



IBM Tivoli Security Compliance Manager (ITSCM)

Component Optimization

AN ENDEAVOR TO PROVIDE ADDITIONAL INFORMATION FOR ITSCM
TOWARDS ENHANCED APPLICATION DEPLOYMENT AND IMPROVED CUSTOMER
SATISFACTION

Version 1.2

Copyright Notice

Copyright IBM Corporation 2001. All rights reserved. May only be used pursuant to a Tivoli Systems Software License Agreement, an IBM Software License Agreement, or Addendum for Tivoli Products to IBM Customer or License Agreement. No part of this publication may be reproduced, transmitted, transcribed, stored in a retrieval system, or translated into any computer language, in any form or by any means, electronic, mechanical, magnetic, optical, chemical, manual, or otherwise, without prior written permission of IBM Corporation. IBM Corporation grants you limited permission to make hardcopy or other reproductions of any machine-readable documentation for your own use, provided that each such reproduction shall carry the IBM Corporation copyright notice. No other rights under copyright are granted without prior written permission of IBM Corporation. The document is not intended for production and is furnished "as is" without warranty of any kind. All warranties on this document are hereby disclaimed, including the warranties of merchantability and fitness for a particular purpose.

U.S. Government Users Restricted Rights —Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corporation.

Trademarks

IBM, the IBM logo, Tivoli, the Tivoli logo, AIX, Cross-Site, NetView, OS/2, Planet Tivoli, RS/6000, Tivoli Certified, Tivoli Enterprise, Tivoli Enterprise Console, Tivoli Ready, and TME are trademarks or registered trademarks of International Business Machines Corporation or Tivoli Systems Inc. in the United States, other countries, or both.

Lotus is a registered trademark of Lotus Development Corporation. Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both. UNIX is a registered trademark of The Open Group in the United States and other countries. C-bus is a trademark of Corollary, Inc. in the United States, other countries, or both. PC Direct is a trademark of Ziff Communications Company in the United States, other countries, or both and is used by IBM Corporation under license. ActionMedia, LANDesk, MMX, Pentium, and ProShare are trademarks of Intel Corporation in the United States, other countries, or both. For a complete list of Intel trademarks, see <http://www.intel.com/sites/corporate/trademark.htm>. SET and the SET Logo are trademarks owned by SET Secure Electronic Transaction LLC. For further information, see <http://www.setco.org/aboutmark.html>. Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries.

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

Notices

References in this publication to Tivoli Systems or IBM products, programs, or services do not imply that they will be available in all countries in which Tivoli Systems or IBM operates. Any reference to these products, programs, or services is not intended to imply that only Tivoli Systems or IBM products, programs, or services can be used. Subject to valid intellectual property or other legally protectable right of Tivoli Systems or IBM, any functionally equivalent product, program, or service can be used instead of the referenced product, program, or service. The evaluation and verification of operation in conjunction with other products, except those expressly designated by Tivoli Systems or IBM, are the responsibility of the user. Tivoli Systems or IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to the IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, New York 10504-1785,

U.S.A.

Table of Contents

INTRODUCTION.....	5
WHAT IS ITSCM?	6
PRODUCT HIGHLIGHTS	6
THE PROBLEM	10
THE SOLUTION.....	10
TUNING GUIDELINES	10
THE LAW OF DIMINISHING RETURNS	10
TUNING JUST FOR THE SAKE OF TUNING.....	10
CONSIDER THE ENTIRE SYSTEM.....	10
ONE PARAMETER AT A TIME	10
TUNE BY LEVELS	11
HARDWARE AS WELL AS SOFTWARE.....	11
UNDERSTAND THE PROBLEM.....	11
PUT FALL-BACK PROCEDURES IN PLACE.....	11
THE TOP TUNING FACTORS.....	11
DB2 INSTALLATION	12
DB2 INSTALLATION SUGGESTIONS.....	22
DB2 TUNING SUGGESTIONS	23
<i>Buffer Pools.....</i>	<i>23</i>
<i>Increasing Buffer pool size.....</i>	<i>23</i>
<i>Placing logs on a separate device.....</i>	<i>23</i>
<i>Increase log file size.....</i>	<i>24</i>
<i>Clear snapshots.....</i>	<i>24</i>
<i>Add indexes as necessary.....</i>	<i>24</i>
<i>Increase maximum number of active applications.....</i>	<i>25</i>
<i>Increase maximum number of open database files per application</i>	<i>25</i>
<i>Enable intra-partition parallelism</i>	<i>25</i>
<i>Increase storage for lock-list.....</i>	<i>25</i>
<i>Increase maximum number of locks.....</i>	<i>26</i>
<i>Additional Tuning.....</i>	<i>26</i>
HINTS AND TIPS	30
<i>General Notes</i>	<i>30</i>
<i>DB2 Installation.....</i>	<i>30</i>
<i>SCM Server Installation.....</i>	<i>30</i>
<i>Installing the SCM Client</i>	<i>32</i>
<i>Installing the SCM Administration Utilities.....</i>	<i>33</i>
<i>Post Installation.....</i>	<i>33</i>
<i>Installing an Email server</i>	<i>33</i>
<i>Configuration.....</i>	<i>33</i>
<i>Implementing a Proxy Agent</i>	<i>37</i>
MISCELLANEOUS	38
<i>Starting and stopping the SCM client.....</i>	<i>38</i>
<i>Using the Build SQL function.....</i>	<i>38</i>
<i>Troubleshooting.....</i>	<i>38</i>
CLIENT.....	38
SERVER.....	41
COLLECTORS.....	43
Admin GUI.....	44
Known Errors	46

