configuration of Tivoli Federated Identity Manager.

Do not add files or subdirectories to ITFIM_WSSM shared library

Do not add files or subdirectories to the ITFIM_WSSM directory. If you add files or subdirectories to the ITFIM_WSSM directory, they might appear in the class path and have adverse effects. The ITFIM_WSSM shared library affects the class path of all applications running on the server. The shared library class path includes the ITFIM_WSSM installation directory.

SSL property settings for the auditing service override other SSL property settings

The SSL properties that are used for an outbound HTTPS connection from an application server to a Common Auditing and Reporting Service server can override the SSL properties on outbound HTTPS requests from other applications deployed in the same WebSphere Application Server.

If the Common Auditing and Reporting Service WebSphere client sets the system properties for the trusted key store, these settings override other SSL clients running in the JVM if they require the use of the same system property settings. For example, if Tivoli Federated Identity Manager is configured to use an HTTPS connection to a Common Auditing and Reporting Service server, the trusted key store specified for that connection will be used on all other HTTPS outbound connections from that WebSphere Application Server. Subsequent HTTPS outbound requests can then fail and return the following exception:

javax.net.ssl.SSLHandshakeException: Received fatal alert: certificate_unknown

To fix this problem on WebSphere Application Server Version 6.0.2, add the CA certificate for other outbound HTTPS requests from the application server to the trusted key store that is used for the HTTPS connection to the Common Auditing and Reporting Service server. For example, if a Web service application running in an application server is configured to make an outbound connection to a Tivoli Federated Identity Manager trust service using HTTPS, add the CA Certificate for that HTTPS connection to the trusted key store used for the HTTPS connection to the Common Auditing and Reporting Service server.

To fix this problem on WebSphere Application Server 6.1, define dynamic outbound endpoint SSL configurations. From the WebSphere Administrative Console, click **Security** -> **SSL certificate and key management** -> **Dynamic outbound endpoint SSL configurations**. This window enables you to define an outbound connection that is based on the target hostname, port, and protocol, and specifies the SSL configuration to be used. You can define unique SSL properties for outbound connections to the Common Auditing and Reporting Service server and any other target systems of outbound SSL requests.

Do not use national language characters in federation names

Tivoli Federated Identity Manager may not function correctly when national language characters are included in the federation name. Using national language characters in federation names is not supported. Federation names are included in

URLs and different Web servers, and browsers support incompatible conventions for converting national language characters into URL components.

Use *only* the following characters in federation names:

- Alphabetic characters: A-Z, a-z
  (Use US-ASCII alphanumeric characters.)
- Numbers: 0-9

Limitations on alias mapping

An alias is created when you federate an service provider account with an associated account on an identity provider. At the end of the federate action, the alias service on the identity provider includes a secSelfAlias (value for user alias given to a partner), and the alias service on the service provider has a secPartnerAlias (value for user alias received from a partner).

Only one account on a service provider can be federated to an account on an identity provider. Multiple accounts on different service providers can be federated to the same account on the identity provider. The secSelfAlias generated for a partner contains, as part of the value, the ProviderId value of the partner (service provider) with which it is associated. The rest of the value of the secSelfAlias is the randomly generated alias. Consequently, if multiple aliases are assigned to the same ProviderId, no mechanism exists to determine which alias should be provided in an assertion generated for consumption by that partner.

For example, assume a user John Smith has accounts jsmith and jsmith1 on Service Provider 1, johns on Service Provider 2, and JohnSmith on the Identity Provider. Both jsmith and johns can be federated with the JohnSmith account, but only jsmith or jsmith1 can be federated with the JohnSmith account. Otherwise, the principal specified in the generated assertion by the identity provider is indeterminate, because there are multiple accounts from the same service provider that are federated with the same identity provider account.

You can update alias values by performing a NameID update. The update changes the secSelfAlias of the identity provider and the secPartnerAlias of the service provider. You can initiate this action at either the identity provider or the service provider.

Workarounds for spaces in path to keytab file for SPNEGO

If WebSphere Application Server is installed into a directory that contains a space in the name, such as the default installation directory on a Windows platform, SPNEGO authentication cannot be configured correctly. This problem applies to the embedded version of WebSphere Application Server and the stand-alone version of WebSphere Application Server.

Two ways to work around this problem are described below.

If SPNEGO authentication is already configured in the Tivoli Federated Identity Manager management console:

After configuring SPNEGO authentication using the Tivoli Federated Identity Manager management console, update the krb5.ini file to change the path to the Kerberos keytab file to a format that does not contain a space (for example, the Windows 8.3 format). The path to the krb5.ini file can be described as:

WAS_HOME/profiles/*WAS_PROFILE*/config/itfim/ *FIM_DOMAIN*/etc/krb5.ini

The krb5.ini file contains a line similar to:

default_keytab_name = C:/Program Files/ IBM/WebSphere/AppServer/profiles/AppSrv01/config/itfim/default/etc/krb5.keytab

Typically, you only need to change "Program Files" to "Progra~1", for example:

default_keytab_name = C:/Progra~1/ IBM/WebSphere/AppServer/profiles/AppSrv01/config/itfim/default/etc/krb5.keytab

After making this update, restart WebSphere Application Server. You need to reapply this change each time you modify the identity provider configuration settings in the Tivoli Federated Identity Manager management console.

If SPNEGO authentication is NOT already configured in the Tivoli Federated Identity Manager management console:

Specify the 8.3 formatted path the krb5.ini.template file *before* you configure SPNEGO authentication in the Tivoli Federated Identity Manager management console. The path might vary depending on your WebSphere Application Server installation directory.

The following directory contains the krb5.ini.template file:

FIM_install_root/etc/krb5.ini.template

Replace the following line in krb5.ini.template:

default_keytab_name = @KEYTAB@

with the fully-qualified 8.3 file name, for example:

default_keytab_name = C:/Progra~1/ IBM/WebSphere/AppServer/profiles/AppSrv01/config/itfim/default/etc/krb5.keytab

Making this change ensures that the file name is preserved each time you configure the identity provider configuration settings in the Tivoli Federated Identity Manager management console

Update permissions to write log files to the log directory

This topic is relevant only for customers using the Tivoli Federated Identity Manager plug-in for IIS on Windows systems.

In non-English locales, log files cannot be added to the *FIM_install_dir*\fimpi\log directory.

Using the cacls.exe command and the locale-specific name for the "NETWORK SERVICE" group, assign the proper permissions so that the log file can be written to the *FIM_install_dir*\fimpi\log directory.

Use the following command:

cacls.exe *FIM_install_dir*\fimpi\log /t /e /g *name_of_network_service_group*:F

where *name_of_network_service_group* is the translated name of the NETWORK SERVICE group for your locale.

Java 2 security cannot be enabled on a Tivoli Federated Identity Manager system

Java 2 security cannot be enabled on a Tivoli Federated Identity Manager system. The Tivoli Federated Identity Manager product will not function, and this configuration is not supported.

Some of the symptoms are:

- An error reporting that the console cannot communicate with the management service.

  ```
  FBTCON257E An error occurred communicating with the Management Service.
  Check the server log files for more information.
  ```

- Errors in the logs indicating a permission issue at startup.

  Stack trace:

  ```
  ava.security.AccessControlException:
  Access denied (java.util.PropertyPermission * read,write)
  at java.security.AccessController.checkPermission(AccessController.java:108)
  ```

- Console Runtime exceptions are:

  ```
  javax.management.RuntimeOperationsException:
  The "name" parameter cannot be null. at
  com.ibm.ws.management.AdminClientImpl.assertObjectNameValid(AdminClientImpl.java:287)
  ```

Note:

- The main purpose of Java 2 security is to protect the WebSphere Application Server container from untrusted Java code. Since the customer trusts IBM there is no technical reason to enable it.
- There is a significant impact on performance when this feature is used. Compare the cost of using this feature to the risk with the code.

Using ServiceName and PortType element to select a specific STS module chain

Depending on the values used within the ServiceName and PortType elements, there can be a mismatch with the following error in the trace:

```
FBTSTM015E Either no configured XPath selected a node from the request, or
the given TokenType or AppliesTo
```

There can also be different requests mapping to the same chain if an "*" is generically used on both fields.

It is necessary that the left side of the colon in the Service Name and Port Type:

- match the NSURI of the WS-Addressing namespace used for PortType and ServiceName elements in the request
- and the right side matches the value

Calling Java classes from XSLT in ITFIM 6.2

ITFIM 6.2 has been developed using OSGI framework model. ITFIM 6.2 has changed the way each module finds other classes. Calling Java from XSLT is relying on an implementation detail that is out of ITFIM control. This constraint is a limitation and is not officially supported. The correct and supported way to use custom Java classes in a mapping rule is to develop a custom module. The module is used in mapping mode in an STS chain.

DirectoryIntegratorSTSModule Configuration Parameters

This topic provides information about setting the following parameters:

- Assembly line handler pool size
- Number of wait threads
- Amount of time for threads to wait for an assembly line handler to become available

The Tivoli Directory Integrator assembly line handlers (ALH) are synchronous and not threadsafe. This limitation is the reason why a pool of them is created. This topic provides some more details about the WebSEAL and WebSphere Application Server configuration.

Set the pool size to be the anticipated number of concurrent clients expected for the system. Setting the pool size to a large number does not make any difference if there are blocking calls that the assembly line is making itself.

