

IBM® Tivoli® Federated Identity Manager
Version 6.2.2

Troubleshooting Guide



IBM® Tivoli® Federated Identity Manager
Version 6.2.2

Troubleshooting Guide



Note

Before using this information and the product it supports, read the information in “Notices” on page 63.

Edition notice

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

© Copyright IBM Corporation 2006, 2011.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

About this publication v

Intended audience	v
Publications	v
IBM Tivoli Federated Identity Manager library	v
Prerequisite publications	vi
Related publications	vi
Accessing terminology online	vi
Accessing publications online	vii
Ordering publications	vii
Accessibility	vii
Tivoli technical training	vii
Support information	vii
Conventions used in this book	viii
Typeface conventions	viii
Operating system-dependent variables and paths	viii

Chapter 1. Troubleshooting and support 1

Chapter 2. Learning more about problem symptoms 3

About troubleshooting	3
About connectivity problems	5
About Tivoli Federated Identity Manager	5
About fixes and updates	6
About messages	7
About performance problems and hangs	7
About traps, crashes, and abends	8

Chapter 3. Troubleshooting checklist for Tivoli Federated Identity Manager. . . . 9

Chapter 4. Known problems and solutions 11

Installation issues	11
Uninstallation issues	13
Tivoli Federated Identity Manager configuration issues	13
Integrated Solutions Console issues	23
Multiple prompts for restarting WebSphere Application Server	23
Incomplete operations after logging off the console	24
Restart WebSphere Application Server after reconfiguring security settings in the management console	24
Stop WebSphere Application Server from the command line after restarting Tivoli Federated Identity Manager from the console	24
Limitation on Module Instances panel	25
WebSphere Application Server console issues	25
WebSphere Application Server must be installed separately for each instance of the management console	25

WebSphere administration console does not open the Runtime or Management application	25
Configuration changes do not propagate to each cluster node before the Tivoli Federated Identity Manager runtime restarts	26
WebSphere console indicates that the management service is not available	27
Querying the Tivoli Federated Identity Manager runtime status	27
Profiles with POST bindings for single sign-on might not work when using WebSphere Application Server 6.1.0.17 or 6.1.0.19.	28
Operational issues	28
Unable to log on to WebSphere when using the VMM Tivoli Access Manager adapter.	33
Deployment issues	34
Error received about deployment operation.	34
Deploying Tivoli Federated Identity Manager in a WebSphere vertical cluster environment	35
Deployment of Tivoli Federated Identity Manager in a cluster requires the discovery of node agents.	35
Customization issues	36

Chapter 5. Fixes 39

Obtaining fixes	39
Receiving fix notifications	39

Chapter 6. Searching knowledge bases 41

Chapter 7. Collecting data 43

Message and trace logs	44
Message logs	44
Trace logs	45
Configuring log settings	46
Configuring message logging	46
Enabling trace logging.	47
Viewing logs	49
Using IBM Support Assistant	50
Using the IBM Support Assistant in graphical mode	50
Using the IBM Support Assistant in console mode	51

Chapter 8. Analyzing data 55

Chapter 9. Contacting support 57

ISA overview	57
IBM software maintenance contracts	58
Determining the business impact	59
Describing a problem	59
Submitting data	59

Notices	63
Index	67

About this publication

IBM® Tivoli® Federated Identity Manager Version 6.2.2 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 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/topic/com.ibm.tivoli.fim.doc_6.2.2/ic/ic-homepage.html

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:

- Tivoli Access Manager Information Center:
<http://publib.boulder.ibm.com/infocenter/tivihelp/v2r1/index.jsp?toc=com.ibm.itame.doc/toc.xml>
- IBM WebSphere Application Server Version 8.0 Information Center:
<http://publib.boulder.ibm.com/infocenter/wasinfo/v8r0/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:

- *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/td/tdprodlist.html>

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 hard copies of some publications.

Many countries provide an online ordering service.

Follow these steps to access this service:

1. Go to <http://www-947.ibm.com/support/entry/portal/Documentation>
2. Select **IBM Publications Center** from **Getting Started**.
3. Select your country from **Select a country/region/language to begin** and click the arrow icon.
4. Follow the instructions for how to order hard copy publications on Welcome to the IBM Publications Center.

If your country does not provide an online ordering service, contact your software account representative to order publications.

Follow these steps to find your local contact:

1. Go to <http://www.ibm.com/planetwide/>
2. Click your country name to display a list of contacts.

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/topic/com.ibm.tivoli.fim.doc_6.2.2/ic/ic-homepage.html.

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. You can also 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 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 must 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. A problem description helps IBM to determine where to start in finding 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. A good problem description 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?" The question might seem straightforward, 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, fail, performance degradation, or incorrect result?
- How does the problem affect the business?

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 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 getting a report from one layer does not mean that the problem originated from 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 trace the problem timeline 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 must look only as far as the first suspicious event that you find in a diagnostic log. However, doing so 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 the preceding questions 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 must 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 significantly affects business, you do not want it to recur. If possible, recreate the problem in a test or development environment, which typically offers you more flexibility and control during your investigation.

- Can the problem be recreated on a test machine?

- Are multiple users or applications encountering the same type of problem?
- Can the problem be recreated 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 the task 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 9.

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_6.2.2/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 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 website 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. Two kinds of cumulative collections of fixes, called fix packs and refresh packs, are published periodically for IBM Tivoli Federated Identity Manager. Fix packs and refresh packs bring users up to the latest maintenance level. Install these update packages as early as possible to prevent problems.

To receive weekly notification of fixes and updates, subscribe to My Support email updates. For more information, see “Receiving fix notifications” on page 39.

The following table describes the characteristics of each fix.

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 might affect.
Fix pack	• A fix pack 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, conduct a regression-test for all of the 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.
Refresh pack	• A refresh pack contains all fixes that have been published since the previous fix pack or refresh pack, and 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 conduct a regression-test on all of the critical functions.

