



# Sun™ Identity Manager 8.0 Tuning, Troubleshooting, and Error Messages

---

Sun Microsystems, Inc.  
4150 Network Circle  
Santa Clara, CA 95054  
U.S.A.

Part No: 820-2962-10

Copyright © 2008 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, U.S.A. All rights reserved.

Sun Microsystems, Inc. has intellectual property rights relating to technology embodied in the product that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more of the U.S. patents listed at <http://www.sun.com/patents> and one or more additional patents or pending patent applications in the U.S. and in other countries.

**THIS PRODUCT CONTAINS CONFIDENTIAL INFORMATION AND TRADE SECRETS OF SUN MICROSYSTEMS, INC. USE, DISCLOSURE OR REPRODUCTION IS PROHIBITED WITHOUT THE PRIOR EXPRESS WRITTEN PERMISSION OF SUN MICROSYSTEMS, INC.**

U.S. Government Rights - Commercial software. Government users are subject to the Sun Microsystems, Inc. standard license agreement and applicable provisions of the FAR and its supplements.

Use is subject to license terms.

This distribution may include materials developed by third parties.

Sun, Sun Microsystems, the Sun logo, Java, Solaris, Sun Java System Identity Manager, Sun Java System Identity Manager Service Provider Edition services, Sun Java System Identity Manager Service Provider Edition software and Sun Identity Manager are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries.

All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon architecture developed by Sun Microsystems, Inc.

UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd.

This product is covered and controlled by U.S. Export Control laws and may be subject to the export or import laws in other countries. Nuclear, missile, chemical biological weapons or nuclear maritime end uses or end users, whether direct or indirect, are strictly prohibited. Export or reexport to countries subject to U.S. embargo or to entities identified on U.S. export exclusion lists, including, but not limited to, the denied persons and specially designated nationals lists is strictly prohibited.

# Contents

<b>List of Tables</b> .....	<b>ix</b>
<b>Preface</b> .....	<b>xi</b>
Intended Audience .....	xi
How This Book Is Organized .....	xii
Conventions Used in This Book .....	xii
Typographic Conventions .....	xii
Symbols .....	xiii
Shell Prompts .....	xiii
Related Documentation and Help .....	xiv
Accessing Sun Resources Online .....	xv
Contacting Sun Technical Support .....	xv
Related Third-Party Web Site References .....	xv
Sun Welcomes Your Comments .....	xvi
<b>Performance Tuning</b> .....	<b>1</b>
Before You Begin .....	2
Intended Audience .....	2
Important Notes .....	2
Related Documentation and Web Sites .....	2
Tuning Roadmap .....	5
Tuning Your Deployment Environment .....	6
Tuning Your Java EE Environment .....	6
Tuning Your Application Server .....	9
Tuning Your Repository Database .....	11
Tuning Identity Manager Performance .....	25
General Performance Tunings .....	26
Tuning Active Sync Adapter Performance .....	27
Tuning Bulk Loading .....	28
Tuning Configurable XML Objects .....	29

Tuning Database Statistics	36
Tuning Data Exporter	37
Tuning the General XML	38
Tuning Identity Manager Service Provider	38
Tuning the Identity Manager Web Interface	38
Tuning Initial Loads	39
Tuning Memory Requirements	39
Tuning Operating System Kernels	40
Tuning Provisioner	40
Tuning Reconciliation	42
Tuning Resource Queries	45
Tuning the Scheduler	46
Tuning Sessions	48
Tuning the Sun Identity Manager Gateway	48
Tuning the Task Bar	51
Debugging Performance Issues	52
Working with Identity Manager Debug Pages	53
Working With Other Debugging Tools	59
<b>Tracing and Troubleshooting</b>	<b>65</b>
Before You Begin	65
Intended Audience	66
Important Notes	66
Before Calling Support	66
Related Documentation and Web Sites	67
Process Overview	68
Tracing Identity Manager Objects and Activities	69
How To Configure Tracing	70
How to View Trace Files	75
Tracing the Identity Manager Server	75
Tracing Adapters	76
Tracing Auditor	76
Tracing Custom Code	77
Tracing Exceptions	78
Tracing Forms	79
Tracing Global XPRESS	80
Tracing PasswordSync	81
Tracing Identity Manager Service Provider Delegated Administration	85
Tracing Reconciliation	85
Tracing the <code>setRepo</code> Command	88
Tracing SPML	88
Tracing the Task Scheduler	91

Tracing Workflows .....	91
Locating Version Information .....	94
Tracing the Identity Manager Gateway Objects and Activities .....	95
How to Configure Tracing .....	95
How to Configure Tracing for the PowerShellExecutor.dll Add-On .....	98
How to Capture Dr. Watson Logs .....	100
Troubleshooting and Fixing Common Problems .....	101
Working with Debugging Tools .....	102
Debugging Errors Displayed in the Browser .....	107
Troubleshooting Adapters .....	107
Troubleshooting Auditor .....	116
Troubleshooting Form Problems .....	117
Troubleshooting the Gateway .....	119
Troubleshooting Java Code Problems .....	120
Troubleshooting Low Memory Conditions .....	120
Troubleshooting PasswordSync Problems .....	122
Troubleshooting Reconciliation Problems .....	124
Troubleshooting Repository Connection Problems .....	125
Troubleshooting Server-Related Problems .....	128
Troubleshooting an SPML Configuration .....	129
Troubleshooting Upgrades .....	129
<b>Errors and Exceptions .....</b>	<b>131</b>
Before You Begin .....	131
Intended Audience .....	131
Important Notes .....	132
Related Documentation and Web Sites .....	132
About Identity Manager Errors and Exceptions .....	133
Where Messages Are Stored .....	134
How Messages Are Displayed .....	135
Error Severity Levels .....	135
Viewing Errors in the System Log Report .....	136
Customizing a Default Error Message .....	139
<b>Index .....</b>	<b>143</b>



# List of Tables

Table 1-1	Related Documentation .....	3
Table 1-2	Useful Web Sites .....	4
Table 1-3	DTrace Commands .....	61
Table 2-1	Useful Web Sites .....	67
Table 2-2	Registry Keys .....	82
Table 2-3	Reconciliation Methods/Classes to Trace .....	86
Table 2-4	Gateway Tracing Command Line Arguments .....	97
Table 2-5	Identity Manager Debug Pages .....	103
Table 2-6	Gateway Command Arguments .....	119
Table 3-1	Useful Web Sites .....	133



# Preface

*Identity Manager Tuning, Troubleshooting, and Error Messages* contains guidelines for tuning Sun Java™ System Identity Manager, describes how to trace and troubleshoot problems, and describes the error messages and exceptions you might encounter as you work with the product.

## Intended Audience

This book is intended for administrators and software developers who implement an integrated identity management and web access platform by using Sun Java System servers and software.

Identity Manager administrators and software developers must understand the following technologies:

- Lightweight Directory Access Protocol (LDAP)
- Java technology
- JavaServer Pages™ (JSP™) technology
- Hypertext Transfer Protocol (HTTP)
- Hypertext Markup Language (HTML)
- Extensible Markup Language (XML)

# How This Book Is Organized

The guide is organized in these sections:

- [Chapter 1, “Performance Tuning,”](#) contains guidelines for tuning Identity Manager performance.
- [Chapter 2, “Tracing and Troubleshooting,”](#) explains how to use trace output to identify and resolve problems.
- [Chapter 3, “Errors and Exceptions,”](#) describes Identity Manager error messages and exceptions.

## Conventions Used in This Book

The tables in this section describe the conventions used in this book.

### Typographic Conventions

The following table describes the typographic changes used in this book.

**Table 1** Typographic Conventions

Typeface	Meaning	Examples
AaBbCc123 (Monospace)	API and language elements, HTML tags, web site URLs, command names, file names, directory path names, onscreen computer output, sample code.	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. % You have mail.
<b>AaBbCc123</b> (Monospace bold)	What you type, when contrasted with onscreen computer output.	% <b>su</b> Password:
<i>AaBbCc123</i> (Italic)	Book titles, new terms, words to be emphasized. A placeholder in a command or path name to be replaced with a real name or value.	Read Chapter 6 in <i>Identity Manager Tuning, Troubleshooting, and Error Messages</i> . These options are called <i>class</i> options. Do <i>not</i> save the file. The file is located in the <i>install-dir/bin</i> directory.

# Symbols

The following table describes the symbol conventions used in this book.

**Table 2** Symbol Conventions

Symbol	Description	Example	Meaning
-	Joins simultaneous multiple keystrokes.	Control-A	Press the Control key while you press the A key.
+	Joins consecutive multiple keystrokes.	Ctrl+A+N	Press the Control key, release it, and then press the subsequent keys.
>	Indicates menu item selection in a graphical user interface.	File > New > Templates	From the File menu, choose New. From the New submenu, choose Templates.

# Shell Prompts

The following table describes the shell prompts that might be used in this book.

**Table 3** Shell Prompts

Shell	Prompt
C shell on UNIX or Linux	<i>machine-name%</i>
C shell superuser on UNIX or Linux	<i>machine-name#</i>
Bourne shell and Korn shell on UNIX or Linux	\$
Bourne shell and Korn shell superuser on UNIX or Linux	#
Windows command line	C: \

# Related Documentation and Help

Sun Microsystems provides additional documentation and information to help you install, use, and configure Identity Manager:

- *Identity Manager Installation*: Step-by-step instructions and reference information to help you install and configure Identity Manager and associated software.
- *Identity Manager Upgrade*: Step-by-step instructions and reference information to help you upgrade and configure Identity Manager and associated software.
- *Identity Manager Administration*: Procedures, tutorials, and examples that describe how to use Identity Manager to provide secure user access to your enterprise information systems.
- *Identity Manager Technical Deployment Overview*: Conceptual overview of the Identity Manager product (including object architectures) with an introduction to basic product components.
- *Identity Manager Deployment Tools*: Reference and procedural information that describes how to use different Identity Manager deployment tool. This information addresses rules and rules libraries, common tasks and processes, dictionary support, and the SOAP-based Web Service interface provided by the Identity Manager server.
- *Identity Manager Workflows, Forms, and Views*: Reference and procedural information that describes how to use the Identity Manager workflows, forms, and views — including information about the tools you must customize these objects.
- *Identity Manager Resources Reference*: Reference and procedural information that describes how to load and synchronize account information from a resource into Sun Java™ System Identity Manager.
- *Identity Manager Service Provider Deployment*: Reference and procedural information that describes how to plan and implement Sun Java™ System Identity Manager Service Provider.
- *Identity Manager Help*: Online guidance and information that offer complete procedural, reference, and terminology information about Identity Manager. You can access help by clicking the Help link from the Identity Manager menu bar. Guidance (field-specific information) is available on key fields.

# Accessing Sun Resources Online

For product downloads, professional services, patches and support, and additional developer information, go to the following:

- Download Center  
<http://www.sun.com/software/download/>
- Professional Services  
<http://www.sun.com/service/sunps/sunone/index.html>
- Sun Enterprise Services, Solaris Patches, and Support  
<http://sunsolve.sun.com/>
- Developer Information  
<http://developers.sun.com/prodtech/index.html>

## Contacting Sun Technical Support

If you have technical questions about this product that are not answered in the product documentation, contact customer support by using one of the following mechanisms:

- The online support web site at <http://www.sun.com/service/online/us>
- The telephone dispatch number associated with your maintenance contract

## Related Third-Party Web Site References

Third-party URLs are referenced in this document and provide additional, related information.

---

**NOTE** Sun is not responsible for the availability of third-party web sites mentioned in this document. Sun does not endorse and is not responsible or liable for any content, advertising, products, or other materials that are available on or through such sites or resources. Sun will not be responsible or liable for any actual or alleged damage or loss caused or alleged to be caused by or in connection with use of or reliance on any such content, goods, or services that are available on or through such sites or resources.

---

# Sun Welcomes Your Comments

Sun is interested in improving its documentation and welcomes your comments and suggestions.

To share your comments, go to <http://docs.sun.com> and click Send Comments. In the online form, provide the document title and part number. The part number is a seven-digit or nine-digit number that can be found on the title page of the book or at the top of the document.

For example, the title of this book is *Sun Java™ System Tuning, Troubleshooting, and Error Messages*, and the part number is 820-2962-10.

# Performance Tuning

You can significantly improve Identity Manager performance across nearly all activities with the proper tuning. In addition to changing settings within the software, you can make performance improvements by tuning the application server, Java Virtual Machine (JVM), hardware, operating system, and network topology.

Additionally, you can use several tools to diagnose and monitor performance. Several of these tools exist within Identity Manager, such as trace and method timers. You can also use other Sun Microsystems and third-party tools to debug performance issues with Identity Manager.

This chapter describes tools, methodologies, and references you can use to improve performance and debug performance-related issues.

---

**NOTE** The tuning process spans many entities and is specific to your deployment environment. This chapter describes different tuning methods for optimizing Identity Manager performance, but these methods are *only intended as guidelines*. You might have to adapt these methods for your deployment environment.

---

This chapter covers the following topics:

- [Before You Begin](#)
- [Tuning Your Deployment Environment](#)
- [Tuning Identity Manager Performance](#)
- [Debugging Performance Issues](#)

# Before You Begin

Review the information in these sections before starting to tune Identity Manager:

- [Intended Audience](#)
- [Important Notes](#)
- [Related Documentation and Web Sites](#)

## Intended Audience

This chapter is intended for application server and database administrators, front line support engineers, and partners who are interested in optimizing Identity Manager performance.

Before tuning Identity Manager, you must

- Be familiar with tuning application servers
- Be familiar with Java 5.0 (required for Identity Manager 8.0)
- Understand performance limitations within your deployment environment
- Be able to identify areas needing performance improvements
- Understand the checklists provided in this chapter

## Important Notes

You should be aware of the following information before trying to tune Identity Manager performance:

- The tuning methods described in this chapter are only provided as *guidelines*. You might have to modify some of these tunings for your deployment.
- Be sure to validate tunings *before* applying changes in a production environment.

## Related Documentation and Web Sites

In addition to the information provided in this chapter, consult the publications and web sites listed in this section for information related to tuning Identity Manager.

## Recommended Reading

See the following publications for information related to performance tuning.

**Table 1-1** Related Documentation

Publication Title	Description
<i>IBM Developer Kit and Runtime Environment, Java Technology Edition, Version 5.0 Diagnostics Guide</i>	Explains how to use AIX JVM to diagnose performance problems. Go to the following web site: <a href="http://download.boulder.ibm.com/ibmdl/pub/software/dw/jdk/diagnosis/diag50.pdf">http://download.boulder.ibm.com/ibmdl/pub/software/dw/jdk/diagnosis/diag50.pdf</a>
<i>Java Tuning White Paper</i>	Contains information, techniques, and pointers related to Java performance tuning. Go to the following web site: <a href="http://java.sun.com/performance/reference/whitepapers/tuning.html">http://java.sun.com/performance/reference/whitepapers/tuning.html</a>
<i>Oracle MetaLink: Note:114671.1: Gathering Statistics for the Cost Based Optimizer</i>	Explains how to use system statistics and Oracle's Cost-Based Optimizer. This publication is available to Oracle Metalink subscribers from the following web site: <a href="https://metalink.oracle.com/">https://metalink.oracle.com/</a> <b>Note:</b> Registration is required.
<i>Solaris Dynamic Tracing Guide</i>	<a href="http://docs.sun.com/app/docs/doc/817-6223">http://docs.sun.com/app/docs/doc/817-6223</a>
<i>Sun Java™ System Application Server Performance Tuning Guide</i>	Describes how to obtain optimal performance from your Sun Java System Application Server. Go to the Sun Microsystems Documentation web site: <a href="http://docs.sun.com">http://docs.sun.com</a>
<i>Tuning Garbage Collection with the 5.0 Java™ Virtual Machine</i>	Describes how to tune your garbage collector application by using JVM. Go to the following web site: <a href="http://java.sun.com/javase/technologies/hotspot/gc/index.jsp">http://java.sun.com/javase/technologies/hotspot/gc/index.jsp</a>
<i>Turbo-charging Java HotSpot Virtual Machine, v1.4.x to Improve the Performance and Scalability of Application Servers</i>	Explains how to download and use the <code>PrintGCStats</code> script and how to collect statistics to derive optimal JVM tunings. Go to the following web site: <a href="http://java.sun.com/developer/technicalArticles/Programming/turbo/#PrintGCStats">http://java.sun.com/developer/technicalArticles/Programming/turbo/#PrintGCStats</a>
<i>Understanding System Statistics</i>	Describes how Oracle's Cost-Based Optimizer (CBO) uses system statistics. Go to the following web site: <a href="http://www.oracle.com/technology/pub/articles/lewis_cbo.html">http://www.oracle.com/technology/pub/articles/lewis_cbo.html</a>
<i>Using JConsole to Monitor Applications</i>	Describes how to use JConsole to monitor applications that run on the Java platform. Go to the following web site: <a href="http://java.sun.com/developer/technicalArticles/J2SE/jconsole.html">http://java.sun.com/developer/technicalArticles/J2SE/jconsole.html</a>

## Useful Web Sites

The following table describes some web sites you might find useful when trying to tune Identity Manager performance.

**Table 1-2** Useful Web Sites

Web Site URL	Description
<a href="http://sunsolve.sun.com/">http://sunsolve.sun.com/</a>	<p>Sun web site containing diagnostic tools, forums, features and articles, security information, and patch contents.</p> <p><b>Note:</b> The information on this site is partitioned into three areas,</p> <ul style="list-style-type: none"> <li>• Internal: Sun employees only</li> <li>• Contract: Available only to customers with contract access</li> <li>• Public: Available to everyone</li> </ul>
<a href="http://forum.java.sun.com/">http://forum.java.sun.com/</a>	Sun Developer Network (SDN) web site where you can browse forums and post questions.
<a href="http://jrat.sourceforge.net/">http://jrat.sourceforge.net/</a>	JRat web site that describes how to use the Java Runtime Analysis Toolkit, an open source performance profiler for the Java platform.
<a href="https://metalink.oracle.com/">https://metalink.oracle.com/</a>	<p>Oracle's internal forum site that contains information about tuning Oracle databases.</p> <p><b>Note:</b> You must be a Oracle Metalink subscriber to access the information provided on this site.</p>
<a href="http://performance.netbeans.org/howto/jvmswitches/index.html">http://performance.netbeans.org/howto/jvmswitches/index.html</a>	NetBeans web site containing information about tuning JVM switches for performance.
<a href="https://sharespace.sun.com/gm/folder-1.11.60181?">https://sharespace.sun.com/gm/folder-1.11.60181?</a>	<p>Identity Manager link on Sun's Share Space.</p> <p><b>Note:</b> You must sign up for a Share Space ID to access information provided on this site.</p>
<a href="http://sharespace.sun.com/gm/document-1.26.2296">http://sharespace.sun.com/gm/document-1.26.2296</a>	<p>Identity Manager FAQ on Sun's Share Space.</p> <p><b>Note:</b> You must sign up for a Share Space ID to access this FAQ.</p>
<a href="http://www.slamd.com/">http://www.slamd.com/</a>	SLAMD Distributed Load Generation Engine web site.
<a href="http://www.opensolaris.org/os/community/dtrace/">http://www.opensolaris.org/os/community/dtrace/</a>	OpenSolaris Community: DTrace web page.
<a href="http://www.solarisinternals.com/">http://www.solarisinternals.com/</a>	Web sites containing information related to tuning Solaris.
<a href="http://docs.sun.com/app/docs/prod/solaris.10">http://docs.sun.com/app/docs/prod/solaris.10</a>	

# Tuning Roadmap

How well your Identity Manager solution performs can depend on the following deployment-specific settings:

- Resource configuration
- How many resources are connecting
- What type of resources are connecting
- How attributes are mapped on the resources
- Exact resource version
- Network topology
- Number (and distribution) of domain controllers
- Number of installed Identity Manager Gateways
- Number of concurrent settings
- Number of concurrent processes (running workflows)
- Number of concurrent users
- Number of concurrent Identity Manager administrators
- Total number of users under management

When you are trying to debug performance problems, start by analyzing and describing the problem. Ask yourself the following questions:

- Where do you see the performance issue? In reconciliation, workflows, custom workflows, GUI page loading, provisioning, Access Review?
- Are you CPU bound, memory bound, resource bound, or network bound?
- Have you examined your configuration (hardware, network, parameters, and so forth) for problems?
- Have you recently changed anything in your deployment environment?
- Have you tried profiling troublesome resources natively to see if the problem is on the resource side and not with Identity Manager?
- What size are the views?
- Are you provisioning to several resources?

- Are resources on slow networks connecting to Identity Manager?
- Are you running additional applications on the server with Identity Manager?
- Do your organizations have a Rule-Driven Members rule?
- Have you run a series of thread dumps to see if there is a consistent pattern?

---

**NOTE** Looking at just a single thread dump can be misleading.

---

- Have you recently turned on tracing?
- Have you checked your JVM garbage collection?
- Have you added organizations that are adding load to memory management?

## Tuning Your Deployment Environment

This section provides information about tuning your deployment environment, including

- [Tuning Your Java EE Environment](#)
- [Tuning Your Application Server](#)
- [Tuning Your Repository Database](#)

## Tuning Your Java EE Environment

This section describes some tuning suggestions you can use to optimize your Java EE environment.

These tuning suggestions were derived from a series of experiments in which a considerable increase in throughputs was observed for the use cases tested. The increases were attributed to JVM sizing and to switches that affected Garbage Collector behavior.

Suggestions for tuning your Java EE environment are organized into the following sections:

- [Tuning Java](#)
- [Tuning the JVM](#)

## Tuning Java

For information, best practices, and examples related to Java performance tuning, see the *Java Tuning White Paper* provided at the following location:

<http://java.sun.com/performance/reference/whitepapers/tuning.html>

## Tuning the JVM

The following tuning scripts were used to derive the tuning suggestions noted in this section. These scripts were added to the `domain.xml` file (located in the domain configuration directory, which is typically `domain-dir/config`) on a Sun Java™ System Application Server.

- **PrintGCStats:** A data mining shell script that collects data from `verbose:gc` logs and displays information such as garbage collector pause times, parameter calculations, and timeline analyses over the application's runtime by sampling the data at user-specified intervals.

For more information about how to use this script and garbage collection statistics to derive optimal JVM tunings, see the following web site:

<http://java.sun.com/developer/technicalArticles/Programming/turbo/#PrintGCStats>

- **PrintGCDetails:** A shell script that can provide more-verbose garbage collection statistics.
- **PrintGCTimeStamps:** A shell script that adds timestamp information to the garbage collector statistics collected by using the `PrintGCDetails` script.

To improve JVM performance

- Be sure to use the required Java version (as noted in the “Supported Software and Environments” section of the *Sun Java™ System Identity Manager Release Notes*) to ensure you are using the most current features, bug fixes, and performance enhancements.
- Be sure you are using a newer version of garbage collection.

Frequently, customers do not remove the older, default garbage collector scheme when installing an application server. Running Identity Manager with an older garbage collector creates many objects, which forces the JVM to constantly collect garbage.

- If you deployed Identity Manager on Sun Java System Application Server,
  - Increase throughput by adding garbage collection elements to the deployed Identity Manager instance `server.xml` file.
  - If you expect a peak load of more than 300 users, try modifying the following settings to increase performance:

---

**NOTE** Sun Java System Application Server exposes tunables that affect the size of various thread pools and connection queues that are maintained by the HTTP container.

By default, most of these tunables are set for a concurrent user load of 300 users or less.

---

- For HTTP listeners configured for the deployed Identity Manager instance, edit the listener definition element in the `server.xml` file and set the number of acceptor threads to the *Number of active CPUs on the host* divided by the *Number of active HTTP listeners on the host*.

For example

```
<http-listener id="http-listener-1" \address="0.0.0.0"
port="80" \acceptor threads="Calculated Acceptor Threads"
...>
```

- Because the static content of most Identity Manager deployments is not projected to change frequently, you can edit the File Cache settings (on the File Cache Configuration page) for static content. Specify a high number (such as the number of seconds in 24 hours) for the maximum age of content within the file cache before the content is reloaded.

To access the File Cache Configuration page, select the File Caching tab on the web-based Admin Console for the HTTP server node. (See the latest *Sun Java™ System Web Server Administrator's Guide* for detailed instructions.)

---

**NOTE** For more information about tuning Java, JConsole, or JVM visit the web sites noted in [Table 1-1](#) and [Table 1-2](#).

---

# Tuning Your Application Server

The following guidelines are provided to help you tune your application server:

- [Tuning a Sun Java System Application Server](#)
- [Tuning a WebSphere Application Server](#)

---

**NOTE** Other than heap size, you can use the default parameter settings for most application servers. You might want to modify the server's heap size, depending on the release being used.

---

## Tuning a Sun Java System Application Server

The “Tuning the Application Server” chapter, in the latest *Sun Java™ System Application Server Performance Tuning Guide*, contains information about tuning a Sun Java™ System Application Server. This publication is available from

<http://docs.sun.com/app/docs>

- In addition, if you are using Sun Java System Application Server 8.2 Enterprise Edition, the following changes solve “concurrent mode failures,” and should give you better and more predictable performance:
  - If you are constantly running old generation collections, review your application's heap footprint and consider increasing the size.

For example, 500 Mbytes is considered a modest size, so increasing this value to 3 Gbytes might improve performance.

With a 2 Gbytes young generation collection, each scavenge promotes about 70 Mbytes. Consider giving this 70 Mbytes at least one more scavenge before promoting. For example, you might need a SurvivorRatio of

$$2 \text{ GB} / 70 \text{ M} \times 2 = 4096 / 70 = 55$$

Where

`-XX:SurvivorRatio=32 -XX:MaxTenuringThreshold=1`

This ratio prevents premature promotion and the added problem of “nepotism,” which can degrade scavenge performance.

- If you specified `-XX:CMSInitiatingOccupancyFraction=60`, and the CMS collections are still starting before they reach that threshold, such as

```
56402.647: [GC [1 CMS-initial-mark: 265707K(1048576K)]
1729560K(3129600K), 3.4141523 secs]
```

Try removing `-XX:CMSInitiatingOccupancy=60` (using the default value of 69%), and add the following line:

```
-XX:UseCMSInitiatingOccupancyOnly
```

- If your young generation collection is 2 Gbytes and old generation collection is 1 Gbyte, this situation might also be causing premature CMS collections. Consider reversing this ratio. Use a 1 Gbyte young generation collection and a 2 Gbytes old generation collection, as follows:

```
-Xms3G -Xmx3G -Xmn1G
```

Also, remove `-XX:NewRatio`. This ratio is redundant when you have explicitly specified young generation and overall heap sizes.

- If you are using a 5uXX version of the JDK, and have excessively long “abortable preclean” cycles, you can use `-XX:CMSMaxAbortablePrecleanLoops=5` as a temporary workaround

---

**NOTE** You might have to adjust this value further.

---

- Add the following line to view more information about garbage collector performance:

```
-XX:+PrintHeapAtGC
```

---

**NOTE** Use this command with caution because it increases how much verbose garbage collector data is produced. Be sure you have enough disk space on which to save the garbage collector output.

---

- If you are using the Sun Fire T2000 Server, large-heap DTLBs (data Translation Look-aside Buffers) can become a scarce resource. Using large pages as follows for the Java heap often helps performance:

```
-XX:+UseLargePages
```

## Tuning a WebSphere Application Server

If you are tuning Identity Manager on an IBM WebSphere® application server, consider limiting how much memory is allocated for the heap because heap memory can affect the memory used by threads.