Similarly, there is no benefit in making the pool size bigger than the maximum number of threads in WebSphere or WebSEAL. The number of wait threads indicate how many threads you are prepared to have blocking in Tivoli Federated Identity Manager waiting for a free ALH to call the Tivoli Directory Integrator. If zero are indicated, threads do not block if an ALH is not immediately available, and returns with an error. If less than zero is indicated, threads block indefinitely. If greater than zero is indicated, they block for the specified number of seconds. Avoid setting the wait thread to a negative number to prevent causing other issues like hanging all the threads in WebSEAL.

Error when using Test Connection for ITFIM Service

When using ITFIM 6.2 on Linux®, you might encounter an issue when creating a domain. After installing Tivoli Federated Identity Manager 6.2 and selecting the Test Connection you might encounter the following error:

```
FBTCON313E  An error occurred while invoking the ITFIM Management Service.
The  Management Service may be unavailable.
```

This error can be due to having several iterations of installation and uninstall in the environment.

To resolve this problem, manually uninstall ITFIM from the environment. All the following commands are run from the wsadmin prompt.

To manually uninstall ITFIM from the environment:

1. Ensure that WebSphere Application Server is started before you run wsadmin.
2. Launch wsadmin using the following command:

   ```
   /<WAS_INSTALL_ROOT>/profiles/<server_name>/bin/wsadmin.sh -user  <user>
   -password <password>
   ```

   for example: `/opt/IBM/WebSphere/AppServer/profiles/AppSrv01/bin/wsadmin.sh -user  wasadmin -password passw0rd`
3. List the applications currently installed on WebSphere Application Server using the following command:

   ```
   wsadmin>$AdminApp list
   ```

Note: The Tivoli Federated Identity Manager console is not be listed here because it is installed as a system application on WebSphere.

4. Uninstall an application.
   - To uninstall an application (such as ITFIMManagementService and ITFIMRuntime) from a standalone WebSphere Application Server, run these commands in order:

     ```
     wsadmin>$AdminApp uninstall <application_name>
     wsadmin>$AdminConfig save
     ```

     The <application_name> can be determined by using the `$AdminApp list` command.
   - To uninstall an application from a WebSphere Application Server cluster, run the same command but from the wsadmin prompt of the deployment manager.
5. A node synchronization is necessary in a cluster in order to propagate configuration changes to the affected node or nodes. By default, this situation occurs periodically, as long as the node can communicate with the deployment manager. You can propagate changes explicitly from the wsadmin prompt of deployment manager using the following command:

   ```
   wsadmin>set Sync1 [$AdminControl completeObjectName
   type=NodeSync,node=myNodeName,*]  wsadmin>$AdminControl invoke $Sync1
   sync
   ```
6. Uninstall the fimconsole from WebSphere Application Server, using the following command:

   ```
   wsadmin>$AdminApp update isclite modulefile {-operation delete
   -contenturi itfim-fimconsole-e.war} wsadmin>$AdminConfig save
   ```
7. After you remove the ITFIMManagementService, ITFIMRuntime, and Fimconsole from WebSphere using wsadmin commands, delete the following directories from your file system:

   ```
   /opt/IBM/FIM  /opt/IBM/WebSphere/AppServer/systemApps/isclite.ear/itfim-
   fimconsole-*.war
   ```

Exception in SystemOut.log during ITFIM server start-up

When starting a node where the ITFIM runtime is deployed, the server start-up can fail and return the following exception:
`com.tivoli.am.fim.rte.config.impl.RuntimeConfigurationImpl setThreadSubject java.lang.NullPointerException`

The full stack trace exception in the Server `SystemOut.log` contains the following information:`13/11/09 13.24.25:502 CET] 00000036 RuntimeConfig I com.tivoli.am.fim.rte.config.impl.RuntimeConfigurationImpl setThreadSubject java.lang.NullPointerException at com.ibm.ws.security.auth.ContextManagerImpl.isWSSubject`

The problem is generated by a wrong ITFIM domain configuration that is using SSL to connect to the Deployment Manager. However, WebSphere Application Server security is not enabled in the WebSphere Application Server cell so traffic to port 8879 (SOAP port for Deployment Manager) is not using SSL.

To resolve this problem, disable the "WebSphere Global Security is enabled". This configuration can be done in the ISC console ITFIM Domain Management Service Endpoint configuration if WebSphere Application Server security is not enabled in the cell.

ITFIM Single Log out not working properly

The issue is caused by an incorrect junction configuration. The best way to configure WebSEAL is by using the tfimcfg.jar tool that creates the junction with the appropriate set-up and makes the needed configuration changes for EAI.

This is how a junction looks like:

```
object show /WebSEAL/ip-ip/FIM/
Name: /WebSEAL/ip-ip/FIM
Description: Object from host saml20sp3.
Type: 16 (Management Object)
Is Policy Attachable: Yes
Extended Attributes
Name: HTTP-Tag-Value
Value(s): ssn=ssn
name=name
email=email
user_session_id=user_session_id
Attached ACL: itfim_saml20ip2_nobody
Attached POP:
Attached AuthzRule:

Effective Extended Attributes
Protected Object Location: /WebSEAL/ip-ip/FIM
Name: HTTP-Tag-Value
Value(s): ssn=ssn
name=name
email=email
user_session_id=user_session_id
Effective ACL: itfim_saml20ip2_nobody
Effective POP:
Effective AuthzRule:
```

Ensure the session ID is being passed by the junction.

itfimcfg.jar also enables EAI for SOAP WebSEAL instance

When configuring WebSEAL SOAP instance using the Tivoli Federated Identity Manager configuration utility, the itfimcfg.jar also enables EAI authentication.

EAI is enabled on SOAP because the same WebSEAL instance can be used for SOAP back channel and single sign-on point of contact. An enhancement was added to WebSEAL to enable it to accept a session termination request directly from a Tivoli Federated Identity Manager server. WebSEAL can accept a session termination request rather than through the PDadmin application programming interface. This feature requires the WebSEAL functionality to be configured for EAI configuration.

An EAI trigger must be set up on all URLs which can receive a log-out request. Since the log-out can be initiated using SOAP, the configuration tool sets an EAI trigger on the SOAP endpoint.

In the case where the SOAP endpoint is hosted on a different WebSEAL server from the one handling the user sessions, Tivoli Federated Identity Manager detects that different servers are being used. Tivoli Federated Identity Manager then returns to using the PDadmin application programming interface to perform the logout. In this case the EAI configuration on the SOAP endpoint is not useful and can be manually disabled.

Console does not show the new cluster member in the Runtime Nodes table

The console does not show the new cluster member in the Runtime Nodes table after it has been added to the cluster.

To display the new cluster member:

1. Edit the software.properties in the `/pkg` subdirectory where Tivoli Federated Identity Manager is installed.
2. Update the `com.tivoli.am.fim.rte.software.serialId` to a different value.

   For example, change:
   `com.tivoli.am.fim.rte.software.serialId=1197396816285` to
   `com.tivoli.am.fim.rte.software.serialId=1197396816286`
3. Log in to the management console.
4. Click **Tivoli Federated Identity Manager** -> **Domain Management** -> **Runtime Node Management**. When the portlet is opened it checks if the `software.properties serialId` has been updated.
5. Click the Deploy Runtime button.

After the deployment is completed, the node will be visible in the table.

LTPA Version 2 BinarySecurityToken ValueType not supported by WebSphere Application Server Versions 6.0.2 and 6.1

If you configure an LTPA Version 2 token module, the Web services layer does not distinguish between LTPA V1 and LTPA V2 if you are using WebSphere Application Server Version 6.0.2 or 6.1. If you send an LTPA token with the QName of `http://www.ibm.com/websphere/appserver/tokentype#LTPAv2` as the value type, the token is rejected in WebSphere Application Server Versions 6.0.2 and 6.1.

WebSphere Application Server Versions 6.0.2 and 6.1 do not distinguish between an LTPA V1 or LTPA V2 token in the Web Services layer. There is only one BinarySecurityToken ValueType supported for LTPA tokens and the QName of the value type is: `http://www.ibm.com/websphere/appserver/tokentype/5.0.2#LTPA`.

As a temporary workaround, you can continue to use an LTPA V1 or LTPA V2 token. However, be sure to set the QName to the following value type no matter what the token contains: `http://www.ibm.com/websphere/appserver/tokentype/5.0.2#LTPA`.

To enable the `http://www.ibm.com/websphere/appserver/tokentype/5.0.2#LTPA` type for LTPAv2 tokens set the following runtime custom property:

- To globally enable the token:
  - Key: ltpa.enable.compat.mode
  - Type: boolean Value: true or false
  - Description: Default is false. If this value is enabled, it ensures that the LTPA STS module issues tokens that are compatible with WAS 6.0.2 and WAS 6.1.
- To enable the token on a single STS Chain:
  - Key: ltpa.enable.compat.mode.[chainId]
  - Example Key: ltpa.enable.compat.mode.[uuid3778696c-0124-1fa7-9b85-be0cd9adb32a]
  - Type: boolean

– Value: true or false
– Description: Default is false. If this value is enabled, it ensures that the LTPA STS module issues tokens that are compatible with WAS 6.0.2 and WAS 6.1. This property enables for the configuration to be set on a specific chain.

Creating a Tivoli Federated Identity Manager domain might require a WebSphere Application Server restart

When creating aTivoli Federated Identity Manager domain (or a connection to a domain), if you specify inaccurate information in the security settings panel, WebSphere Application Server might have to be restarted. If you enter correct data and theTivoli Federated Identity Manager console successfully connects to the management service (use **Test Connection** to test the connection), you do not need to reconnect to WebSphere Application Server. If the Tivoli Federated Identity Manager console cannot connect to the Management Service, even if correct security information is supplied, then you need to restart WebSphere Application Server.