<i>FAQs (Frequently Asked Questions)</i>	47
APPENDIX	52
ADDITIONAL TOOLS	52
<i>db2look</i>	52

Table of Figures

FIGURE 1, IBM TIVOLI SECURITY COMPLIANCE MANAGER (ITSCM) HIGHLIGHTS.....	7
FIGURE 2, IBM TIVOLI SECURITY COMPLIANCE MANAGER (ITSCM) HARDWARE REQUIREMENTS.....	8
FIGURE 3, IBM TIVOLI SECURITY COMPLIANCE MANAGER (ITSCM) SOFTWARE REQUIREMENTS.....	9
FIGURE 4, CLIENT.PREF.....	34
FIGURE 5, SERVER.INI	36

Introduction

You are reading a document about IBM Tivoli Security Compliance Manager (ITSCM), a software product that allows the creation of security policies and their regular monitoring for compliance and vulnerabilities. Essentially, ITSCM performs two functions:

1. Collects data from multiple computer systems and stores that data in a central repository.
2. Provides reporting capabilities to show adherence to, or violation of defined security policies.

Operating as a data collection service, IBM Tivoli Security Compliance Manager will gather and store a wide variety of information from multiple computer systems. This information can include any data located on the system, such as operating system version information, software patch levels, registry information, file and directory information, application information, and other security-related data. Functioning as a monitoring and reporting service, Tivoli Security Compliance Manager will extract the data collected, analyze the data for conformity, and produce reports that can reveal adherence to internal and industry-standard security policies.¹

This data collection function, into a DB2 UDB Relational Database Management System (RDBMS), often results in a condition where system resources are exhausted, or unavailable, creating a situation that requires immediate intervention.

As a monitoring and reporting service, IBM Tivoli Security Compliance Manager can extract the data collected, analyze the data for vulnerabilities, and produce reports that can reveal adherence to internal and industry-standard security policies.

Examples of such security and privacy requirements in the United States include:

- Gramm-Leach-Bliley Act (GLBA), also known as the Financial Services Modernization Act
- Sarbanes-Oxley Act (SOA)
- Health Insurance Portability and Accountability Act (HIPAA)

The purpose of this document is to address the tuning requirements of DB2 UDB in the context of IBM Tivoli Security Compliance Manager (ITSCM) and the proper handling of additional components, thereby attempting to alleviate the conditions mentioned above.

To this end, a correct implementation, and/or successful tuning, of the DB2 UDB RDBMS used for data collection of ITSCM information promises a more stable

¹ Levoy Wheeler, IBM Tivoli Security Compliance Manager Version 5.1: Deployment and Tuning Guide, 30 April, 2004.

deployment, greater uptime, and a clearer path to an evident Return-On-Investment (ROI).

What is ITSCM?

IBM Tivoli Security Compliance Manager is a new security policy compliance product that acts as an early warning system by identifying security vulnerabilities and security policy violations for small, medium and large businesses. Tivoli Security Compliance Manager helps organizations define consistent security policies and monitor compliance of these defined security policies. Security policies can be based on both internal security requirements and industry-standard security policies.

Product Highlights

- Automates scans of servers and desktop systems which can help reduce the cost and time associated with manual security checks
- Provides reports to security officers and compliance auditors with detailed information about the security health of the business so they can take the appropriate steps to make individual systems and departments compliant
- Identifies software security vulnerabilities prior to costly damage being inflicted by security incidents
- Improves business operations and helps to increase efficiencies through automation and centralization
- Assists in addressing compliance issues in regulations and standards by automating compliance tasks, monitoring correspondence, reducing human error, and taming compliance costs
- Integrates with Tivoli's automated security management tools to help mediate security policy violations and risks

Tivoli Security Compliance Manager is designed for customers who are concerned with compliance and regulations as well as those concerned with identifying and fixing security weaknesses and vulnerabilities in their IT infrastructure. In assisting with compliance issues, the software comes with pre-defined recommended security policies and can be customized to fit specific regulatory, industry or corporate security policies. The software helps provide a more efficient way to gather and manage information about the health of systems, and enterprises can use that data to help them comply with regulations like the Health Insurance Portability and Accountability Act (HIPAA).