If many threads are created simultaneously and the heap size increases, the application's space limit can be quickly impacted and the following error results:

```
JVMCI015:OutOfMemoryError
```

## Tuning Your Repository Database

Identity Manager relies on the repository database to store and manage its identity and configuration data. For this reason, database performance can greatly influence Identity Manager's performance.

---

**NOTE** Detailed information about performance tuning and managing databases is beyond the scope of this publication, because this information is dataset-specific and vendor-specific. In addition, customer database administrators (DBAs) should already be experts on their own databases.

---

This section characterizes the Identity Manager application and provides general information about the nature of Identity Manager data and its typical usage patterns to help you plan, tune, and manage your databases. This information is organized into the following sections:

- [Repository Table Types](#)
- [XML Columns](#)
- [Data Classes](#)
- [Object IDs](#)
- [Prepared Statements](#)
- [Character Sets and Encodings](#)
- [General Tuning Guidelines](#)
- [Vendor-Specific Database Tuning Guidelines](#)

## Repository Table Types

The Identity Manager repository contains three types of tables, and each table has slightly different usage characteristics. Information about these tables is organized into the following sections:

- [Attribute Tables](#)
- [Change Tables](#)
- [Object Tables](#)

### *Attribute Tables*

Attribute tables enable you to query for predefined single-valued or multivalued object attributes.

For most object types, stored attributes are hard-coded.

---

**NOTE** The `User` and `Role` object types are exceptions to this rule. The inline attributes that are stored in the object table for `User` and `Role` are configurable, so you can configure additional custom attributes as queryable.

---

When you search for objects based on attribute conditions, Identity Manager accesses attribute tables in joins with the corresponding object tables. Some form of join (such as a `JOIN`, an `EXISTS` predicate, or a `SUB-SELECT`) occurs for each attribute condition.

The number of rows in the attribute table are proportional to the number of rows in the corresponding object table. The values distribution might exhibit skew, where multivalued attributes have a row per value and some objects might have more attributes or more attribute values than others. Typically, there is a many-to-one relation between rows in the attribute table and rows in the object table.

Attribute tables have `ATTR` in the table name.

### *Change Tables*

Identity Manager uses a change table to track changes made to a corresponding object table. These table sizes are proportional to the rate of object change, but the tables are not expected to grow without bound. Identity Manager automatically truncates change tables.

Change tables can be highly volatile because the lifetime of a row is relatively short and new rows can be created frequently.

Access to a change table is typically performed by a range scan on the time-stamp field.

Change tables have `CHANGE` in the table name.

### *Object Tables*

The Identity Manager repository uses object tables to hold serialized data objects, such as Large Objects (LOBs). Object tables can also hold commonly queried, single-valued object attributes.

For most object types, stored attributes are hard-coded.

---

**NOTE** The `User` and `Role` object types are exceptions to this rule. The inline attributes that are stored in the object table are configurable, and you can configure additional custom attributes as queryable for `User` and `Role`.

---

The number of rows in an object table equals the number of objects being stored. The number of objects stored in each object table depends on which object types are being stored in the table. Some object types are numerous, while other types are few.

Generally, Identity Manager accesses an object table by object ID or name, though Identity Manager can also access the table by using one of the attributes stored in the table. Object IDs and names are unique across a single object type, but attribute values are not unique or evenly distributed. Some attributes have many values, while other attributes have relatively few values. In addition, several object types can expose the same attribute. An attribute may have many values for one object type and few values for another object type. The uneven distribution of values might cause an uneven distribution of index pages, which is a condition known as *skew*.

Object tables are tables that *do not* have `ATTR` or `CHANGE` suffixes in the table name.

## XML Columns

Every object table contains an XML column, which is used to store each serialized object — with the exception of the LOG table-set. Certain LOG table-set optional attributes are stored in the XML column if these attributes are present. For example, if digital signing is enabled.

## Data Classes

You can roughly divide Identity Manager data into a number of classes that exhibit similar properties with respect to access patterns, cardinality, lifetime, volatility, and so forth. Each of the following classes corresponds to a set of tables in the repository.

- [User Data](#)
- [Role Data](#)
- [Account Data](#)
- [Compliance Violation Data](#)
- [Entitlement Data](#)
- [Organization Data](#)
- [Task Data](#)
- [Configuration Data](#)
- [Export Queue Data](#)
- [Log Data](#)

### *User Data*

User data consists of user objects.

You can expect this data to grow quite large because there is an object for each managed identity. After an initial population phase, you can expect a proportionally small number of creates because the majority of operations will be updates to existing objects.

User objects are generally long-lived and they are removed at a relatively low rate.

User data is stored in USEROBJ, USERATTR, and USERCHANGE tables.

### *Role Data*

Role data consists of `Role` objects, including Roles subtypes such as Business Roles, IT Roles, Applications, and Assets.

Role data is similar to organization data, and these objects are relatively static after a customer deploys Identity Manager.

---

**NOTE** An exception to the preceding statement is a deployment that is integrated with an external source containing an authoritative set of roles. One integration style might be to feed role changes into Identity Manager, which causes Identity Manager Role data to be more volatile.

---

Generally, the number of role objects is small when compared to the number of identity objects such as users (assuming that multiple users share each role), but this depends on how each enterprise defines its roles.

Role data is stored in the `ROLEOBJ`, `ROLEATTR`, and `ROLECHANGE` tables.

### *Account Data*

Account data solely consists of account objects in the Account Index.

As with user data, you expect account data to become rather large, with an object for each known resource account. Account objects are generally long-lived, removed at a relatively low rate, and after initial population, are created infrequently. Unless you frequently add or remove native accounts, or enable native change detection, account objects modifications occur infrequently.

Identity Manager stores account data in `ACCOUNT`, `ACCTATTR`, and `ACCTCHANGE` tables.

### *Compliance Violation Data*

Compliance Violation data contains violation records that indicate when the evaluation of an Audit Policy failed. These violation records exist until the same Audit Policy is evaluated against the same User and policy passes. Violation records are created, modified, or deleted as part of an Audit Policy Scan or as part of an Access Review.

The number of violation records is proportional to the number of Audit Policies that are used in scans and the number of Users. An installation with 5000 users and 10 Audit Policies might have 500 violation records ( $5000 \times 10 \times 0.01$ ), where the 0.01 multiplier depends on how strict the policies are and how user accounts are changed.

Identity Manager stores Compliance Violation records in `OBJECT`, `ATTRIBUTE`, and `OBJCHANGE` tables.

### *Entitlement Data*

Entitlement data predominately consists of user entitlement objects, which are only created if you are doing compliance access reviews.

Entitlement records are created in large batches, modified slowly (days) after initial creation, and are then untouched. These records are deleted after an Access Review is deleted.

Identity Manager stores entitlement data in `ENTITLE`, `ENTATTR`, and `ENTCHANGE` tables.

### *Organization Data*

Organization data consists of object group or organization objects.

Object group data is similar to configuration data, and this data is relatively static after being deployed. Generally, the number of objects is small (one for each defined organization) when compared to task objects or to identity objects such as users or accounts, but the number can become large compared to other configuration objects.

Organization data is stored in `ORG`, `ORGATTR`, and `ORGCHANGE` tables.

### *Task Data*

Task data consists of objects that are related to tasks and workflows, including state and result data.

The data contained in these tables is short-lived compared to other classes because objects are created, modified, and deleted at a high rate. The volume of data in this table is proportional to the amount of activity on the system.

Task data is stored in `TASK`, `TASKATTR`, and `TASKCHANGE` tables.

### *Configuration Data*

Configuration data consists of objects related to Identity Manager system configuration, such as forms, roles, and rules.

Generally, Configuration data is:

- Relatively small compared to other classes
- Only expected to change during deployment and upgrade, and changes occur in large batches
- Not expected to change much after being deployed

Identity Manager stores configuration data in `ATTRIBUTE`, `OBJCHANGE`, and `OBJECT` tables.

### *Export Queue Data*

If you enable Data Exporting, some records are queued inside Identity Manager until the export task writes those records to the Data Warehouse. The number of records that are queued is a function of Data Exporting configuration and the export interval for all queued types. The following data types are queued by default, and all other data types are not:

- ResourceAccount
- WorkflowActivity
- TaskInstance
- WorkItem

The number of records in these tables will grow until the export task drains the queue. The current table size is visible through a JMX Bean.

Records added to this table are never modified. These records are written during other Identity Manager activities, such as Reconciliation, Provisioning, and Workflow Execution. When the Data Exporter `export` task runs, the task drains the table.

Identity Manager stores Export Queue Data records in `QUEUE`, `QATTR`, and `QCHANGE` tables.

### *Log Data*

Log data consists of audit and error log objects. Log data is write-once only, so you can create new audit and error log objects, but you cannot modify these objects.

Log data is long-lived and can potentially become very large because you can only purge log data by explicit request. Access to log data frequently relies on attributes that are stored in the object table instead of in the attribute table. Both the distribution of attribute values and queries against the log specifically depend on how you are using Identity Manager.

For example, the distribution of attribute values in the log tables depends on

- What kind of changes are made
- Which Identity Manager interfaces was used to make the changes
- Which types of objects were changed

The pattern of queries against the log table also depends on which Identity Manager reports, which custom reports, or which external data mining queries a customer runs against the log table.

Identity Manager stores audit log records in `LOG` and `LOGATTR` tables and error log records in `SYSLOG` and `SLOGATTR` tables. This data does not have corresponding change tables.

## Object IDs

Identity Manager generates globally unique identifiers (GUIDs) for objects by using the `VMID` class provided in the JDK.

These GUID values exhibit a property that gets sorted by their string representations, based on order in which the objects are created. For example, when you create new objects with Identity Manager, the newer objects have object IDs that are greater than the older objects. Consequently, when Identity Manager inserts new objects into the database, the index based on object ID can encounter contention for the same block or blocks.

## Prepared Statements

Generally, Identity Manager uses prepared statements for activities (such as inserting and updating database rows), but does not use prepared statements for queries.

If you are using Oracle, this behavior can create issues with the library cache. In particular, the large number of statements versions can cause contention on the library cache latch.

To address this contention, change Oracle's `CURSOR_SHARING` parameter value from `EXACT` to `SIMILAR`. Changing this value causes Oracle to replace literals in SQL statements with bind variables, thereby reducing the number of versions.

## Character Sets and Encodings

Identity Manager does not restrict which encoding the database uses because Identity Manager is a Java application that generally reads and writes character data rather than bytes.

Identity Manager only requires that the data is sent and returned correctly. For example, the data does not become corrupted when written or re-read. Use an encoding that supports multibyte characters and is appropriate for the customer's data. Generally, UTF-8 encoding is sufficient, but enterprises with a large number of true multibyte characters, such as Asian or Arabic, might prefer UTF-16.

Most database administrators prefer to use an encoding that supports multibyte characters because

- Their deployments often grow to support international characters
- Their end users cut-and-paste from Microsoft application's text containing characters that look like ASCII but are actually multibyte, such as em dashes (–).

## General Tuning Guidelines

This section describes some general guidelines for tuning a repository database:

- DBAs must frequently run statistics to monitor what is happening with the repository database.
- If you are using a data source, set the `connectionPoolDisable` attribute to `true` in the `RepositoryConfiguration` object:

```
<RepositoryConfiguration connectionPoolDisable='true'>
```

Setting this attribute to `true` disables automatic internal connection pooling in the Identity Manager repository so you avoid having two connection pools (one for Identity Manager and one for your application server).

- You can edit the `RepositoryConfiguration` object to enhance the performance of searches against specific, single-valued attributes. For example, an administrator might edit this object to add an extended attribute, such as `employeeID`, that is used to search for Users or as a correlation key.

The default `RepositoryConfiguration` object looks like the following example:

---

**NOTE**      The ellipses represent XML attributes that are not relevant here.

---

```
<RepositoryConfiguration ... >
  <TypeDataStore Type="User" ... attr1="MemberObjectGroups",
  attr2="lastname" attr3="firstname" attr4="" attr5="">
  </TypeDataStore>
</RepositoryConfiguration>
```

Each of the `attr1`, `attr2`, `attr3`, `attr4`, and `attr5` XML attributes specifies a single-valued attribute to be copied into the `waveset.userobj` table. The `waveset.userobj` table can contain up to five inline attributes. The attribute value named by `attr1` in `RepositoryConfiguration` will be copied into the "attr1" database column in this table.

---

**NOTE** In previous releases, you specified inline attributes in the `UserUIConfig` object. As of Identity Manager 8.0, inline attributes are stored in the base object table for a `Type` (rather than as rows in the child attribute table).

---

Using inline attributes improves the performance of repository queries against those attributes. (Because inline attributes reside in the main "object" table, queries against inline attributes are faster than those against non-inline attributes, which are stored in the child "attribute" table. A query condition against a non-inline attribute requires a "join" to the attribute table.)

By default, Identity Manager inlines `MemberObjectGroups`, `lastname`, and `firstname`.

- You can add two more attributes to enable faster searching, as long as those attributes are queryable.

For example, if your deployment contains an `employeeID` extended attribute, inlining that attribute will improve the performance of repository searches against that attribute.

- If you do not need `lastname` or `firstname`, you can remove them or replace them with other attributes.
- *Do not* remove `MemberObjectGroups`. Identity Manager uses this attribute internally to speed up authorization checks.

---

**NOTE** For more information about which object types are stored in each set of tables, see "[Data Classes](#)" on page 14.

---

## Vendor-Specific Database Tuning Guidelines

This section describes some vendor-specific guidelines for tuning Oracle and SQL Server repository databases.

---

**NOTE** Currently, MySQL databases are only supported in development and for demonstrations.

---

### *Oracle Databases*

This section describes guidelines for tuning Oracle repository databases:

- The Identity Manager application does not require Oracle database features or options.
- If you are using an Oracle repository database and Service Provider or Identity Manager, you might encounter problems with object table fragmentation because Identity Manager uses `LONG` (rather than `LOB`), data types by default. Using `LONG` data types can result in large amounts of “unallocated” extent space, which cannot be made into usable space.

To mitigate this problem

- Take `EXPORT` dumps of the Object table and reimport them to free up unallocated extent space. After importing, you must stop and restart the database.
- Use `LOB` data types and DataDirect Technologies’ Merant drivers, which provide a standard `LOB` implementation for Oracle.
- Use Locally Managed Tablespaces (LMTs), which offer automatic free space management. LMTs are available in Oracle 8.1.5.
- Identity Manager does not require Oracle `init.ora` parameter settings for SGA sizing, buffer sizing, `open_cursors`, `processes`, and so forth.
- Identity Manager's repository is a general-purpose database; but in general, Identity Manager's repository is best described as an object database.

Of the Identity Manager tables, the `TASK` table-set comes closest to having transaction-processing characteristics. The `LOG` and `SYSLOG` table-sets are also exceptional because these tables do not store serialized objects.

---

**NOTE** See “[Repository Table Types](#)” on page 12 and “[Data Classes](#)” on page 14 for descriptions of the tables, the object types stored in each table, and the general access pattern for each table.

---

- If you have performance issues with Oracle, check for issues related to poor query plans being chosen for what Identity Manager expects to be relatively efficient queries.

For example, Identity Manager is configured to perform a full table-scan when an index is available for use. These issues are often visible in Automated Workload Repository (AWR) reports provided in the SQL by the `buffer gets` table. You can also view issues in the Enterprise Manager tool.

Performance problems typically appear to be the result of bad or missing database table statistics. Addressing this problem improves performance for both the database and Identity Manager.

The following articles (available from Oracle) are a good source of information about the cost-based optimizer (CBO) in Oracle:

- *Understanding System Statistics*
- *Oracle MetaLink: Note:114671.1: Gathering Statistics for the Cost Based Optimizer*
- *Cost Control: Inside the Oracle Optimizer*

You might also investigate using SQL Profiles, which are another method for choosing the best query plans. You can use the SQL Advisor within Enterprise Manager to create these profiles when you identify poorly performing SQL.

- If you detect unexpected growth in the Oracle redo log, you might have workflows that are caught in an infinite loop with a Manual Action. The loop causes constant updates to the repository, which in turn causes the size of the `TaskInstances` to grow substantially. The workflow errors are caused by improper handling of `WF_ACTION_TIMEOUT` and users closing their browser in the middle of a workflow.

To prevent problematic workflows, preview each Manual Action before a production launch and verify the following:

- Have you set a timeout?
- Have you created appropriate transition logic to handle a timeout for the activity with the Manual Action?
- Is the manual action using the exposed variables tag when there is a large amount of data in the TaskInstance?
- Frequently, you can significantly improve Identity Manager performance if you change the `CURSOR_SHARING` parameter value from `EXACT` to `SIMILAR`.

Identity Manager uses prepared statements for some activities (such as inserting and updating database rows), but does not use these statements for most queries.

When you use Oracle, this behavior can cause issues with the library cache. In particular, the large number of statements versions can create contention on the library cache latch. Changing `CURSOR_SHARING` to `SIMILAR` causes Oracle to replace literals in SQL statements with bind variables, which greatly reduces the number of versions.

---

**NOTE** See [“Prepared Statements” on page 18](#) for more information.

---

### *SQL Server Databases*

Some customers who used a SQL Server 2000 database as a repository reported that as concurrency increased, SQL Server 2000 reported *deadlocking* problems that were related to SQL Server's internal use of pessimistic locking (primarily lock escalation).

These deadlock errors display in the following format:

```
com.waveset.util.IOException:
==> com.microsoft.sqlserver.jdbc.SQLServerException: Transaction (Process
ID 51) was deadlocked on lock | communication buffer resources with another
process and has been chosen as the deadlock victim. Rerun the transaction.
```

To prevent or address deadlocking problems,

1. Use the SQL Server 2005 database.
2. Configure the `READ_COMMITTED_SNAPSHOT` parameter. Format the command as follows:

```
ALTER DATABASE waveset SET READ_COMMITTED_SNAPSHOT ON
```

Enabling the `READ_COMMITTED_SNAPSHOT` parameter

- Removes contention during the execution of `SELECT` statements that can cause blocks, which greatly reduces the potential for deadlocks internal to SQL Server.
- Prevents uncommitted data from being read and guarantees that `SELECT` statements receive a consistent view of committed data.

---

**NOTE** Go to following URL for more information about the `READ_COMMITTED_SNAPSHOT` parameter:

<http://msdn2.microsoft.com/en-us/library/ms188277.aspx>

---

# Tuning Identity Manager Performance

Suggestions for optimizing Identity Manager's performance are organized into the following areas:

- [General Performance Tunings](#)
- [Tuning Active Sync Adapter Performance](#)
- [Tuning Bulk Loading](#)
- [Tuning Configurable XML Objects](#)
- [Tuning Database Statistics](#)
- [Tuning Data Exporter](#)
- [Tuning the General XML](#)
- [Tuning Identity Manager Service Provider](#)
- [Tuning the Identity Manager Web Interface](#)
- [Tuning Initial Loads](#)
- [Tuning Memory Requirements](#)
- [Tuning Operating System Kernels](#)
- [Tuning Provisioner](#)
- [Tuning Reconciliation](#)
- [Tuning Resource Queries](#)
- [Tuning the Scheduler](#)
- [Tuning Sessions](#)
- [Tuning the Sun Identity Manager Gateway](#)
- [Tuning the Task Bar](#)

## General Performance Tunings

In general, you can optimize Identity Manager performance if you

- Turn off tracing (such as Java class, userform, and workflow tracing). Tracing can add substantial overhead.
- Run Identity Manager’s built-in Audit Log Maintenance Task and System Log Maintenance Task to configure log record expirations. Log records can grow without bound, so use these tasks to prevent the repository database from running out of space. For information, see *Identity Manager Administration*.
- Check the README file in Identity Manager updates (formerly called *service packs* or *installation packs*) to see if any performance improvements have been made to the product. If so, schedule an upgrade.
- Consider the performance impact when fetching data from one or more remote systems, including the Identity Manager repository.
- Increase the number of application server instances running Identity Manager, either on the same server or by adding servers, and use a load-balancing tool to distribute the requests between instances.
- Keep the size of files referenced in a binary attribute as small as possible. Loading extremely large graphics files, for example, can decrease Identity Manager performance.
- Write robust and readable XML code that minimizes duplication (for example, refactored), that uses memory efficiently, and mitigates the impact to overall system performance.
- Configure Identity Manager system monitoring to track events in real-time.

You can view these events in dashboard graphs to quickly assess system resources, spot abnormalities, understand historical performance trends (based on time of day, day of week, etc.), and interactively isolate problems before looking at audit logs. Dashboards do not provide as much detail as audit logs, but can provide hints about where to look for problems in the logs.

For more information about dashboards, see the “Reporting” chapter in *Identity Manager Administration*.

# Tuning Active Sync Adapter Performance

Because synchronization is a background task, how you configure an Active Sync adapter can affect server performance.

Use the Resources list to manage Active Sync adapters. Choose an Active Sync adapter and access start, stop, and status refresh control actions from the *Synchronization* section of the Resource Actions list.

To improve Active Sync adapter performance,

- Evaluate and adjust polling intervals based on the type of activity being performed.

The polling interval determines when the Active Sync adapter will start processing new information. For example, if the adapter reads in a large list of users from a database and updates these users in Identity Manager each time, you could run this process in the early morning every day. Some adapters have a quick search for new items to process and can be set to run every minute.

- Edit the synchronization file for the resource to specify the host where the adapters will run.

You can configure Active Sync adapters requiring more memory and CPU cycles to run on dedicated servers to help load balance the systems.

- If you have the appropriate administrator capability, you can change Active Sync resources to disable, manually start, or automatically start Active Sync adapters.

When you set an adapter to automatic, the adapter restarts when the application server starts. When you start an adapter, it runs immediately and executes at the specified polling interval. When you stop an adapter, it stops the next time the adapter checks for the stop flag.

- Adjust the level of detail captured by the synchronization logs.

Synchronization logs capture information about the resource that is currently processing. Each resource has its own log file, path, and log level. The amount of detail captured by the adapter log depends on the specified logging level. You specify these values in the Logging section of the Synchronization Policy for the appropriate user type (Identity Manager or Service Provider).

## Tuning Bulk Loading

To improve performance during bulk load operations,

- Simplify default workflows to improve processing time (especially for bulk processing actions such as Active Sync, bulk actions, and reconciliation) by removing the callout to the Approval subprocess.
- Keep user forms that are assigned to administrators as simple as possible. For example,
  - When creating a form for data loading, remove any code that is designed to display data.
  - When using bulk add actions, be sure your CSV file defines basic attributes such as `firstname` and `lastname`. You can then remove these attributes from the administrator's user form.

---

**NOTE** Do not modify the default forms provided with Identity Manager. Instead, make a copy of the form, give the copy a unique name, and modify the renamed copy. This approach prevents your customized forms from being overwritten during upgrades and product updates.

See the “Identity Manager Forms” chapter in *Identity Manager Workflows, Forms, and Views* for more information about creating and editing forms.

---

- Implement the following features in deployment environments where you have an NIS (Network Information Service) implemented:
  - Add an account attribute named `user_make_nis` to the schema map and use this attribute in your reconciliation or other bulk provisioning workflow. Specifying this attribute causes the system to bypass the step of connecting to the NIS database after each user update on the resource.
  - To write the changes to the NIS database after provisioning has completed, create a ResourceAction named `NIS_password_make` in the workflow.

## Tuning Configurable XML Objects

Configurable XML objects offer a broad spectrum of user interface specifications that enable you to define how data is presented to users for different tasks and automate complex business processes. However, this same flexibility can affect efficiency, performance, and reliability.

This section describes some guidelines for tuning Identity Manager's configurable XML objects, which consist of forms, rules, and workflows. The information is organized into the following sections:

- [Tuning Forms](#)
- [Tuning Rules](#)
- [Tuning Workflows](#)

### Tuning Forms

You can use Identity Manager forms to define interfaces to interact with views or variable contexts in an executing task. Forms also provide an execution context for business and transformation logic on a set of data elements. Although you can create very powerful, dynamic forms that perform a variety of tasks, reducing the complexity of the form increases efficiency.

The following sections describe some methods for improving the performance of your customized forms

- [Optimizing New Forms](#)
- [Optimizing Admin Forms](#)
- [Optimizing End User Forms](#)
- [Optimizing Expressions in Form Fields](#)

### *Optimizing New Forms*

When designing new Identity Manager forms, system integrators can optimize a form's performance by

- Performing “expensive” queries only one time, wherever possible. To minimize these queries,
  - Use `<Field>` `<Default>` elements to execute and store query results.
  - Use field names to reference values in later fields.
  - For custom tasks
    - Calculate the value in the task before `ManualAction`, then store that value in a task variable.
    - Use `variables.tmpVar` to reference variables in the form.
    - Use `<setvar name='tempVar' />` to clear the variable after `ManualAction`.
- Use `<defvar>` for calculations that are performed for the initial display and with each refresh.

### *Optimizing Admin Forms*

To improve the performance of Admin forms:

- Specify `TargetResources` that only fetch specific resources for editing. (See [“Tuning Workflows” on page 32](#) for more information.)
- Use `cacheList` and `cacheTimeout` caching parameters for objects that change infrequently if you are working with `FormUtil.getResourceObjects` or `FormUtil.listResourceObjects`.
- Store the results of time-consuming calculations and fetches in `<Field>` elements and evaluate in the `<Default>` expression to help ensure an operation occurs only one time.
- Use `update.constraints` to limit which resources are fetched at runtime (see [Dynamic Tabbed User Form](#))
- Use background approval (`ManualAction` with different owners and one-second timeouts) for faster page submissions.
- Be aware that Identity Manager refreshes *all fields* defined on *all panels* of a `Tab Panel Form` when the page reloads, regardless of which panel is selected.

### *Optimizing End User Forms*

To improve the performance of End User forms:

- Use `TargetResources` to limit view checkouts to just those resource accounts of interest, which reduces fetch time for view and the memory consumed by `TaskInstance` and `WorkItems`
- Consider using `Session.getObject(Type, name)` to return a `WSUser` if just the view properties and attributes of the Identity Manager user object are of interest (useful for managing multiple deferred task triggers).
- Be aware that End User tasks typically have more `WorkItems` than Provisioning tasks, so End User tasks are especially susceptible to `WorkItem` size.
- Consider using temporary generic objects for “view” editing that is constructed on view checkout then merged back into a full view for check in.
- Consider using scalable forms instead of the default Create and Edit User interfaces.

When you use the default User forms to edit a user, Identity Manager fetches the resources owned by that user the moment you start editing the user’s account. In deployment environments where users have accounts on many resources, this potentially time-intensive operation can result in performance degradation.

### *Optimizing Expressions in Form Fields*

Some activities performed in forms call resources that are external to Identity Manager. Accessing these resources can affect Identity Manager performance, especially if the results contain long lists of values, such as compiling a list of groups or email distribution lists.

To improve performance during these calls, follow the guidelines in “Using a Java Class to Obtain Field Data” in *Identity Manager Workflows, Forms, and Views*.

Also, avoid using JavaScript in performance-critical expressions such as `<Disable>` expressions. Short XPRESS expressions are easier to debug if you use the built-in tracing facilities. Use JavaScript for complex logic in workflow actions.

If a form is slow to display, you can use the `debug/Show_Timings.jsp` page to determine the problem. Look for calls to `Formconvert.convertField()`. This method shows how long each field took to compute its value.