Integrated Solutions Console issues

The following issues and solutions are related to the use of the Tivoli Federated Identity Manager management console, which operates as a plug-in to the IBM Integrated Solutions Console.

Multiple prompts for restarting WebSphere Application Server

After performing an operation in the console, you might be prompted to restart WebSphere Application Server multiple times.

About this task

If you restart WebSphere Application Server when prompted but you do not allow enough time for the restart to complete before you try to resume using the console, you might be prompted to restart the WebSphere server again.

Procedure

1. Wait a sufficient amount of time after restarting WebSphere Application Server to allow the restart to complete. You can check the WebSphere Application Server logs to determine when the restart is complete.
2. When the restart completes, continue using the console.

Incomplete operations after logging off the console

If you log off the console before an operation completes, an incomplete result might occur.

About this task

Several operations, such as creating a federation, require the completion of several panels of information in the console. If you log off prior to clicking the **Finish** button on the final panel, the operation in progress is performed with the data currently entered. As a result, the operation, such a definition of a federation, might be incomplete. The following steps can prevent incomplete operations.

Procedure

1. Finish the operation in progress.

2. Then, log off from the console. If you must log off from the console in the middle of an operation:
   a. Log in to the console.
   b. Verify that the operation completed.
   c. If it did not, modify any values, as necessary.

Restart WebSphere Application Server after reconfiguring security settings in the management console

You must restart WebSphere Application Server if you set new security values in the management console.

The Tivoli Federated Identity Manager management console operates as a plug-in to the IBM Integrated Solutions Console, which operates as an application in WebSphere Application Server. WebSphere Application Server can optionally be configured to enable WebSphere Application Server security. When WebSphere Application Server security is enabled, the IBM Integrated Solutions Console and the Tivoli Federated Identity Manager management console must be configured to use the appropriate WebSphere Application Server security values (for example, the location of encryption files and the passwords needed to access the files).

Stop WebSphere Application Server from the command line after restarting Tivoli Federated Identity Manager from the console

This topic applies only to Windows systems. If you select restart in the Tivoli Federated Identity Manager management console, you must stop WebSphere Application Server service using the command line.

On Windows systems where WebSphere Application Server is running as a service, clicking Restart in the Tivoli Federated Identity Manager management console causes the service, as seen through the Windows Services panel, to lose contact with WebSphere Application Server. The restart of WebSphere Application Server is successful; however, the Services panel might not indicate that the service was restarted.

If this occurs, you cannot stop the WebSphere Application Server service through the Services panel, and Tivoli Federated Identity Manager management console can lose the connection to WebSphere Application Server. To stop the WebSphere Application Server service, use the command line to stop the service. Refer to the WebSphere Application Server Information Center for details on how to stop the service using the command line.

To work around this problem, do the following:

1. Restart WebSphere Application Server from the Tivoli Federated Identity Manager management console.
2. Log off the Tivoli Federated Identity Manager management console.
3. Log in to the Tivoli Federated Identity Manager management console.

As a more permanent workaround, you can install WebSphere Application Server on Windows without registering it as a service. When WebSphere Application Server is not registered as a service, the Tivoli Federated Identity Manager management console can restart it successfully by issuing the appropriate commands.

Limitation on Module Instances panel

When configuring the Trust Service in Tivoli Federated Identity Manager, the Module Instances panel uses the language of the locale of the server.

The Tivoli Federated Identity Manager domain is created based on the locale of the server. The language used in the server is based on the locale of the server. The language used on the server is implemented regardless of the language used on the browser connecting to the server.

For example, an English browser connects to a non- English Tivoli Federated Identity Manager console. The language seen on the Module Instances panel is that of the non-English language.

WebSphere Application Server console issues

The following issues and solutions are related to the use of the WebSphere Application Server administration console.

WebSphere Application Server must be installed separately for each instance of the management console

You can only install one management console component on each WebSphere Application Server installation.

The IBM Tivoli Federated Identity Manager management console component can only be installed to a single profile of a WebSphere Application Server installation. If you need to install the management console component multiple times on a single machine, you must create a new installation of WebSphere Application Server for each instance of the management console.

WebSphere administration console does not display the Runtime or Management application

If the WebSphere Application Server administration console does not display the Runtime or Management application, log out and log back in.

About this task

Click the Logout link on the console to logout correctly. Closing the Browser does not correctly logout a session. If logging out and logging back in does not clear the problem, use the following steps:

Procedure

1. Look for the appropriate EAR files:

Component	File name
Management Service	ITFIMManagementService.ear
Runtime	ITFIMRuntime.ear

Depending on the WebSphere Application Server environment, these files are in one of the following locations:

Server configuration	Default path
Single server	$WAS_PROFILE_HOME/installedApps/*cell_name/node_name/*

Server configuration	Default path
Cluster	$WAS_PROFILE_HOME/installedApps/cell_name

2. If you do not locate the EAR files, you might have an installation problem. Refer to the *Tivoli Federated Identity Manager Installation Guide*.

Configuration changes do not propagate to each cluster node before the Federated Identity Manager runtime restarts

The Federated Identity Manager runtime restarts before the specified configuration changes are propagated to each node in the cluster.

After making a Federated Identity Manager configuration change using the WebSphere Administrative Console, the following message and selection button are displayed:

```
FBTCON197W Recent configuration changes need to be
 reloaded to the Tivoli Federated Identity Manager runtime...

Load configuration changes to Tivoli Federated Identity Manager runtime
```

When you select the button, the Federated Identity Manager runtime begins restarting on each application server in the cluster.

Normally, the configuration changes should take only a second or two to update each Federated Identity Manager runtime. However, it is possible that, due to network delays, you click the button and start the runtimes before all changes have been propagated to each node in the WebSphere cell. If this occurs, you can start the Federated Identity Manager runtime again from the WebSphere Administrative console by selecting the following sequence of options: **Tivoli Federated Identity Manager**-> **Domain Management**-> **Runtime Node Management**-> **Reload Configuration**.

WebSphere console indicates that the management service is not available

The WebSphere Application Server administration console might incorrectly indicate that the IBM Tivoli Federated Identity Manager management service is not available.

About this task

The IBM Tivoli Federated Identity Manager management service runs as a WebSphere application. The Enterprise Applications page of the WebSphere Application Server administration console is used to view the availability of WebSphere applications. In a WebSphere Application Server Network Deployment (cluster) environment, the Enterprise Applications page might display the status of the management service as unavailable even though the application is running. The status is indicated by a gray icon with a slash through it.

Procedure

1. To verify the application's status, stop and start the application using one of the following methods:
   - Use the wsadmin command line interface.
   - Use the WebSphere Application Server console.

2. Check the console to see if the management service is displayed as available.

Querying the Tivoli Federated Identity Manager runtime status

It is not possible to query the status of the IBM Tivoli Federated Identity Manager runtime from the eWAS console. The following wsadmin commands show how to query the IBM Tivoli Federated Identity Manager runtime's status as well as how to start and stop the IBM Tivoli Federated Identity Manager runtime from the command line. These commands assume the WebSphere Application Server instance is named "server1".

- Determine whether IBM Tivoli Federated Identity Manager runtime is installed ("ITFIMRuntime" appears if it is):

  ```
  wsadmin>$AdminApp list
  ```

- Check whether IBM Tivoli Federated Identity Manager runtime is running (if no output is returned the Runtime is not running):

  ```
  wsadmin>$AdminControl queryNames
  type=Application,process=server1,name=ITFIMRuntime,*
  ```

- Stop IBM Tivoli Federated Identity Manager runtime:

  ```
  wsadmin>set appManager [$AdminControl queryNames
  type=ApplicationManager,process=server1,*]
  wsadmin>$AdminControl invoke $appManager stopApplication ITFIMRuntime
  ```

- Start IBM Tivoli Federated Identity Manager runtime:

  ```
  wsadmin>set appManager [$AdminControl queryNames
  type=ApplicationManager,process=server1,*]
  wsadmin>$AdminControl invoke $appManager startApplication ITFIMRuntime
  ```

Operational issues

This topic describes issues related to system operation.

Federated Single Sign-on (SSO) fails although user is defined in the point of contact directory

When using Tivoli Access Manager, a single sign-on attempt can fail if the WebSphere Application Server is specified as the point of contact server and the Tivoli Access Manager authorization server (pdacld) is not synchronized with the WebSphere Application Server user directory.

If you specify WebSphere Application Server as the point of contact, and the directory is not synchronized with the Tivoli Access Manager user registry, a service provider, while attempting to login a user during a single sign-on operation, will still call the Tivoli Access Manager authorization server (pdacld) to retrieve the user principal (for example, for a SAML11 federation). Although the user is defined in the WebSphere Application Server user directory, if Tivoli Access Manager is using a different directory and the user is not in that directory, the SSO attempt will fail with the following message:

```
HPDIA0202W   An unknown user name was presented to Access Manager.
```

Workaround:

When creating a domain for Federated Single Sign-on, the wizard prompts whether you want to configure into a Tivoli Access Manager environment. Specify to configure into a Tivoli Access Manager environment only if you are using

WebSEAL as the point-of-contact server, or if the Tivoli Access Manager user registry is synchronized with the specified point-of-contact user directory.

Intermittent failures in validating messages that are received

Intermittent failures in validating received messages can occur in the Federated Identity Manager environment if fix IY93387 is not installed for the JVM.