Tivoli Security Compliance Manager audits systems and applications for vulnerabilities and identifies violations against security policies. The software uses a red, yellow, green, stoplight-style warning system to alert security administrators and users, whether systems comply with security policy. For example, the software could send reports or an email to administrators showing which systems have passwords set properly, which have anti-viruses signature files are up-to-date, and which have dangerous or unnecessary services running. The software can help identify security vulnerabilities across AIX, Solaris, HP-UX, Windows, Linux and Linux on zSeries.

ITSCM Component Optimization

Figure 1, IBM Tivoli Security Compliance Manager (ITSCM) Highlights

Features	Advantages	Benefits
Centralized administration	Centralizes the policy definition and scanning of systems	Consistent security audits across the organization reducing human error
Automation	Automates security scans of servers and desktop systems which can help reduce the cost and time associated with manual security checks	Helps reduce costs by providing a more effective and efficient way to perform routine tasks, which allows administrators to focus on more strategic functions
Security vulnerability scans	Helps to secure business systems	Helps to identify software security vulnerabilities prior to costly damage being inflicted by security incidents
Compliance	Security policy compliance can assist in collecting and managing data about security efforts and systems that companies use as part of their overall regulatory compliance initiatives	Addresses compliance issues in regulations and standards by automating compliance tasks, monitoring correspondence, reducing human error, and taming compliance costs
Auditing and reporting mechanisms	Enables administrators to produce reports with detailed information about the security health of the business so they can take the appropriate steps to make individual systems and departments compliant	Quickly produce reports for audits and ensuring regulatory compliance

ITSCM Component Optimization

Figure 2, IBM Tivoli Security Compliance Manager (ITSCM) Hardware Requirements

Processor and memory requirements for IBM Tivoli Security Compliance Manager Server			
Type of ITSCM deployment		CPU RQMTS	RAM RQMTS
Small	(1-500 clients)	1	512MB RAM
Medium	(501-2500 clients)	2	512MB RAM
Large	(2501-10000 clients)	2-4	2-4GB RAM
Server package requires 5MB to install			
Disk and memory requirements for client and collectors			
<i>Disk and memory requirements for ITSM Client</i>			
Client platform	Disk Requirements for installation directory	Disk Requirements for temporary directory	RAM RQMTS
AIX	64MB	45MB	75MB RAM
HP-UX	64MB	6MB	75MB RAM
Linux	64MB	46MB	75MB RAM
Solaris	64MB	65MB	75MB RAM
Windows	64MB	44MB	75MB RAM
<i>Note: The HP-UX platform values in the table are much smaller due to the JRE being packaged in the OS</i>			
Disk and memory requirements for Administration Utilities			
<i>Disk and memory requirements for ITSM Administration Console</i>			
Administration Console Platform	Disk Requirements for installation directory	Disk Requirements for temporary directory	RAM RQMTS
Windows	64MB	42MB	128MB RAM Min 256MB RAM Rec

ITSCM Component Optimization

Figure 3, IBM Tivoli Security Compliance Manager (ITSCM) Software Requirements

Tivoli Security Compliance Manager requires:

IBM DB2 UDB version 7.2 or 8.1

Supported Operating Systems:

The following information are the supported Operating Systems for ITSCM Server, Client, and Administration Console:

IBM Tivoli Security Compliance Manager (ITSCM) Server

Operating System	Level	Patch/Maintenance Level
AIX	5.1, 5.2	None required
Windows 2000	Server	Latest Fix Pack
Solaris	2.8, 2.9	Latest Fix Pack
SuSe Linux Enterprise Server	8	Latest Fix Pack

IBM Tivoli Security Compliance Manager (ITSCM) Client

Operating System	Level	Patch/Maintenance Level
AIX	5.1, 5.2	Latest cumulative patches
HP-UX	11.0, 11i	Latest cumulative patches
Red Hat Linux	6.2, 7.0, 7.1, 7.2, 8.0, 9.0	Latest cumulative patches
Solaris	2.6, 2.7, 2.8, 2.9	Latest cumulative patches
Windows NT	4.0 Server, 4.0 Workstation	Latest service pack and security roll up package
Windows 2000	Server, Advanced Server, Professional	Latest service pack and security roll up package
Windows 2003	Server Standard Edition, Enterprise Edition	Latest cumulative patches
Red Hat Enterprise Linux (EL)	2.1	Latest cumulative patches
Red Hat EL Advanced Server	3.0	Latest cumulative patches
Red Hat EL for zSeries	3.0	Latest cumulative patches
Red Hat EL Advanced Server	2.1	Latest cumulative patches
SuSe Linux	7.0	Latest cumulative patches
SuSe Linux Enterprise Server (ES)	8	Latest cumulative patches
SuSe Linux ES for zSeries	8	Latest cumulative patches
SuSe Linux ES for iSeries/pSeries	8	Latest cumulative patches