## Tuning Rules

You use Identity Manager rules to encapsulate constants and XPRESS logic that can be reused in forms, workflows, and other configurable components in the product.

When writing rules, use the following guidelines (as applicable) to obtain optimal performance:

- Use static declarations to return a constant value.
- Use `<defvar>`s to implement algorithms with temporary values for incremented values or for values that are referenced only one time.
- Use `putmap`, `setlist`, or `setvar` methods for complex or expensive calculations whose value must be returned multiple times. Be sure to eventually set the value to `<null>`.

## Tuning Workflows

You customize Identity Manager workflows to facilitate and automate complex business processes with various human and electronic touchpoints.

You can use the following methods to improve custom workflow performance:

- Simplify default workflows to improve processing time (especially for bulk processing actions such as Active Sync, bulk actions, and reconciliation) by removing the callout to the Approval subprocess.
- Ensure that no infinite loops exist in your workflows. In particular, be sure break flags are updated and properly checked in the loops that exist in approval subprocesses.
- Put fetched objects into a variable for use later if you must contact the repository for the same object multiple times.

Using a variable is necessary because Identity Manager does not cache all objects.

- Specify `TargetResources` options in `WorkflowServers checkoutView` to restrict the number of resources that are queried for account information.

The following example shows how to restrict the number of resources being queried for account information.

```

<Argument name='TargetResources'>
  <list>
    <string>resource name[| #]</string>
  </list>
</Argument>

```

---

**NOTE** In the preceding example, [| #] is an optional parameter that you can use when more than one account exists on a particular resource. In most cases, the resource name is sufficient.

---

- Clear unnecessary view variables left by forms, especially large maps and lists. For example

```
<setvar name='myLargeList'></null></setvar>
```

View is copied multiple times in a `TaskInstance` object, so large views greatly increase the size of `TaskInstances` and corresponding `TaskResults`.

- Use `resultLimit` (in seconds) in the `TaskDefinition`, or set this option during task execution to quickly dispose of completed tasks. Large numbers of `TaskInstances` impact
  - How `taskResults.jsp` in footers and some JSPs in the Administrator interface are displayed
  - How JSP tasks are displayed
  - Querying `TaskInstances` for task renaming
  - Database size
- Set the following options as needed:
  - **delete** (*preferred option*): Causes an older `TaskInstance` of the same name to be deleted before the new task begins execution.
  - **wait**: Suspends the current `TaskInstance` until the older `TaskInstance` is deleted or expired due to reaching its `resultLimit`.
  - **rename**: Inserts a time stamp into the `TaskInstance` name to avoid naming collisions.
  - **terminate**: Deletes an older `TaskInstance` of the same name. Any currently executing `TaskInstances` of the same name are terminated.

## Tuning WorkItems (ManualActions)

The number and size of `WorkItems` (indicated by `ManualActions` in a workflow) can affect memory and system performance *significantly*.

By default, Identity Manager copies an entire workflow context into a `WorkItem`, then writes the workflow context back out after submission.

To improve performance for `ManualActions` and `WorkItems`:

- Reduce the size of `WorkItems`

By default, `ManualAction` creates a `WorkItem`, then copies each variable in the task context into `WorkItem.variables`. Limiting task context variables prevents overwrites from parallel approvals.

- Use `ExposedVariables` to limit which variables are copied back into `WorkItem`. For example

```
<ExposedVariables><List><String>user ...
```

- Use `EditableVariables` to limit the variables assimilated back into the executing task from `WorkItem`. For example

```
<EditableVariables><List><String>user ...
```

Remember to include an approval flag, a form button value, and the actual approver's name.

- Change the confirmation page and background processing to improve user interface response time:
  - Create a confirmation `ManualAction` or background `ManualAction`, owned by another user such as `Configurator`.
  - Set `timeout='-5'` (5 seconds) and `ignoreTimeout='true'` to prevent an error message if a user submits an action after the task is executed and the `WorkItems` are deleted.
- Optimize memory use by setting large attribute values (such as value maps and lists) to `<null>` on submission or instantiate them as `Activity-scoped` variables that quickly pass out of scope.
- Shorten the lifetime of finished tasks
  - Prevent dead-end tasks by ensuring that each `WorkItem` specifies a `Timeout` and the workflow anticipates a `Timeout` for each `WorkItem`.

- Consider using `resultLimit` and `resultOption` options in the `TaskDefinition` to determine how the Scheduler handles a task after the task completes.
  - Use `resultLimit` to control how many seconds a task is allowed to live after the task has completed. The default is zero (0), which means the task instance will be deleted immediately after task completion.
  - Use `resultOption` to specify what action to take when repeated instances of a task are started (such as `wait`, `delete`, `rename`, or `terminate`). Default is `delete`.

If you want to immediately delete tasks that complete successfully, but you also want to keep tasks containing errors long enough to debug, you can conditionally delete finished tasks.

Set a `resultLimit` in the `TaskDefinition` to a sufficient time period to debug issues. You can set `resultLimit` to zero (or a small value) if no errors are reported at runtime (such as `WF_ACTION_ERROR` is `<null/>`) after a `WorkflowServices` call.

- Evaluate and fix poorly scoped variables. Scope variables according to where they are declared:
  - **Global variables** are values that must be used across many activities (such as the case owner, view) and as approval flags in subprocesses.
 

If a variable is declared as an element of `<TaskDefinition>`, scope the variable globally. If a variable is declared `external`, its value is resolved up the call stack.
  - **Activity variables** of expensive values (such as those variables requiring a resource fetch or storing a large list or map of values) that may be referenced in a `WorkItem`.

If a variable is declared as an element of `<Activity>`, ensure the variable is visible to actions and transition elements in `Activity`.

---

**NOTE** Use `<Activity>` variable values in Forms, not workflows, to avoid copying values on `WorkItem` creates (only for Identity Manager 2005Q3M1 SP1 and later).

---

- **Activity variables** are values used in transition logic.

If a variable is declared as an element of `<Action>`, the variable should pass out of scope on completion of the action. Action variables are typically used in `WorkflowApplication` invocations to “change” the names of variables set by the application (such as `View > User`).

- Do not specify synchronous execution (`syncExec='true'`) for the last page in a wizard workflow.

If set to `true`, the task will run to completion when the user executes the task. The interface will hang until the task completes or encounters another stopping point (such as another `ManualAction`).

- Remove unnecessary approval checks.

---

**NOTE** For Active Sync, use a streamlined provisioning task in place of the system-specified provisioning task specified by `viewOptions.Process`.

---

- Do not modify the provisioning tasks provided with Identity Manager.

You must create a new task, then identify that task in the form and in the process mappings configuration (unless the task is not listed).

## Tuning Database Statistics

As a database administrator, you should frequently run statistics to monitor your repository database.

Performance problems are often caused by bad or missing database table statistics. Fixing this problem improves performance for both the database and Identity Manager performance.

See the following Oracle articles for more information. Website locations are provided in [Table 1-1](#):

- *Understanding System Statistics*
- *Oracle MetaLink: Note:114671.1: Gathering Statistics for the Cost Based Optimizer*
- *Cost Control: Inside the Oracle Optimizer*

Also consider using SQL Profiles, which is another method for choosing the best query plans. You can use the SQL Advisor within Enterprise Manager to create these profiles when you identify poorly performing SQL.

## Tuning Data Exporter

Data Exporter enables you to export new, changed, or deleted Identity Manager data to an external repository that is suitable for reporting or analytic work. The actual exporting of data is done in batches, where each type of data to be exported is able to specify its own export cycle. The data to be exported comes from the Identity Manager repository and, depending on the length of the export cycle and the amount of changed data, the volume of exported data can be large.

Some Identity Manager data types are queued into a special table for later export. Specifically, `WorkflowActivity` and `ResourceAccount` data is queued because this data is not persisted otherwise. Any persisted data type can also be queued if the warehouse needs to see all changes to the type, or if the type has a lifecycle that does not correspond to the export cycle, such as `TaskInstance` and `WorkItem` data.

To maximize performance, only queue and export the types of data that you require in the warehouse. Data exporting is disabled by default, but if you enable data exporting, it exports all data types. Turn off any data types that you do not need.

When the export task exports data, the task attempts to complete the export as quickly as possible, using multiple threads to achieve as much throughput as possible. Depending on the I/O speed of the Identity Manager repository and the warehouse, the export task can fully utilize the processors on the Identity Manager server, which causes any interactive performance to suffer. Ideally, the export should occur on a machine dedicated to that task or at least occur during periods when there is no interactive activity on the machine.

The export task supports the following tuning parameters:

- Queue read block size
- Queue write block size
- Queue drain thread count

The drain thread count is the most important throughput. If there is a large number of records in the queue table, increasing the number of threads (up to 24) tends to increase throughput. However, if the queue is dominated by one type of record, fewer drain threads might actually be faster. The export task attempts to divide the queue table contents into as many sets as there are threads allocated, and give each thread a set to drain. Note that these threads are in addition to the drain threads that are draining the other repository tables.

## Tuning the General XML

You can usually optimize the general XML by using static `XMLObject` declarations wherever possible. For example, use

- `<List>` instead of `<list>`
- `<String>` instead of `<s>`
- `<Map><MapEntry . . .></Map>` instead of `<map>`

Also, depending on the context, you might have to wrap objects instead of using the `<o></o>` element.

## Tuning Identity Manager Service Provider

You can use Identity Manager dashboard graphs to quickly assess the current system, spot abnormalities, and understand historical trends (such as concurrent users or resource operations over a time period) for Identity Manager Service Provider.

---

**NOTE** Service Provider does not have an Administrative interface. You use Identity Manager's Administrative interface to perform almost all administrative tasks (such as viewing dashboard graphs).

---

For more information about tuning Sun Java™ System Identity Manager Service Provider (Service Provider) see the latest version of *Identity Manager Service Provider Deployment*, which is available from the following location:

<http://docs.sun.com/app/docs/prod/ident.mgr#hic>

## Tuning the Identity Manager Web Interface

When you are working with the Identity Manager Web Interface, you can optimize performance by using the OpenSPML toolkit that is co-packaged with Identity Manager.

---

**NOTE** Using the `openspml.jar` file from the <http://openspml.org/> web site might cause memory leaks.

---

## Tuning Initial Loads

To improve performance during a large, initial user load:

- Disable all Audit Events. Audit Logging can add several records per operation, making future audit reports perform more slowly.

To disable Audit Events from the Identity Manager Administrator interface,

- a. Select Configure > Audit.
  - b. On the Audit Configuration page, deselect the Enable auditing box, and click Save.
- Disable the list cache. The list cache keeps lists of users in frequently accessed organizations in memory. To maintain these lists, the list cache searches for and checks all newly created users.

To disable the list cache, you can shutdown the web server or change the `ChangeNotifier.updateSearchIntervalCount` property (on the `debug/Show_WSProp.jsp` debug page) to **0**.

- Clear the current list cache on the `debug/Clear_List_Cache.jsp` page.
- Ensure the workflow being used to process the users does not contain approvals.
- Use alternative load methods, such as
  - Splitting the load and running the data in zones
  - Using bulk loads, which are much faster
  - Loading from a file
- Disable Data Exporter for the `WorkflowActivity` type.

## Tuning Memory Requirements

You must determine your memory needs and set values in your application server's JVM by adding maximum and minimum heap size to the Java command line. For example

```
java -Xmx512M -Xms512M
```

To improve performance

- Set the maximum and minimum heap size values to the same size.
- Depending on your specific implementation, you might want to increase these values if you run reconciliation.
- For performance tuning purposes, you may also set the following in the `waveset.property` file:

```
max.post.memory.size value
```

---

**NOTE** The `max.post.memory.size` specifies the maximum number of bytes that a posted file (for example, by using an `HTML FileSelect` control) might contain without being spooled to the disk. For cases where you do not have permission to write to temp files, increase the `max.post.memory.size` to avoid having to spool to the disk. The default value is 8 Kbytes.

---

---

**NOTE** For additional information about system requirements, see the *Identity Manager Release Notes*.

---

## Tuning Operating System Kernels

For information about tuning Solaris and Linux operating system kernels, see the “Tuning the Operating System” chapter in the *Sun Java System Application Server Enterprise Edition 8.2 Performance Tuning Guide*.

For information about tuning Oracle operating system kernels, see the product documentation provided with your Oracle system.

## Tuning Provisioner

Network latency tends to be a popular cause for performance issues when dealing with view provisioning. Tracing individual resource adapters can help you determine what is causing performance problems.

To improve provisioner performance

- Set `provisioner.maxThreads` in the `Waveset.properties` file to control the number of simultaneous account provisioning threads where a thread is started for each resource operation.

Generally, you can achieve optimal performance by setting this value to 10. Specifying a value greater than 20 significantly degrades the provisioner's performance.

- Configure quota settings in the `Waveset.properties` file to control the number of concurrent operations (such as reprovisioning) a user can execute for a specific task. Increasing the number of concurrent actions can help more operations complete faster, but trying to process too many actions at once might cause bottlenecks.

You can create configuration sets on a per-pool basis. For example, if you create configuration A, configuration B, and configuration C, when you create a `TaskDefinition` (workflow), you can assign a specific pool configuration to the workflow from the configurations that you defined.

The following example shows the quota settings that limit user *bob* to running one reprovisioning task at a time:

#### Code Example 1-1 Limiting Concurrent Operations

```
Quota.poolNames=ReProvision,Provision
Quota.pool.ReProvision.defaultLimit=1
Quota.pool.ReProvision.unlimitedItems=Configurator
Quota.pool.ReProvision.items=bob,jan,ted
Quota.pool.ReProvision.item.bob.limit=1
```

To enforce the task quota, reference `poolName` in a `TaskDefinition`. The format is as follows:

```
<TaskDefinition ... quotaName='{poolName}'...>
```

Most users start only one task at a time. For proxy administrators who perform reconciliation or Active Sync tasks, set the task quota higher.

---

**NOTE** Avoid using the `Configurator` user for reconciliation and Active Sync tasks. The `Configurator` has access to unlimited tasks and can monopolize available resources, which adversely affects concurrent processes.

---

## Tuning Reconciliation

The Reconciler is the Identity Manager component that performs reconciliation. This section suggests methods for improving Reconciler performance, including

- [General Tuning Suggestions](#)
- [Tuning the Reconciler Server Settings](#)
- [Tuning Reconciliation for Multiple Resources](#)

### General Tuning Suggestions

In general, you can improve Reconciler performance if you

- Avoid using the `Configurator` user for reconciliation tasks. The `Configurator` has access to unlimited tasks and can monopolize available resources, which adversely affects concurrent processes.

Instead, use a streamlined, minimal user for reconciliation and Active Sync tasks. Because the subject executing the task is serialized as part of the task, a minimal user takes less space, or overhead, for each task and update in the repository.

- Use information on the Reconciler status page (`debug/Show_Reconciler.jsp`) to decide which settings to adjust based on queue sizes, available system resources, and performance benchmarks. Be aware that these settings are dependent on the environment.
- Use the System Memory Summary page (`debug/Show_Memory.jsp`) to see how much total and free memory is available. Reconciliation is a memory-intensive function, and you can use this information to determine whether there is sufficient memory allocated to the JVM. You can also use this page to launch garbage collection or to clear unused memory in the JVM for investigating heap usage.
- When you assign user forms to proxy administrators who are performing reconciliations, keep the user forms as simple as possible and only use essential fields. Depending on the schema map, including a field that calculates the `waveset.organization` attribute is generally sufficient.

---

#### NOTE

Administrators who need to view or edit the Identity Manager schema for Users or Roles must be in the `IDM Schema Configuration AdminGroup` and must have the `IDM Schema Configuration capability`.

---

- Use per-account workflows judiciously. The reconciliation process does not start provisioning tasks for performance reasons by default.

If you must use a per-account workflow task, edit the reconciliation policy to limit the Reconciler's automatic responses to events of interest only. (See the Situation area of the Edit Reconciliation Policy page.)

## Tuning the Reconciler Server Settings

Although the default settings are usually adequate, you can sometimes improve Reconciler performance if you adjust the following settings on the Edit Server Settings page:

- **Parallel Resource Limit:** Specifies the maximum number of resource threads that the Reconciler can process in parallel.  
  
Resource threads allocate work items to worker threads, so if you add additional resource threads, you might also have to increase the maximum number of worker threads.
- **Minimum Worker Threads:** Specifies the number of processing threads that the Reconciler always keeps open.
- **Maximum Worker Threads:** Specifies the maximum number of processing threads that the Reconciler can use. The Reconciler starts only as many threads as the workload requires, which places a limit on that number. Worker threads automatically close if they are idle for a short duration.

During idle times, the threads stop if they have no work to do, but only down to the minimum number of threads specified. As the load increases, the Reconciler adds more threads until the maximum number of threads is reached. The Reconciler never has less than the minimum number of threads or more than the maximum.

Generally, more threads allow more concurrency; but at some point, too many threads can put too much load on the machine or just do not provide additional benefit.

---

**NOTE**      Recommending generic, optimal settings is not possible because deployments are so different. Reconciler settings must be adjusted differently for each deployment environment.

---

To change the Reconciler server settings:

1. Log into the Administrator interface.
2. Select the Configure > Servers > Reconciler tabs.
3. When the Edit Server Settings page displays, adjust the settings as necessary.

---

**NOTE** See “Configuring Identity Manager Server Settings” in *Identity Manager Administration* for more information.

---

## Tuning Reconciliation for Multiple Resources

If you are configuring reconciliation for multiple resources in Identity Manager, you have several options:

- All of the resources on the same server, all at the same time.

This option is the most efficient from Identity Manager’s perspective, but if you have many resources (for example more than 20), you are likely to experience Java resource issues.

- All of the resources on the same server, each at a different time.

This option is easier on Java resource loading, but puts a significant burden on your schedule configuration.

- Each resource on a different server, all at the same time.

This option minimizes elapsed time, but increases the number of servers.

An ideal solution does not exist for this configuration because deployments are so different. You might have to mix and match these options to find an acceptable solution for your deployment.

Preparing a usage survey, based on the business reasons behind this functionality might help you decide how to proceed. Ask yourself these questions:

- Why are you reconciling these resources?
- Do you have the same the goal for each of these resources?
- Are each of these resources equally important (or critical)?
- Must all resources be reconciled on the same schedule, or can you spread out the reconciliations?
- How often must each resource be reconciled?

Also, remember that the reconciliation server does not have to be one of the pools that handles web traffic. You can add a server that you never interact with directly because this server exists solely for transaction processing. Having a server dedicated to transaction processing might make the first option more attractive for very large systems.

## Tuning Resource Queries

---

**NOTE** Network latency tends to be a popular cause of performance issues during view provisioning. Tracing individual resource adapters can help you determine what is causing performance problems.

---

You can improve resource query performance if you use `FormUtil.getResourceObjects` to implement the query.

Use one of the following methods to cache query results:

- `getResourceObjects(Session session, String objectType, String resID, Map options, String cacheList, String cacheTimeout, String cacheIfExists)`
- `getResourceObjects(String subjectString, String objectType, String resId, Map options, String cacheList, String cacheTimeout, String clearCacheIfExists)`

- 
- NOTE**
- Set `cacheTimeout` in milliseconds.
  - Restrict searches to specific `searchContext`, if applicable.
  - Return the minimum number of attributes in `options.searchAttrsToGet`
-

## Tuning the Scheduler

The Scheduler component controls task scheduling in Identity Manager. This section suggests methods for improving Scheduler performance, including:

- [General Tuning Suggestions](#)
- [Tuning the Scheduler Server Settings](#)

### General Tuning Suggestions

The following `TaskDefinition` options determine how the Scheduler handles tasks after they are completed:

- **resultLimit:** Controls how many seconds a task is allowed to live after the task has completed. The default setting varies for different tasks. A setting of zero immediately removes tasks after completion.
- **resultOption:** Controls what action is taken when repeated instances of a task are started. The default setting is `delete`, which removes extra task instances.

These default settings are designed to optimize memory by shortening the lifetime of finished Scheduler tasks. Unless there is a compelling reason to change these settings, leave the default values set.

If you want to immediately delete tasks that completed successfully, but you also want to keep tasks containing errors long enough to debug, you can

- Conditionally delete finished tasks by setting the `resultLimit` value to a time period that is sufficient for debugging issues.
- Set the `resultLimit` to zero (or a small value) if no errors (such as `WF_ACTION_ERROR` is `<null/>`) are reported at runtime after a `WorkflowServices` call.

## Tuning the Scheduler Server Settings

You can sometimes improve Scheduler performance by adjusting the following settings on the Edit Server Settings page:

- **Maximum Concurrent Tasks:** Specify the maximum number of tasks that the Scheduler can run at one time.

When more tasks are ready to run than the Maximum Concurrent Tasks setting allows, the extra tasks must wait until there is room available or they are run on another server.

If too many tasks are being swapped out of memory and sharing CPU time, the overhead slows down performance. Alternatively, setting the maximum too low results in idle time. The Scheduler checks for available tasks every minute, so a waiting task waits at least a minute before being run.

The default Maximum Concurrent Tasks setting (100) is usually adequate. You can decide whether to adjust this setting up or down based on which tasks are being run in the deployment and by profiling the runtime behavior after the deployment is otherwise complete.

In some cases, you may want to suspend or disable the Scheduler. For example, if you want a server dedicated to handling the End User interface, disabling the Scheduler will prevent tasks from running on that server. The server would only serve the End User interface pages and store launched tasks for other servers to execute.

- **Task Restrictions:** Specify the set of tasks that can execute on the server.

The Task Restrictions setting can provide a finer granularity of control over what tasks are allowed to run on a server. You can restrict tasks individually or through the server settings.

---

**NOTE**      Recommending generic, optimal settings is not possible because deployments are so different. Scheduler settings must be adjusted differently for each deployment environment.

---

To change the Scheduler server settings:

1. Log into the Administrator interface.
2. Select the Configure > Servers > Scheduler tabs.
3. When the Edit Server Settings page displays, adjust the settings as necessary.

---

**NOTE** See “Configuring Identity Manager Server Settings” in *Identity Manager Administration* for more information.

---

## Tuning Sessions

Identity Manager maintains a least recently used (LRU) cache of authenticated sessions for use by authenticated users. By using existing authenticated sessions, you can speed up repository access for objects and actions that require a session.

To optimize the authentication pool size, change the `session.userPoolSize` value in the `Waveset.properties` file to the maximum number of expected, concurrent user sessions on the server.

## Tuning the Sun Identity Manager Gateway

The Sun Identity Manager Gateway generates a thread for each connection, and uses a different pool for each unique combination of resource type, Gateway host, and Gateway port. The Gateway checks for idle connections every five minutes and when a connection has been idle for 60 seconds, the Gateway closes and removes that connection from the pool.

When the Gateway receives a request,

- If there are no idle connections in the corresponding pool, the Gateway creates a new connection.
- If an idle connection exists in the pool, the Gateway retrieves and reuses that connection.

You must configure the maximum number of connections on the resource, and you must configure these connections the same way for all resources of the same type, that are using the same Gateway. For that resource type, the first connection made to the Gateway on a given host and port uses that resource's maximum connections value.

---

**NOTE** When you change the maximum number of connections on a resource, you must start and stop the server for the change to take effect.

---

The following example shows how connections, requests, and Gateway threads are related:

If you set the maximum number of connections to ten on an Active Directory resource, and you are using two Identity Manager servers, then you can have up to 20 simultaneous connections (ten from each Identity Manager server) to the Gateway for that Active Directory resource. The Gateway can have ten simultaneous requests outstanding from each server, and the Gateway processes each request on a different thread. When the number of simultaneous requests exceeds the maximum number of Gateway connections, additional requests are queued until the Gateway completes a request and returns the connection to the pool.

---

**NOTE** Although the Gateway code is multithreaded, this characteristic does not address the APIs or services being used by the Gateway. For Active Directory, the Gateway uses the ADSI interface provided by Microsoft. No investigation has been done to determine whether this interface handles Gateway requests in parallel.

---

Other methods for improving Gateway performance, include:

- Locating the Gateway near (from a network connectivity perspective) the domain controllers of the managed domain.
- Increasing the block size on a Gateway resource can increase throughput during reconciliation or load operations.

Increased throughput results have been noted for basic reconciliations with no custom workflows and in which there are no attribute reconciliations are being performed. Initially, the Gateway does consume more system memory, but this memory is eventually released.

Be aware that there is a diminishing return. At some point, larger block sizes do not result in proportionately increased performance. For example, the following data shows the speed observed for a Load from Resource of 10,000 users from an Active Directory resource. Also, the peak memory usage for the Gateway process during the load is included.

Block Setting	Users Created Per Hour	Peak Gateway Memory Usage
100	500	20M
200	250	25M
500	9690	60M
1000	10044	92M

- For Exchange Server 2007, the `PowerShellExecutor` performs actions for Exchange Server 2007. You can modify the following registry settings to change the behavior of the `PowerShellExecutor` inside the gateway:

---

**NOTE** Both settings can have a large impact on the behavior and memory usage of the gateway. Changes to these parameters should only be considered after careful testing.

---

- `powerShellTimeout`

**Content:** Timeout for Powershell actions (registry type `REG_DWORD`)

**Default:** 60000 ms (1 minute)

When the `powerShellTimeout` setting times out, any `RunSpace` actions are interrupted and cancelled to prevent runaway actions in the PowerShell environment that cause the gateway to become unresponsive.

Decreasing the `powerShellTimeout` value to a small value can prematurely cancel actions, and can prevent the `RunSpace` initialization from finishing correctly. Observed start-up times for the first `RunSpace` in the pool range from 2-5 seconds.

The `powerShellTimeout` value is read-only on start-up, and you cannot change without restarting the gateway.

- o `runSpacePoolSize`

**Content:** Number of RunSpaces in the pool (registry type REG\_DWORD)

**Default:** 5

**Minimum:** 5

**Maximum:** 25

The number of RunSpaces in the pool allow for parallel execution of PowerShell actions by the gateway. One provisioning action or update of a user in Exchange 2007 can result in multiple PowerShell actions being executed.

A started RunSpace can consume a large amount of memory. For the first RunSpace, the typical size is approximately 40 MB, subsequent RunSpaces normally use between 10-20 MB.

---

**NOTE** The preceding figures can differ in specific environments and are only given as *guidelines*, so be careful when changing this value.

---

The `runSpacePoolSize` value is read-only on start-up, and you cannot change the pool size value without restarting the gateway.

## Tuning the Task Bar

The Administrative interface task bar displays links to previously performed provisioning tasks, which causes the interface to render more slowly when there are a large number of tasks.

To improve interface performance, remove the `taskResults.jsp` link from all JSPs by deleting the `<List>...</List>` element from the `UserUIConfig` object.

The following example shows `<List>...</List>` entries within `<TaskBarPages>`:

**Code Example 1-2**    Modifying the UserUIConfig Object

```
<TaskBarPages>
  <List>
    <String>account/list.jsp</String>
    <String>account/find.jsp</String>
    <String>account/dofindexisting.jsp</String>
    <String>account/resourceReprovision.jsp</String>
    <String>task/newresults.jsp</String>
    <String>home/index.jsp</String>
  </List>
</TaskBarPages>
```

## Debugging Performance Issues

This section describes the different Identity Manager and third-party debugging tools you can use to debug performance issues. The information is organized into the following sections:

- [Working with Identity Manager Debug Pages](#)
- [Working With Other Debugging Tools](#)

# Working with Identity Manager Debug Pages

---

**NOTE** Tracing affects system performance. To help ensure optimal performance, specify the minimum tracing level or turn tracing off after debugging the system.