Errors in the trace file indicate a failure to validate the received XML. The following trace entry is an example:

```
[9/7/07 17:05:27:807 EST] 0000002d KessServiceJk 3 \
com.tivoli.am.fim.kess.service.jks.worker.impl.\
    KessServiceJksWorkerImpl validateXML The certificate\
      retrieved for xml signature validation was: [
[
  Version: V3
  Subject: CN=Centrelink Test Cert :7100000475, \
    O=Centrelink, L=CANBERRA, ST=ACT, C=AU
  Signature Algorithm: SHA1withRSA, OID = \
    1.2.840.113549.1.1.5

  Key:  IBMJCE RSA Public Key:
modulus:
141898594998619915567596662789190627967501703539604746432 69\
648282818399447366631175233283931353522781074184183075861 42\
873259432347557013082642779335744236865727648453965872003 30\
994380979671978481017289673087451941781181786204778418108 67\
011457005652616004635395030912776952386403821852272253222 70\
78677140534763
public exponent:
65537

  Validity: [From: Wed Nov 01 14:28:26 EST 2006,
              To: Sun Mar 30 11:46:30 EST 2008]
  Issuer: CN=TEST SecureNet Health OCA, O=SecureNet Limited,\
 C=AU
  SerialNumber: [1162351707]
<...removed for brevity ...>
[9/7/07 17:05:27:813 EST] 0000002d KessServiceJk 3 \
  com.tivoli.am.fim.kess.service.jks.worker.impl.\
  KessServiceJksWorkerImpl validateXML
Signature was NOT valid on the XML document.
Core Validity: false
SignedInfo: false
Msg: SignatureValue mismatched.
```

This problem is resolved by APAR IY93387, and fixed in the JVM 1.4.2 SR9 and JVM 5.0 SR4.

The Java SDK 1.4.2 SR9 fix level can be downloaded from the following Web site:

http://www-1.ibm.com/support/docview.wss?rs=180&uid=swg24011104

The Java SDK 1.5 SR5 fix level can be downloaded from the following Web site:

http://www-1.ibm.com/support/docview.wss?rs=180&uid=swg24015843

Disabling replay validation detection in a passticket

A timestamp is embedded within a passticket, but the time value interval is only granular to a full second. If two passtickets are generated for the same object (user, target app, secret-key) within one second, then the two passtickets will be identical, that is, the passtickets will look to the validator like a "replay attack." To manage

this problem, RACF® allows "disable replay detection," and this APAR enables IBM Tivoli Federated Identity Manager to support this functionality. To disable replay, you can set either or both of the following custom runtime properties:

```
passticket.disable.replay.check.[chainid_uuid]=true

passticket.disable.replay.check=true
```

here *chainid_uuid* is the value of the chain UUID. For example:

```
passticket.disable.replay.check.[uuideb42e428-011b-1ebc-a0cb-
9e6c4b35c1c7]=true
```

To determine the value of Chain UUID, in the administration console select **Trust Service Chains**-> **Select Action**, then select **Show Chain ID in column in table**. This action selection causes a new column to appear in the table that displays the unique Chain ID.

Unable to log in to WebSphere when using the VMM Tivoli Access Manager adapter

Symptom: Unable to log in as WebSphere administrator or unable to stop the WebSphere Application Server. An error message indicates a problem with initializing or contacting the Virtual Member Management (VMM) Tivoli Access Manager adapter registry. This error can occur when WebSphere Application Server security is enabled, and the adapter is configured to use WebSphere Federated Repositories.

About this task

The VMM Tivoli Access Manager Adapter uses the Tivoli Access Manager supported registries. This usage requires that the Tivoli Access Manager registry is available whenever the adapter is in use. The Tivoli Access Manager configuration file that is required by the VMM Tivoli Access Manager Adapter must also exist. The configuration file must be located in the exact location that you specified when you configured the adapter.

The problem with WebSphere login can occur if one of the following conditions is true:

- The configuration property for the VMM Tivoli Access Manager Adapter is deleted
- The Tivoli Access Manager registry server is not running.
- Tivoli Access Manager is configured.

When this problem occurs, you must recover access to the WebSphere Application Server.

Note: First, check the WebSphere Application Server log and determine if the underlying Tivoli Access Manager registry server is not running. If that is the case, restart the registry server, and verify that you can reach it from the WebSphere environment. If you are successful, you do not need to complete the task steps in this topic.

The following steps describe how to restore access to WebSphere Application Server. The steps apply only to configurations where two conditions are true:

1. WebSphere Application Server Federated Repositories is configured to contain more than one adapter. Note that an adapter is needed for each configured registry type.
2. The VMM Tivoli Access Manager adapter is one of the configured adapters in the WebSphere Application Server Federated Repositories.

Procedure

1. Stop the WebSphere Application Server.

   If you cannot log in, you might need to use an operating system command to halt the process.
2. Connect to WebSphere Application Server:

   ```
   wsadmin -conntype none
   ```
3. If there is a Base Entry configured for the VMM Tivoli Access Manager Adapter, delete it from the Realm.

   For example, if the Realm name is `defaultWIMFileBasedRealm`, and the adapter base entry is `o=ibm,c=us`:

   ```
   $AdminTask deleteIdMgrRealmBaseEntry {-name defaultWIMFileBasedRealm
   -baseEntry o=ibm,c=us}
   ```
4. Delete the VMM Tivoli Access Manager Adapter repository entry.

   For example, if the VMM Tivoli Access Manager Adapter registry ID is `TAMRegistryAdapter`:

   ```
   $AdminTask deleteIdMgrRepository {-id TAMRegistryAdapter}
   ```
5. Save the configuration.

   ```
   $AdminConfig save
   ```
6. Restart the WebSphere Application Server server.
7. If necessary, reconfigure the VMM Tivoli Access Manager Adapter.

Deployment issues

The following issues and solutions are related to deployment operations in IBM Tivoli Federated Identity Manager.

Deploying Tivoli Federated Identity Manager fails on z/OS

When deploying Tivoli Federated Identity Manager on z/OS, the actions hang and fail after a delay.

A limitation of the z/OS platform can cause Tivoli Federated Identity Manager actions to hang and fail. The failure can occur with the deployment of the Federated Identity Manager runtime. You can diagnose the failure by examining the WebSphere Application Server log file and looking for a WARNING message such as the following:

```
Trace: 2008/02/20 15:30:48.909 01 t=9BE748 c=UNK key=P8 \
    (13007002)
   ThreadId: 00000044
   FunctionName: com.ibm.ws.runtime.component.\
    ThreadMonitorImpl
   SourceId: com.ibm.ws.runtime.component.ThreadMonitorImpl
   Category: WARNING
   ExtendedMessage: BBOO0221W: WSVR0605W: Thread "WebSphere:\
    ORB.thread.pool t=009c22b8" (00000022) has been active \
    for 181010 milliseconds and may be hung.  There is/are 1\
    thread(s) in total in the server that may be hung.
```

To resolve this problem, define a WebSphere Application Server environment variable that increases an essential thread pool size.

To define the environment variable for a standalone application server in the administrative console, browse to **Servers** -> **Application servers** -> *server_name* -> **Server Infrastructure** -> **Administration** -> **Custom properties**. Add the private_bboo_internal_work_thread_pool_size property with a value of 5.

To define the environment variable for a network deployment configuration in the administrative console, browse to **System Administration** -> **Deployment manager** -> **Administration services** -> **Custom properties**. Add the private_bboo_internal_work_thread_pool_size property with a value of 5.

Restart the WebSphere Application Service server that has had the environment changed. To verify that the new value has taken effect, when the server starts, look for this message in the output of the server:

```
BBOM0001I private_bboo_internal_work_thread_pool_size: 5
```

Note: This failure has been reported only during the deployment of the Federated Identity Manager runtime, and the value of 5 has resolved the issue. However, if similar error messages are generated while performing other Federated Identity Manager activities, the pool size environment variable might need to be increased even further.

Error received about deployment operation

FBTCON137E: An error occurred during the deploy operation.

This message is a generic description of any deployment failure. You can receive this error even when the operation is successful but the operation took longer than the specified SOAP request timeout value.

After creating the domain and clicking **Deploy** from the Runtime Node Management Panel, you might receive this message. To validate the deployment, perform the following steps:

1. Close the Runtime Node Management panel.
2. Open the Runtime Node Management panel.

You should notice that the runtime shows as deployed with a check mark in the status column. You can now configure the runtime.

SOAP timing out

You can prevent the SOAP Client Connector from timing out during an operation by increasing the value of the `com.ibm.SOAP.requestTimeout` parameter to be 45 seconds longer than the time that it took for the operation to complete. To determine the appropriate timeout setting, perform the following operations:

1. Disable the timeout value by setting the value of the `com.ibm.SOAP.requestTimeout` parameter to 0.
2. Deploy the application and record how long it took.
3. Set the timeout value by setting the value of the `com.ibm.SOAP.requestTimeout` parameter to be the recorded time plus 45.

For example, if it took 3 minutes and 45 seconds to deploy the application, set `com.ibm.SOAP.requestTimeout` to 270 (4 minutes and 30 seconds).

To set the value of the com.ibm.SOAP.requestTimeout parameter, perform the following steps:

1. Stop the WebSphere Application Sever node where the administration console is installed.
2. Edit the $ISC/AppServer/properties/soap.client.props file.
3. Modify the value of the com.ibm.SOAP.requestTimeout parameter.
4. Save and close the file.
5. Start the WebSphere Application Sever node where the administration console is installed.