About messages

You can often resolve the problem stated in the warning or error message you get from IBM Tivoli Federated Identity Manager 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 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 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 processor, I/O, or memory.
- Query or other workload processing is using more resource than is expected or available.
- One system is performing better than another.
- A query, application, or system hangs.

Hangs can be 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 hung. In many cases, the query might be complex, and the system might also be heavily used at the time, so the system has not hung. But, the system might be slow to respond. Also, during a system shutdown, a significant buildup of activity can result in most or all commands seem to hang.

Aside from characterizing the problem correctly in terms of what the symptoms are, and where the symptoms are observed, you need several other pieces of information to put the problem in context. Symptoms typically range from slowness, too much resource used, and so on. The symptoms are typically found in a query, application, system resource, and so on.

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. Historical data helps you to focus on changes in system behavior and then focus on why these changes were introduced. You also must 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 if the system has started to handle a larger workload. You can also check 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 happen 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 must be addressed.

4. If the problem is isolated to one component, does one particular activity cause 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 does the problem occur in random?

Determine if any common functions are involved. The function involvement might indicate 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. The exception 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 must provide the trace log for analysis.

Generate trace files only when IBM Software Support asks you to do so, although Tivoli Federated Identity Manager can generate trace logs on demand. See “Trace logs” on page 45 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 39 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 41.
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 44 and “Trace logs” on page 45 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*.
 - IBM Tivoli Access Manager for e-business
See the *IBM Tivoli Access Manager for e-business Troubleshooting 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 13.
8. Does the problem occur when you are trying to run the administration console?
See “Integrated Solutions Console issues” on page 23 and “WebSphere Application Server console issues” on page 25.
9. Does the problem involve a deployment failure?
See “Error received about deployment operation” on page 34.
10. If you cannot resolve the problem in the preceding steps, gather 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 while configuring a specific function or performing 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. See Chapter 7, "Collecting data," on page 43 for details.

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.

Tivoli 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 correct JVM.

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

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

To stop and restart vold:

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

To manually put the CD:

1. Stop vold with the `/etc/init.d/volmgt stop` command.
2. Find your CD device with the `ls -la /dev/sr0` command.
This command produces the `c#t#d#s2` device for the CD.
3. Use the `mount -F hsfs /dev/c#t#d#s2mount_point` command to mount the device.
4. Run the Tivoli Federated Identity Manager installation.
5. Remove the device after you install Tivoli Federated Identity Manager and restart vold with the following commands:
`umount mount_point/etc/init.d/volmgt start`

Tivoli Federated Identity Manager generated installation response file does not work properly on a silent installation

The installation response file that is generated by using the `-record-options` installation option contains two extra `licenseAccepted` parameters that are set to false. The extra `licenseAccepted` parameters prevent the generated installation response file from working properly during a silent installation.

As a workaround, you must edit the generated installer response file and delete the extra instances of `-G licenseAccepted=false`. Ensure that there is only *one* `licenseAccepted` parameter left, and that it is set to true. For example: `-G licenseAccepted=true`.

Tivoli Federated Identity Manager command-line interface 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. Use the Tivoli Federated Identity Manager CLI commands to manually stop and then restart the WebSphere Application Server before you begin.

Fix pack installation script fails due to SOAP port mismatch

The fix pack installation of the Tivoli Federated Identity Manager runtime must connect to a WebSphere Application Server SOAP port to deploy the runtime. The fix pack installer acquires its SOAP port value from the following line in the /<-installation-directory>/etc/fim.appservers.properties file of the Tivoli Federated Identity Manager instance being patched:

```
was.soap.port=8880
```

OR

```
ewas.soap.port=8880
```

This value is set in the file when the Tivoli Federated Identity Manager instance is installed.

For the connection to be successful, the WebSphere Application Server instance to which it is being deployed must still be using that SOAP port. If it is not, then the Tivoli Federated Identity Manager fix pack installation fails in the WebSphere Application Server UPDI and the error is reported as:

Prerequisite checking has failed. Click Back to select a different package, or click Cancel to exit.

Associated failure messages are:

The WebSphere server does not seem to be listening in host localhost port 8881 as specified in /opt/IBM/FIM/etc/fim.appservers.properties. Make sure the server is running and that the specified port and host are correct.

If the specified port is different than the actual SOAP port used, change the value in the fim.appservers.properties file to work with the port being used by WebSphere Application Server. Then, reapply the fix pack.

The FIM Installer fails if a '\$' is used in a WebSphere Application Server JKS or in the password of a P12 file

It is a permanent restriction in the Tivoli Federated Identity Manager 6.1.1. The same code is used by the fix pack installer so the same restriction applies to fix pack installations.

The workaround for this restriction is documented in technote #1307656 entitled A \$ symbol in a truststore or keystore password prevents installation. It is publicly available on the Tivoli Federated Identity Manager support site.

IBM Tivoli Federated Identity Manager installation fails on Windows Server 2008 R2 when using the graphical user interface mode

Installing Tivoli Federated Identity Manager by using the graphical user interface on Windows Server 2008 R2 fails. You might encounter the following error message: "The installer is unable to run in graphical mode. Try running the installer with the -console or -silent flag". This issue is specific to using Windows 7 or Windows Server 2008 R2.

To resolve the installation issue:

1. Right-click **Properties** on the `install.exe` file in the Tivoli Federated Identity Manager folder of the installation image for Tivoli Federated Identity Manager. The **Properties** window opens.
2. Click the **Compatibility** tab on the Properties window.
3. Click **Change settings for all users** in the **Compatibility** tab. A new **Properties** window opens.
4. Select the **Run the program in compatibility mode** option and select **Windows Server 2008 (Service Pack 1)** mode.
5. Click **OK** at the bottom of the window and click **OK** on the remaining window.

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.

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 show 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 keystore, the settings override other SSL clients running in JVM if they must use 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 keystore specified for that connection is 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 keystore that is used for the HTTPS connection to the Common Auditing and Reporting Service server.