---

This section provides instructions for accessing the Identity Manager debug pages describes how to use these pages to identify and debug Identity Manager performance issues. See the following sections for information:

- [Accessing the Debug Pages](#)
- [Control Timings \(callTimer.jsp\)](#)
- [Edit Trace Configuration \(Show\\_Trace.jsp\)](#)
- [Host Connection Pool \(Show\\_ConnectionPools.jsp\)](#)
- [List Cache Cleared \(Clear\\_XMLParser\\_Cache.jsp\)](#)
- [Method Timings \(Show\\_Timings.jsp\)](#)
- [Object Size Summary \(Show\\_Sizes.jsp\)](#)
- [Provisioning Threads for Administrator Configurator \(Show\\_Provisioning.jsp\)](#)
- [System Cache Summary \(Show\\_CacheSummary.jsp\)](#)
- [System Memory Summary \(Show\\_Memory.jsp\)](#)
- [System Properties \(SysInfo.jsp\)](#)
- [System Threads \(Show\\_Threads.jsp\)](#)
- [User Session Pool Cleared \(Clear\\_User\\_Cache.jsp\)](#)
- [Waveset Properties \(Show\\_WSProp.jsp\)](#)

## Accessing the Debug Pages

---

**NOTE** You must have the *Debug*, *Security Administrator*, or the *Waveset Administrator* capabilities to access and execute operations from the Identity Manager Debug pages. Administrators and the Configurator are assigned this capability by default.

If you do not have the Debug capability, an error message results.

---

To access the Identity Manager debug pages:

1. Open a browser and log in to the Identity Manager Administrative interface.
2. Type the following URL into the browser:

```
http://host:port/idm/debug
```

Where:

- *host* is the application server on which you are running Identity Manager.
  - *port* is the number of the TCP port on which the server is listening.
3. When the System Settings page displays, type the `.jsp` file name for the debug page you want to open. For example

```
http://host:port/idm/debug/pageName.jsp
```

---

**NOTE** Some debugging utilities are not linked from the System Settings page, but you can use them to enhance your ability to gather data for product performance and usability. For a complete list of debug pages, open a command window and list the contents of the `idm/debug` directory.

---

### Control Timings (`callTimer.jsp`)

Use the Control Timings page to collect and view call timer statistics for different methods. You can use this information to track bottlenecks to specific methods and invoked APIs. You can also use options on the Call Timings page to import or export call timer metrics.

---

**NOTE** Call timing statistics are only collected while trace is enabled.

---

To view call timer statistics

1. Open the Control Timing page, and click Start Timing & Tracing to enable trace and timing.
2. To stop the timing, click Stop Timing & Tracing or click Stop Timing.

The page redisplay and populates the Show Timings table with a list of methods for which statistics are available and the methods' aggregate call timer statistics (*not broken down by caller*). This table contains the following information:

- Method name (Click a method name to see which methods it calls.)
  - Total time
  - Average time
  - Minimum time
  - Maximum time
  - Total calls
  - Total errors
3. To clear the list, click Clear Timing.

---

**NOTE** You can also use the `callTimer` command to collect call timer data from the Console. This command is useful when you are debugging performance issues during an upgrade or in other situations where Identity Manager is not running on the application server.

---

### Edit Trace Configuration (`Show_Trace.jsp`)

Use the Edit Trace Configuration page to enable and configure tracing for the Java classes provided with your Identity Manager installation.

Specifically, you can use this page to configure the following trace settings:

- Choose methods, classes, or packages to trace and specify the level of trace you want to capture.
- Send trace information to a file or to standard output.
- Specify the maximum number of trace files to be stored and the maximum size for each file.
- Specify how dates and times are formatted in the trace output file.

- Specify the maximum number of methods to be cached.
- Indicate how to write data to the trace file.

Write data to the trace file as the data is generated or queue the data and then write it to the file.

### Host Connection Pool (`Show_ConnectionPools.jsp`)

If you are not using a data source, you can use the Host Connection Pool page to view connection pool statistics. These statistics include the pool version, how many connections were created, how many are active, how many connections are in the pool, how many requests were serviced from the pool, and how many connections were destroyed.

You can also use the Host Connection Pool page to view a summary of the connection pools used to manage connections to the Gateway. You can use this information to investigate low-memory conditions.

### List Cache Cleared (`Clear_XMLParser_Cache.jsp`)

Use the List Cache Cleared page to clear recently used XML parsers from the cache and to investigate low memory conditions.

### Method Timings (`Show_Timings.jsp`)

Use the Method Timings page to quickly detect and assess hotspots at a method level. The following information is gathered from Identity Manager methods and displayed on the Method Timings page:

- Method names
- How many times the methods were called
- How many times the methods exited with an error status
- Average time consumed by the methods
- Minimum and maximum times consumed by invocations of each method

The Method Timings page also contains a table with the following links (click these links to view additional information):

- **Details:** Shows call stack information.
- **History:** Shows a graph of call duration versus time for the most recent calls.

- **History data:** Shows a list of the most recent calls, showing what time the call was made and the duration of the call.

Identity Manager does not keep stack history by default. To keep stack history and to control its depth, edit `Waveset.properties` and look at the `MethodTimer` keys.

---

**NOTE** The Clear ALL option on the Method Timings page clears all results. This option is enabled by default.

---

## Object Size Summary (`Show_Sizes.jsp`)

Use the Object Size Summary page to detect problematically large objects that can affect your system.

The Object Size Summary page shows information about the size of objects (in characters) stored in the repository. These objects are listed by type, along with the total number of objects of each type, and the objects' total combined size, average size, maximum size, and minimum size.

Click entries in the Type column to view additional size information about each specific object type. For example, click Configuration to view the ID, name, and size of the largest configuration objects in the repository.

You can also access this size information from the console command line.

1. Open the console.
2. At the command prompt, type:

```
showSizes [<type> [<limit>]]
```

---

**NOTE** For upgrades, existing objects will report a size of 0 until they have been updated or otherwise refreshed.

---

## Provisioning Threads for Administrator Configurator (`Show_Provisioning.jsp`)

Use the Provisioning Threads for Administrator Configurator to view a summary of the provisioning threads in use by the system (a subset of the information available in `Show_Threads.jsp`).

---

**NOTE** Looking at just a single thread dump can be misleading.

---

## System Cache Summary (`Show_CacheSummary.jsp`)

Use the System Cache Summary page to view information about the following items to help you investigate low-memory conditions:

- Administrator-associated object caches
- System object cache
- User log-in sessions
- XML parser cache

## System Memory Summary (`Show_Memory.jsp`)

Use the System Memory Summary page to view how much total and free memory you have available in Mbytes. When you are using memory-intensive functionality such as Reconciliation, this information can help you determine whether there is sufficient memory allocated to the JVM.

You can also use this page to launch garbage collection or to clear unused memory in the JVM for investigating heap usage.

## System Properties (`SysInfo.jsp`)

This page also provides information about your environment, including software versions, paths and environmental variables.

## System Threads (`Show_Threads.jsp`)

Use the System Threads page to see which processes are running so you can verify that automated processes (such as reconciliation or Active Sync) are running.

This page includes information about the process type, process name, its priority, if the process is a daemon and if the process is still alive (running).

---

**NOTE**      Looking at just a single thread dump can be misleading.

---

## User Session Pool Cleared (`Clear_User_Cache.jsp`)

Use the Session Pool Clearer page to clear all of the cached sessions for users who have recently logged in and to investigate low memory conditions.

## Waveset Properties (Show\_WSProp.jsp)

Use the Waveset Properties page to view and *temporarily* edit properties in the `Waveset.properties` file. You can test different property settings for a particular server on which the `Waveset.properties` file resides without having to restart the server to pick up the changes. The edited property settings only remain in effect until the next time you restart the server.

## XML Resource Adapter Caches Flushed and Cleared (Clear\_XMLResourceAdapter\_Cache.jsp)

Use the XML Resource Adapter Caches Flushed and Cleared page to clear test XML resource adapters from the cache and to investigate low memory conditions.

# Working With Other Debugging Tools

You can use the following Sun Microsystems' and third-party tools to identify potential performance bottlenecks — especially if your deployment uses custom Java classes:

- [Identity Manager Profiler](#)
- [DTrace](#)
- [JConsole](#)
- [JRat](#)

## Identity Manager Profiler

Identity Manager provides a Profiler utility to help you troubleshoot performance problems in your deployment.

Customized forms, Java, rules, workflows, and XPRESS can cause performance and scale problems. The Profiler profiles how much time is spent in these different areas, enabling you to determine whether these forms, Java, rules, workflows, or XPRESS objects are contributing to performance and scale problems and, if so, which parts of these objects are causing the problems.

---

**NOTE** For more information about JProfiler, see “Working with the Identity Manager Profiler” in the *Identity Manager 8.0 Release Notes*.

---

## DTrace

DTrace is a dynamic tracing framework for the Solaris 10 operating environment that enables you to monitor JVM activity.

DTrace provides more than 30,000 probes and uses integrated user- and kernel-level tracing to give you a view into your production system. You can also trace arbitrary data and expressions by using the D language (similar to C or awk).

The DTrace facility

- Includes special support for monitoring the JVM, and enables you to watch your whole system and span outside the JVM
- Is easiest to use with Java 6 because probes are built into the JVM
- Works with Java 1.4 and Java 5, but you must download JVM PI or JVM TI agents from the following location:

<https://solaris10-dtrace-vm-agents.dev.java.net/>

The following example shows how to write a DTrace script:

### Code Example 1-3 Example DTrace Script

```
#!/usr/sbin/dtrace -Zs
#pragma D option quiet
hotspot$1:::
{
    printf("%s\n", probename);
}
```

In this example, you would replace *\$1* with the first argument to the script, which is the PID of the Java process you want to monitor. For example

```
# ./all-jvm-probes.d 1234
```

Use the following commands to enable different DTrace probes:

**Table 1-3** DTrace Commands

Command	Description
<code>-XX:+DTraceMonitorProbes</code>	Enables JVM support in Java 6 (patches for 1.4 and Java 5)
<code>-XX:+ExtendedDTraceProbes</code>	Provides the following information: <ul style="list-style-type: none"> <li>• JVM startup (begin and end) and shutdown</li> <li>• Thread starting and stopping</li> <li>• Class loading and unloading</li> <li>• Garbage collection (several options available)</li> <li>• JIT compilation begin and end</li> <li>• Compiled method loading and unloading</li> <li>• Monitor contention, wait and notify</li> <li>• Method entry, method return, and object allocation</li> </ul>
<code>/usr/sbin/dtrace -n 'hotspot*:::'</code>	Enables all JVM probes for all Java processes on the system
<code>/usr/sbin/dtrace -n 'hotspot1234:::'</code>	Enables all JVM probes for only the Java process with PID <i>1234</i>
<code>/usr/sbin/dtrace -n 'hotspot1234:::gc-begin'</code>	Enables only the probe that starts when garbage collection for process <i>1234</i> begins

**NOTE** Because DTrace causes additional work in the system, enabling this facility affects system performance. The effect is often negligible, but can become substantial if you enable many probes with costly enablings.

Instructions for minimizing the performance effect of DTrace are provided in the “Performance Considerations” chapter of the *Solaris Dynamic Tracing Guide*. You can view this publication from the following location:

<http://docs.sun.com/app/docs/doc/817-6223>

For more information about DTrace, see `/usr/demo/dtrace` and `man dtrace`.

## JConsole

The Java Monitoring and Management Console (*JConsole*) is a Java Management Extension (*JMX*) technology-compliant graphical management tool that is co-packaged with JDK 5 (and later). JConsole connects to a running JVM and gathers information from the JVM MBeans in the connected JMX agent.

Specifically, you can use JConsole to

- Detect low memory and deadlocks
  - JConsole accesses the memory system, memory pools, and MBeans garbage collector to provide information about memory use, including memory consumption, memory pools, and garbage collection statistics.
- Enable or disable garbage collection
- Enable or disable verbose tracing
- Monitor local and remote applications
- Monitor and manage MBeans including current heap memory use, non-heap memory use, and how many objects are pending for finalization
- View information about performance, resource consumption, and server statistics
- View information about operating system resources (Sun's platform extension), such as
  - CPU process time
  - How much total and free physical memory is available
  - The amount of committed virtual memory (how much virtual memory is guaranteed to be available to the running process)
  - How much total and free swap space is available
  - The number of open file descriptions (*UNIX only*)
- View summary information about the JVM and monitored values, threads running on the application, and loaded classes

---

**NOTE** For more information about using JConsole to monitor applications on the Java platform, see the following Sun Developer Network (SDN) article:

<http://java.sun.com/developer/technicalArticles/J2SE/jconsole.html>

---

Identity Manager supplies some JMX Mbeans that provide information about:

- Identity Manager Server Cluster
- Data Exporter
- Scheduler

## JRat

You can use the Java Runtime Analysis Toolkit (JRat), an open-source performance profiler for the Java platform, to identify potential performance bottlenecks, especially if your deployment uses custom Java classes. JRat monitors your application's execution and persists the application's performance measurements.

For example, if you have a custom workflow for provisioning, you can use JRat to see which classes are being invoked and how much time is required to run your workflow compared to the default Identity Manager provisioning workflow.

---

**NOTE** For more information about JRat, see the following web site:

<http://jrat.sourceforge.net/>

---



# Tracing and Troubleshooting

This chapter explains how you can use tracing to help fix errors, solve performance issues, and understand how things flow within Identity Manager. This chapter also describes methods for troubleshooting problems that might occur with different Identity Manager components.

The information in this chapter is organized into the following sections:

- [Before You Begin](#)
- [Tracing Identity Manager Objects and Activities](#)
- [Tracing the Identity Manager Gateway Objects and Activities](#)
- [Troubleshooting and Fixing Common Problems](#)

## Before You Begin

Review the information in the following sections before tracing or troubleshooting Identity Manager:

- [Intended Audience](#)
- [Important Notes](#)
- [Before Calling Support](#)
- [Related Documentation and Web Sites](#)

## Intended Audience

This chapter is intended for application server and database administrators, front line support engineers, and partners who are responsible for maintaining Identity Manager in a deployment environment.

Before troubleshooting problems with Identity Manager, you must

- Be familiar with Java 5.0 (required for Identity Manager 8.0)
- Understand the components you are trying to troubleshoot

## Important Notes

You should be aware of the following information before trying to troubleshoot Identity Manager:

- Tracing affects system performance. To ensure optimal performance, use the minimum tracing level needed or turn off tracing after debugging the systems.
- Do not enable tracing for the `com.waveset` class. The `com.waveset` class is verbose and has many classes, so tracing this class might cause your server to hang.
- Do not configure Identity Manager to trace at the advanced level unless instructed to do so by Sun Support.
- If you set `exception.trace` in the `Waveset.properties` file to debug a specific problem, do not use it for an extended period of time and be sure you disable it after debugging. Setting `exception.trace` in the `Waveset.properties` file can significantly affect system performance.

## Before Calling Support

Before you call Sun Support for assistance with a problem, consider the following suggestions for narrowing the problem area:

- Try using web browser search tools to research the problem
- Use resource-specific support options when available. Your problem might be related to a resource and fixed using a resource patch.
- Remove Identity Manager from the equation by using an appropriate client tool when possible. For example, using a Java LDAP browser.

## Related Documentation and Web Sites

In addition to the information provided in this chapter, consult the publications and web sites listed in this section for information related to tracing and troubleshooting Identity Manager.

### Recommended Reading

See the following publications for information related to tracing and troubleshooting Identity Manager:

- See “Testing Your Customized Form” section in *Identity Manager Workflows, Forms, and Views* for the recommended method of tracing XPRESS functions.
- See the article titled, *Using JConsole to Monitor Applications*, for information about using JConsole to monitor applications that run on the Java platform. This article is available from the following location:

<http://java.sun.com/developer/technicalArticles/J2SE/jconsole.html>

### Useful Web Sites

The following table describes web sites related to troubleshooting Identity Manager.

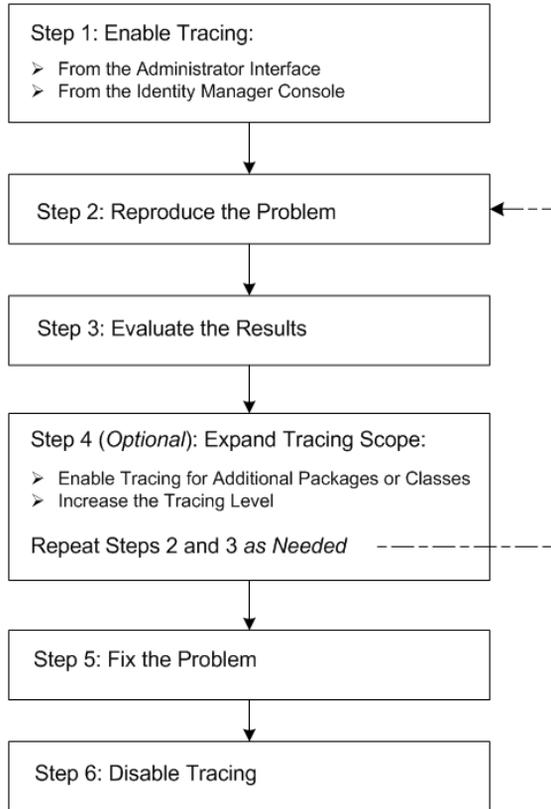
**Table 2-1** Useful Web Sites

Web Site URL	Description
<a href="http://sunsolve.sun.com/">http://sunsolve.sun.com/</a>	Sun Microsystems' web site containing diagnostic tools, forums, features and articles, security information, and patch contents. <b>Note:</b> The information on this site is partitioned into three areas, <ul style="list-style-type: none"> <li>• <b>Internal:</b> Sun employees only</li> <li>• <b>Contract:</b> Available only to customers with contract access</li> <li>• <b>Public:</b> Available to everyone</li> </ul>
<a href="http://www.sun.com/service/online/us">http://www.sun.com/service/online/us</a>	Sun Microsystems' Technical Support web site
<a href="https://sharespace.sun.com/gm/folder-1.11.60181?">https://sharespace.sun.com/gm/folder-1.11.60181?</a>	Identity Manager folder on Sun Microsystems' Share Space, which contains Identity Manager's FAQ, links to forums, features and articles, and more. <b>Note:</b> You must sign up for a Share Space ID to access information provided on this site.
<a href="https://identitymanageride.dev.java.net">https://identitymanageride.dev.java.net</a>	Open source Identity Manager Integrated Development Environment (Identity Manager IDE) project. Includes instructions for installing and configuring Identity Manager IDE.

# Process Overview

Generally, you can identify and resolve problems in your deployment if you follow these steps:

**Figure 2-1** Tracing and Troubleshooting a Problem



See [“Tracing Identity Manager Objects and Activities” on page 69](#) for information about how to enable tracing for different Identity Manager objects and activities.

See [“Troubleshooting and Fixing Common Problems” on page 101](#) for information about how to troubleshoot problems that commonly occur for those objects and activities.

# Tracing Identity Manager Objects and Activities

Trace output can be very helpful when you are trying to identify and resolve problems, or when you are developing custom resource adapters.

This section describes how to enable tracing for a variety of Identity Manager objects and activities. The information is organized as follows:

- [How To Configure Tracing](#)
- [How to View Trace Files](#)
- [Tracing the Identity Manager Server](#)
- [Tracing Adapters](#)
- [Tracing Custom Code](#)
- [Tracing Exceptions](#)
- [Tracing Forms](#)
- [Tracing Global XPRESS](#)
- [Tracing PasswordSync](#)
- [Tracing Identity Manager Service Provider Delegated Administration](#)
- [Tracing Reconciliation](#)
- [Tracing the setRepo Command](#)
- [Tracing SPML](#)
- [Tracing the Task Scheduler](#)
- [Tracing Workflows](#)
- [Locating Version Information](#)

---

**NOTE** Tracing affects system performance. To ensure optimal performance, specify the minimum tracing level or turn tracing off after debugging the system.

---

## How To Configure Tracing

You can enable and configure tracing from several locations within Identity Manager. Instructions are provided in the following sections:

- [From the System Settings Page](#)
- [From Individual Debug Pages](#)
- [From the Identity Manager Console](#)

### From the System Settings Page

The System Settings page is the primary Identity Manager Debug page. From this page, you can

- View and edit objects in the repository
- Clear caches
- Set specialized traces
- Reload the `Waveset.properties` file

### *Editing the Global Trace Configuration Object*

Use the following steps to view and edit the default trace configuration object from the Identity Manager System Settings page:

1. Open a browser and log in to the Identity Manager Administrative interface.
2. Type the following URL into the browser:

```
http://host:port/idm/debug
```

Where:

- *host* is the local server on which you are running Identity Manager.
- *port* is the number of the TCP port on which the server is listening.

3. When the System Settings page displays, click Show Trace (located near the bottom of the page) to manage a single trace configuration, where you can create, modify, or delete up to ten trace settings.

---

**NOTE** The remaining instructions assume you have only one trace configuration.

To manage multiple trace configurations, click the Show Trace List button instead. When the Trace Configuration page displays, select the configuration name to edit the current settings.

Identity Manager supplies one configuration, called *Global*, by default. However, if you have multiple servers in an Identity Manager instance, defining different configurations for particular servers might be useful. If the name of a trace configuration matches the name of the current host, the host configuration overrides the Global configuration.

---

4. On the Edit Trace Configuration page, select the Trace Enabled box to enable tracing.

---

**NOTE** Deselecting this box stops tracing, but keeps your configuration. You can turn tracing on and off without having to remember and retype the classes you were tracing.

---

5. Specify which classes, packages, or methods to trace by typing their names into the table. For example:
  - To trace all classes in the `waveset.repository` package, enter `com.waveset.repository`.
  - To trace the `AbstractDataStore` class in the `waveset.repository` package, enter `com.waveset.repository.AbstractDataStore`.
  - To trace the `list` method in the `AbstractDataStore` class in the `waveset` package, enter `com.waveset.repository.AbstractDataStore#list`.

---

**NOTE** Do not enable trace for the `com.waveset` class. The `com.waveset` class is verbose and has many classes, so tracing this class might cause your server to hang.

---

6. In the same table, choose a Method/Class tracing level from the Level menu. Each level captures different types of information, as described in the following table:

Trace Level	Description
0	Minimum debugging output, traces exceptions and error information only
1	Trace Level 0 events, plus entries and exits for public methods
2	Trace Level 1 events, plus entries and exits for non-public methods
3	Trace Level 2 events, plus decision points and significant variables
4	Maximum debugging output

---

**NOTE** Method/Class tracing produces a predictable, but possibly very large volume of trace output. Try to be as specific as possible when specifying methods and classes to trace.

---

7. (Optional) To enable subcall tracing, choose a level from the Subcall Trace menu, which uses the same trace numbering scale described in the previous step.

---

**NOTE**

- The default Subcall Tracing level is *none*, which disables subcall tracing on a per-method or per-class basis.
- Subcall Tracing levels are independent of the Method/Class tracing levels you specify for [Step 6](#).

---

When you enable Subcall Tracing for a particular method that supports subcall tracing, you are automatically setting the tracing level for methods that are *called by* this method. Subcall Tracing enables you to produce a short, but detailed burst of trace output that is triggered by the entry into a specified method and terminated by the method's exit.

For example, assume you created a trace configuration setting for the `com.waveset.adapter.NewRes#init` method, set Method/Class tracing to level one and set Subcall tracing to level three. Also, assume that the `init` method calls two other methods:

- `NewRes#subcallA`
- `NewRes#subcallB`

When the `init` method runs, the `com.waveset.adapter.NewRes#init` method produces trace output at level one until reaching `subcallA`. When `subcallA` begins executing, the trace level changes to three and continues at that level until `subcallA` exits. The `com.waveset.adapter.NewRes#init` method returns to the `init` method and restores the trace level to one. Later, when `init` calls `subcallB`, there is another burst of level three trace detail that lasts until `subcallB` exits. Finally, when `init` exits, level one tracing stops.

8. Send the trace results to a specified file location or to `stdout`.

If you choose output to a file, the Trace File field displays. Use this field to specify an alternate location and file name for trace output files. The default location and file name is

`path_to_idm_install\export\pipeline\config\WSTrace.log`.

9. Specify the maximum number of trace files to store (default is 2).
10. Specify the maximum size for each file (default is 512K).
11. Specify whether to write the trace output file as generated (synchronously) or to queue the data and then write it to the trace file (asynchronously).
12. Save your changes.

### *To Create a New Trace Configuration Object*

Use the following steps to create a new trace configuration object:

1. Decide which package or method you want to trace.  
Generally, you specify a resource adapter name or use information that was revealed in an error message.
2. Log in to the Identity Manager Administrator interface and open the System Settings page as described on [page 70](#).
3. On the Systems Setting page, click Show Trace List.
4. When the Identity Manager Trace Configuration page displays, click New.

5. On the Edit Trace Configuration page, choose one of the following options from the Trace Configuration menu:
  - **Global:** Select to enable tracing for all servers.
  - *Server\_name:* Select a server name to enable tracing for a particular server.
6. Select the Trace Enabled box to enable tracing for this object and configure the remaining parameters on this page as described in [“Editing the Global Trace Configuration Object” on page 70](#).
7. Save your changes.

## From Individual Debug Pages

Use the following steps to enable or disable tracing from an Identity Manager Debug page:

1. Open a browser and log in to the Identity Manager Administrative interface.
2. Type the following URL into the browser:

```
http://host:port/idm/debug/pageName.jsp
```

Where:

- *host* is the local server on which you are running Identity Manager.
- *port* is the number of the TCP port on which the server is listening.
- *pageName.jsp* is the particular Debug page you want to open.

For example, to trace adapter classes and methods for a custom adapter, type `http://host:port/idm/debug/Show_Trace.jsp` to open the Edit Trace Configuration page.

## From the Identity Manager Console

Use the following steps to enable tracing from the Identity Manager console:

1. Set `$WSHOME`.

For example, to set the variable to the default installation directory:

```
set WSHOME=C:\Program Files\tomcat\webapps\idm
```

2. Open the Identity Manager console from the `bin` directory with the command:

```
lh console
```

- From the console, type `trace` to see a detailed summary of available trace options, including `enable` and `disable`. Use the following syntax:

```
trace [ -s server ] $subcommand
```

## How to View Trace Files

By default, Identity Manager sends tracing information to a file named `WSTrace#.log` that is stored in the `path_to_idm_install\export\pipeline\config\` directory. If necessary, you can specify alternate file names and locations when you configure tracing for an object.

Each log file is numbered sequentially, up to the maximum number of files specified on the Edit Trace Configuration page. For example, if you specify a maximum of three files, the files are named `WSTrace1.log`, `WSTrace2.log`, and `WSTrace3.log`.

To view these log files,

- If you are sending trace information to a file, open the trace file from the specified location. For example,
 

```
path_to_idm_install\export\pipeline\config\WSTrace2.log
```
- If you are sending trace information to `stdout`, open your application server's `stdout` file in a text editor to view the trace output logs.

## Tracing the Identity Manager Server

Identity Manager is a Java-based product whose executables consist of Java classes grouped as packages. When the code is implemented, many classes can supply trace output.

Tracing the Identity Manager server can provide helpful information such as where a server is failing, having problems, or not running. You can use the Identity Manager Debug pages to enable package-level and method-level tracing in a running Identity Manager server.