Deploying Tivoli Federated Identity Manager in a WebSphere vertical cluster environment

When deploying Tivoli Federated Identity Manager in a vertical cluster environment, deployment in the first WebSphere Application Server instance is successful, but deployment fails in the next server instance.

A WebSphere Application Server vertical cluster environment is created by hosting multiple application server instances on the same physical machine (node). All instances of the WebSphere Application Server use the same JVM environment.

If you attempt to deploy the Tivoli Federated Identity Manager server into multiple instances of WebSphere Application Server on the same system, the deployment fails against the second server instance because it is already configured in the JVM.

To bypass the problem, back up and remove the following files before configuring the IBM Tivoli Federated Identity Manager server into the second instance of WebSphere Application Server:

WAS_HOME/java/jre/PolicyDirector/PD.properties

WAS_HOME/java/jre/PolicyDirector/PDCA.ks

Repeat this procedure each time you deploy the Tivoli Federated Identity Manager server into a separate instance of WebSphere Application Server on the same node.

Deployment of Tivoli Federated Identity Manager in a cluster requires the discovery of node agents

In a WebSphere Application Server cluster deployment, the Federated Identity Manager runtime requires the discovery of a node agent to properly initialize.

Although the Federated Identity Manager runtime requires the discovery of a node agent to properly initialize, a node agent might not be available to the managed node server at startup. The following list includes key reasons why the node agent is not discovered upon startup:

Node agent was not started
The node agent was not started before the managed nodes were started. Node agents must be started before the managed nodes are started or they will not be discovered by applications at startup.

Performance issues
The Federated Identity Manager runtime uses a node agent discovery event to reinitialize its components. If the system or network is under

heavy load or response latencies are large for other reasons, this event might not be issued by WebSphere Application Server. As a result, the Federated Identity Manager runtime will be left in an improperly initialized state.

The following scenarios describe how a node agent can be detected:

- The Federated Identity Manager runtime is configured without error using the administrative console, or the runtime node is displayed as both deployed and configured using the Runtime Node Management page.
- With Federated Identity Management tracing enabled, you can find an entry that is logged by the com.tivoli.am.fim.fedmgr2.servlet.SSOPSServlet handleNotification() method that states:

  ```
  received mbean notification: <notification_message>
  ```

  where *notification_message* indicates a com.ibm.websphere.management.NotificationConstants event type.

Workaround:

If the node agent cannot be detected, reinitialize the Federated Identity Manager runtime by restarting the ITFIMRuntime EAR, or by restarting the managed node. Ensure that both the node agent and the deployment manager (dmgr) are running before attempting to start or restart the Federated Identity Manager runtime.

Customization issues

The following issues and solutions are related to the customization of Tivoli Federated Identity Manager, such as developing or modifying plug-ins.

Resolving ClassDefNotFoundErrors or ClassNotFoundExceptions

If you receive a class-not-found or class-definition-not-found error on a class that is provided by the WebSphere J2EE container, such as the ClassDefNotFoundErrors and the ClassNotFoundExceptions, ensure that the missing package of the class is listed in the Export-Package stanza of the com.tivoli.am.fim.osgi.connector's MANIFEST.MF file. You can change the MANIFEST.MF file that is located in $WAS_PROFILE/config/itfim/com.tivoli.am.fim.osgi.connector_6.2.1/META-INF/. Note that when involving methods in the WebSphere J2EE container, you might need to switch the context class loader.

Debugging OSGi bundle resolution and startup problems

Federated Identity Manager Eclipse runtimes can be launched in debug mode, which opens a port to the OSGi console (osgiConsole). You can use telnet to connect to the osgiConsole. The osgiConsole, enables you to query for the status of all the bundles in the Federated Identity Manager plug-ins directory and to perform life-cycle management operations on the bundles, for example, start, stop, install, and uninstall. Perform the following steps to launch a Federated Identity Manager Eclipse runtime in debug mode:

1. Navigate to directory where the launch.ini file is installed for the com.tivoli.am.fim.war.management.war:

   WAS_PROFILE/installedApps/*node*/ITFIMManagementService.ear/*com.tivoli.am.fim.war.management.war*/WEB-INF/eclipse/launch.ini

   or navigate to the com.tivoli.am.fim.war.runtime.war:

WAS_PROFILE/installedApps/*node*/ITFIMRuntime.ear/*com.tivoli.am.fim.war.runtime.war*/WEB-INF/eclipse/launch.ini

2. Edit the launch.ini file and specify a value for the property osgi.console.port (for example, osgi.console.port=8888), then save the file. If you are enabling debugging for both the management service and runtime on the same machine, ensure that you use different port values.
3. Restart the ITFIMManagementService EAR or the ITFIMRuntime EAR.
4. From a command prompt (or shell), run telnet localhost 8888 to access the OSGi console for the particular Federated Identity Manager Eclipse runtime. You should see an osgi> prompt after you establish a telnet connection to that port.
5. From the osgi> prompt, run the ss command to list the bundles in the Eclipse runtime and the status of each bundle. Run help to view a list of possible command options.
6. Find the bundle you introduced and the number listed next to it. If the bundle's status is INSTALLED and not RESOLVED or STARTED, you can try to run the command start *number* to print a Java exception stack trace that is associated with starting your customized bundle.
7. To exit from the OSGi console, run disconnect to back out from the port.

Locating the temp directory for OSGi error logs

To debug a Federated Identity Manager OSGi error, you might need to check the OSGi error logs. When Federated Identity Manager is started, the plug-ins, configuration, and features directories are copied from the configuration repository to a location under the WebSphere Application Server temp directory where the Federated Identity Manager OSGi instance is launched. This location is usually under either of the following directories:

WAS_APPSERVER/profiles/*profilename*/temp/*node*/*server*/ITFIMManagementService/com.tivoli.am.fim.war.management.war/itfim/

WAS_APPSERVER/profiles/*profilename*/temp/*node*/*server*/ITFIMRuntime/com.tivoli.am.fim.war.runtime.war/itfim/

The error logs from OSGi are located in the itfim/configuration directory under the temp directory.

Invoking a WebSphere API method from a custom module

This topic describes how to use the J2EEContainerAction in a custom module to invoke a WebSphere API method (such as JNDI lookup, RMI lookup, or SOAP MessageFactory).

A callback mechanism is available that allows a portion of Federated Identity Manager plug-in code to execute as if it was running in the J2EE container. This mechanism is necessary because some method calls rely on resources that are visible only to the class loaders in the J2EE container. Examples of method calls currently in Federated Identity Manager include: JNDI lookup, RMI endpoint lookup, SOAP MessageFactory creation, and JAAS login.

To use this mechanism, first create a class that implements com.tivoli.am.fim.osgi.J2EEContainerAction, which contains a run() method:

```
J2EEContainerAction myAction = new J2EEContainerAction(){
 public Object run() throws Exception{
  //do something that needs to run within a J2EE container
 }
};
```

Next, invoke your action class using com.tivoli.am.fim.osgi.J2EEContainerFactory:

```
J2EEContainerFactory.runInJ2EEContainer(myAction);
```

J2EEContainerFactory handles the contextclassloader switching from the Eclipse runtime to the J2EE container. It then calls the run() method in your action class, and then switches the contextclassloader back to the original. The following three commonly used J2EEContainActions are already defined, and you can simply reuse them:

```
com.tivoli.am.fim.j2eeactions.CreateMessageFactoryAction
\\creates a SOAP Message Factory
com.tivoli.am.fim.j2eeactions.JndiLookupAction
\\performs a JNDI lookup in WAS
com.tivoli.am.fim.j2eeactions.RmiLookupAction
\\performs a RMI lookup in WAS
```

Chapter 5. Fixes

You can obtain fixes from the product support Web site and sign-up for notifications of product support information, including fixes.

Obtaining fixes

A product fix might be available to resolve your problem. You can determine what fixes are available for Tivoli Federated Identity Manager by checking the product support Web site:

Procedure

1. Go to the IBM Software Support Web site for Tivoli Federated Identity Manager: http://www.ibm.com/software/sysmgmt/products/support/IBMTivoliFederatedIdentityManager.html. A list of most recent fixes is listed in the Download section of the page.
2. Click the name of a fix to read the description and optionally download the fix.

Receiving fix notifications

To receive e-mail notifications about fixes and other news about IBM products, follow these steps:

Procedure

1. Go to the IBM Software Support Web site for Tivoli Federated Identity Manager: http://www.ibm.com/software/sysmgmt/products/support/IBMTivoliFederatedIdentityManager.html.
2. Click **My Support** in the upper-right corner of the page. A sign-in page is displayed.
3. If you have already registered, skip to the next step. If you have not registered, click **Register now** to establish your user ID and password.
4. Sign in to **My support**.
5. Click the **Edit profile** tab.
6. Select **Software** → **Security** → **Access** in the fields that are displayed.
7. Select **IBM Tivoli Federated Identity Manager** from the list of products displayed.
8. Click **Add products**.
9. To enable e-mail notification, click **Subscribe to email** at the top of the page.
10. From the list, click **Software**.
11. Select the check boxes that best describe the e-mail notifications that you would like to receive.
12. Click **Update**.
13. Sign out of the session by clicking **Sign out** or click **Go to my personalized page** to see your personalized support page.

Chapter 6. Searching knowledge bases

IBM Support for Tivoli Federated Identity Manager maintains a knowledge base of technical articles, problems, and workarounds.