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 by using HTTPS, add the CA Certificate for that HTTPS connection to the trusted keystore used for the HTTPS connection to the Common Auditing and Reporting Service server.

To fix this problem on WebSphere Application Server 6.1, define the dynamic outbound endpoint SSL configurations.

From the WebSphere administrative console, select **Security > SSL certificate and key management > Dynamic outbound endpoint SSL configurations**.

Define an outbound connection that is based on the target host name, 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 might 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

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.

The following section describes the two ways to work around this problem.

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

After configuring SPNEGO authentication in 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 must only change Program Filesto 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 must 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 qualified 8.3 file name, for example:

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

The update 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 with 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.

Use the **cacls.exe** command, and the locale-specific name for the "NETWORK SERVICE" group, to assign the permissions for the log file to 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 might 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 effect 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.

The following settings must match in the Service Name and Port Type:

- The left side of the colon must match the NSURI of the WS-Addressing namespace used for PortType and ServiceName elements in the request.
- The right side of the colon must match the values used within the ServiceName and PortType elements.

Calling Java classes from XSLT in ITFIM 6.2

Tivoli Federated Identity Manager 6.2 has been developed using the OSGI framework model.

Tivoli Federated Identity Manager 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 block in Tivoli Federated Identity Manager waiting for a free ALH to call the Tivoli Directory Integrator.

If zero is 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. Start the wsadmin with 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 with the following command:

```
wsadmin>$AdminApp list
```

Note: The Tivoli Federated Identity Manager console is not 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 stand-alone 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 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 with 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, with the following command:

```
wsadmin>$AdminApp update isclite modulefile {-operation delete
-contenturi itfim-fimconsole-e.war} wsadmin>$AdminConfig save
```
7. After you use wsadmin commands to remove the ITFIMManagementService, ITFIMRuntime and Fimconsole from WebSphere, 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** option. 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 logout not working properly

The issue is caused by an incorrect junction configuration.

The issue occurs in the following conditions:

- WebSEAL is the point of contact.
- The method that is used to invalidate sessions requires Tivoli Federated Identity Manager to know the *session ID* of the session that it is invalidating.

In order for WebSEAL to pass the session ID as an attribute, the junction that points to Tivoli Federated Identity Manager must be configured accordingly.

Use the tfimcfg.jar tool to configure WebSEAL in creating the junction with the required configuration to have WebSEAL pass the session ID.

An example of a junction:

```
object show /WebSEAL/ip-ip/FIM/
Name: /WebSEAL/ip-ip/FIM
Description: Object from host sam120sp3.
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

The `itfimcfg.jar` also enables EAI authentication when you use the Tivoli Federated Identity Manager configuration utility to configure the WebSEAL SOAP instance.

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 you can use SOAP to initiate log-out, the configuration tool sets an EAI trigger on the SOAP endpoint.

If 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 returns to using the PDadmin application programming interface to log out unless the environment has been enabled to use Session Management Server (SMS). In the case where SMS is not being used 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 show the new cluster member:

1. Edit the `software.properties` in the `/pkg` subdirectory where Tivoli Federated Identity Manager is installed.
2. Update `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 on to the management console.
4. Click **Tivoli Federated Identity Manager > Domain Management > Runtime Node Management**. When the portlet opens, it checks if the `software.properties serialId` has been updated.
5. Click **Deploy Runtime**.

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 compatible with WebSphere Application Server 6.0.2 and WebSphere Application Server 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 compatible with WebSphere Application Server 6.0.2 and WebSphere Application Server 6.1. This property enables the configuration to be set on a specific chain.

Tivoli Federated Identity Manager console cannot connect to the Management Service during domain creation

The connection problem might be caused by any of the following reasons:

- The Management Service is not running.
- The Management Service is not initialized properly.
- The information provided in the WebSphere Security Settings panel is not correct.

To resolve the problem, you must complete this procedure:

1. Ensure that the Tivoli Federated Identity Manager Management Service is installed and is running.

If the Management Service is not running, the following error message is shown:

For a remote connection:

FBTCON313E

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

For a local connection:

FBTCON166E

An error was encountered while retrieving environment settings. Check the environment settings and try again.

To run the Management Service application, complete these steps:

- a. From the administration console, select **Applications > Application Types > WebSphere enterprise applications**.
 - b. Select **ITFIMManagementService**.
 - c. Click **Start**.
2. Ensure that the information in the WebSphere Security panel of the Domain wizard is correct.

If the information is inaccurate, the following error message is shown:

FBTCON314E

While contacting the ITFIM Management Service, a WebSphere AdminClient connector could not be created to the Management Service's application server with the given host and port.

3. Restart the WebSphere Application Server.

The tfimcfg.jar tool fails to configure Tivoli Access Manager when FIPS is enabled

An error occurs when Tivoli Federated Identity Manager is accessing the WebSphere Application Server environment to query the InfoService for the configured federation information. The problem occurs because the tfimcfg tool tries to use the SSL protocol that is not supported on the FIPS enabled environment.

Tivoli Federated Identity Manager provides the option of running in a FIPS compliant environment. After FIPS is enabled, the TLS SSL connection factory must be used when running the tfimcfg.jar tool. Specify the **-sslfactory TLS** command parameter when running **tfimcfg** to configure the Tivoli Access Manager environment. For example:

```
java -jar tfimcfg.jar -action tamconfig -cfgfile
/opt/pdweb/etc/webseald.conf -sslfactory TLS
```

SAML 2.0 bearer subject confirmation data processing not conformant

The Recipient attribute of the SubjectConfirmationData is *not* set when issuing an assertion in the following conditions:

- The SAML 2.0 assertion is generated with a bearer subject confirmation method.
- No Claims element is supplied in the RST.

As a workaround, set an Attribute in the ContextAttributes section of the STSUniversalUser in your mapping rule to override the Recipient attribute in SubjectConfirmationData for bearer subject confirmation method. This solution is available for SAML 2.0 identity provider scenarios, or other STS scenarios which issue SAML 2.0 assertions.

An example of this attribute looks like:

```
https://saml.some.receipient.com
```

Another issue is that the assertion validation fails when:

- The bearer subject confirmation method is present.
- The Recipient attribute does not match the assertion consumer service URL in claims.

When validating SAML 2.0 assertions, the Recipient attribute is only validated for bearer subject confirmation methods when the assertion is being validated as part of a single sign-on flow. One exception to this is if you set a runtime custom property:

```
SAML2.AlwaysValidateBearerSubjectConfirmationData = true
```

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

You might be prompted to restart the WebSphere server again if you did not allocate enough time for the restart to complete.

Procedure

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

Incomplete operations after logging off the console

If you log off from 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 before you click **Finish** from 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. Log off from the console. If you must log off from the console in the middle of an operation:
 - a. Log on to the console.
 - b. Verify that the operation is complete.
 - c. If the operation is not complete, modify the necessary values.

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.

You can also configure the 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 that is required 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, use the command line to stop the WebSphere Application Server.

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 issue occurs, you cannot stop the WebSphere Application Server service through the Services panel. Tivoli Federated Identity Manager management console

can also lose the connection to WebSphere Application Server. To stop the WebSphere Application Server service, use the command line to stop the service. See the WebSphere Application Server information center for details on how to stop the service with 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 from the Tivoli Federated Identity Manager management console.
3. Log on 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

Read the issues and solutions about the WebSphere Application Server administration console in this section.

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

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

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

WebSphere administration console does not open the Runtime or Management application

If the WebSphere Application Server administration console does not open the Runtime or Management application, log off from the console and log on to the console again..

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/
Cluster	\$WAS_PROFILE_HOME/installedApps/ cell_name

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

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

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

After making a Tivoli Federated Identity Manager configuration change by using the WebSphere Administrative Console, the following message and option are shown:

```
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 **Load configuration changes to Tivoli Federated Identity Manager runtime**, the Tivoli Federated Identity Manager runtime begins restarting on each application server in the cluster.

Normally, the configuration changes take only a second or two to update each Tivoli Federated Identity Manager runtime. However, it is possible that, due to network delays, you click the option and start the runtimes before all changes have been propagated to each node in the WebSphere cell.

If this scenario occurs, you can start the Tivoli Federated Identity Manager runtime again from the WebSphere Administrative console. 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 show the status of the management service as unavailable. The service might become unavailable even if application is running. The status is indicated by a gray icon with a slash through it.

Procedure

1. Use one of the following methods to verify the status of the application, stop, and start the application:
 - Use the **wsadmin** command-line interface.
 - Use the WebSphere Application Server console.
2. Check the console to see if the management service shows as available.

Querying the Tivoli Federated Identity Manager runtime status

You can use certain wsadmin commands to query the IBM 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 status of the IBM Tivoli Federated Identity Manager runtime, and 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" shows 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
```

Profiles with POST bindings for single sign-on might not work when using WebSphere Application Server 6.1.0.17 or 6.1.0.19

WebSphere v6.1.0.17 and v6.1.0.21 has a defect that causes Tivoli Federated Identity Manager flows that use a POST profile to fail.

For details on addressing this issue, see technote #1326460: IBM Tivoli Federated Identity Manager 6.1.1 and 6.2 POST profile SSO might fail with a SRVE0216E error.

Operational issues

Tivoli Federated Identity Manager might encounter issues related to system operation. Learn about identified issues and how to address them.

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 might fail under the following circumstances:

- If the WebSphere Application Server is specified as the point of contact server, and
- 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 log in a user during a single sign-on operation, still calls 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 fails with the 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.

Configure the 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 Tivoli 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 17:05:27:813 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 website:

<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 website:

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

Disabling the replay validation detection in PassTicket validation

A PassTicket includes a timestamp that 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 are identical. The content of the PassTickets are slightly different, but the value of the timestamps are the same.

An error occurs during token validation when a PassTicket with a timestamp that has the same value as a previous request is called to the STS. This behaviour is considered to be a replay request.

This error message is shown on the response that indicates that the PassTicket has been replayed:

FBTSTS024E

The given Username token was replayed.

The given Username token was verified before and now it is being reused.

This server's configuration does not allow Username tokens to be reused.

The only way to work around this issue is to disable the replay validation detection with the understanding that this enables applications to reuse the same PassTicket.

Administrators must consider the security risks involved in disabling the replay validation detection.

To manage this problem, use the disable replay detection functionality which RACF® supports. To disable replay validation, 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 of the custom chain that uses the PassTicket STS module instance.

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 > Show Chain ID in column in table**. This action selection adds a new column in the table that shows the unique Chain ID.

StringIndexOutOfBoundsException is thrown when adding a SAML 2.0 service provider as a partner

This exception occurs if the parameter `ValidateKeyIdentifier` does not contain an underscore, and the parameter `signatureKeyAlias` and `signatureKeystoreName` are not specified. This exception also occurs if the parameter `EncryptionKeyIdentifier` does not contain an underscore, and the parameter `encryptionKeyAlias` and `encryptionKeystore` are not specified.

SAML STS modules calculate the wrong validity period of assertion

An error during the SAML assertion timestamps validation occurs if you set the custom runtime property to `saml.use.legacy.clockskew.default = true`.

To fix this error, change the value of the property
`tosaml.use.legacy.clockskew.default = false`.

Note: Tivoli Federated Identity Manager 6.2.2, by default, does not use a clock skew when validating SAML assertion timestamps. It instead uses the local clock of the run time.

Tivoli Federated Identity Manager alias service does not work with Oracle database

Attempts to use Oracle database for Tivoli Federated Identity Manager alias service shows errors like:

```
com.ibm.ws.ejbpersistence.utilpm.PersistenceManagerException:
```

```
PMGR1012E: The current backend id DB2UDBNT_V8_1, does not match the  
datasource connected to.
```

To fix this error, you must follow additional steps after installing the fix pack for Tivoli Federated Identity Manager Version 6.2.1 (on UNIX-based system).

Configuration steps to use Oracle database for Tivoli Federated Identity Manager alias service:

1. Create a backup of the `itfim.ear` file with the commands:

```
cd FIM_INSTALL_DIR/pkg/release  
cp itfim.ear itfim-orig.ear
```
2. Modify the EAR to be Oracle-aware for deployment with the commands:

```
mkdir /tmp/work  
rm FIM_INSTALL_DIR/pkg/release/itfim.ear
```

```

WEBSPPHERE_INSTALL_DIR/AppServer/bin/ejbdeploy.sh FIM_INSTALL_DIR/pkg/
release/itfim-orig.ear /tmp/work FIM_INSTALL_DIR/pkg/release/itfim-
oracle.ear -dbschema FIMAliasesSchema -dbname FIMAliases -dbvendor
ORACLE_V10G -trace
cp FIM_INSTALL_DIR/pkg/release/itfim-oracle.ear FIM_INSTALL_DIR/pkg/
release/itfim.ear
rm -rf /tmp/work

```

3. A new FIM_INSTALL_DIR/pkg/release/itfim.ear is now available for deployment to work with Oracle. Use a text editor to update the file FIM_INSTALL_DIR/pkg/software.properties to change the property com.tivoli.am.fim.rte.software.serialId to a different value, for example, increment.
4. Use the Tivoli Federated Identity Manager console to select **Tivoli Federated Identity Manager > Domain Management > Runtime Node Management**. A message indicating that a new runtime is available for deployment shows. Use the console to deploy the new runtime.
5. Restart the WebSphere process where the runtime is deployed.

For more details about setting up Oracle database and configuring the WebSphere Application Server to use JDBC, see the technote entitled "After installing FP3 for FIM 6.2, Oracle db still cannot be used for FIM Alias Service."

Note: If you receive subsequent fix packs or anything that alters the deployed itfim.ear, you must do the configuration steps again.

Missing the closing macro delimiter "@" from the consent.html template

The closing macro delimiter, @OPTIONAL_ATTRIBUTE@ is missing from the consent.html template causing the label field for the optional attributes not to render correctly.

To modify the consent.html template:

1. Open the page template <fim_install_root>/pages/<lang>/openid/consent.html
2. Look for <label for="chk_@OPTIONAL_ATTRIBUTE">@OPTIONAL_ATTRIBUTE@</label>
 and change it to:
 <label for="chk_@OPTIONAL_ATTRIBUTE">OPTIONAL_ATTRIBUTE@</label>

3. Use the console to Publish Pages and Reload Configurations for the change to take effect.

Generating debug statements for identity mapping in XSLT rules

When authoring XSLT rules for identity mapping, there is no mechanism to log or trace statements for debugging purposes. you can use an extension to generate debugging statements to XSLT rules.

To start debug statements for identity mapping, add entries in the XSLT rules with the syntax:

```

<xsl:variable name="variablename" select="mapping-ext:traceString('debug
string')">

```

To enable the trace configuration, set the trace level to fine for the class `com.tivoli.am.fim.trustserver.sts.utilities.IDMappingExtUtils`. The output is typically rendered in `/opt/IBM/WebSphere/AppServer/profiles/<profile-name>/logs/<server-name>`.

Recipient checking is not performed correctly by the SAML browser post support

If a SAML 1.x service provider must accept a `samlp:Response` that does not contain a Recipient attribute, this runtime custom property can be used at a Tivoli Federated Identity Manager identity provider that is issuing a SAML response:

```
SAML.AllowNoRecipient=true
```

Typically the Recipient is a required attribute so there is no need to set this runtime custom property. It is only offered for uncommon backwards compatibility use cases.

The IBM Tivoli Federated Identity Manager IVCred STS Module does not validate IVCRED tokens that correspond to an unauthenticated user

When using a WebSEAL junction configured to invoke the Tivoli Federated Identity Manager STS to generate a token, the IVCred sts module will generate a `NullPointerException` if the endpoint accessed has an unauthenticated ACL. When the endpoint has an authenticated ACL, the credential received by Tivoli Federated Identity Manager corresponds to an unauthenticated user.

The IVCred STS Module has been enabled to consume and validate IVCred tokens that correspond to an unauthenticated user.

There are two modes of operation:

- The sts module will generate an error if a token received corresponds to an unauthenticated user. This is the default setting. The error is:
`FBTSTS015E The IV-Cred binary token is invalid or not present.`
- The IVCred STS Module can be configured to map the unauthenticated user token to a special user account that can be configured. The user account selected should be considered as a low entitlements or guest account.

The IVCred STS module adds an unauthenticated user name to the universal user structure.

To enable this mode, add the custom property:

```
ivcred.unauthenticated.user.name=myusername
```

Where `myusername` is the user name value to use for mapping.

Additional properties can also be provided to describe the user account to map to:

`ivcred.unauthenticated.user.registry.id` - used to include the registry id of the account.

`ivcred.unauthenticated.user.uuid` - used to indicate the unique id for the user account

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

Symptom: Unable to log on 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: You need an adapter 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 on, use an operating system command to stop 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.
7. If necessary, reconfigure the VMM Tivoli Access Manager Adapter.

Deployment issues

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

Error received about deployment operation

FBTCON137E: An error occurred during the deployment 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.

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

SOAP timing out

Increase the value of the `com.ibm.SOAP.requestTimeout` parameter to be 45 seconds longer than the time it takes for the operation to complete. Doing so prevents the SOAP Client Connector from timing out during an operation.

To determine the appropriate timeout setting, perform the following steps:

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 Server 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 Server node where the administration console is installed.

Deploying Tivoli Federated Identity Manager in a WebSphere vertical cluster environment

The deployment against the second server instance fails when Tivoli Federated Identity Manager is deployed in a vertical cluster environment.

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 Tivoli 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, although the Tivoli Federated Identity Manager runtime requires the discovery of a node agent to properly initialize.

The following list includes the 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, otherwise they are not detected by applications at startup.

Performance issues

The Tivoli Federated Identity Manager runtime uses a node agent discovery event to reinitialize its components. This event might not be issued by WebSphere Application Server under any of the following circumstances:

- If the system or network is under heavy load
- If the response latencies are large for other reasons

As a result, the Tivoli Federated Identity Manager runtime is left in an improperly initialized state.

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

- The Tivoli Federated Identity Manager runtime is configured without error through the administrative console. Or, the runtime node shows as both deployed and configured in the Runtime Node Management page.
- When the Federated Identity Management tracing is 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 Tivoli 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) run before attempting to start or restart the Tivoli 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

Put the missing package of the class in the Export-Package stanza of the MANIFEST.MF file in the `com.tivoli.am.fim.osgi.connector` if you receive a `class-not-found` or `class-definition-not-found` error on a WebSphere J2EE container class. Example of the error messages: `ClassDefNotFoundErrors` and `ClassNotFoundExceptions`.

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: Switch the context class loader when involving methods in the WebSphere J2EE container.

Debugging OSGi bundle resolution and startup problems

You can run Tivoli Federated Identity Manager Eclipse runtimes in debug mode, which opens a port to the OSGi console (`osgiConsole`). You can use telnet to connect to the `osgiConsole`.

You can also do the following through the `osgiConsole`:

- Query the status of all the bundles in the Tivoli Federated Identity Manager plug-ins directory.

- Perform life cycle management operations on the bundles. For example, start, stop, install, and uninstall.

Follow these steps to start a Tivoli Federated Identity Manager Eclipse runtime in debug mode:

1. Go to directory where the `launch.ini` file is installed for `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 go to `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 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 Tivoli Federated Identity Manager Eclipse runtime. You can 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 status is **INSTALLED**, and not **RESOLVED** or **STARTED**, run the command **start number**. Running the **startnumber** command prints 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 Tivoli Federated Identity Manager OSGi error, check the OSGi error logs. When the Tivoli Federated Identity Manager starts, the following files are copied from the configuration repository to a location under the WebSphere Application Server temp directory

- The plug-ins
- Configuration
- Features directories

It is the same location where the Tivoli Federated Identity Manager OSGi instance started.

The 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

You can use the `J2EEContainerAction` in a custom module to start a WebSphere API method (such as JNDI lookup, RMI lookup, or SOAP MessageFactory).

A callback mechanism can make a portion of Tivoli Federated Identity Manager plug-in code run 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 Tivoli 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, start your action class by 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 `J2EEContainerActions` are already defined, and you can 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 website and sign-up for notifications of product support information, including fixes.

Obtaining fixes

A product fix might be available to solve your problem. Check the product support site to determine what fixes are available for Tivoli Federated Identity Manager.

Procedure

1. Go to the IBM Software Support site for Tivoli Federated Identity Manager at <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. You can also download the fix information.

Receiving fix notifications

You can receive email notifications about fixes and other news about IBM products.

Procedure

1. Go to the IBM Software Support 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 opens.
3. If you have already registered, go 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 show.
7. Select **IBM Tivoli Federated Identity Manager** from the list of products.
8. Click **Add products**.
9. To use email 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 email notifications that you 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 documentation, problems, and workaround.

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. See “ISA overview” on page 57 for more information.

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

Procedure

1. Go to the IBM Software Support site for Tivoli Federated Identity Manager:
<http://www.ibm.com/software/sysmgmt/products/support/IBMTivoliFederatedIdentityManager.html>.
2. Under **Solve a problem**, click any of the following options:
 - **Technotes**, which lists information about the product by document title.
 - **APARs**, which lists known problems according to Authorized Program Analysis Report (APAR) numbers.
3. Optionally, you can search for specific terms, error codes, or APARs by using the **Search** field on the product support page. You can also browse through the Technotes or APARs pages.

Chapter 7. Collecting data

There are other ways to solve a problem aside from troubleshooting the symptoms. Another way of solving a problem is collecting more diagnostic data.

Before you begin to collect data for a problem report, install and run the IBM Support Assistant. This troubleshooting tool includes a console where you can 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 about IBM Support Assistant, see “ISA overview” on page 57 for details.

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 the problem 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 must collect and the sources from which you collect that data, depends on the type of problem 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 in 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 must 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 must collect additional data, such as message and trace information. See “Message and trace logs” on page 44 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.

See 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 must be aware of, but does not necessarily require any action.

Error messages

Indicate that a condition has occurred that requires an 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. See the WebSphere Application Server information center for more information.