---

**NOTE** Configure Identity Manager to trace at this advanced level *only* if instructed by Sun Support.

---

## Tracing Adapters

You can use trace information to verify that an adapter has started, to verify that all setting changes were saved, and to diagnose problems with the adapter.

When you enable tracing for an adapter, you must identify the methods that you want to trace, as follows:

```
com.waveset.adapter.sample.MyResourceAdapter
```

You must also provide calls to create log entries for any new methods in your custom adapter so the adapter can write its resource settings to a log file. In some cases, you can specify additional logging parameters for the resource instance, such as:

- Maximum Log Archives
- Maximum Active Log Age
- Log File Path
- Maximum Log File Size
- Log Level

- 
- NOTE**
- To further debug the synchronization process, you can configure synchronization logging for the adapter. (Instructions are provided in the Identity Manager online help.)
  - Tracing custom Java code can also be useful when you are writing your own resource adapter. See [“Tracing Custom Code” on page 77](#) for more information.
  - For more information about tracing specific adapter methods, see the *Identity Manager Resources Reference*.
- 

## Tracing Auditor

You can trace the following methods to troubleshoot issues with Identity Auditor:

- `com.sun.idm.auditor.policy` to trace issues with Audit Scans.
- `com.sun.idm.auditor.accessreview` to trace issues with Access Reviews.
- `com.sun.idm.auditor.report` to trace issues with Audit Reports.
- `com.sun.idm.auditor.view` to trace issues with Auditor Views.

To enable tracing,

1. Open a browser and log in to the Administrator interface.
2. Select Configure > Servers.
3. When the Configure Servers page displays, click the server name in the Server column to edit the settings.
4. On the Edit Server Settings page, click the Scheduler tab.
5. Select the Tracing Enabled box to activate Scheduler debug tracing and write the results to `stdout`.
6. Save your changes.

You can add the following hidden flags by modifying

## Tracing Custom Code

You can use the Identity Manager tracing facility in custom-written code to provide seamless integration when troubleshooting your deployment.

Use the following classes to provide this tracing facility:

- **`com.sun.idm.logging.trace.Trace`** objects are used to interface with the tracing subsystem.
- **`com.sun.idm.logging.trace.TraceManager`** is a factory for these objects.

You can also use these `com.sun.idm.logging.trace` facilities to record trace output. See the Javadoc for more information about these facilities and classes.

---

**NOTE** Be aware that the `com.sun.idm.logging.trace` facilities were *not* available prior to the Identity Manager 6.0 SP1 release. For earlier versions of Identity Manager, you must use the `com.waveset.util.Trace` class instead.

---

## Tracing Exceptions

Exception logs are stack traces that you can view from the Identity Manager Debug pages or from the `config/Waveset.properties` file. Trace data does not include all exceptions by default, but exception logging can be an important and informative troubleshooting tool.

Use one of the following methods to enable exception logging:

- Open the Waveset Properties page (`debug/Show_WSProp.jsp`) in the Identity Manager Administrator interface. Locate the `exception.trace` key and change the Value to `true`.
- Open the `config/Waveset.properties` file in a text editor, and change the `exception.trace` key value to `true`. The following figure shows an example exception trace.

**Figure 2-2** Exception Trace Example

```

WavesetException: Validation errors detected in form.
con.waveset.exception.FormValidation: Validation errors detected in form.
  at con.waveset.util.WavesetException.checkBreakpoint(WavesetException.java:513)
  at con.waveset.util.WavesetException.<init>(WavesetException.java:114)
  at con.waveset.exception.FormValidation.<init>(FormValidation.java:39)
  at con.waveset.object.ViewMaster.runExpansions(ViewMaster.java:1004)
  at con.waveset.server.ViewMaster.runExpansions(ViewMaster.java:648)
  at con.waveset.view.UserViewer.runExpansions(UserViewer.java:1438)
  at con.waveset.view.UserViewer.checkInView(UserViewer.java:1133)
  at con.waveset.object.ViewMaster.checkInView(ViewMaster.java:747)
  at con.waveset.session.LocalSession.checkInView(LocalSession.java:611)
  at con.waveset.ui.util.GenericViewSource.checkInView(GenericViewSource.java:522)
  at con.waveset.ui.util.GenericEditForm.process(GenericEditForm.java:613)
  at org.apache.jsp.modify_jsp_jspService.modify_jsp.java:395)
  at org.apache.jasper.runtime.HttpJspBase.service(HttpJspBase.java:92)
  at javax.servlet.http.HttpServlet.service(HttpServlet.java:809)
  at org.apache.jasper.servlet.JspServletWrapper.service(JspServletWrapper.java:162)
  at org.apache.jasper.servlet.JspServlet.serviceJspFile(JspServlet.java:240)
  at org.apache.jasper.servlet.JspServlet.service(JspServlet.java:187)
  at javax.servlet.http.HttpServlet.service(HttpServlet.java:809)
  at org.apache.catalina.core.ApplicationFilterChain.internalDoFilter(ApplicationFilterChain.java:200)
  at org.apache.catalina.core.ApplicationFilterChain.doFilter(ApplicationFilterChain.java:146)
  at org.apache.catalina.core.StandardWrapperValve.invoke(StandardWrapperValve.java:209)
  at org.apache.catalina.core.StandardPipeline$StandardPipelineValveContext.invokeNext(StandardPipeline.java:596)
  at org.apache.catalina.core.StandardPipeline.invoke(StandardPipeline.java:433)
  at org.apache.catalina.core.ContainerBase.invoke(ContainerBase.java:948)
  at org.apache.catalina.core.StandardContextValve.invoke(StandardContextValve.java:144)
  at org.apache.catalina.core.StandardPipeline$StandardPipelineValveContext.invokeNext(StandardPipeline.java:596)
  at org.apache.catalina.core.StandardPipeline.invoke(StandardPipeline.java:433)
  at org.apache.catalina.core.ContainerBase.invoke(ContainerBase.java:948)
  at org.apache.catalina.core.StandardHostValve.invoke(StandardHostValve.java:133)
  at org.apache.catalina.core.StandardPipeline$StandardPipelineValveContext.invokeNext(StandardPipeline.java:596)
  at org.apache.catalina.valves.ErrorDispatcherValve.invoke(ErrorDispatcherValve.java:118)
  at org.apache.catalina.core.StandardPipeline$StandardPipelineValveContext.invokeNext(StandardPipeline.java:594)
  at org.apache.catalina.valves.ErrorReportValve.invoke(ErrorReportValve.java:116)
  at org.apache.catalina.core.StandardPipeline$StandardPipelineValveContext.invokeNext(StandardPipeline.java:594)
  at org.apache.catalina.core.StandardPipeline.invoke(StandardPipeline.java:433)
  at org.apache.catalina.core.ContainerBase.invoke(ContainerBase.java:948)
  at org.apache.catalina.core.StandardEngineValve.invoke(StandardEngineValve.java:127)
  at org.apache.catalina.core.StandardPipeline$StandardPipelineValveContext.invokeNext(StandardPipeline.java:596)
  at org.apache.catalina.core.StandardPipeline.invoke(StandardPipeline.java:433)
  at org.apache.catalina.core.ContainerBase.invoke(ContainerBase.java:948)
  at org.apache.coyote.tomcat4.CoyoteHandler.service(CoyoteHandler.java:152)
  at org.apache.coyote.http11.Http11Processor.process(Http11Processor.java:799)
  at org.apache.coyote.http11.Http11Protocol$Http11ConnectionHandler.processConnection(Http11Protocol.java:705)
  at org.apache.tomcat.util.net.TcpWorkerThread.runIt(PoolTcpEndpoint.java:577)
  at org.apache.tomcat.util.threads.ThreadPool$ControlRunnable.run(ThreadPool.java:683)
  at java.lang.Thread.run(Thread.java:595)

```

Identity Manager sends exception logs to `stdout` on the web application instance, which is often the application server console.

---

**NOTE** When you are finished, disable exception logging to stop unnecessary output to the application server logs. To disable exception logging, set the `exception.trace` key value back to **false**.

---

## Tracing Forms

You can enable tracing to troubleshoot edited forms and to check for expression statement errors within your form fields.

Use either of the following methods to enable tracing:

- Open the Waveset Properties page (`debug/Show_WSPProp.jsp`) in the Identity Manager Administrator interface. Locate the `form.trace` key and change the Value to **true**.
- Open the `config/Waveset.properties` file in a text editor, and change the `form.trace` key value to **true**.

Identity Manager reports any problems with form expression syntax to standard output.

---

**NOTE** The `form.trace` key is disabled by default because it produces trace information for every field on every page, including the Accounts List page, which affects system performance. Consider using a more targeted form and field tracing method.

When you are finished troubleshooting your forms, remember to disable tracing by changing the `form.trace` key value back to **false**.

---

Global XPRESS tracing might also be helpful while you are developing and updating forms and form processes. Although Global XPRESS tracing produces a large amount of output that affects system performance, this tracing method shows XPRESS output and might expose where problems are happening in your form.

For more information, see

- [“Tracing Global XPRESS” on page 80](#)
- [“Testing Your Customized Form” in \*Identity Manager Workflows, Forms, and Views\*](#)

## Tracing Global XPRESS

While not generally recommended, you can use global XPRESS tracing to trace any and all XPRESS code, wherever the code is located. For example, you can trace XPRESS code in forms, views, and workflows. The resulting trace shows XPRESS output that can expose potential problems.

---

**NOTE** XPRESS tracing is disabled by default because it produces a large amount of output, which affects system performance.

See “Testing Your Customized Form” in *Identity Manager Workflows, Forms, and Views* for more information about tracing XPRESS functions.

---

Use the following steps to enable global XPRESS tracing:

1. Open a command window.
2. Change directories to `config/Waveset.properties` in the default Identity Manager installation directory.
3. Open the `config/Waveset.properties` file and edit the `xpress.trace` line to read:

```
xpress.trace=true
```

4. Save the `Waveset.properties` file.
5. Restart your application server or reload the `Waveset.properties` file from the Identity Manager debug pages.
6. Replicate the XPRESS trace output to a file by adding this line to the `Waveset.properties` file:

```
xpress.traceFile=FileName.txt
```

XPRESS tracing is written to the JVM standard output.

## Tracing PasswordSync

This section describes how to enable tracing for PasswordSync and how to configure tracing in Direct access or JMS modes.

### To Enable Trace for PasswordSync

You can use the following methods to configure tracing for Identity Manager's PasswordSync feature:

- [Using the PasswordSync Configuration Tool](#)
- [Editing the Registry Keys](#)

#### *Using the PasswordSync Configuration Tool*

This section describes how to configure tracing from the PasswordSync Configuration tool Trace tab.

---

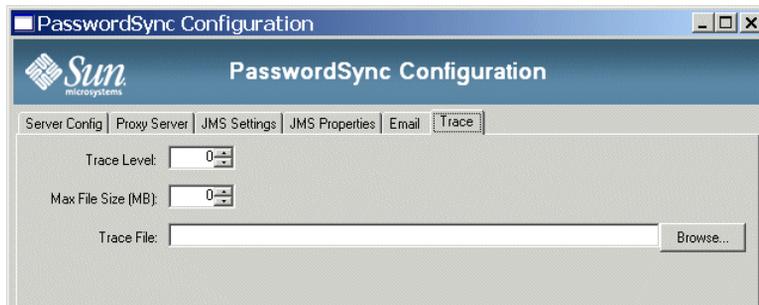
**NOTE** See Chapter 10, "PasswordSync," in *Identity Manager Administration* for PasswordSync configuration instructions.

The first time you run the configuration tool, the wizard does not allow you to configure tracing. Subsequently, when you run the configuration tool, the wizard displays a Trace tab where you can configure tracing.

---

The following figure shows the PasswordSync Configuration tool Trace tab.

**Figure 2-3** Trace Tab in the PasswordSync Configuration Tool



- Use the Trace Level field to specify the level of detail you want PasswordSync to provide when writing to the trace log. A value of 0 turns tracing off, while a value of 4 shows full detail.
- Use the Max File Size field to specify a maximum size for the log file.

When the trace file exceeds the size specified in the Max File Size (MB) field, PasswordSync starts a new trace file and appends .bk to the old trace file name. For example, if your trace level is set to 100 Mbytes, and your trace file writes to C:\logs\pwicsvc.log, when the trace file exceeds 100 Mbytes PasswordSync renames the file to C:\logs\pwicsvc.log.bk. PasswordSync then creates a new C:\logs\pwicsvc.log file where new trace file messages are written.

- Use the Trace File field to specify a location for the PasswordSync trace file.

### *Editing the Registry Keys*

To enable additional PasswordSync configuration settings, edit the following PasswordSync registry keys using the PasswordSync configuration tool.

---

**NOTE** Using PasswordSync configuration tool is the safest method for editing PasswordSync registry keys. Editing these keys directly in the Windows Registry is discouraged.

---

**Table 2-2** Registry Keys

Key Name	Type	Description
allowInvalidCerts	REG_DWORD	<p>If this registry key is set, the client tolerates expired certificates and certificates that have an invalid CN or hostname.</p> <p>You might find this setting useful when debugging test environments in which most certificates are produced from invalid certificate authorities (CAs).</p> <p>These settings only apply if SSL is being used.</p> <p>The default is 0.</p> <p>If set to 1, this key sets the following flags on the .NET client:</p> <ul style="list-style-type: none"> <li>• SECURITY_FLAG_IGNORE_UNKNOWN_CA</li> <li>• INTERNET_FLAG_IGNORE_CERT_CN_INVALID</li> <li>• INTERNET_FLAG_IGNORE_CERT_DATE_INVALID</li> </ul>

**Table 2-2** Registry Keys (*Continued*)

Key Name	Type	Description
dumpFilebase	REG_SZ	<p>Enables Windows to generate a dump file if the PasswordSync DLL displays an exception.</p> <p>Set the registry key to the fully qualified directory path where you want to write the memory dump. For example: <code>c:\temp</code></p> <p>Set the registry value to write the memory dump each time Identity Manager catches an exception during password processing.</p> <p><b>Note:</b> On Windows 2000 server (any service pack), you also must install in the configured directory <code>DbgHelp.dll</code>, which is available from Microsoft. The minimum release version for the <code>DbgHelp.dll</code> file must be 5.1. Download the <code>DbgHelp.dll</code> file here:</p> <p><a href="http://www.microsoft.com/whdc/DevTools/Debugging/default.mspx">http://www.microsoft.com/whdc/DevTools/Debugging/default.mspx</a></p> <p>If <code>DbgHelp.dll</code> is not installed, no dump files will be generated on Windows 2000.</p> <p>The format for dump file names is</p> <p><code>lhpwic-YYYYMMDD-HH:mm-xxxxx.dmp</code></p> <p>In this name, <code>YYYYMMDD</code> will be the date of the dump, <code>HH:mm</code> is the time of the dump (24-hour clock), and <code>xxxxx</code> is the thread number of the application.</p> <p><b>You must manually remove dump files!</b> Dump files range in size from 20 MB to more than 100 MB, depending on the size of the Windows Local Security Authority Subsystem (LSASS) process. Over time, systems with limited disk space could fill up if these dump files are not removed.</p>
installdir	REG_SZ	The directory where the PasswordSync application is installed.

The PasswordSync registry keys are located in the following location:

`HKEY_LOCAL_MACHINE\SOFTWARE\Waveset\Lighthouse>PasswordSync`

Other keys are present in this location.

## To Collect Logs for the Different Modes

PasswordSync trace logs are the same, whether you are using a direct access mode or JMS mode configuration. However, these trace logs might only provide partial information. You must configure different classes for each configuration to collect logs on the server side, as described in the following sections.

### *Tracing in Direct Mode*

When using PasswordSync with a direct access mode configuration, the trace logs show failures, but not all logged failures are real failures. For example, in some circumstances the view check-in takes a long time, which shows as a failure in the log. You must trace on the server side to see this information.

In Direct mode, PasswordSync talks to the servlet that generates the view to be checked into the repository. You can trace the `com.waveset.rpc.GenericMessageHandler` class at level 4 to view all phases of password synchronization — from receiving the password change to the response generated and returned to the servlet. Level 4 is the only level that supplies enough detail for troubleshooting.

### *Tracing in JMS Mode*

When using PasswordSync with a JMS mode configuration, the logs only show successful or failed deliveries to the JMS server. From this point on, you must rely on server side logs. JMS tracing is a little more complex.

You can trace the `com.waveset.rpc.PasswordSyncHandler` class at level 4 to convert the messages generated by the PasswordSync dll into a JMS message and add those messages to the JMS queue. Limited tracing is available in this class, and only level 4 provides enough information to help with troubleshooting.

If PasswordSync successfully delivers the JMS message to the JMS queue, the tracing will not help you find the cause of a problem. The next, and final step is to trace the JMS adapter. See the *Identity Manager Resources Reference* for instructions.

## Tracing Identity Manager Service Provider Delegated Administration

Enabling Identity Manager tracing at Method/Class Level 2 for the following classes allows you to trace authorization flows when listing or accessing Identity Manager Service Provider (Service Provider) users and when AdminRoles are dynamically assigned when Service Provider users log in.

- Trace `com.sun.idm.idmx.view.IDMXBrowseViewer` when searching for Service Provider users.
- Trace `com.sun.idm.idmx.view.IDMXUserViewer` when creating, editing, or deleting Service Provider users.
- Trace `com.sun.idm.idmx.api.IDMXServerUtils` when using both of the preceding classes.
- Trace `com.waveset.security.authn.WSSPELoginModule` when you are logging in as a Service Provider user.

---

**NOTE** You configure Service Provider tracing from the Identity Manager debug pages. If necessary, review [“Tracing the Identity Manager Server” on page 75](#) for instructions.

---

## Tracing Reconciliation

If you are having problems with a reconciliation task, you can use the standard tracing facility on `com.waveset.task.Reconciler` to trace the Reconciler.

Use either of the following methods to enable tracing:

- Open the Waveset Properties page (`debug/Show_WSPProp.jsp`) in the Identity Manager Administrator interface. Locate the `exception.trace` key and change the Value to **true**.
- Open the `config/Waveset.properties` file in a text editor, and change the `exception.trace` key value to **true**.

You can also enable tracing from the System Settings page and trace the following reconciliation methods, at the Method/Class trace Levels noted, to view useful debugging information:

**Table 2-3** Reconciliation Methods/Classes to Trace

Trace this Method/Class	At this Trace Level	To See
<code>com.waveset.recon.ReconTask\$ResourceThreadwaitForLighthouseWorkItems#</code>	4	How many users queued and processed during the Identity Manager user examination process
<code>com.waveset.recon.ReconTask\$ResourceThread#waitForReconcileWorkItems</code>	4	How many reconciles and responses queued and processed for the specified resource
<code>com.waveset.recon.ReconTask\$ResourceThread#examineResource</code>	3	How many accounts read from the resource
<code>com.waveset.recon.ReconTask\$ResourceThread#queueAccountReconciles</code>	3	Detailed information about each account read from the resource such as <code>accountId</code> , <code>accountGUID</code> , <code>accountDisabled</code>
<code>com.waveset.recon.ReconTask\$ResourceThread#examineLighthouse</code>	3	How many Identity Manager users who claim to own an account on the reconciled resource queued
<code>com.waveset.recon.ReconTask\$WorkerThread#findClaimants</code>	3	All Identity Manager users who claim to have an account on the resource
<code>com.waveset.recon.ReconTask\$WorkerThread#correlateUsers</code>	4	A list of correlated/confirmed users
<code>com.waveset.recon.ReconTask\$WorkerThread#confirmPossibleOwners</code>	3	A list of all confirmed owners of resource accounts
<code>com.waveset.recon.ReconTask\$WorkerThread#reconcileAccount</code>	2	The individual account being reconciled
<code>com.waveset.recon.ReconTask\$WorkerThread#performResponse</code>	2	The individual account and user during response
<code>com.waveset.recon.ReconTask\$WorkerThread#respondOrRequeue</code>	4	The error message if there is a problem performing the previous response method for an account
<code>com.waveset.recon.ReconTask\$WorkerThread#applyResponse</code>	3	The response list that is being applied
<code>com.waveset.recon.ReconTask\$WorkerThread#failUserExamination</code>	2	All user examination requests that failed with the error
<code>com.waveset.recon.ReconTask\$WorkerThread#failUserResponses</code>	2	All user response requests that failed with the error

**Table 2-3** Reconciliation Methods/Classes to Trace

<b>Trace this Method/Class</b>	<b>At this Trace Level</b>	<b>To See</b>
<code>com.waveset.recon.AccountContext#processAttributeWorkflow</code>	3	The attribute changes and the formatted changes during the launch of the attribute change workflow
<code>com.waveset.recon.Response#perform</code>	4	The user, response, and the user's resource info XML before and after the modifications took place
<code>com.waveset.recon.Response#perform</code>	4	A summary of the response action on the user
<code>com.waveset.recon.Response#createNewUserFromAccount</code>	4	Full details of the response action on the user
<code>com.waveset.recon.ReconUtil#getRuleState</code>	3	Full User view with the user's attribute during rule processing
<code>com.waveset.recon.ReconUtil#evaluateCorrelationRule</code>	3	The value of the correlation rule state and result for examination
<code>com.waveset.recon.ReconUtil#evaluateConfirmationRule</code>	4	The value of the confirmation rule state and result for examination
<code>com.waveset.recon.ReconUtil#getCorrelatedUsers</code>	4	A list of correlated users matching the specified rule result
<code>com.waveset.recon.ReconUtil#confirmPotentialOwners</code>	3	A list of users who have been confirmed via the confirmation rule
<code>com.waveset.recon.ReconUtil#getIfExistsAccountIndexEntry</code>	3	To output info related to examination of the account index for a specified entry
<code>com.waveset.recon.ReconUtil#launchWorkflow</code>	3	The task instance and task definition information when launched
<code>com.waveset.recon.ReconUtil#deleteAccountIndex</code>	2	User information to be deleted from the account index
<code>com.waveset.recon.ReconUtil#indexFoundAccount</code>	3	The account and situation recorded during a create or update of the index for an account known to exist
<code>com.waveset.recon.ReconUtil#indexMissingAccount</code>	3	The account and situation recorded during a create or update of the index for an account NOT known to exist
<code>com.waveset.recon.ReconUtil#listAccountsIndexSaysExist</code>	3	Account information that the index says exists
<code>com.waveset.recon.UserContext#aquireRepoLock</code>	2	The user who is being locked for update
<code>com.waveset.recon.UserContext#releaseRepoLock</code>	2	The user who is being unlocked in the repository

---

**NOTE** Remember, the higher the tracing level, the larger the trace file. Also tracing all of these methods at the same time will create a *very* large trace file.

When you are finished troubleshooting, remember to disable tracing by setting the `exception.trace` key value back to **false**.

---

## Tracing the `setRepo` Command

If you see errors while you are using the `setRepo` command to configure the Identity Manager repository, use the following flags to isolate and debug the problem:

```
-Dtrace.enabled=true
-Dtrace.level.com.waveset.repository.AbstractDataStore=2
-Dtrace.level.com.waveset.repository.DefaultTypeHandler=4
// Use one of the following based on your repository type
-Dtrace.level.com.waveset.repository.OracleDataStore=4
-Dtrace.level.com.waveset.repository.SqlServerDataStore=4
-Dtrace.level.com.waveset.repository.MysqlDataStore=4
-Dtrace.level.com.waveset.repository.DB2DataStore=4
```

Identity Manager sends output from the `setRepo` command to the default `$WSHOME/config/WSTrace.log` file.

## Tracing SPML

This section describes methods for enabling trace for SPML version 1.0 and SPML version 2.0.

### To Enable Tracing for SPML 1.0

SPML 1.0 provides the following options for turning on trace output so you can log Identity Manager's SPML traffic and diagnose problems.

**Method 1: Enable the `setTrace` Method**

You can use the `setTrace` method, provided by the `SpmIClient` and `LighthouseClient` classes, to enable tracing for SPML 1.0.

When you enable this `setTrace` method, the XML for the request sent by the client and the XML for the response received from the server are printed to the *client* console as they are sent and received.

The `setTrace` method takes a Boolean argument. For example

```
SpmIClient client = new SpmIClient();
client.setURL("http://localhost:8080/idm/spml");
client.setTrace(true);
```

**Method 2: Pass the `trace` Operational Attribute**

You can enable tracing for an individual SPML RPC request by passing a `trace` operational attribute to the RPC request on the server side.

Tracing occurs during servlet initialization, and it controls how information is output for the RPC traffic of a servlet handling SPML version 1.0 requests. For example, the `trace` prints the raw XML that is sent back and forth on whatever the `System.out` is for that servlet (which is a function of the Application container). For example:

```
AddRequest ar = new AddRequest();
ar.setOperationalAttribute("trace", "true");
```

When you use the `trace` attribute, how the attribute affects server operation is vendor-specific. Currently, Identity Manager prints the raw request and response data to the *server* console, which is useful if the client application is not associated with a console window.

For more information consult your OpenSPML Toolkit product documentation.

## To Enable Tracing for SPML 2.0

SPML 2.0 provides the following options for turning on trace output so you can log Identity Manager's SPML traffic and diagnose problems.

### *Method 1: Initializing the SOAP `rpcrouter` Servlet*

You can enable tracing by initializing the SOAP `rpcrouter` servlet, which controls the output of RPC traffic information for the servlet handling SPML requests. The `rpcrouter` servlet takes an *<init parameter>* that enables SOAP tracing on the server side. The servlet's initialization logic checks for a `trace` configuration parameter and then, if the parameter's value is `true`, prints the raw request and response data to the console.

For example, servlet prints the raw XML that is sent back and forth to whatever `System.out` has been specified for that servlet. (Where `System.out` is a function of the application container.)

---

**NOTE** The SOAP `rpcrouter` servlet is a third-party, open source `org.openspml.server.SOAPRouter` class from the SPML organization. For more information about this servlet, refer to your OpenSPML Toolkit documentation.

---

### *Method 2: Pass the `trace` Operational Attribute*

As for SPML 1.0, you can enable tracing for an individual SPML RPC request by passing a `trace` operational attribute to the RPC request on the server side.

```
AddRequest ar = new AddRequest();
ar.setOperationalAttribute("trace", "true");
```

Tracing controls what information is output for the RPC traffic of a servlet handling SPML 2.0 requests.

When you use the `trace` attribute, how that attribute affects server operation is vendor-specific. Currently, Identity Manager prints the raw XML request and response data to the *server* console, which is useful if the client application is not associated with a console window.

For more information consult your OpenSPML toolkit product documentation.

## Tracing the Task Scheduler

If a Scheduler task is having problems, you can use the standard tracing facility on `com.waveset.task.Scheduler` to trace the task scheduler. The output shows detailed information about task scheduling.

To enable tracing,

1. Open a browser and log in to the Administrator interface.
2. Select Configure > Servers.
3. When the Configure Servers page displays, click the server name in the Server column to edit the settings.
4. On the Edit Server Settings page, click the Scheduler tab.
5. Select the Tracing Enabled box to activate Scheduler debug tracing and write the results to `stdout`.
6. Save your changes.

- 
- NOTE**
- In a clustered environment, tracing occurs on each instance.
  - See [“Tracing the Identity Manager Server” on page 75](#) for more information about defining and editing trace configuration objects, and about viewing trace files.
- 

## Tracing Workflows

Enabling workflow tracing can help you resolve problems with workflow activities and understand the workflow process.