About this task

IBM Support Assistant includes a hierarchical search tool to help you focus your search for information related to a specific product, platform, or issue; refer to "ISA Overview" on page 55 for more information.

The following procedure describes how to perform a manual search for information.

Procedure

1. Go to the IBM Software Support Web site for Tivoli Federated Identity Manager: http://www.ibm.com/software/sysmgmt/products/support/IBMTivoliFederatedIdentityManager.html
2. Under **Solve a problem**, click either:
   - **Technotes**, which lists information about the product by article title.
   - **APARs**, which lists known problems by their Authorized Program Analysis Report numbers.
3. Optionally, you can search for specific terms, error codes, or APARs using the Search field on the product support page or the Technotes or APARs pages.

Chapter 7. Collecting data

Sometimes you cannot solve a problem simply by troubleshooting the symptoms. In such cases, you need to collect more diagnostic data.

Before you begin to collect data for a problem report, it is strongly recommended that you install and run the IBM Support Assistant. This troubleshooting tool includes a console that enables you submit an online problem management record (PMR). As part of the process, information specific to your system, environment and product is gathered into a file used by IBM Software Support. For more information on IBM Support Assistant, refer to "ISA Overview" on page 55.

Collecting data early, even before opening a problem management record (PMR), can help you to answer the following questions:

1. Do the symptoms match any known problems?
2. If so, has a fix or workaround been published?
3. Is this a non-defect-oriented problem that can be identified and resolved without a code fix?
4. Where does the problem originate?

The diagnostic data that you need to collect and the sources from which you collect that data are dependent on the type of problem that you are investigating. For example, if you are investigating a potential disk error in an AIX environment, one critical source of diagnostic data is the output from an errpt command.

For help identifying the component from which the problem originates, follow the questions in the troubleshooting checklist for Tivoli Federated Identity Manager.

Collecting general data

When you submit a problem to IBM Software Support, there is a base set of information that you typically need to provide details about the affected system or systems, such as:

- Version of Tivoli Federated Identity Manager and patch levels on affected systems
- Operating system name and version
- General details about the structure of your environment, such as number of servers and software installed, domains and federations configured, and so on

Collecting problem-specific data

For specific symptoms, or for problems in a specific part of the product, you might need to collect additional data, such as message and trace information. Refer to the topics in this chapter for more information. After you collect the appropriate diagnostic data, you can attempt to analyze the data yourself or you can provide it to IBM Software Support.

Message and trace logs

Tivoli Federated Identity Manager message and trace logs are managed and stored by WebSphere Application Server.

Refer to the troubleshooting topics in the WebSphere Application Server information center for detailed information about logs and logging.

Message logs

Message logs are text files in which the operations of the system are recorded.

The following types of messages are recorded by default:

Informational messages
Indicate conditions that are worthy of noting but that do not require you to take any precautions or perform an action.

Warning messages
Indicate that a condition has been detected that you should be aware of, but does not necessarily require that you take any action.

Error messages
Indicate that a condition has occurred that requires you to take action.

Message log files

All Tivoli Federated Identity Manager messages are logged in the following default WebSphere Application Server message logs.

Table 2. Message logs

Log	Default file name	Content
JVM Logs	SystemOut.log	Messages in text format for the application server instance.
IBM Service Log	activity.log	Messages in binary Common Base Event format for the application server installation. **Note:** Tools for viewing this format are provided with WebSphere Application Server. Refer to the WebSphere Application Server information center for more information.

Using the WebSphere Application Server administrative console, you can configure some settings of the logs, such as the location, name, maximum size of the log files and the levels of severity that you want to log (such as Warning and Severe). For more information, refer to "Configuring log settings" on page 44.

Message log locations

By default, the message logs are located in the following directories.

Table 3. Application server default message log locations

Log	Path
JVM Logs	**UNIX and Linux:** /opt/IBM/WebSphere/AppServer/profiles/*profile_name*/logs/*server_name*/SystemOut.log **z/OS:** /usr/lpp/WebSphere/V6R0/profiles/*profile_name*/logs/*server_name*/SystemOut.log **Windows:** C:\Program Files\IBM\WebSphere\AppServer\profiles*profile_name*\ logs*server_name*\SystemOut.log
IBM Service Log	**UNIX and Linux:** /opt/IBM/WebSphere/AppServer/profiles/*profile_name*/logs/activity.log **z/OS:** /usr/lpp/WebSphere/V6R0/profiles/*profile_name*/logs/activity.log **Windows:** C:\Program Files\IBM\WebSphere\AppServer\profiles*profile_name*\ logs*server_name*\activity.log

Console message logs are saved in the message log directories of the WebSphere Application Server node where the administration console is installed.

Trace logs

Trace logging, or tracing, provides IBM Software Support personnel with additional information relating to the condition of the system at the time a problem occurred.

In contrast to message logs, in which records are made of noteworthy events that have occurred, trace logs capture transient information about the current operating environment when a component or application fails to operate as intended. Trace logs are available in English only.

Trace logging is not enabled by default because in some circumstances it can cause large amounts of data to be collected in a short amount of time and might result in significant performance degradation. Therefore, you should enable trace logging only at the direction of IBM Software Support personnel. Refer to for more information.

Trace log entries can provide the following level of detail:

Fine Minimal detail.

Finer Moderate detail.

Finest Maximum (verbose) detail.

Trace log file

If tracing is enabled for an application server, Tivoli Federated Identity Manager trace information is logged in the following default WebSphere Application Server trace log.

Table 4. Trace log

Default log name	Default file name	Content
Diagnostic Trace	trace.log	Trace information in text format for the application server instance.

Using the WebSphere Application Server administrative console, you can configure some settings of the logs, such as the location, name, maximum size of the log files and the level of detail that you want to log (such as Fine, Finer, Finest). For more information, refer to "Configuring log settings."

Trace log locations

By default, the trace log is located in the following directories.

Table 5. Application server default trace log locations

Log	Path
Diagnostic Trace	**UNIX and Linux:** /opt/IBM/WebSphere/AppServer/profiles/*profile_name*/logs/*server_name*/trace.log **z/OS:** /usr/lpp/WebSphere/V6R0/profiles/*profile_name*/logs/*server_name*/trace.log **Windows:** C:\Program Files\IBM\WebSphere\AppServer\profiles*profile_name*\logs*server_name*\trace.log

Console trace logs are saved in the trace log directories of the WebSphere Application Server node where the administration console is installed.

Configuring log settings

Settings for message and trace logs can be configured using the administration console. Message logging is enabled by default. Trace logging should be enabled only at the direction of IBM Support personnel.

Configuring message logging

Message logging to the JVM log and the IBM Service log is enabled by default. Both logs are configured to log messages for all Tivoli Federated Identity Manager components of all severity levels. You can modify the names, location, file size, and severity level to be logged.

Configuring the JVM log

You can modify the file name, location, file format, file size, logging start and stop times, number of logs to keep, and severity level to be logged in the JVM Log.

About this task

The JVM log, also referred to as SystemOut.log, is a standard WebSphere Application Server log used for messages. For detailed information, refer to the JVM log topics in the WebSphere Application Server information center.

Procedure

1. Start the WebSphere Application Server administrative console and log in, if necessary.
2. Click **Troubleshooting** → **Logs and Trace** to open the Logging and Tracing page.
3. Click the name of the server that you want to configure, for example, server1.
4. Click **JVM Logs** to view the configuration options.
5. Select the **Configuration** tab.
6. Scroll through the panel to display the attributes to configure.
7. Change the appropriate configuration attributes and click **Apply**.
8. Save your configuration changes.

Configuring the IBM Service log

The IBM Service log is enabled by default. You can change this setting or modify the names, location, file size, and severity level to be logged in the log.

About this task

The service log, also referred to as the activity.log, is a standard WebSphere Application Server log used for messages. For detailed information about the log, refer to the service log topics in the WebSphere Application Server information center.

Procedure

1. Start the WebSphere Application Server administrative console and log in, if necessary.
2. Click **Troubleshooting** → **Logs and Trace** to open the Logging and Tracing page.
3. Click the name of the server that you want to configure, for example, server1.
4. Click **IBM Service Logs** to view the configuration options.
5. Select or clear the **Enable service log** box to enable or disable logging. The service log is enabled by default.
6. Set the name for the service log in the **File Name** field. The default name is activity.log. If the name is changed, the run time requires write access to the new file, and the file must use the .log extension.
7. Specify the number of megabytes to which the file can grow in the **Maximum File Size** field. When the file reaches this size, it wraps, replacing the oldest data with the newest data.
8. Click **Apply** to save the configuration changes.
9. Restart the server for the configuration changes to take effect.

Enabling trace logging

You can enable trace logging at server startup or on a running server.

About this task

Note: To maintain system performance, you should enable trace logging only at the direction of IBM Support personnel.

Enabling trace at server startup

Trace logging can be enabled at server startup.

About this task

The trace log is a standard WebSphere Application Server log used for trace information. For detailed information about the log, refer to the WebSphere Application Server information center.