You can use the WebSphere Application Server administrative console to configure the following log settings:

- location
- name
- maximum size of the log files
- levels of severity that you want to log, such as Warning and Severe

See “Configuring log settings” on page 46 for more information.

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/ <i>profile_name</i> /logs/ <i>server_name</i> /SystemOut.log Windows: C:\Program Files\IBM\WebSphere\AppServer\profiles\ <i>profile_name</i> \ logs\ <i>server_name</i> \SystemOut.log
IBM Service Log	UNIX and Linux: /opt/IBM/WebSphere/AppServer/profiles/ <i>profile_name</i> /logs/ <i>activity.log</i> Windows: C:\Program Files\IBM\WebSphere\AppServer\profiles\ <i>profile_name</i> \ logs\ <i>server_name</i> \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.

Trace logs capture transient information about the current operating environment when a component or application fails to operate as intended. Trace logs are different from message logs because message logs records only the occurrence of noteworthy events. Trace logs are available only in English.

Trace logging is not enabled by default. In some circumstances, trace logging might cause large amounts of data to be collected in a short amount of time which results in significant performance degradation. Therefore, enable the trace logging only at the direction of IBM Software Support personnel. See “Configuring log settings” on page 46 for details.

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). See “Configuring log settings” for more information.

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/ <i>profile_name</i> /logs/ <i>server_name</i> /trace.log Windows: C:\Program Files\IBM\WebSphere\AppServer\profiles\ <i>profile_name</i> \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

You can use the administration console to configure the settings for message and trace logs. Message logging is enabled by default. Trace logging must 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 called as SystemOut.log, is a standard WebSphere Application Server log used for messages. For detailed information, see the JVM log topics in the WebSphere Application Server information center.

Procedure

1. Start the WebSphere Application Server administrative console and log on, 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 show the attributes to configure.
7. Change the appropriate configuration attributes.
8. Click **Apply**.
9. 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 called as the activity.log, is a standard WebSphere Application Server log used for messages. For detailed information about the log, see the service log topics in the WebSphere Application Server information center.

Procedure

1. Start the WebSphere Application Server administrative console and log on, 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 runtime 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. 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, see the WebSphere Application Server information center.

Procedure

1. Start the WebSphere Application Server administrative console and log on, 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, see 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 required 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 the **All Components** option has been enabled, you turn it off and then enable specific components.
10. Click a component or group name. See "Trace components" on page 49 for a list of Tivoli Federated Identity Manager components.

Note: If the selected server is not running, it does not show 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, see the WebSphere Application Server information center.
12. Click **OK**.
13. Click **Save** to save your changes. You must restart WebSphere Application Server for the change to take effect.

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, see the WebSphere Application Server information center.

Procedure

1. Start the WebSphere Application Server administrative console and log on, 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 required state.
8. Configure the trace output if you want to change the existing one.
9. Click **Apply**.

Trace components

Trace 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 in this topic.

The following table describes a partial list of the components you can specify for trace logging.

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 Note: Liberty protocol is being deprecated in the Tivoli Federated Identity Manager 6.2.2 release.
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.ws federation	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

You can view the JVM logs with any of the following options:

- Use the WebSphere Application Server administrative console, which also supports viewing from a remote machine
- Use a text editor on the machine where the log files are stored.

See the **Viewing JVM logs** section 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 must be dumped to a file before it can 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

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 prompt, 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 show all problem types. Find your problem type and select it.

4. Supply a file name for the data collection .zip file.

You can use any file name. The tool automatically appends the .zip file extension. For example, if you enter the file name 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 finish 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 file name in the **Output Filename/Path** box.

The tool appends the server host name and current timestamp to the file name that you entered.

7. Send the .zip file to IBM Support

You can choose FTP or HTTPS for file transfer.

Note: 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. You can use the tool to 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 later.
 - 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 <i>response.txt</i>
UNIX or Linux	runISALiteConsole.sh -record <i>response.txt</i>

You can specify your own file name 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. Use the response file to run the tool.

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

Operating system type	Command
Windows	<code>runISALiteConsole.bat response.txt</code>
UNIX or Linux	<code>runISALiteConsole.sh response.txt</code>

Note:

- 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, 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 must determine how that data can help you to resolve your 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 from those sources. For example, if the problem is related to installation, start your analysis with the installation log files, if any. This action helps narrow down the problem, instead of starting with the general product or operating system log files.
- Understand how the various pieces of data relate to each other. For example, if the data spans to more than one system, keep your data organized to know which pieces of data come from which sources.
- Use timestamps to confirm that each piece of diagnostic data is relevant to the timing of the problem.

Note: Data from different sources can have different timestamp formats. Understand the sequence of the different elements in each timestamp format so you can tell when the different events occur.

The specific method of analysis is unique to each data source. 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, go through the data to trace the root cause of the problem.

To investigate 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:** You can search for information about a problem and collect the logged information required to troubleshoot a problem by using the IBM Support Assistant. You can use it to open a problem management report (PMR) online. It is also helpful to use the IBM Support Assistant to report a problem to IBM Software Support. You need your user account ID and password to submit a PMR through the IBM Support Assistant. The IBM Support Assistant client is included on your product CD and updates can be downloaded from the following website: <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://www14.software.ibm.com/webapp/set2/sas/f/handbook/contacts.html>, and click the name of your geographic region.

If the problem that you submit is for a software defect, missing content, 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 website 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 effect 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

Use multiple filters to focus your search to access troubleshooting repository information quickly. 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 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. Use the service feature 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 save it 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 website:
<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. Click your IBM Support Assistant desktop icon to start. Select the **User Guide** tab for information about performing the various available tasks.

IBM software maintenance contracts

Ensure that your company has an active maintenance contract before you submit a problem to IBM Software Support. Make sure that you are also 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://www14.software.ibm.com/webapp/set2/sas/f/handbook/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 must 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 <i>critical</i> 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 <i>significant</i> business impact: The program is usable, but it is severely limited.
Severity 3	The problem has <i>some</i> business impact: The program is usable, but less significant features (not critical to operations) are unavailable.
Severity 4	The problem has <i>minimal</i> 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 change anything in the system? For example, did you change anything in 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. See the “ISA overview” on page 57 for more information about IBM Support Assistant.

FTP (EcuRep)

Use the FTP service called EcuRep to submit data files to IBM. 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 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)

Use FTP to transfer your files and follow these steps:

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

ESR tool

Registered users who are on an authorized caller list can use the Electronic Service Request (ESR) tool to submit diagnostic data. Use the ESR tool 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 give 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.
1623-14, Shimotsuruma, Yamato-shi
Kanagawa 242-8502 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 might 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 document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

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 measurement 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 is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to 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. 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 IBM's application programming interfaces.

If you are viewing this information in softcopy form, the photographs and color illustrations might not be displayed.

Trademarks

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

Adobe, 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.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency which is now part of the Office of Government Commerce.

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.

ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.

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



Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

Linear Tape-Open, LTO, the LTO Logo, Ultrium, and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

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

Index

A

- abends
 - overview 8
- accessibility vii
- activity.log
 - configuring 47
 - IBM service log default name 44
- Authorized Program Analysis Report (APAR)
 - searching 41

C

- cacls.exe command
 - using 16
- class loading errors
 - ClassDefNotFoundErrors 36
 - ClassNotFoundExceptions 36
- CLI will not work after installation 12
- cluster
 - deploying vertical cluster 35
- command line interface does not work after installation 12
- commands
 - calcs.exe 43
 - errpt 43
 - wsadmin 43
- configuration changes
 - troubleshooting 26
- configuration problems
 - ITFIM_WSSM directory 14
 - national language characters 15
 - spaces in keytab filepath 15
 - SSL property settings 14
- connectivity
 - troubleshooting 5
- connectivity problems
 - overview 5
- console issues
 - configuration changes 26
 - incomplete operations 24
 - starting WebSphere Application Server 24
 - stopping WebSphere Application Server 24
- conventions
 - typeface viii
- crashes
 - overview 8
- customization
 - OSGi bundle resolution 36
 - OSGi error logs 37
 - WebSphere API method 38

D

- data
 - collecting for a problem 43
- data files
 - packaging using EcuRept 59

- deployment issues
 - overview 34
- deployment operation
 - error 34
- deployment problems
 - node agent 35
 - Simple Object Access Protocol (SOAP)
 - timing out 34
 - WebSphere vertical cluster 35
- diagnostic data
 - submitting 60
- directory names, notation viii

E

- EAR files
 - locating 26
- EcuRep service
 - using 60
- education
 - See Tivoli technical training
- Electronic Service Request (ESR) tool
 - using 60
- environment variables, notation viii
- error messages
 - overview 7
- errpt command
 - using 43

F

- federation names
 - using valid characters 15
- fix notifications
 - obtaining 39
 - receiving 39
- fix pack
 - overview 6
- fixes
 - locating 6
 - obtaining 39
 - overview 39

I

- IBM Service log
 - configuring 47
 - path 45
- IBM Software Support
 - contacting 57
- IBM Support Assistant
 - overview 57
 - submitting data to 59
- IBM Support Assistant (ISA)
 - using 50
- IBM Support AssistantUsage
 - using 50
- installation failure on Solaris 11
- integrated solutions console
 - troubleshooting 23

- ITFIM_WSSM directory
 - adding files 14
 - adding subdirectories 14
- ITFIMManagementService.ear
 - using management service file 26
- ITFIMRuntime.ear
 - using runtime file 26
- IY93387 fix
 - validating message failure error 28

J

- JAR files
 - troubleshooting 13
- Java Virtual Machine (JVM)
 - fixing error IY93387 28
- Java Virtual Maching (JVM) log
 - configuring 47
- JVM log
 - filepath 45

K

- keytab file
 - troubleshooting 15
- knowledge bases
 - searching 41

L

- log data
 - analyzing 55
- log files
 - permissions 14
 - writing 16
- log settings
 - configuring 46
- logs
 - analyzing data 55
 - enabling trace at server startup 48
 - enabling trace on running server 48
 - file names 45
 - locations 46
 - message types 44
 - tracing 46
 - viewing 49

M

- maintenance contracts
 - overview 58
- management application
 - troubleshooting 26
- management console
 - troubleshooting multiple instances 25
- management service
 - troubleshooting 27

- message logs
 - configuring 46
 - overview 44
- messages
 - errors 7
 - types 44

N

- national language characters
 - using in federation names 15
- node agent
 - troubleshooting discovery of 35
- notation
 - environment variables viii
 - path names viii
 - typeface viii

O

- operation
 - validating message failure 28
- operations
 - troubleshooting 24
- ordering publications vii
- OSGi bundle resolution
 - debugging 36
- OSGi error logs
 - locating 37
- osgiConsole (OSGi)
 - debugging 36
- osgiConsole (OSGi) error logs
 - locating directory 37

P

- path names, notation viii
- performance problems
 - overview 7
 - troubleshooting 7
- problems
 - collecting general data 43
 - collecting specific data 43
 - describing 59
 - learning about 3
 - obtaining symptoms 3
 - reporting 59
 - severity levels 59
 - submitting data to IBM 59
- product fixes
 - overview 6
- product updates
 - overview 6
- publications
 - ordering vii
 - related vi

R

- refresh pack
 - overview 6
- runtime status
 - querying 27

S

- secure sockets layer (SSL) property settings
 - configuring 14
- security settings
 - reconfiguring 24
- Simple and Protected GSSAPI Negotiation Mechanism (SPNEGO)
 - keytab filepath spaces 15
- Simple Object Access Protocol (SOAP)
 - timing out 34
- software maintenance contracts
 - overview 58
- software problem resolution
 - searching 39
- software support
 - contacting 57
- Solaris installation failure 11
- solutions
 - overview 11
- SPNEGO
 - configuring spaces in keytab filepath 15
- SSL property settings
 - configuring for auditing service 14
 - overriding other property settings 14
- support
 - contacting 57
- support handbook
 - overview 58
- SystemOut.log
 - message logs 44

T

- technotes
 - searching 41
- Tivoli technical training vii
- trace component names
 - understanding 49
- trace components
 - enabling 49
- trace logs
 - enabling 47
 - overview 45
 - using file names 46
- training, Tivoli technical vii
- traps
 - overview 8
- troubleshooting
 - checklist 1
 - installation 11, 12
 - obtaining symptoms 3
 - process 1
 - uninstallation 13
- typeface conventions viii

V

- variables, notation for viii

W

- WebSphere API
 - invoking method 38

- WebSphere Application Server
 - console issues 24
 - displaying Runtime or Management application 26
 - installing for multiple consoles 25
 - restarting after reconfiguring security settings 24
 - stopping from the command line 24
 - troubleshooting management service 27
 - troubleshooting multiple instances of management console 25
- WebSphere cluster
 - node agent not discovered 35
- WebSphere vertical cluster
 - deploying 35
- wsadmin command
 - using 27



Printed in USA

GC27-2715-01