- 
- NOTE**
- In a clustered environment, tracing occurs on each instance.
  - To debug workflows in a multiple server deployment environment, consider shutting down all but one server. If all servers are running, you cannot determine which server in the environment is executing the workflow, which causes troubleshooting problems.
-

## To Enable Tracing for Workflows

Use the following steps to enable workflow tracing:

1. Open a browser and log in to the Identity Manager Administrator interface.
2. From the System Setting page, choose Configuration from the List Objects Type menu.
3. Click the List Objects button.
4. When the List Objects of type: Configuration page displays, scroll down the list of objects to locate the System Configuration object and click the edit link.
5. When the Checkout Object: Configuration, #ID#CONFIGURATION:SYSTEMCONFIGURATION page displays, you can edit any of the following workflow options in the SystemConfiguration object:

---

**NOTE** Typically, you enable only one option, but it is possible to enable more than one option at a time.

These attributes are not dependent on each other. You can turn on one type of trace while the other types are turned off.

---

- Specify `workflow.consoleTrace=true` to redirect workflow trace messages to the application server console, which can be useful when workflows are terminating due to a fatal exception because this attribute prints more trace output than `workflow.fileTrace`. (Default value is false.)
- Specify `workflow.fileTrace=PathToFile` to redirect workflow trace messages to a file that is easy to read. This attribute does not have a value by default.

```
<Attribute name='fileTrace' />
```

Add a value tag to the `workflow.fileTrace` attribute and, using Unix-style forward slashes, enter the path to a log file. Identity Manager stores relative pathnames relative to the application server's installation directory. For example,

### On Windows:

```
<Attribute name='fileTrace' value='C:\mydir\workflow.txt' />
```

### On Solaris/Unix:

```
<Attribute name='fileTrace' value='/mydir/workflow.txt' />
```

- Specify `workflow.traceLevel=tracingLevel` to specify the level of workflow tracing you want to see.
  - Specifying no value or a value of 0 generates the most tracing detail.
  - Specifying a value of 1, suppresses tracing of variable and argument values that are `GenericObjects`. `GenericObject` values are frequently large and passed among subprocesses, so their trace can be redundant.
- Specify `workflow.Trace=true` to trace workflow processing. You must restart the application server to start tracing with this option. Identity Manager stores the trace results in the task's `WavesetResult` object. Use this tracing option when file system access is unavailable. (Default value is `false`.)
- Trace messages in the task's `WavesetResult`. (Default value is 1.)

---

**NOTE** Specifying `workflow.Trace=true` appends trace messages into one long, unformatted string that is difficult to read. Use this option only when you do not have access to the file system.

---



---

**NOTE** With the first two options, you might lose some of the workflow trace if a fatal exception occurs.

---

6. Save the `SystemConfiguration` object.
7. Restart your application server, or reload the `SystemConfiguration` object from the Identity Manager debug area.

## To Enable Tracing for a Designated Workflow

Use one of the following methods to enable tracing for a designated workflow:

- **Editing the WFProcess definition:** To enable trace unconditionally for a particular process, edit the XML of the `TaskDefinition` object by adding the `trace='console'` attribute to the `<WFProcess>` element.
- **Editing the workflow variable:** Gives you more control over the timing of tracing by using a workflow variable. The workflow engine will look for a top-level workflow variable named `trace`.

The following example shows how to trace a workflow variable:

**Code Example 2-1** Tracing a Workflow Variable

```
<Variable name='trace'>
  <cond><eq><ref>accountId</ref><s>jfaux</s></eq>
    <s>workflowTrace.txt</s>
  </cond>
</Variable>
```

The trace variable turns tracing on only if the workflow is operating on a user named `jfaux`. You could also specify `trace` in a form field to control tracing interactively.

In this example, the trace output is written to the `workflowTrace.txt` file.

## Locating Version Information

You can use one of the following methods to get information about the Identity Manager version you are currently using:

- Hover your cursor above the Help button in the upper right corner of the Identity Manager application window. A pop-up is displayed that contains the version number.
- Open the `Waveset.properties` file, and check the information provided at the top of the file.

Use the Identity Manager System Properties page (`/debug/SysInfo.jsp`) to get information about which JVM versions, and other system software versions, you are currently using.

# Tracing the Identity Manager Gateway Objects and Activities

This section describes how to trace objects and activities in Sun Identity Manager Gateway, the information is organized as follows:

- [How to Configure Tracing](#)
- [How to Configure Tracing for the PowerShellExecutor.dll Add-On](#)
- [How to Capture Dr. Watson Logs](#)

- 
- NOTE**
- When viewing or editing a Gateway trace file, use Notepad to avoid file restrictions.
  - When you start the Gateway, the program appends new trace entries to the trace file instead of deleting entries. To locate the point at which the Gateway trace entries begin, look for a Gateway version string.
  - The Gateway version is output in the trace automatically when you start the Gateway. You can also type `gateway -v` from the command line to get the version.
- 

## How to Configure Tracing

You can enable tracing from the Gateway Debug page (`Gateway.jsp`) or from the command line to debug problems with Windows accounts on Identity Manager.

Instructions are provided in the following sections:

- [From the Gateway Debug Page](#)
- [From the Command Line](#)

### From the Gateway Debug Page

Enable tracing from the Gateway Debug page (`Gateway.jsp`) if you cannot access the Gateway. You can specify and retrieve Gateway trace files from this debug page.

Use the following steps to enable tracing:

1. Log in to the Identity Manager Administrator interface.
2. Type the following URL in to your browser to open the Gateway Debug page:  
`http://host:port/idm/debug/Gateway.jsp`
3. Choose a resource to trace from the Gateway Resource list.
4. Click the following buttons to modify the existing settings:
  - **Get Version:** Returns the Gateway version and the operating system of the machine on which you are running the Gateway.
  - **Get Trace File:** Returns the contents of the trace file.
  - **Get Trace Parameters:** Returns the path of the trace file, the trace level, and the maximum size of the trace file.
  - **Set Trace Parameters:** See [“To Create a New Trace Configuration Object” on page 73](#) for information about these options.
  - **Get Loaded Modules:** Returns the load addresses of modules (DLLs) being used by the Gateway.

---

**NOTE** The Get Loaded Modules list consists of load addresses, followed by module names and only includes loaded modules. The list does not include delay-loaded modules that have not been called.

The Get Loaded Modules option only supports Active Directory and Domino.

---

## From the Command Line

Enabling trace from the command line is useful if you want a wider range of options.

To enable tracing:

1. Open a command window.
2. Start the Gateway, specifying the necessary trace command arguments.

See the following table for a description of these command line arguments.

**Table 2-4** Gateway Tracing Command Line Arguments

Argument	Description
-f	Specify the path to the trace file
-l	Specify the level of tracing: <ul style="list-style-type: none"> <li>• <b>Level 0:</b> Disables tracing. (Default)</li> <li>• <b>Level 1:</b> Traces the flow of control between components and generally defines a <i>low-detail</i> trace point that includes entry and exit from high-level functional methods.</li> <li>• <b>Level 2:</b> Generally defines a <i>medium-detail</i> trace point that includes entry and exit from every method, and information and data trace points for high-level functional methods. Level 2 adds the flow of control within each component, major decision points, and items of information.</li> <li>• <b>Level 3:</b> Generally defines a <i>high-detail</i> trace point that includes entry and exit from every method, information and data trace points for high-level functional methods, and significant subroutines. Level 3 adds lower-level decision points and items of information.</li> <li>• <b>Level 4:</b> Generally defines a <i>hyper-detail</i> trace point that includes everything traced in the other trace levels. Level 4 traces at a very low level and provides a level of detail that is seldom needed but might be useful in characterizing complex behaviors of some components. <b>NOTE:</b> Not all components support level 4.</li> </ul> <p>Trivial methods, such as getters and setters, generally do not have entry or exit trace points because they add overhead.</p>
-m	Specify the maximum trace file size in kilobytes <p>When the trace file reaches <code>-m</code> Kbytes, Identity Manager closes the current trace file, deletes any existing back-up files, renames the current trace file to the name specified by the <code>-f</code> argument with <code>.bk</code> appended, and opens a new trace file with the <code>-f</code> argument name.</p> <p>For example, if you specified <code>-f beebble.trc</code> on the command line, the following two files result after <code>-m</code> Kbytes are recorded:</p> <pre>beebble.trc.bk beebble.trc</pre> <p>Where <code>beebble.trc</code> contains the most recent traces.</p>

Usage: gateway -f name -l -m

For example:

```
cd %WSHOME%\bin\winnt
```

```
gateway -d -p 11319 -f %CD%\gateway.trc -l 2 -m 500
```

The preceding invocation starts the Gateway with the following characteristics:

- **-d**: Use regular application (not a service)
- **-p 11319**: Use port 11319

You must configure this port for Gateway resources from the Identity Manager resource configuration. For example, for an Active Directory resource

- **-f %CD%\gateway.trc**: Directory to which the trace output is written. Identity Manager writes the trace output to a text file in this directory.
- **-l 2**: Output level 2 of Gateway tracing.
- **-m**: Maximum size in Kilobytes of trace log file.

---

**NOTE** If specified, Identity Manager saves **-f**, **-l**, and **-m** values in the registry, so the next time you run Gateway from the command line or as a service, the same values are used.

Identity Manager sends the Gateway trace output to the console *and* to a trace file.

---

## How to Configure Tracing for the PowerShellExecutor.dll Add-On

The PowerShellExecutor.dll is an add-on that implements communication between the gateway and Microsoft PowerShell. The PowerShell is used to manage Exchange Server 2007 accounts. This add-on cannot share tracing facilities with the rest of the gateway and provides a similar stand-alone tracing facility as the rest of the gateway.

The trace configuration for the PowerShellExecutor is stored in the same registry key as the other gateway registry keys:

```
HKEY_LOCAL_MACHINE\Software\Waveset\Lighthouse\Gateway
```

This base key is created when you configure tracing through the Identity Manager debug pages or when you start the gateway with trace command arguments.

On shut down, the gateway writes the current `PowerShellExecutor` settings for the tracing to the registry. These settings include:

- `traceFileName`

**Content:** File name for the trace output (registry type `REG_SZ`)

**Default** " "

Name of the trace file to generate for the `PowerShellExecutor` tracing. Where the name:

- Can be a fully qualified path, including the filename.
- Cannot end in a slash (\)

The full path, except the file, provided in the `traceFileName` must exist.

If configured, log rotation adds a timestamp, in the form "`yyyyMMddHHmmss`," to the configured filename after rotation, when the file is no longer active.

- `traceLevel`

**Content:** Trace level (registry type `REG_DWORD`)

**Default:** 0 (no tracing)

**Allowed:** 0–4

This key is shared with the rest of the gateway. The whole gateway always provides tracing at the same level.

- `traceMaxSize`

**Content:** Maximum file size in bytes (registry type `REG_DWORD` or `REG_QWORD`)

**Default:** 100000 bytes

**Minimum:** 100000 bytes

Tracing text is written as UTF8 encoded text with a byte order mark to make it portable to other systems.

- `traceMaxFiles`

**Content:** Number of trace files (registry type REG\_DWORD)

**Default:** 2

**Minimum:** 1

This setting controls the number of trace files to keep on the system. Setting the maximum number of files to keep to **1**, causes the file to be overwritten when the maximum size is reached. The oldest file, based on last write time, is removed when the maximum number of files is reached.

- `traceConfigInterval`

**Content:** Time out in milliseconds (registry type REG\_DWORD)

**Default:** 300000 ms (5 minutes)

**Minimum:** 60000 ms (1 minute)

All trace settings are re-read from the registry based on this timeout value. In a production environment, consider setting this value to a large value, such as 24 hours, to minimize overhead.

## How to Capture Dr. Watson Logs

If the Gateway encounters a serious problem and exits abnormally, you can send the resulting Dr. Watson logs to Sun Support for analysis.

---

**NOTE** You must have administrator privileges on the system to view these logs.

---

To capture a Dr. Watson log:

1. Open the Windows Event Viewer.
2. Open the application log.
3. Look for an event with `DrWatson` source.
4. Open the event to view detailed information.
5. Ensure that the Bytes option is selected for Data.
6. Right-click in the display dialog and choose Select all from the menu.

7. Type **Ctrl-C** to copy the information.
8. Paste the information into Notepad and save the file.
9. Send the file in an email to Sun Support with a detailed description of the problem. Be sure to indicate which version of the Identity Manager Gateway you are running.

## Troubleshooting and Fixing Common Problems

Use the information provided in the following sections to help diagnose and fix problems you might encounter as you work with Identity Manager:

- [Working with Debugging Tools](#)
- [Debugging Errors Displayed in the Browser](#)
- [Troubleshooting Adapters](#)
- [Troubleshooting Form Problems](#)
- [Troubleshooting the Gateway](#)
- [Troubleshooting Java Code Problems](#)
- [Troubleshooting Low Memory Conditions](#)
- [Troubleshooting PasswordSync Problems](#)
- [Troubleshooting Reconciliation Problems](#)
- [Troubleshooting Repository Connection Problems](#)
- [Troubleshooting Server-Related Problems](#)
- [Troubleshooting an SPML Configuration](#)
- [Troubleshooting Upgrades](#)

---

**NOTE** For additional troubleshooting information, including the Identity Manager FAQ, visit the following location:

[https://sharespace.sun.com/gm/document-1.26.2296/IdM\\_FAQ.html?](https://sharespace.sun.com/gm/document-1.26.2296/IdM_FAQ.html?)

You must sign up for a Share Space ID to access information provided on this site.

---

## Working with Debugging Tools

You can use several different debugging tools to help identify and fix problems in your Identity Manager deployment. These tools include:

- [Identity Manager Debug Pages](#)
- [Identity Manager IDE](#)
- [Identity Manager System Monitoring](#)
- [Adapter Logs](#)
- [DTrace](#)
- [JConsole](#)

### Identity Manager Debug Pages

---

**NOTE** You must have the *Debug* capability to access and execute operations from the Identity Manager Debug pages. Administrators and the Configurator are assigned this capability by default.

If you do not have the Debug capability, an error message results.

---

To access individual Identity Manager debug pages:

1. Open a browser and log in to the Identity Manager Administrative interface.
2. Type the following URL into the browser to open the System Settings page:

`http://host:port/idm/debug`

Where:

- *host* is the local server on which you are running Identity Manager.
- *port* is the number of the TCP port on which the server is listening.

From this page, you can enable or disable tracing for various Identity Manager activities and objects and use the information displayed on these pages to troubleshoot problems in your deployment.

Some Debug pages are not linked to the System Settings page, and you must type the page's `.jsp` file name into the browser to open the page, as follows:

`http://host:port/idm/debug/pageName.jsp`

Where *pageName.jsp* is the particular Debug page you want to open.

The following table describes the most commonly used Debug pages and their actual .jsp file names.

**Table 2-5** Identity Manager Debug Pages

Page Name	Use This Page to
Control Timings ( <code>callTimer.jsp</code> )	<p>Collect and view call timer statistics for different methods. You can use this information to track bottlenecks to specific methods and invoked APIs. You can also use options on the Call Timings page to import or export call timer metrics.</p> <ul style="list-style-type: none"> <li>• How long to fetch initial User view (with no resources) during a scan</li> <li>• How long to refresh initial User view (including resources) during a scan</li> <li>• How long to evaluate policy on the User view</li> <li>• How long each user scan takes (including User view fetch, policy evaluation, and so forth)</li> <li>• How long to fetch a list of users for an access scan</li> <li>• How long to evaluate the attestation rule in access review</li> </ul>
Edit Trace Configuration ( <code>Show_Trace.jsp</code> )	<p>Enable and configure trace settings for the Java classes provided with your Identity Manager installation. You can specify</p> <ul style="list-style-type: none"> <li>• Which methods, classes, or packages to trace and the level of trace.</li> <li>• Whether to send trace information to a file or to standard output and how dates and times are formatted in the trace output file.</li> <li>• Maximum number of trace files to store and the maximum size of each file.</li> <li>• Specify the maximum number of methods to be cached.</li> <li>• Indicate how to write data to the trace file and whether to write data to the trace file as data is generated or to queue and then write the data to a file.</li> </ul>
Host Connection Pool ( <code>Show_ConnectionPools.jsp</code> )	<p>View connection pool statistics (if you are not using a data source), including pool version, how many connections were created, how many are active, how many connections are in the pool, how many requests were serviced from the pool, and how many connections were corrupted.</p> <p>You can also use the Host Connection Pool page to view a summary of the connection pools used to manage connections to the Gateway. You can use this information to investigate low-memory conditions.</p>
List Cache Cleared ( <code>Clear_XMLParser_Cache.jsp</code> )	<p>Clear recently used XML parsers from the cache and investigate low memory conditions.</p>

**Table 2-5** Identity Manager Debug Pages (Continued)

Page Name	Use This Page to
Method Timings ( <a href="#">Show_Timings.jsp</a> )	<p>Detect and assess hotspots at a method level. Use this page to gather information from Identity Manager methods, including:</p> <ul style="list-style-type: none"> <li>• Method names</li> <li>• How many times the methods were called</li> <li>• How many times the methods exited with an error status</li> <li>• Average time consumed by the methods</li> <li>• Minimum and maximum times consumed by invocations of each method</li> </ul>
Object Size Summary ( <a href="#">Show_Sizes.jsp</a> )	<p>Detect problematically large objects that can affect your system. This page shows the size of objects (in characters) stored in the repository, including the objects' total combined size, average size, maximum size, and minimum size. Click entries in the Type column to view the ID, name, and size of the largest configuration objects in the repository.</p>
Provisioning Threads for Administrator Configurator ( <a href="#">Show_Provisioning.jsp</a> )	<p>View a summary of provisioning threads in use by the system (a subset of the information available in <a href="#">Show_Threads.jsp</a>).</p>
System Cache Summary ( <a href="#">Show_CacheSummary.jsp</a> )	<p>View the following information to investigate low-memory conditions:</p> <ul style="list-style-type: none"> <li>• Administrator-associated object caches</li> <li>• System object cache</li> <li>• User login sessions</li> <li>• XML parser cache</li> </ul>
System Memory Summary ( <a href="#">Show_Memory.jsp</a> )	<p>View how much total and free memory is available (in MB) when you are using memory-intensive functionality, such as reconciliation, to help determine whether there is sufficient memory allocated to the JVM. You can also use this page to launch garbage collection or to clear unused memory in the JVM for investigating heap usage.</p>
System Properties ( <a href="#">SysInfo.jsp</a> )	<p>View information about your environment.</p>
System Threads ( <a href="#">Show_Threads.jsp</a> )	<p>View which processes are running to verify that automated processes are running. Includes information about the process type, process name, priority, if the process is a daemon, and if the process is still alive (running).</p>
User Session Pool Cleared ( <a href="#">Clear_User_Cache.jsp</a> )	<p>Use the Session Pool Clearer page to investigate low memory conditions.</p>
Waveset Properties ( <a href="#">Show_WSProp.jsp</a> )	<p>View and <i>temporarily</i> edit properties in the <code>Waveset.properties</code> file. Edited property settings remain in effect only until the next server restart.</p>
XML Resource Adapter Caches Flushed and Cleared ( <a href="#">Clear_XMLResourceAdapter_Cache.jsp</a> )	<p>Clear test XML resource adapters from the cache and use to investigate low memory conditions.</p>

For a more-complete list of debug pages, open a command window and list the contents of the `idm/debug` directory.

See “[Working with Identity Manager Debug Pages](#)” on page 53 for more information about these Debug pages.

## Identity Manager IDE

The Sun™ Identity Manager Integrated Development Environment (Identity Manager IDE) is Java application that enables you to view, customize, and debug Sun™ Identity Manager (Identity Manager) objects in your deployment.

Specifically, the Identity Manager IDE provides a graphical Debugger that you can use to debug Identity Manager forms, rules, and workflows. You can use this Debugger to set breakpoints and watches, step through code, examine and modify variables, examine classes and the callstack, follow threads, and run multiple sessions.

Instructions for installing and configuring the Identity Manager Integrated Development Environment (Identity Manager IDE) are now provided on <https://identitymanageride.dev.java.net>.

## Identity Manager System Monitoring

You can configure Identity Manager system monitoring to track system events. System monitoring collects and aggregates statistics at various levels to present a real-time view of system events, based on your specifications.

Viewing this information in dashboard graphs enables you to quickly assess system resources, view abnormalities, understand historical performance trends, and interactively isolate problems before looking at audit logs. Although dashboards do not provide as much detail as audit logs, dashboards can indicate where to look for problems in the logs.

For more information about dashboards and system monitoring, see the “Reporting” chapter in *Identity Manager Administration*.

## Adapter Logs

Adapter logs capture information about the adapter that is currently processing. You can use this information to monitor the adapter’s progress and to diagnose and debug adapter problems.

---

**NOTE** You must enable tracing and identify the methods for which tracing is requested before any logging can occur. Also, your customized adapter must include calls that create log entries for new methods.

---

Nearly every adapter has its own log file, path, and log level. You can specify the level of detail captured by the adapter log, along with these other values in the Logging section of the Synchronization Policy for the appropriate Identity Manager or Service Provider user type.

For more information about using adapter log files as a debugging tool, see [“Troubleshooting Adapters” on page 107](#).

## DTrace

DTrace is a comprehensive, dynamic tracing framework for the Solaris operating environment. DTrace provides more than 30,000 probes into your production system and integrates user- and kernel-level tracing. You can use DTrace to monitor JVM activity. This facility also allows you to use the D language (similar to C or awk) to trace arbitrary data and expressions.

## JConsole

The Java Monitoring and Management Console (*JConsole*) is a Java Management Extension (*JMX*) technology-compliant graphical management tool bundled with JDK 5 (and later). JConsole connects to a running JVM and gathers information from the JVM MBeans in the connected JMX agent.

For example, you can use JConsole to

- Detect low memory
- Enable or disable garbage collection
- Enable or disable verbose tracing
- Detect deadlocks
- Control Identity Manager log levels
- Access information about operating systems resources (Sun's platform extension).
- Monitor and manage MBeans
- View information about the JVM and monitored values, threads running on the application, and class loading

---

**NOTE** For more information about JConsole, see the article titled, *“Using JConsole to Monitor Applications.”* You can view this article from the following location:

<http://java.sun.com/developer/technicalArticles/J2SE/jconsole.html>

---

## Debugging Errors Displayed in the Browser

If a red error message displays in the Identity Manager interface after you have performed an action, you might be able to view more complete information and further analyze the error by viewing and saving the page source.

- To view the page source if you are using Internet Explorer, select View > Source from the menu bar.
- To view the page source if you are using Netscape, select View > Page Source from the menu bar.

If you still need help resolving the problem,

1. View the page source, and then select File > Save to save the file to your system.
2. Locate the error in your saved file.
3. Send the error information, the URL from the page where the problem occurred, and a description of the problem in an email to Sun Support for resolution assistance.

## Troubleshooting Adapters

To troubleshoot an adapter, review the adapter's log file. Almost all adapters write their resource settings to a log file, and you can use this information to confirm that the adapter started and that all setting changes have been saved.

---

**NOTE** You must enable tracing and identify the methods for which tracing is requested before any logging can occur. Also, your custom adapter must include calls that create log entries for new methods.

If necessary, review [“Tracing the Identity Manager Server” on page 75](#) for instructions about how to enable tracing.

---

Most adapter log files are located in the \$WSHOME/config directory and they are named WSTrace1.log.

Active Sync-enabled adapters that make log calls to the ActiveSyncUtil instance create a log file or set of log files in a directory specified by the Log File Path resource attribute. Be sure to check these log files for additional Active Sync-related log entries.

The information in this section is organized as follows:

- [To Debug an Adapter](#)
- [To Debug LoginConfig Changes](#)
- [To Debug Adapter Connection Problems](#)

## To Debug an Adapter

Follow these general steps to debug your custom adapter.

1. Create a test program for your adapter, and be sure this Java file performs the following basic functions:
  - Create a new resource
  - Create a user
  - Get a user
  - Update a user
  - Delete a user
  - Perform create, get, update and delete operations on multiple users

---

**NOTE** A sample test file (`SkeletonResourceTests.java`) is provided in the `/REF` directory on your installation CD.

---

2. Set an appropriate logging level for debugging.

For example, for the first debugging pass, increase the logging level to 4 (maximum debugging output), set the log file path, and specify a maximum file size.

When you start the adapter, all of the resource settings are written to the log file. You can use this information to validate that the adapter started and that all setting changes were saved.

3. Compile and test your adapter.
  - To compile the test program, open a command window and enter the `javac -d . test/filename.java` command. This command creates the class file in the appropriate `com/waveset/adapter/test` directory.

- To test your new adapter using this class file, be sure that your compiled adapter is in the `com/waveset/adapter` directory and use the following command to run the adapter:

```
java -D waveset.home=path com.waveset.adapter.test.MyResourceAdapter
```

4. Create an HTML help file for your resource.

---

**NOTE**

- Example help files are supplied in the `idm.jar` file located in the `com/waveset/msgcat/help/resources` directory.
- See *Identity Manager Workflows, Forms, and Views* for information about how to include online help with the application.

---

5. (For Active Sync-enabled adapters only) To reset synchronization on the last resource, delete the `XmlData SYNC_resourceName` object.

6. Read the error log and modify the adapter.

7. Reset the logging level.

For example, specifying Level 2 debugging yields information about the adapter settings and any errors, but limits the amount of log detail to a manageable level.

8. Before starting Identity Manager, you must identify the new adapter in the `$WSHOME/config/Waveset.properties` file by placing the adapter name under the `resource.adapters` entry or Identity Manager cannot recognize the adapter.

9. Install your adapter and its associated help file into Identity Manager.

---

**NOTE** Before Identity Manager can recognize an instance of a new adapter in the display, you must create a new resource of that type from the List Resource page.

From this page, select `New > new adapter` and use the Resource Wizard to create the new adapter.

---

10. Use Identity Manager to create a resource and a user on that resource.

---

**TIP** When troubleshooting an Active Sync-enabled adapter, if you edit the `XmlData SYNC_resourceName` object to remove the `MapEntry` for the Active Sync synchronization process from the Debug page, the adapter starts over from the first detected change.

If you used the IAPI event, you must set the `Property()` method to store synchronization state for the resource, such as a last change processed value. Setting this method is very useful for troubleshooting adapters. You can set the adapter to run and ignore past changes. Subsequently, you can modify the adapter and see your change results in the adapter log file.

---

If your resource is an Active Sync resource, you might see additional information if you enable logging on the resource edit page. Set the logging level (0-4) and the file path where the log file will be written (as `resource_name.log`).

11. (For Active Sync-enabled adapters only) Restart synchronization for the last resource.