Procedure

1. Start the WebSphere Application Server administrative console and log in, if necessary.
2. Click **Troubleshooting** → **Logs and Trace** to open the Logging and Tracing page.
3. Click the name of the server that you want to configure, for example, server1.
4. Click **Diagnostic Trace** to view the configuration options.
5. Click the **Configuration** tab.
6. Select or clear the **Enable log** box to enable or disable logging. The trace log is disabled by default.
7. Complete the configuration as instructed by IBM Support personnel. For additional information about the configuration settings, refer to the troubleshooting and trace log topics in the WebSphere Application Server information center.
8. Click **Apply** to save your configuration changes.
9. To enter a trace string to set the trace specification to the desired state:
   a. Click **Troubleshooting** → **Logs and Trace** to open the Logging and Tracing page.
   b. Click the server you want to configure, for example, server1.
   c. Click **Change Log Detail Levels**.
   d. If **All Components** has been enabled, you might want to turn it off and then enable specific components.
10. Click a component or group name. Refer to "Trace components" on page 47 for a list of Tivoli Federated Identity Manager components. **Note:** If the selected server is not running, it will not be displayed in the list.
11. Enter a trace string in the trace string box. For example, to specify tracing for only the trust server, enter: `*=info: com.tivoli.am.fim.trustserver.*=all` For more information about trace strings, refer to the WebSphere Application Server information center.
12. Click **OK**.
13. Click **Save** to save your changes. You might need to restart WebSphere Application Server for the change to become effective.

Enabling trace on a running server

Trace logging can be enabled on a running server.

About this task

The trace log is a standard WebSphere Application Server log used for trace information. For detailed information about the log, refer to the WebSphere Application Server information center.

Procedure

1. Start the WebSphere Application Server administrative console and log in, if necessary.
2. Click **Troubleshooting → Logs and Trace** to open the Logging and Tracing page.
3. Click the name of the server that you want to configure, for example, server1.
4. Click **Diagnostic Trace**.
5. Click the **Runtime** tab.
6. Select the **Save runtime changes to configuration as well** box if you want to write your changes back to the server configuration.
7. Change the existing trace state by changing the trace specification to the desired state.
8. Configure the trace output if a change from the existing one is desired.
9. Click **Apply**.

Trace components

The following table describes a partial list of the components you can specify for trace logging. **Note:** Enable tracing for specific components only as directed by IBM Support personnel. Depending on the symptoms of your problem, IBM Support personnel might suggest that you enable tracing for additional components that are not described here.

Table 6. Trace component names and descriptions

Trace component	Component description
com.tivoli.am.fim	All Tivoli Federated Identity Manager components
com.tivoli.am.fim.audit	Audit
com.tivoli.am.fim.fedmgr2	Single sign-on protocol service (SPS)
com.tivoli.am.fim.kess	Key service
com.tivoli.am.fim.liberty	Liberty single sign-on protocol
com.tivoli.am.fim.management	Console management
com.tivoli.am.fim.mgmt	Console management
com.tivoli.am.fim.saml	SAML single sign-on protocol
com.tivoli.am.fim.saml20	SAML 2.0 single sign-on protocol
com.tivoli.am.fim.soap	SOAP connections
com.tivoli.am.fim.sps	Single sign-on protocol service framework
com.tivoli.am.fim.trust	Trust service client
com.tivoli.am.fim.trustserver	Trust service
com.tivoli.am.fim.wsfederation	WS-Federation single sign-on protocol
com.tivoli.am.fim.wssm	Web services security management

Viewing logs

The format of the logs determines how they can be viewed.

JVM logs

To view the JVM logs, you can use the WebSphere Application Server administrative console, which supports viewing from a remote machine, or use a text editor on the machine where the log files are stored. Search on "viewing JVM logs" in the WebSphere Application Server information center for more information.

IBM service logs

The service logs are written in binary format. To view the log, you can use tools that are part of WebSphere Application Server. Search on "viewing the service log" in the WebSphere Application Server information center for more information.

Trace logs

Trace data is generated as plain text in basic, advanced or log analyzer format. On an application server, trace data can be directed to a file or an in-memory circular buffer. If the circular buffer is used, the data would need to be dumped to a file before it could be viewed.

On an application client or stand-alone process, trace data can be directed to a file or to the process console window. Search on "trace output" in the WebSphere Application Server information center for more information.

Using IBM Support Assistant

The IBM® Support Assistant Lite for Tivoli Federated Identity Manager tool aids troubleshooting of Tivoli Federated Identity Manager. Use the tool to automatically collect problem data.

You must install the Tivoli Federated Identity Manager plug-in for IBM Support Assistant as part of the Tivoli Federated Identity Manager installation. If you did not specify the IBM Support Assistant component when installing the product, install it now.

To use the tool, see:

- "Using the IBM Support Assistant in graphical mode"
- "Using the IBM Support Assistant in console mode" on page 49

Using the IBM Support Assistant in graphical mode

You can use a graphical user interface to collect data with IBM Support Assistant.

About this task

To access the graphical user interface, run a script from the command line.

Procedure

1. Ensure that your Java environment is configured correctly:
   a. Verify that your Java runtime environment is at level 1.4.2 or higher.
   b. Determine if the location of the Java runtime environment is included in your PATH environment setting. If the location is not included in your path, set the variable `JAVA_HOME` to point to the Java runtime environment.

Table 7. Specifying JAVA_HOME for your environment

Operating system	Sample command
Windows	For example, if you have a Java Development Kit installed at `c:\jre1.4.2`, use the command: `SET JAVA_HOME=c:\jre1.4.2`
UNIX or Linux	For example, if you are using the bash shell and you have a Java Development Kit installed at `/opt/jre142`, use the command: `export JAVA_HOME=/opt/jre142`

2. Start the IBM Support Assistant tool:

   Open a command window, and change directory to the ISAlite installation directory. The ISAlite installation directory is the location where you uncompressed the `TFIMISALite.zip` file. Enter the command for your environment:

Table 8. Running IBM Support Assistant

Operating system type	Command
Windows	runISALite.bat
UNIX or Linux	runISALite.sh **Note:** Ensure that the script is executable. If necessary, use the following command to change the file permissions: `chmod 755 runISALite.sh`

   The IBM Support Assistant now starts a graphical user interface.

3. In the Problem Type window, select a problem type.

   Expand the folders to display all problem types. Find your problem type and select it.

4. Supply a filename for the data collection ZIP file.

   You can use any filename. The tool automatically appends the ZIP file extension. For example, if you enter the filename `Install_problem`, the file is named `Install_problem.zip`.

5. Click **Collect Data**.

   The collection script runs and prompts you for additional information. The information can include configuration information or, the sequence of events leading to the problem. The script might also prompt you for preferences for data collection.

   When the scripts finishes collecting the setup information, it collects the necessary data. The tool creates a ZIP file that you can send to IBM Support.

6. When prompted, enter a filename in the **Output Filename/Path** box.

   The tool appends the server hostname and current timestamp to the filename that you entered.

7. Send the ZIP file to IBM Support

   You can choose FTP or HTTPS for file transfer. Note that FTP is unencrypted and HTTPS is encrypted.

Using the IBM Support Assistant in console mode

You can collect data with IBM Support Assistant in console mode.

About this task

Console mode provides command-line control of the IBM Support Assistant Lite collection scripts. The tool lets you record your responses from a console-mode session in a response file. You can then use the response file to drive subsequent executions of the same collection script.

Procedure

1. Ensure that your Java environment is configured correctly:
   a. Verify that your Java runtime environment is at level 1.4.2 or higher.
   b. Determine if the location of the Java runtime environment is included in your PATH environment setting. If the location is not included in your path, set the variable `JAVA_HOME` to point to the Java runtime environment.

Table 9. Specifying JAVA_HOME for your environment

Operating system	Sample command
Windows	For example, if you have a Java Development Kit installed at `c:\jre1.4.2`, use the command: `SET JAVA_HOME=c:\jre1.4.2`
UNIX or Linux	For example, if you are using the bash shell and you have a Java Development Kit installed at `/opt/jre142`, use the command: `export JAVA_HOME=/opt/jre142`

2. Start the IBM Support Assistant tool:

   Open a command window, and change directory to the ISAlite installation directory. The ISAlite installation directory is the location where you uncompressed the `TFIMISALite.zip` file. Enter the command for your environment:

Table 10. Running IBM Support Assistant

Operating system type	Command
Windows	runISALiteConsole.bat
UNIX or Linux	runISALiteConsole.sh **Note:** Ensure that the script is executable. If necessary, use the following command to change the file permissions: `chmod 755 runISALite.sh`

   The IBM Support Assistant now starts in console mode.
3. Create a response file.

Table 11. Syntax for recording data input for IBM Support Assistant

Operating system type	Command
Windows	runISALiteConsole.bat -record *response.txt*
UNIX or Linux	runISALiteConsole.sh -record *response.txt*

   You can specify your own filename for *response.txt*.

   When running in this mode, you supply data input during an interactive session. The tool records your responses into the file that you specify.
4. Run the tool using the response file.

Table 12. Syntax for using IBM Support Assistant with a response file

Operating system type	**Command**
Windows	runISALiteConsole.bat *response.txt*
UNIX or Linux	runISALiteConsole.sh *response.txt*

Notes®:

- The response file is a plain text file. You can edit it to modify values as needed. For example, you can use the file on another computer after adjusting the response file values to reflect settings for the local computer.
- Remember that sensitive information, such as user names and passwords, might be stored in the response file. Manage the file carefully, to prevent unauthorized access to important information.
- Some data collection sessions require interaction with the user, and thus are not suitable for the silent collection option. For example, IBM Support might ask you to reproduce a problem during data collection, in order to collect log and trace files. In this case, silent collection cannot record and reproduce all steps.

Chapter 8. Analyzing data

After you collect data from multiple sources, you need to determine how that data can help you to resolve your particular problem.