## To Debug LoginConfig Changes

To debug LoginConfig-related changes to your adapter, you must

1. Enable trace for the selected files and trace the following classes at Method/Class Level 1 trace:
  - `com.waveset.security.authn.WSResourceLoginModule`
  - `com.waveset.session.LocalSession`
  - `com.waveset.session.SessionFactory`
  - `com.waveset.ui.LoginHelper`
  - `com.waveset.ui.web.common.ContinueLoginForm`
  - `com.waveset.ui.web.common.LoginForm`

2. Test Single Sign-On (SSO) pass-through authentication login through Telnet as follows:
  - a. After correctly configuring the SSO login module, telnet directly to the http port and send an http request to login.jsp.
  - b. Paste the following request, which contains an SSO login module that looks for the sm\_user HTTP header, into your telnet session:

```
HEAD /idm/login.jsp HTTP/1.0
Accept: text/plain,text/html,*/*
Accept-Language: en-us
Accept-Encoding: gzip, deflate
User-Agent: Mozilla/4.0
Host: LOCALHOST
sm_user: Configurator
```

A trace displays to indicate that your user has logged in correctly. For example,

#### Code Example 2-2 Sample Output

```
2003.07.08 14:14:16.837 Thread-7 WSResourceLoginModule#checkForAuthenticatedResourceInfo()
Found authenticated resource accountId, 'Configurator@Netegrity SiteMinder' on Identity
Manager user 'Configurator'. null null 2003.07.08 14:14:16.837 Thread-7
WSResourceLoginModule#checkForAuthenticatedResourceInfo()
Exit null null 2003.07.08 14:14:16.837 Thread-7 WSResourceLoginModule#login()
Exit, return code = true null null 2003.07.08 14:14:16.847
Thread-7 LocalSession#login() Login succeeded via Netegrity SiteMinder null null
2003.07.08 14:14:16.847 Thread-7 LocalSession#login() Overall authentication
succeeded null null 2003.07.08 14:14:16.897 Thread-7 LocalSession#checkIfUserDisabled()
Entry null null 2003.07.08 14:14:16.897 Thread-7 LocalSession#checkIfUserDisabled() Exit
null null 2003.07.08 14:14:16.927 Thread-7 LocalSession#login() Exit null null
```

## To Debug Adapter Connection Problems

This section describes methods for debugging some common adapter connection problems. The information is organized as follows:

- [Adapter Authentication Problems](#)
- [Active Sync Adapter Problems](#)
- [Domino Gateway Adapter Problems](#)

- [Mainframe Host Adapter Problems](#)
- [PeopleSoft Adapter Problems](#)
- [SAP Adapter Problems](#)
- [Unix Adapter Problems](#)

---

**NOTE** Generally, you can identify adapter connection issues by tracing the adapter class `com.waveset.adapter.adapter_classname`. For example:

```
com.waveset.adapter.ADSIResourceAdapter
```

If necessary, review instructions for enabling trace in [“Tracing the Identity Manager Server” on page 75](#).

---

### *Adapter Authentication Problems*

Some common authentication problems include

- Missing authentication properties

You must include a property name for the specified DataSource type in the set of names output in the trace.

- An Identity Manager user is missing a matching resource account

If authentication succeeds for the resource adapter, but an exception occurs indicating that no Identity Manager user could be found with a matching resource account ID, be sure that the resource `accountId` associated with the user is the same as the `accountId` returned by your resource adapter’s `authenticate` method.

To verify the Identity Manager user’s resource `accountId`, review Identity Manager’s Edit Trace Configuration page (`debug/Show_Trace.jsp`). If a mismatch exists, change the content of the name being returned by your `authenticate` method or change your resource’s ID template. The template must generate a resource `accountId` that matches the `accountId` being returned by the `authenticate` method.

### *Active Sync Adapter Problems*

The most common problems with custom Active Sync adapters are form-related. These errors generally occur because you have not provided necessary information, such as password or email information, in a required field.

Identity Manager prints form validation errors to the adapter log after the final XML of the view. For example,

```
20030414 17:23:57.469: result from submit (blank means no errors):  
20030414 17:23:57.509: Validation error: missing required field password
```

Identity Manager also prints all messages to the adapter log. These messages include account creation and update times, adapter errors, and a summary of the schema map data.

Active Sync resource adapters store information about the last change processed in the `SYNC.resourceName XMLData` object.

### *Domino Gateway Adapter Problems*

Following are some common Domino gateway and adapter configuration errors and instructions for fixing these problems:

- If an error message states the 'New Domino Gateway' resource is not accessible and the connection is refused, try stopping and restarting the Identity Manager Gateway.
- If an error message states no ID file name is specified and the path to the userID file is set incorrectly, specify a target location for the userID file and edit the resource adapter to set this attribute to a correct path. Typically, the target location for the userID file is the directory into which you installed the Gateway.
- If an error message states you are not authorized to use the server, you have not set the correct access permissions for the ID file. Specify the correct permissions for this file and retry.

### *Mainframe Host Adapter Problems*

When RACF, ACF2, or TopSecret host adapters fail to reuse or cache connections, users are forced to log in frequently, which negatively impacts performance. Generally, the cache timeout setting causes this problem.

To check the cache timeout setting, trace Identity Manager's adapter connection pool as follows:

1. From Identity Manager's Edit Configuration Object page, trace the `com.waveset.adapter.HostConnPool#reapConnections` method at level 4.  
If necessary, review instructions for enabling trace in ["Tracing the Identity Manager Server" on page 75](#).
2. Capture trace for a sufficiently long period of time (*at least 30-60 minutes*), while the adapter performs operations.
3. Review the trace output in the application server `stdout` or trace file and look for `Info reaping connection` entries.

If this entry occurs more than once every 30 minutes, you have a good indication that connections are being timed out unnecessarily.

To resolve this problem, increase the `Idle Timeout` resource attribute value to prevent connections from being reaped too frequently. The `Idle Timeout` attribute controls how long a connection remains idle before the connection is logged out. The default value is 90 seconds, which causes new log ins to occur frequently.

Ideally, specify a value that is greater than the average idle time for your deployment environment. For example, adjust the `Idle Timeout` attribute to 30 minutes (1800000 milliseconds) or more.

### PeopleSoft Adapter Problems

This section describes methods for troubleshooting the following PeopleSoft adapter problems:

- Executing a “Test Connection” from the PeopleSoft resource adapter causes a failure with a generic exception.
  - a. Open the `Waveset.properties` file and set `exception.trace=true`.
  - b. Retry the test connection, and if you see the following results:

```
FormState: expansion complete
java.lang.NullPointerException: PSProperties not loaded from file
WavesetException:
WavesetException:
==> com.waveset.util.WavesetException:
FormState: derivation
```

- c. Log in to the PeopleSoft web interface to verify that you are using the correct UID and password.
- The PeopleSoft application server logs are not showing log-in attempts.

This problem generally occurs because you did not use the `psjoa.jar` file supplied with the PeopleSoft installation to which you are connecting. (See the *Identity Manager Resources Reference* for more information about the PeopleSoft resource adapter).

### SAP Adapter Problems

This section contains information about debugging some common problems with SAP, SAP Enterprise Portal, or SAP HR Active Sync adapters.

- If an error results when you try to test the connection from an SAP or SAP HR Active Sync adapter to the SAP system, open a command window and run this command from your installation directory `WEB-INF/lib`:

```
java -jar sapjco.jar
```

The `sapjco.jar` command shows which version of the SAP Java Connector (JCO) is installed and whether the adapter is installed correctly. The command also returns the JNI platform-dependent and the RFC libraries that communicate with the SAP system.

If these platform-dependent libraries are not found, consult the SAP documentation to find out how to correctly install the SAP Java Connector.

### *Unix Adapter Problems*

This section contains information about debugging some common problems with Unix adapters.

- If you see timeout errors when provisioning to a UNIX resource adapter, you can determine where the provisioning process is failing by tracing the `com.waveset.adapter.ScriptedConnection` method at Method/Class level 4, which gives you the maximum logging output.
- When becoming `root` to perform administrative commands, if the resource adapter executes a `su root` command instead of a `su - root` command, the environment does not inherit any of the custom environment variables defined for `root`; including any custom prompts (environment variable of `PS1`).

When configuring a Unix adapter, you can determine which prompt to enter into the Root Shell Prompt field as follows:

- a. Telnet or ssh to the system as the user you specified in the Login User field.
- b. After typing the password and logging in, type **su root** without a dash and press return.
- c. Type the root password.
- d. The next prompt displayed is the prompt you must enter into the Root Shell Prompt field.

## Troubleshooting Auditor

You can trace the following methods to troubleshoot issues with Identity Auditor:

- `com.sun.idm.auditor.policy` to trace issues with Audit Scans.
- `com.sun.idm.auditor.accessreview` to trace issues with Access Reviews.
- `com.sun.idm.auditor.report` to trace issues with Audit Reports.
- `com.sun.idm.auditor.view` to trace issues with Auditor Views.

In addition, you can set the following hidden flags by modifying Forms or TaskDefinitions:

- **Audit Policy Scan Task - maxThreads (int):** Number of concurrent threads used. (Defaults to 5)
- **Audit Policy Scan Task - userLock (int):** Time (in milliseconds) to wait to acquire a lock on the User object. (Defaults to 5000)
- **Audit Policy Scan Task - scanDelay (int):** Time (in milliseconds) to delay between launching a new scan thread. (Defaults to 0, no delay)
- **Audit Policy Scan Task - clearuserLocks (boolean):** If true, the scanner attempts to clear the user lock before scanning.

In addition, the Show Timings page (`/debug/callTimer.jsp`) provides the following information:

- How long it takes to fetch the initial User view, with no resources, during the scan
- How long it takes to refresh the User view, including resources, during the scan
- How long it takes to evaluate the policy on the User view
- How long each user scan takes, including User view fetch, policy evaluation, and so forth
- How long it takes to fetch a list of users for access review
- How long it takes to evaluate the attestation rule in access review

## Troubleshooting Form Problems

This section describes some common form problems and how to fix these problems.

- You added a `<script>` tag to your user form, but see the following exception when trying to access the user form:

```
java.lang.NoClassDefFoundError: org/mozilla/javascript/NativeScript
```

You must put the `WEB-INF/lib/javascript.jar` file into the application server's classpath. For example,

- If you installed Tomcat as a service (Windows only), edit the following registry key:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tomcat\
Parameters\JVM Option Number 0
```

Append the path to your jar file to the end of the `JVM Option Number 0` key string value, wherever it is located. For example,

```
;C:\tomcat\lib\javascript.jar
```

- If you just started Tomcat from the command line, move the jar file into the `tomcat/lib` directory.
- You customized a Tabbed User Form, but when you try accessing the form through User Creation, you see the following exception:

```
com.waveset.util.WavesetException: Maximum form stack depth exceeded
```

You are limited to pushing 1000 hard-coded elements on the stack and you exceeded this limit.

If you added a check to detect `Field` containers that include a `FieldRef` or `Field` that matches the name of the `Field` container, you might have inadvertently created a circular reference on the form. To fix this problem, change the `FieldRef` or `Field` name.

- You used a `MultiSelect` field in a form for LDAP users, assigned groups to a user in this field, edited the user again, and then see the same group on both sides of the `MultiSelect`. For example,

- Left side value: `cn=Group1,dc=test,dc=Com`

- Right side value: `cn=Group1,dc=test,dc=com`

In this example, the LDAP `baseContext` resource field is set to `dc=test,dc=com` and the LDAP groups are listed as `dc=test,dc=Com`, which causes a problem because LDAP is not case sensitive, but the `MultiSelect` widget is case-sensitive.

To alleviate this problem, change the `baseContext` on the LDAP resource to match the case in your LDAP resource `dc=test,dc=Com` or use the `<upcase>` XPRESS function to uppercase both the left and right side displays.

## Troubleshooting the Gateway

When troubleshooting the Sun Identity Manager Gateway, it is often useful to run the Gateway from the command line. Using command line options allows you to input a wider range of start-up options, which includes starting the Gateway as a normal application instead of a service and running the Gateway on a different port.

---

**NOTE** You must kill the Identity Manager Gateway as a service before running it from the command line. For example, type

```
gateway.exe -k
```

---

The following table describes the Gateway command line arguments.

**Table 2-6** Gateway Command Arguments

Argument	Description
-i	Install this program as an NT service, with specified startup
-r	Remove this program from the Service Manager
-s	Start the service
-k	Kill the service
-t	Set start-up for an existing service
-d	Debug, and run as a regular application
-p	Specify a TCP/IP port number (Default is 9278)
-f	Specify the path to the trace file
-l	Specify the level of tracing (Default is 0, no information)
-m	Specify the maximum trace file size in kilobytes
-v	Display the version

Usage: gateway -i n -r -s -k -t n -d -p n -f name -l n -m n -v

---

**NOTE**

- The -d and -s options are mutually exclusive.
- See [“From the Command Line” on page 96](#) for more information about the Gateway tracing levels.

---

You can also use the Identity Manager Gateway Debug page (`debug/Gateway.jsp`) to troubleshoot the Gateway. See [“How to Configure Tracing” on page 95](#) for more information.

## Troubleshooting Java Code Problems

If you have the basic Java programming skills required to work with Identity Manager, you should be able to diagnose and resolve most Java code problems.

However, a fairly common problem occurs where someone opens a connection to the database but does not close the connection properly. If you do not close the connection properly, performance issues result.

## Troubleshooting Low Memory Conditions

This section describes tools that you can use to investigate low memory conditions, including:

- [From the Identity Manager Debug Pages](#)
- [From JConsole](#)

### From the Identity Manager Debug Pages

---

**NOTE** You must have the *Debug* capability to access and execute operations from the Identity Manager Debug pages. Administrators and the Configurator are assigned this capability by default.

If you do not have the Debug capability, an error message results.

---

You can open the following Identity Manager Debug pages from the Administrator interface to monitor how much memory is being used by your system:

- **Host Connection Pool** page (`debug/Show_ConnectionPools.jsp`): View a summary of connection pool statistics (if you are not using a data source), including the pool version, how many connections were created, how many are active, how many connections are in the pool, how many requests were serviced from the pool, and how many connections were destroyed.

You can also use the Host Connection Pool page to view a summary of the connection pools used to manage connections to the Gateway. You can use this information to investigate low-memory conditions.

- **List Cache Cleared** (`debug/Clear_XMLParser_Cache.jsp`): Clear the cache of recently used XML parsers.
- **Private collection pool** (`debug/Show_JDBC.jsp`): View a summary of the cache of JDBC connections being used by the repository and some resource adapters.

- **System Memory Summary page** (`debug/Show_Memory.jsp`): View the used and total memory in the system. You must click the Garbage Collect button to get the most current used memory value.
- **System Memory Summary page** (`debug/Show_Memory2.jsp`): View an updated `Show_Memory.jsp` page that allow you to clear all unused memory in the JVM so you can investigate heap usage.
- **User Session Pool Cleared** (`Clear_User_Cache.jsp`): Clear the cached sessions for recently logged in users.
- **XML Resource Adapter Caches Flushed and Cleared** (`Clear_XMLResourceAdapter_Cache.jsp`): Clear the cache of the test XML resource adapter.

## From JConsole

Use the Java Monitoring and Management Console (*JConsole*) to detect low memory and deadlocks. JConsole is a Java Management Extension (*JMX*) technology-compliant graphical management tool that is co-packaged with JDK 5 (and later).

JConsole accesses the memory system, memory pools, and MBeans garbage collector to provide information about memory use such as memory consumption, memory pools, and garbage collection statistics. In addition, You can use JConsole to monitor MBeans for information about current heap memory use and non-heap memory use.

---

**NOTE** For information about using JConsole to monitor applications that run on the Java platform, see *Using JConsole to Monitor Applications*. This publication is available from the following web site:

<http://java.sun.com/developer/technicalArticles/J2SE/jconsole.html>

---

## Troubleshooting PasswordSync Problems

When you are trying to troubleshoot problems with PasswordSync, review the following logs for information:

- **PasswordSync Error Logs:** PasswordSync writes all failures to the Windows Event Viewer. (For more information about Event Viewer, see Windows' Help.) The source name for error log entries is *PasswordSync*.
- **PasswordSync Trace Logs:** PasswordSync writes all trace logs to the file location specified when you configured tracing. See [“Using the PasswordSync Configuration Tool” on page 81](#).

Some common PasswordSync problems and solutions include

- PasswordSync is not propagating password changes from the Windows server to Identity Manager.

PasswordSync relies on the registry settings when creating a connection from Active Directory to the Identity Manager Server. During startup, PasswordSync reads the registry and processes the settings, but PasswordSync does not perform any checks to see if it can create a connection.

The following example shows a registry entry for a PasswordSync server. The example includes the default registry setting values, but does not show all of the settings used by PasswordSync.

```
[HKEY_LOCAL_MACHINE\SOFTWARE\Waveset\Lighthouse>PasswordSync]
"clientTimeout"=dword:00000384
"lhuser"="config"
"endpointURL"="http://10.10.10.10:8080/idmboot/servlet/rpcrouter2"
"tracefile"="?"
"tracelevel"=dword:00000000
"proxyPort"="8080"
"securityFlags"=dword:00000000
"requestTimeout"=dword:00007530
"proxyServer"="?"
"lhcred"="rsVtQZpa5Ys="
"soapClientTimeout"=dword:00002710
"installdir"="C:\Program Files\Sun\Identity Manager>PasswordSync\"
"allowInvalidCerts"=dword:00000000
"icptFile"="C:\WINNT\system32\lhfiqw.tmp"
"connectionFlags"=dword:00000000
"tracemaxKB"=dword:00002710
```

If you have not enabled tracing at an appropriate level, PasswordSync does not log connection failures in much detail. To see more detailed trace information, edit the PasswordSync registry settings as described on [page 82](#). Specify `tracelevel 4` to output the maximum trace information, and change the `tracefile` value to point to a writable file. For example:

```
"tracelevel"=dword:00000004
"tracefile"="C:\\Program
Files\\Sun\\IdentityManager\\PasswordSync\\pwicsvc.log"
```

After setting these values, you must restart the PasswordSync service. The start-up routine logs the values read from the registry in the trace file. For each intercepted password change, PasswordSync logs the actions taken to push the password to Identity Manager, which includes creating the SOAP connection.

- If a connection fails during creation, you might encounter the following situations:

---

**NOTE** Each of these situations has its own error code and set of log entries. Identity Manager removes the date, time stamp, and process number from these entries to keep them short.

---

- An incorrect or unreachable URL error that occurs when the server cannot be reached, is not running, or does not reply with a correct SOAP response. Check that PasswordSync can access the server and page.
  - Be sure the server is running and that you have configured your firewalls and routers correctly.
  - Check the application server to be sure it is running, and that PasswordSync can connect to the endpointURL without the application path. If PasswordSync can does not return a page or an error, the application server is not running.
  - Check the SOAP response by opening the endpointURL in a standard browser. If you do not see the `ERROR: org.opensaml.server.SOAPRouter: GET is unsupported message`, see if the servlet is compiling and available.
- An incorrect user name error generally occurs when the `userID` stored in the `lhuser` entry is incorrect. Use the `Configure.exe` utility to replace the user or replace the `lhuser` registry key value with a valid `userID`.

- An incorrect password error generally occurs when the password stored in the `lhcred` entry is not correct when used in combination with the `userID` stored in `lhuser`. Use the `Configure.exe` utility to replace the password, but do not manually edit the `lhcred` registry key.
- A garbage in the password entry error generally occurs when the registry key is corrupted and or when the registry key is manually edited, which causes garbage in the password entry.
- This situation causes the process to hang in `RAEncryptor::Decrypt3DES` and `PasswordSync` cannot decrypt the entry. Use the `Configure.exe` utility to replace the password.

## Troubleshooting Reconciliation Problems

When you are trying to troubleshoot problems with a reconciliation task, review the Reconciliation Status Debug page (`debug/Show_Reconciler.jsp`) to see what the resource threads are working on.

Some common reconciliation problems include

- Reconciliation does not start
  - Open the System Threads Debug page (`debug/Show_Threads.jsp`) to see if the `Reconcilerserver_name` task is running. If the task is not running,
    - a. Open the System Settings page and click Trace Scheduler.  
Running the Scheduler restarts the `Reconcilerserver_name` task.
    - b. Check the System Threads page again, and if the `Reconcilerserver_name` task has not restarted, the Scheduler might be hung.
    - c. Restart your application server.
- Reconciliation fails for a certain user object
  - If the object exists in Identity Manager, try editing the object directly. If the Identity Manager account does not exist, load the one account from the resource.
  - Verify the reconciliation user has the proper access rights.
  - Try extracting the object to a file using the same code path as reconciliation.
  - Open the System Memory Summary Debug page (`debug/Show_Memory.jsp` or `debug/Show_Memory2.jsp`) and verify you have enough free memory.

# Troubleshooting Repository Connection Problems

Identity Manager's `lh` commands are very useful when you are troubleshooting connection problems. These commands use Identity Manager's web application installation, but remove the application server from the equation.

This section describes the following

- [Using `lh` Commands to Debug Problems](#)
- [Testing DataSource Connections](#)

## Using `lh` Commands to Debug Problems

This section describes how to use the `lh` commands; starting with using the more basic commands and progressing to using commands that exercise most of Identity Manager.

- [Using `lh setRepo -c -n`](#)
- [Using `lh setRepo -c -v`](#)
- [Using `setRepo`](#)
- [Using `lh console`](#)

After becoming familiar with these debugging tools, you can develop your own variations for using these `lh` commands.

*Using* `lh setRepo -c -n`

Use the `lh setRepo -c -n` command to perform the most basic connection test, which allows you to examine the current repository location *without connecting*. You can use this command to verify that parameters, such as URL and JDBC driver, are correct.

- If the connection is successful, you can read the `ServerRepository.xml` bootstrap file.
- If the connection fails, try to solve this failure first. A decryption error is the most common cause of this failure. For example, you might have a J2EE mismatch or classpath conflict.

*Using* `lh setRepo -c -v`

Use the `lh setRepo -c -v` command to *connect to* and examine the current repository location. (The `-v` provides verbose output.) You can use this command to exercise almost all of the Repository code without requiring the Identity Manager server.

- If the connection is successful, then you are successfully connected to the current repository location.
- If the connection fails, try to solve this problem first. Solving your connection problem can be very helpful in resolving DNS, firewall, or remote connection privilege problems.

---

**NOTE** For more information, see [“Testing DataSource Connections” on page 127.](#)

---

*Using* `setRepo`

Use the `setRepo` command throughout the debugging process, to specify a new repository location or to set the repository to the same location.

You can use this command to confirm that all of the necessary components, such as the JAR files, are in place. The `setRepo` command also lets you vary connection information, such as `userid` and `password`, to debug table ownership or privilege problems.

*Using* `lh console`

Use this command to actually start an Identity Manager Server using the JAR files in the `WEB-INF/lib` and the classes in `WEB-INF/classes` under `WSHOME`. The `lh console` command uses your Identity Manager installation environment and actually starts the Identity Manager application, but *removes* the application server from the equation.

- If the connection is successful, the problem is specific to the application server environment or configuration.
- If the connection fails, review the failure messages.
  - If the connection failure is the same as the application server failure, you have reproduced this failure with significantly fewer variables.
  - If the failure appears to be different from the application server failure, try fixing the Identity Manager connection problem first because there are fewer variables and more of these variables are under Identity Manager control.

## Testing DataSource Connections

If you are testing a DataSource connection, the `lh setRepo -c` command might fail.

This failure is especially likely if you configured Identity Manager to use the application server's database connectivity service or the application server's directory service. These services often work only in the environment that a running application server provides to a web application.

Initially, approach the DataSource configuration you want in a step-by-step manner. Once you are comfortable with these steps, you can adapt your approach to suit your needs.

1. Try using a direct JDBC DriverManager connection, such as a non-DataSource connection, that bypasses the application server's database connectivity service.
2. Use a DataSource, but store the DataSource object in a directory service *other than* application server's directory service.

---

**NOTE** If you have no other directory service available, you can download a free directory service, including the reference implementation of JNDI that uses only the local file system.

---

If these steps work, you have localized the problem to the application server.

Then, if useful, you can add the application server's database connectivity service or the application server's directory service, whichever service works outside of the environment that the application server provides to web applications.

## Troubleshooting Server-Related Problems

You can analyze your application server logs for fatal errors and other server-related problems.

To troubleshoot server problems, use the application server's Administrative Console to increase the logging level for each module. For more information, see the product documentation supplied with your server.

Most application servers have a standard location for standard out files (`stdout`) and standard error files (`stderr`) for the JVM running the application server. To analyze your application server logs, locate the logs directory or the log files specified for your Identity Manager application server.

- Open the `stdout` file to view minor messages and other tracings.
- Open the `stderr` file to view fatal and critical exceptions.

---

**NOTE** You will see Identity Manager start up and shut down the messages in this trace output.

---

This section explains how to prevent or fix common server-related problems:

- Beginning with Identity Manager version 7.1, a sealing violation exception occurs in the application server log when you use Identity Manager with Oracle 10g on Sun™ Application Server Enterprise Edition 8.2.

This exception generally occurs if you are using more than one Java Archive file (JAR file) containing Oracle JDBC drivers.

To prevent this problem, be sure the `CLASSPATH` contains only one JDBC driver JAR file, and that you use the `ojdbc14_g.jar` provided with Oracle 10g. In addition, you must use the `ojdbc14_g.jar` provided by the Oracle 10g installation to ensure correct operation.

## Troubleshooting an SPML Configuration

To test an SPML configuration:

1. Open the Connect page and click Test.

A dialog indicating that the connection was successful pops up.

2. Open the Schema page and click Submit.

The system displays a tree view of the schemas supported by the Identity Manager server.

If you cannot establish a successful connection

- Double-check the URL you entered.
- If the error you receive contains phrases such as “no response” or “connection refused,” then the problem is most likely the host or port used in the connection URL.
- If the error suggests that a connection was made, but the web application or servlet could not be located, the problem is most likely in the `WEB-INF/web.xml` file. See [“Deployment Descriptor” on page 43](#) for more information.

## Troubleshooting Upgrades

If you encounter problems during the upgrade, check the upgrade log files located in the `$WSHOME/patches/logs` directory. The file names for the logs are based on a timestamp and the stage of the upgrade.



# Errors and Exceptions

This chapter describes the error and exception messages generated by Identity Manager. The information is organized as follows:

- [Before You Begin](#)
- [About Identity Manager Errors and Exceptions](#)
- [Viewing Errors in the System Log Report](#)
- [Customizing a Default Error Message](#)

## Before You Begin

Review the following sections before you start working with Identity Manager error and exception messages:

- [Intended Audience](#)
- [Important Notes](#)
- [Related Documentation and Web Sites](#)

## Intended Audience

This chapter is intended for system administrators and deployers who need additional information about the errors and exceptions generated by Identity Manager.

You must have a background in programming and experience working with XML and Java.

## Important Notes

Be sure to read the following information before working with Identity Manager error and exception messages:

- Examples in this chapter use a locale (*xx\_XX locale*) that was devised for example purposes only.
- When you interpret error messages, be aware of the following:
  - Some messages have different keys but display the same error message text.
  - Some messages are used by multiple components.
  - Exceptions are generally listed by exception type, component, or both.
  - Some exceptions are caused by internal programming errors that cannot be viewed by Identity Manager users.
  - Some exceptions are simple wrappers and their parameters are the entire exception. For example, Identity Manager exceptions wrap resource messages, adapter code, and multi-part exceptions, such as password policy violations.
- If you are using parameterized messages with single or double quotes in the message, you must use an additional single or double quote to escape the message. (You will see only one quote symbol when the message is output to the system.)

For example, the following message begins and ends with two *single* quote marks:

```
''0}''
```

- If you need help diagnosing problems, contact Sun Technical Support:  
<http://www.sun.com/service/online/us>

## Related Documentation and Web Sites

In addition to the information provided in this chapter, consult the publications and web sites listed in this section for information related to tuning Identity Manager for your deployment.

## Recommended Reading

See the following publications for information related to Identity Manager error and exception messages:

- For information about creating a customized message catalog, see *Identity Manager Technical Deployment Overview* for instructions.
- For information about creating or editing System Log reports, see *Identity Manager Administration*.

## Useful Web Sites

The following table describes some web sites you might find useful:

**Table 3-1** Useful Web Sites

Web Site URL	Description
<a href="http://sharespace.sun.com/gm/document-1.26.2296">http://sharespace.sun.com/gm/document-1.26.2296</a>	Identity Manager FAQ on Sun's Share Space <b>Note:</b> You must sign up for a Share Space ID to access this FAQ.
<a href="http://sunsolve.sun.com/">http://sunsolve.sun.com/</a>	Sun web site containing diagnostic tools, forums, features and articles, security information, and patch contents <b>Note:</b> The information on this site is partitioned into three areas, <ul style="list-style-type: none"> <li>• Internal: Sun employees only</li> <li>• Contract: Available only to customers with contract access</li> <li>• Public: Available to everyone</li> </ul>
<a href="http://www.sun.com/service/online/us">http://www.sun.com/service/online/us</a>	Sun Technical Support web site

# About Identity Manager Errors and Exceptions

This section describes the Identity Manager error and exception message system and the components that can generate errors and exceptions.

The information is organized into the following sections:

- [Where Messages Are Stored](#)
- [How Messages Are Displayed](#)
- [Error Severity Levels](#)

## Where Messages Are Stored

Error messages are stored as follows:

- Identity Manager messages are stored in a `WPMessages.properties` file in the `idmcommon.jar` file and in an `RAMessages.properties` in the `idmadapter.jar`.
  - The `WPMessages.properties` file is located in `project_directory/waveset/idm/common/src/com/waveset/msgcat`
  - The `RAMessages.properties` file is located in `project_directory/waveset/idm/adapter/src/com/waveset/adapter`
- Identity Manager Service Provider uses a standard Java message resource bundle for displaying strings in the user interface. This resource bundle is normally included in the `idmspe.jar` file and must be extracted if you plan to customize message strings.