To analyze the data, take the following actions:

- Determine which data sources are most likely to contain information about the problem, and start your analysis there. For example, if the problem is related to installation, start your analysis with the installation log files (if any), rather than starting with the general product or operating system log files.
- Have a clear understanding of how the various pieces of data relate to each other. For example, if the data spans more than one system, keep your data well organized so that you know which pieces of data come from which sources.
- Confirm that each piece of diagnostic data is relevant to the timing of the problem by checking timestamps. Note that data from different sources can have different timestamp formats; be sure to understand the sequence of the different elements in each timestamp format so that you can tell when the different events occurred.

The specific method of analysis is unique to each data source, but one tip that is applicable to most traces and log files is to start by identifying the point in the data where the problem occurs. After you identify that point, you can work backward in time through the data in order to unravel the root cause of the problem.

If you are investigating a problem for which you have comparative data for a working and non-working environment, start by comparing the operating system and product configuration details for each environment.

Chapter 9. Contacting support

IBM Software Support provides assistance with product defects.

You can contact IBM Software Support in the following ways:

- **IBM Support Assistant:** This Web application enables you to search for information on a problem, collect the logged information required to troubleshoot a problem, and open a problem management report (PMR) online. The IBM Support Assistant is the recommended method for reporting a problem to IBM Software Support. You need your user account ID and password to submit a PMR using IBM Support Assistant. The IBM Support Assistant client is included on your product CD and updates can be downloaded from the following Web site:http://www.ibm.com/software/support/isa/
- **Online:** Go to the **Submit** and **track problems** tab on the IBM Software Support site at http://www.ibm.com/software/support/probsub.html. Type your information into the appropriate problem-submission tool.
- **By phone:** For the phone number to call in your country, go to the Contacts page of the *IBM Software Support Handbook* at http://techsupport.services.ibm.com/guides/contacts.html, and click the name of your geographic region.

If the problem that you submit is for a software defect or for missing or inaccurate documentation, IBM Software Support creates an Authorized Program Analysis Report (APAR). The APAR describes the problem in detail. Whenever possible, IBM Software Support provides a workaround that you can implement until the APAR is resolved and a fix is delivered. IBM publishes resolved APARs on the Software Support Web site daily, so that other users who experience the same problem can benefit from the same resolution.

Before you submit a problem to IBM Software support, answer these questions:

1. Do you have an active IBM software maintenance contract?
2. Do you understand the business impact of the problem?
3. Can you describe the problem?

ISA Overview

IBM Support Assistant simplifies the process of researching and reporting on software problems.

IBM Support Assistant provides quick access to support-related information along with serviceability tools for problem determination. IBM Support Assistant consists of three tools:

Search

Access troubleshooting repository information quickly, using multiple filters to focus your search. The concurrent search tool spans the bulk of IBM documentation and returns results that are categorized by source for easy review.

Product information links

These self help links include:

- Product support pages

- Product home pages
- Product troubleshooting guides
- Product education road maps and the IBM Education Assistant
- Product recommended updates
- Product news groups and forums

Service feature

The service feature is an automated system collector and symptom-based collector. The system collector gathers general information from your operating system, registry, and other sources. The symptom-based collector gathers specific product information relating to a particular problem that you are having. The service feature also enables you to automatically set tracing to help IBM support in the data gathering process.

The service feature also enables you to submit a problem online to IBM Software Support. Enter your entitlement information once and have it saved for future sessions. You can then create a problem report for IBM and attach the gathered data in the collector file.

IBM Support Assistant provides a complete online user guide to assist you in the setup and use of the tool. The following steps describe the basic procedure for setting up your system to use IBM Support Assistant:

1. Install the IBM Support Assistant tool from your product CD, or from the following Web site:

   http://www.ibm.com/software/support/isa/
2. In an IBM software repository, locate the IBM Support Assistant plug-in for the version of WebSphere Application Server that you are running. The Tivoli Federated Identity Manager plug-in uses the WebSphere Application Server plug-in as the base code.

   Note: Be patient downloading the WebSphere Application Server plug-in; it can take a long time depending on network traffic and the availability of system resources.
3. Install the WebSphere Application Server plug-in into the \plugin subdirectory of the IBM Support Assistant installation directory.
4. Locate the IBM Support Assistant plug-in for Tivoli Federated Identity Manager on the product CD or in an IBM software repository.
5. Install the IBM Support Assistant plug-in into the \plugin subdirectory of the IBM Support Assistant installation directory.
6. Start the IBM Support Assistant using your desktop icon. Select the **User Guide** tab for information on performing the various available tasks.

IBM software maintenance contracts

Before you submit a problem to IBM Software Support, ensure that your company has an active maintenance contract, and that you are authorized to submit problems to IBM.

If you are not sure what type of software maintenance contract you need, call 1-800-IBMSERV (1-800-426-7378) in the United States. From other countries, go to the Contacts page of the *IBM Software Support Handbook* at http://techsupport.services.ibm.com/guides/contacts.html, and click the name of your geographic region for phone numbers of people who provide support for your location.

Determining the business impact

When you submit a problem to IBM, you are asked to supply a severity level. Therefore, you need to understand and assess the business impact of the problem that you are reporting.

Use the following criteria:

Table 13. Severity levels

Severity 1	The problem has a *critical* business impact: You are unable to use the program, resulting in a critical impact on operations. This condition requires an immediate solution.
Severity 2	This problem has a *significant* business impact: The program is usable, but it is severely limited.
Severity 3	The problem has *some* business impact: The program is usable, but less significant features (not critical to operations) are unavailable.
Severity 4	The problem has *minimal* business impact: The problem causes little impact on operations or a reasonable circumvention to the problem was implemented.

Describing a problem

When describing a problem to IBM, be as specific as possible. Include all relevant background information so that IBM Software Support specialists can help you solve the problem efficiently.

To save time, know the answers to these questions:

- What software versions were you running when the problem occurred?
- Do you have logs, traces, and messages that are related to the problem symptoms?
- Can you re-create the problem? If so, what steps do you perform to re-create the problem?
- Did you make any changes to the system? For example, did you make changes to the hardware, operating system, networking software, or other system components?
- Are you currently using a workaround for the problem? If so, be prepared to describe the workaround when you report the problem.

Submitting data

You can send diagnostic data, such as log files and configuration files, to IBM Software Support.

Use one of the following methods:

- IBM Support Assistant
- FTP (EcuRep)
- ESR tool

IBM Support Assistant

IBM Support Assistant includes a service feature which has an automated system collector and a symptom-based collector. The system collector gathers general

information from your operating system, registry, and other sources. The symptom-based collector gathers specific product information relating to a particular problem that you are having. The service feature also enables you to automatically set tracing to help IBM support in the data gathering process. Refer to "ISA Overview" on page 55 for more information on IBM Support Assistant.

FTP (EcuRep)

To submit files using the FTP service called EcuRep, package the data files that you collected into ZIP or TAR format, and name the package according to your Problem Management Record (PMR) identifier. Your file must use the following naming convention in order to be correctly associated with the PMR:

```
xxxxx.bbb.ccc.yyy.yyy
```

where:

Table 14. File naming convention

xxxxx	PMR number
bbb	Branch, from the PMR identifier
ccc	Country code, from the PMR identifier
yyy.yyy	File type (ZIP or TAR format)

To transfer your files using FTP, complete these steps:

1. Using an FTP utility, connect to the emea.ibm.com server (for example, ftp.emea.ibm.com).
2. Log in as anonymous, and enter your e-mail address as your password.
3. Change directories to toibm (for example, cd toibm).
4. Change to one of the platform-specific subdirectories: aix, cae, hw, linux, lotus, mvs, os2, os400, swm, tivoli, unix, vm, vse, and windows.
5. Change to binary (bin) mode (for example, bin).
6. Put your file on the server. You can send but not update files on the FTP server; therefore, any subsequent time that you need to change the file, you need to create a new file with a unique name.

For more information about the EcuRep service, see IBM EMEA Centralized Customer Data Store Service at http://www.ibm.com/de/support/ecurep/index.html.

If your product runs in a z/OS environment and you want to compress your data sets, you can use the TRSMAIN utility, which you can download from the following Web page: ftp://ftp.software.ibm.com/s390/mvs/tools/packlib.

ESR tool

Registered users who are on an authorized caller list can submit diagnostic data using the Electronic Service Request (ESR) tool. The ESR tool enables you to submit and manage Problem Management Records (PMRs) on demand, 24 hours a day, seven days a week, 365 days a year.

To submit data using ESR, complete these steps:

1. Sign onto ESR.

2. On the Welcome page, enter your PMR number in the **Enter a report number** field, and click **Go**.
3. Scroll down to the **Attach Relevant File** field.
4. Click **Browse** to locate the log, trace, or other diagnostic file that you want to submit to IBM Software Support.
5. Click **Submit**. Your file is transferred to IBM Software Support through FTP, and it is associated with your PMR.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd
3-3-12, Roppongi, Minato-ku, Tokyo 106-8711 Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:
INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
2Z4A/101
11400 Burnet Road
Austin, TX 78758 U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not

been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Trademarks

IBM, the IBM logo, and ibm.com® are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at http://www.ibm.com/legal/copytrade.shtml

Adobe, Acrobat, PostScript® and all Adobe-based trademarks are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, other countries, or both.

Intel®, Intel logo, Intel Inside®, Intel Inside logo, Intel® Centrino®, Intel Centrino logo, Celeron®, Intel® Xeon®, Intel SpeedStep®, Itanium®, and Pentium® are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft®, Windows, Windows NT®, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Java and all Java-based trademarks are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Index

Printed in USA

GC27-2715-00