```
$ cd $WSHOME/WEB-INF/classes
$ jar xvf ../lib/idmspe.jar com/sun/idm/idmx/msgcat/IDMXMessages.properties
```

- Identity Auditor messages are stored in the `AUMessages.properties` file, which is located in `project_directory/waveset/idm/auditor/src/com/sun/idm/auditor/msgcat/AUMessages.properties`
- The Data Exporter's default implementation includes its own message bundle, called `WICMessages.properties`. The `WICMessages.properties` file is located in the `exporter.jar` file in

```
com/sun/idm/warehouse/msgcat/WICMessages.properties
```

---

**NOTE** You can create a customized message catalog to add messages or to modify messages provided with the system.

See “Customizing Message Catalogs” in *Identity Manager Technical Deployment Overview* for instructions.

---

## How Messages Are Displayed

For easy identification, Identity Manager displays page-level error and exception messages along the top of the page as boxed text with a unique error icon.

**Figure 3-1** Example Login Authentication Error




---

**NOTE** In the Identity Manager messaging system, items displayed as {0} or {1} represent parameters that are supplied by code. For example, The file is larger than the maximum supported length of {0} bytes.

In this exception, the code replaces the {0} with the parameter value representing the maximum number of bytes.

---

## Error Severity Levels

Within Identity Manager, error severities are defined as follows:

- **Fatal:** A severe error that causes your system to crash, resulting in the loss or corruption of unsaved data.
- **Error:** A severe error that might cause the loss or corruption of unsaved data. Immediate action must be taken to prevent losing data.
- **Warning:** Action must be taken at some stage to prevent a severe error from occurring in the future.
- **Info:** An informative message, usually describing server activity. No action is necessary.

---

**NOTE** You can check the Identity Manager System Log for more information about an error or exception message. (See [“Viewing Errors in the System Log Report”](#) on page 136.)

---

# Viewing Errors in the System Log Report

System Log reports can provide information about errors generated by Identity Manager. The System Log report consists of the error's timestamp, severity, server name, component name, error code or ID, stack trace (structure of the execution stack at that point in the program's life), and error text information.

You can use Identity Manager's Administrator interface or command-line interface to run and view System Log reports.

---

**NOTE** Instructions for creating and editing System Log reports are provided in *Identity Manager Administration*.

---

## To Run a System Log Report From the Administrator Interface

Perform the following steps to run a System Log report from the Administrator interface:

1. Log in to the Identity Manager Administrator interface.
2. Select Reports > Run Reports to open the Run Reports page.
3. Locate the appropriate System Log Report entry in the Report Type column, and then click the Run button in that same row.

The Report Results page displays, listing the system messages that were reported during the specified interval. For example, [Figure 3-2 on page 137](#) depicts information about two system messages.

## Figure 3-2 Example Report Results Page Report Results

### System Messages from the previous week

Monday, February 25, 2008 3:00:20 PM CST

Lists all system messages reported during the past 7 days

Number of records reported: 2

▼ TimeStamp	Event	Severity	Server	Component	Error Code	Message
<a href="#">Thursday, February 21, 2008 3:09:06 PM CST</a>	N/A	Fatal	idmvm042	Server	SECURITY	Security Violation: Incoming HttpServletRequest considered invalid by CSRFGUARD from address: 129.147.62.21
<a href="#">Tuesday, February 19, 2008 7:21:15 PM CST</a>	SV-0220-012115	Error	idmvm042	Server	ERROR	An error occurred starting the connection: io exception: The Network Adapter could not establish the connection ==> java.sql.SQLException: io exception: The Network Adapter could not establish the connection

#### 4. The Report Results table shows the following information:

- **Timestamp:** Shows the day, date, and time the error occurred.

Click the Timestamp links to view detailed information about that System Log record. For example, if you clicked the first Timestamp link shown in [Figure 3-2](#), the following information displays.

### Figure 3-3 Partial System Log Record Detail

#### System Log Record Details

Timestamp	Tuesday, February 19, 2008 7:21:15 PM CST
Event	SV-0220-012115
Server	idmvm042
Severity	Error
Component	Server
Error Code	ERROR
Message	<p>An error occurred starting the connection: io exception: The Network Adapter could not establish the connection ==&gt; java.sql.SQLException: io exception: The Network Adapter could not establish the connection</p> <p>com.waveset.util.WavesetException: An error occurred starting the connection: io exception: The Network Adapter could not establish the connection java.sql.SQLException: io exception: The Network Adapter could not establish the connection ==&gt; java.sql.SQLException: io exception: The Network Adapter could not establish the connection</p> <p>at com.waveset.adapter.OracleERPResourceAdapter.makeConnection(OracleERPResourceAdapter.java:797)</p> <p>at com.waveset.adapter.OracleERPResourceAdapter.startConnection(OracleERPResourceAdapter.java:484)</p> <p>at com.waveset.adapter.OracleERPResourceAdapter.getUser(OracleERPResourceAdapter.java:2736)</p> <p>at com.waveset.adapter.ResourceAdapterProxy.getUser(ResourceAdapterProxy.java:902)</p> <p>at com.waveset.provision.FetchContext.doFetchContext(fetchContext.java:368)</p> <p>at com.waveset.provision.FetchContext.processOpFetchContext(fetchContext.java:230)</p> <p>at com.waveset.provision.ThreadContext.processContext(ThreadContext.java:348)</p> <p>at com.waveset.provision.ThreadContext.launchThreads(ThreadContext.java:258)</p> <p>at com.waveset.provision.Provisioner.fetchAccountsList(Provisioner.java:2345)</p> <p>at com.waveset.provision.Provisioner.fetchAccounts(Provisioner.java:2790)</p> <p>at com.waveset.view.UserViewer.assembleView(UserViewer.java:858)</p> <p>at com.waveset.view.UserViewer.checkoutView(UserViewer.java:754)</p> <p>at com.waveset.object.ViewMaster.checkoutView(ViewMaster.java:630)</p> <p>at com.waveset.session.LocalSession.checkoutView(LocalSession.java:681)</p> <p>at com.waveset.util.GenericViewSource.getView(GenericViewSource.java:447)</p> <p>at org.apache.jsp.account.modify_jsp._jspService(modify_jsp.java:428)</p> <p>at org.apache.jasper.runtime.HttpJspBase.service(HttpJspBase.java:94)</p> <p>at javax.servlet.http.HttpServlet.service(HttpServlet.java:302)</p> <p>at org.apache.jasper.servlet.JspServletWrapper.service(JspServletWrapper.java:324)</p> <p>at org.apache.jasper.servlet.JspServlet.service(JspServlet.java:292)</p> <p>at org.apache.jasper.servlet.JspServlet.service(JspServlet.java:238)</p> <p>at javax.servlet.http.HttpServlet.service(HttpServlet.java:302)</p> <p>at org.apache.catalina.core.ApplicationFilterChain.internalDoFilter(ApplicationFilterChain.java:237)</p> <p>at org.apache.catalina.core.ApplicationFilterChain.doFilter(ApplicationFilterChain.java:157)</p> <p>at com.sun.idm.profiler.instrumentation.RequestTimingFilter.doFilter(RequestTimingFilter.java:81)</p> <p>at org.apache.catalina.core.ApplicationFilterChain.internalDoFilter(ApplicationFilterChain.java:188)</p> <p>at org.apache.catalina.core.ApplicationFilterChain.doFilter(ApplicationFilterChain.java:157)</p> <p>at org.apache.catalina.core.StandardWrapperValve.invoke(StandardWrapperValve.java:214)</p> <p>at org.apache.catalina.core.StandardValveContext.invokeNext(StandardValveContext.java:104)</p> <p>at org.apache.catalina.core.StandardPipeline.invoke(StandardPipeline.java:520)</p>

- **Event:** Identifies the syslog ID of the target entry (when applicable).
- **Severity:** Shows the severity level of the error:
  - **Fatal:** A severe error that causes the system to crash, resulting in the loss or corruption of unsaved data.
  - **Error:** A severe error that might cause the loss or corruption of unsaved data. Immediate action must be taken to prevent losing data.
  - **Warning:** Action must be taken at some stage to prevent a severe error from occurring in the future.
  - **Info:** An informative message, usually describing server activity. No action is necessary.
- **Server:** Identifies the server on which the error occurred.
- **Component:** Identifies the system component that generated the error.
- **Error Code:** Shows the error code associated with that error.
- **Message:** Shows the actual error message text.

## To Run a System Log Report From the Command-Line Interface

---

**NOTE** These instructions assume you are familiar with the Identity Manager command-line interface and `lh` commands. For more information, read “Appendix A: lh Reference” in *Identity Manager Administration*.

---

Perform the following steps to run and view a System Log report from the command line:

1. Open a command window.
2. Change directories to the default Identity Manager installation directory.
3. At the prompt, type the `lh syslog [options]` command.

Use these *options* to include or exclude information:

- `-d Number`: Show records for the previous number of days (default=1)
- `-F`: Show only records with fatal severity level

- `-E` : Show only records with error severity level or higher
- `-i logid`: Show only records with a specified Syslog ID  
Syslog IDs are displayed on some error messages and reference a specific System Log entry.
- `-W` : Show only records with warning severity level or higher (default)
- `-X`: Include reported cause of error, if available

## Customizing a Default Error Message

---

**NOTE** To add message catalog entries or to modify entries provided with the system, you must create a customized message catalog. See *Identity Manager Technical Deployment Overview* for instructions.

---

Identity Manager's default error messages are stored in a `WPMessages.properties` file in the `idmcommon.jar` file and in an `RAMessages.properties` file in the `idmadapter.jar`.

- The `WPMessages.properties` file is located in  
`project_directory/waveset/idm/common/src/com/waveset/msgcat`
- The `RAMessages.properties` file is located in  
`project_directory/waveset/idm/adapter/src/com/waveset/adapter`

The Identity Manager Service Provider's default error messages are located in the `IDMXMessages.properties` file.

You can customize these default error messages by modifying attributes in the `ErrorUIConfig` object provided with that message.

To modify the `ErrorUIConfig` object:

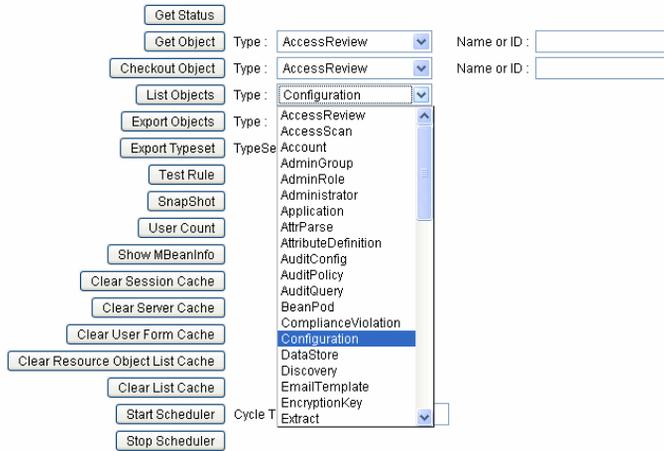
1. Log in to Identity Manager Administrator interface.
2. Open the System Settings page by typing `http://host:port/idm/debug` in to your browser.

3. Locate the Type menu, located next to the List Objects button. Choose Configuration from the menu.

**Figure 3-4** List Objects Type Menu

### System Settings

Click a button to effect a system change.



4. Click the List Objects button.
5. On the List Objects of type: Configuration page, click the `ErrorUIConfig` edit link.

The following example shows the XML for an `ErrorUIConfig` object on Checkout Object: Configuration page:

**Code Example 3-1** Example `ErrorUIConfig` XML Source

```
<?xml version='1.0' encoding='UTF-8'?>
<!DOCTYPE Configuration PUBLIC 'waveset.dtd' 'waveset.dtd'>
<!-- MemberObjectGroups="#ID#Top" extensionClass="GenericObject"
id="#ID#9787BA467E01441B:178655:1156195C008:-7FFE"
lastModifier="com.waveset.object.ErrorUIConfig" name="ErrorUIConfig"-->
<Configuration id='#ID#9787BA467E01441B:178655:1156195C008:-7FFE' name='ErrorUIConfig'
lock='Configurator#1200600790328' creator='com.waveset.object.ErrorUIConfig'
createDate='1191343145328' lastModifier='com.waveset.object.ErrorUIConfig'
lastModDate='1191343145328'>
  <Extension>
    <Object>
      <Attribute name='Enabled'>
        <Boolean>true</Boolean>
      </Attribute>
      <Attribute name='ErrorMsgID' value='UI_ERROR_DEFAULT_FATAL_MESSAGE' />
    </Object>
  </Extension>
  <MemberObjectGroups>
    <ObjectRef type='ObjectGroup' id='#ID#Top' name='Top' />
  </MemberObjectGroups>
</Configuration>
```

6. You can modify the following `ErrorUIConfig` object attributes:
- **Enabled:** Controls whether the message is enabled (`true`) or disabled (`false`).

---

**NOTE** The ability to disable this attribute is provided for backward-compatibility; however, disabling this attribute will result in cryptic messages and you will not see the typical extended messages in the System Log.

---

- **ErrorMsgID:** Identifies the error message to be displayed.

Change the `ErrorMsgID` attribute value to provide the message text you want displayed.

The current setting for this attribute is as follows, and it references the `UI_ERROR_DEFAULT_FATAL_MESSAGE` message in the message catalog:

```
<Attribute name='ErrorMsgID' value='UI_ERROR_DEFAULT_FATAL_MESSAGE' />
```

---

**NOTE** If you are using parameterized messages with single or double quotes in the message, you must use an additional single or double quote to escape the message. (You will see only one quote symbol when the message is output to the system.)

For example, the following message begins and ends with two *single* quote marks:

```
''{0}''
```

---

7. When you are finished, click Save to save your changes.

---

**NOTE** If you have difficulty creating a default error message, contact your Administrator or Administrator hotline.

---

# Index

## SYMBOLS

`$WSHOME/patches` directory 129

## A

account data 15

accountIds, verifying 112

Active Sync adapters

logs 105

performance tuning 27

adapters

debugging 108

troubleshooting authentication 112

troubleshooting authentication problems 112

tuning 27

Admin forms, optimizing 30

allowInvalidCerts registry key 82

application servers

analyzing logs 128

garbage collection 7

increasing instances 26

memory 39

PeopleSoft 115

redirecting workflow trace messages to  
console 92

testing DataSource connections 127

troubleshooting 128

troubleshooting repository problems 125

tuning 8, 9, 39

tuning IBM WebSphere® 11

tuning Sun Java System Application Server 7, 8, 9

viewing trace output files 75, 114

attribute tables 12

attributes

account 28

binary 26

connectionPoolDisable 19

modifying ErrorUIConfig object 139

optimizing memory 34

querying for predefined object 12

using inline 20

XML 20

Audit Log Maintenance Task 26

authentication

optimizing pool size 48

properties, missing 112

testing pass-through 111

troubleshooting adapter problems 112

## B

binary attributes 26

## C

callTimer command 55

callTimer.jsp debug page 54, 103, 117

change tables 13

## Section D

- characters sets and encoding 18
- Clear\_List\_Cache.jsp debug page 39
- Clear\_User\_Cache.jsp debug page 58
- Clear\_XMLParser\_Cache.jsp debug page 56
- Clear\_XMLResourceAdapter\_Cache.jsp debug page 59
- Compliance Violation data 15
- concurrent
  - mode failures 9
  - operations, limiting 41
  - processes 42
  - tasks 47
  - user loads 8
  - users 38
- configuration data 16
- connectionPoolDisable attribute 19
- Control Timings debug page 54, 103
- customizing
  - default error messages 139
  - Service Provider message strings 134

## D

- data classes 14
- data mining 7
- data Translation Look-aside Buffers (DTLBs) 10
- database repository, tuning 11
- database statistics, tuning 36
- DataSource connections, testing 127
- debug pages
  - accessing 54
  - Clear\_List\_Cache.jsp 39
  - Clear\_User\_Cache.jsp 58
  - Clear\_XMLParser\_Cache.jsp 56
  - Clear\_XMLResourceAdapter\_Cache.jsp 59
  - Control Timings 54, 103
  - debugging performance issues 53
  - Show Timings 117
  - Show\_CacheSummary.jsp 58
  - Show\_ConnectionPools.jsp 56
  - Show\_Memory.jsp 58
  - Show\_Provisioning.jsp 57
  - Show\_Sizes.jsp 57

- Show\_Timings.jsp 56
- Show\_Trace.jsp 55
- Show\_WSProp.jsp 39, 59
- SysInfo.jsp 58
- debugging
  - adapters 108
  - errors displayed through browser 107
  - LoginConfig changes 110
  - PasswordSync 122
  - performance issues 53
- default error messages, customizing 139
- defvarmethod 32
- deployment-specific settings for tuning 5
- directories
  - \$WSHOME/patches 129
  - logs 129
  - patches 129
- documentation, related
  - errors/exceptions 132
  - performance tuning 2
  - troubleshooting/tracing 67
- domain.xml file 7
- Dr. Watson logs 100
- DTrace facility
  - commands 61
  - description/purpose 60, 106
  - enabling probes 61
  - writing scripts 60

## E

- editing
  - Identity Manager schema for Users/Roles 42
  - registry keys 82
  - RepositoryConfiguration object 19
- enabling
  - methods 89
  - setTrace method 89
  - SOAP tracing 90
  - trace 110
- encoding and character sets 18
- End User forms, optimizing 31
- entitlement data 16

- environment, Java EE 6
- errors
  - debugging 107
  - important notes 132
  - severity levels 135
  - viewing 136
- ErrorUIConfig object, modifying 139
- example locale 132
- exception logging
  - description 78
  - disabling 79
  - enabling 78
- exceptions
  - important notes 132
  - tracing 78
  - wrappers 132
- experience requirements
  - for performance tuning 2
  - for troubleshooting 66
  - for working with errors/exceptions 131
- export queue data 17

## F

- File Cache Configuration page 8
- forms, optimizing
  - Admin forms 30
  - End User forms 31
  - form field expressions 31
  - new forms 30

## G

- garbage collection 7
- garbage collector
  - detecting low memory/deadlocks 62
  - tuning JVM performance 7, 10
- gateway tracing 95
- general XML, optimizing 38
- getResourceObjects 45

- global XPRESS tracing 80
- GUIDs 18

## H

- heap size 39
- HTTP listeners 8
- HTTP requests 111

## I

- IBM WebSphere® application server, tuning 11
- Identity Manager
  - and Identity Manager IDE 105
  - errors and exceptions 131
  - forms performance, optimizing 29
  - gateway tracing 95
  - general performance, optimizing 26
  - resource query performance, optimizing 45
  - session performance, optimizing 48
  - taskbar performance, optimizing 51
  - workflows, optimizing 32
- Identity Manager Integrated Development Environment. *See Identity Manager IDE*
- Identity Manager schema, editing 42
- Identity Manager Service Provider
  - errors and exceptions 134
- idmadapter.jar 139
- idmcommon.jar file 134, 139
- idmspe.jar file 134
- IDMXMessages.properties file 139
- important notes
  - for errors/exceptions 132
  - for performance tuning 2
  - for troubleshooting/tracing 66
- inline attributes 20
- installdir 83
- instances, application server 26

## J

- Java
  - test programs 108
  - tuning 7
- Java EE environment 6
- JVM, tuning 7

## L

- Large Objects (LOBs) 13
- limiting concurrent operations 41
- LOB data types 21
- locally managed tablespaces (LMTs) 21
- log data 17
- logging SPML traffic 89, 90
- logs
  - analyzing application server 128
- logs directory 129

## M

- ManualActions, tuning 34
- memory requirements 39
- message catalogs
  - adding/editing entries 139
  - creating customized 133, 134
- messages
  - parameterized 132, 142
  - where stored 134
- methodology, tuning 5
- methods
  - caching query results 45
  - enabling setTrace 89
  - improving form performance 29
  - improving Gateway performance 49
  - reconciliation 86
  - tuning rules 32
- modifying `ErrorUIConfig` objects 139

## O

- object IDs 18
- object tables 13
- operations, limiting concurrent 41
- optimizing
  - Admin forms 30
  - authentication pool size 48
  - End User forms 31
  - form field expressions 31
  - memory 34
  - new forms 30
- Oracle database repository 21
- organization data 16
- output, trace 89, 90

## P

- parameterized messages 132, 142
- pass-through authentication, testing 111
- PasswordSync
  - debugging 122
  - registry keys 82
- patches directory 129
- PeopleSoft application servers, troubleshooting 115
- per-account workflows 43
- performance
  - improving JVM 7
  - provisioner 41
- performance tuning
  - experience requirements 2
  - important notes 2
  - Java EE environment 6
  - optimizing Identity Manager 25
  - optimizing the database repository 11
  - recommended reading 3
  - related documentation 2
  - XML 38
- performance, provisioner 40
- predefined object attributes 12
- prepared statements 18
- PrintGCDetails script 7
- PrintGCStats script 3, 7

- PrintGCTimeStamps script 7
- probes, enabling DTrace 61
- processes, concurrent 42
- properties
  - authentication 112
  - missing authentication 112
- provisioner performance 40, 41
- putmap method 32

## R

- RAMessages.properties file 134, 139
- recommended reading 3
- reconciliation
  - methods 86
  - performance 42
  - tracing 85
  - troubleshooting 124
  - tuning 42
- registry keys, PasswordSync 82
- related documentation
  - errors/exceptions 132
  - performance tuning 2
  - troubleshooting/tracing 67
- related web sites 4, 67, 133
- repository databases
  - tuning 11, 19
  - tuning Oracle 21
  - tuning SQL Server 23
- repository tables
  - attribute 12
  - change 13
  - object 13
- RepositoryConfiguration object 19
- requests
  - HTTP 111
  - SPML RPC 89, 90
- requirements, experience
  - for performance tuning 2
  - for troubleshooting 66
  - for working with errors/exceptions 131
- requirements, memory 39
- resource queries, tuning 45

- resource query, Identity Manager 45
- role data 15
- rules, tuning 32
- running System Log reports 136

## S

- scheduler, tracing 91
- scripts
  - DTrace 60
  - PrintGCDetails 7
  - PrintGCStats 3, 7
  - PrintGCTimeStamps 7
- server.xml file 8
- servers
  - connection settings 54, 70, 74, 102
  - tracing 74, 75
  - troubleshooting 128
  - tuning 7, 9
- sessions, Identity Manager 48
- setlist method 32
- setTrace method, enabling 89
- setvar method 32
- severity levels 135
- Show Timings debug page 117
- Show\_CacheSummary.jsp debug page 58
- Show\_ConnectionPools.jsp debug page 56
- Show\_Memory.jsp debug page 58
- Show\_Provisioning.jsp debug page 57
- Show\_Sizes.jsp debug page 57
- Show\_Timings.jsp debug page 56
- Show\_Trace.jsp debug page 55
- Show\_WSProp.jsp debug page 39, 59
- Single Sign-On (SSO) 111
- SOAP, tracing 90
- Solaris
  - patches xv
  - support xv
- special considerations
  - for errors/exceptions 132
  - for performance tuning 2
  - for troubleshooting/tracing 66

## Section T

- SPML
  - tracing messages 88
  - troubleshooting 129
- SPML requests
  - RPC 89, 90
- SQL Server databases, tuning 23
- statistics, running 19
- Sun Java System Application Server 8
- Sun Java System Identity Manager. *See* Identity Manager
- Sun JVM. *See* JVM
- support
  - contacting Technical Support xv
  - Solaris xv
- synchronization logs 27
- SysInfo.jsp debug page 58
- System Log Maintenance Task 26
- System Log reports 136

## T

- table names
  - attribute 12
  - change 13
  - object 13
- tables
  - attribute 12
  - change 13
  - object 13
- task bar, Identity Manager 51
- task bar, tuning 51
- task data 16
- task scheduler tracing 91
- tasks, specifying maximum concurrent 47
- testing adapters 108
- trace configuration 73
- trace file 75
- trace output files, viewing 75, 114
- tracing
  - application server logs 128
  - debugging error display 107
  - Dr. Watson logs 100
  - enabling 90, 110
  - enabling for SPML 89, 90
  - exceptions 78
  - gateway 95
  - Identity Manager server 75
  - important notes 66
  - reconciliation 85
  - related documentation 67
  - SPML messages 88
  - task scheduler 91
  - workflows 91
  - XPRESS 80
- tracing method 69
- troubleshooting
  - application servers 128
  - experience requirements 66
  - important notes 66
  - PeopleSoft application servers 115
  - reconciliation 124
  - related documentation 67
- tuning
  - application servers 8
  - database statistics 36
  - deployment-specific settings 5
  - general performance tunings 26
  - IBM WebSphere® application server 11
  - Java 7
  - ManualActions 34
  - methodology 5
  - Provisioner 40
  - reconciliation 42
  - repository database 19
  - resource queries 45
  - roadmap 5
  - rules 32
  - session performance 48
  - Sun Java System Application Server 7, 8, 9
  - task bar 51
  - workflows 32
  - WorkItems 34

**U**

- upgrade log files [129](#)
- user data [14](#)
- user loads, concurrent [8](#)
- users, concurrent [38](#)

**V**

- viewing
  - errors [136](#)
  - trace output files [75](#), [114](#)

**W**

- web sites, useful [4](#), [67](#), [133](#)
- workflow trace messages [92](#)
- workflows
  - per-account [43](#)
  - tracing [91](#)
  - tuning [32](#)
- WorkItems, tuning [34](#)
- WPMessages.properties file [134](#), [139](#)
- wrappers, exception [132](#)
- writing DTrace scripts [60](#)
- WSHOME [107](#), [109](#)

**X**

- XML columns [14](#)
- XML, optimizing [38](#)
- XPRESS tracing [80](#)