IBM Tivoli Security Compliance Manager (ITSCM) Administration Console

Operating System	Level	Patch/Maintenance Level
Windows 2000	Professional	Latest service pack and security roll up package
Windows XP	Professional	Latest service pack and security roll up package

The Problem

The problem, as defined by Tivoli Support and Field Service personnel, is the default settings of a typical ITSCM implementation which, in all but small deployments, are not sufficient enough to deal with an ever-increasing data stream. This fact coupled with limited knowledge, adequate training, or skilled personnel needed to tune a Relational Database Management System (RDBMS) often results in frantic calls to remote support to resolve the situation.

The Solution

The solution is to provide adequate information that will allow anyone to appropriately tune an RDBMS utilized in conjunction with ITSCM resulting in a smooth running application and a faster Return-On-Investment.

Tuning Guidelines

The law of diminishing Returns

The greatest performance improvements usually result from the initial efforts at tuning.

Tuning just for the sake of tuning

Only adjust identified limitations. If non-associated resources are tuned, these will have little or no effect on response time until the major constraints are relieved. Major resource barriers are the ones that significantly impact performance the most.

Consider the entire system

Never tune in isolation. This means that tuning should not be done without considering the entire system.

One parameter at a time

Don't change more than one parameter at a time. Even if you're sure that all the changes will be beneficial, you will have no way of evaluating how much each change contributed. Also, the trade-offs cannot be judged effectively. Every time a parameter is adjusted to improve one area, it almost always affects at least one other area that may not be considered. By changing only one at a time, this allows you to have a benchmark to evaluate whether the change does what you want.

Tune by levels

For the same reasons that you should only change one parameter at a time, tune one level of your system at a time. You can use the following list of levels within a system as a guide:

- Hardware
- Operating System
- Application Server
- Database Manager
- SQL Statements
- Application Programs

Hardware as well as software

Some performance problems may be corrected by applying service either to your hardware, or to your software, or to both. Do not spend excessive time monitoring and tuning your system when simply applying service may make it unnecessary.

Understand the problem

Even if it seems that additional storage or processor power could immediately improve performance, take the time to understand where your bottlenecks are. You may spend money on additional disk storage only to find that you do not have the processing power or the channels to exploit it.

Put fall-back procedures in place

As noted earlier, some tuning can cause unexpected performance results. If this leads to poorer performance, it should be reversed and alternative tuning tried. If the former setup is saved in such a manner that it can be simply recalled, the backing out of the incorrect information becomes much simpler.

The top tuning factors

1. Make sure there is enough disk space
2. Buffer pools should use 50-75% of available memory for high-end production servers
3. Runstats should be performed on all tables
4. Use the Configuration Advisor to configure the database for the application environment
5. Logging should be to separate high-speed disks. This is done via the NEWLOGPATH configuration parameter.
6. Commit often to increase concurrency.
7. To avoid sort overflows, increase SORTHEAP.
8. Table space should be SMS for the system catalog temporary or DMS for the rest.
9. Use parameter markers for repeated SQL statements.


DB2 Installation

DB2 must be installed prior to installing Tivoli Security Compliance Manager Server. Consider making a separate file system for the DB2 database that is to be used by the server. This might be necessary if the server encounters DB2 errors indicating that there is not enough space to store Tivoli Security Compliance Manager Data. Creating a separate file system does not remove the possibility of filling up a file system but it does reduce the impact to other applications that are using DB2 and using the file system that DB2 uses by default. If the file system has an error, it will be easier to isolate the problem to Tivoli Security Compliance Manager. – A DB2 8.1 database may either be on the server machine, or on a remote machine. In order to use a DB2 8.1 database on a remote machine, you will need to place a copy of the db2java.zip and the db2jcc.jar files onto your IBM Tivoli Security Compliance Manager Server machine. These files must be located in the same directory on the IBM Tivoli Security Compliance Manager server. You will need to provide the fully-qualified directory path to the db2java.zip file during install.

A DB2 7.2 database may either be on the server machine, or on a remote machine. In order to use a DB2 7.2 database on a remote machine, you will need to place a copy of the db2java.zip file onto your IBM Tivoli Security Compliance Manager Server machine. You will need to provide the fully-qualified directory path to the db2java.zip file during install.²


A step-by-step default DB2 installation scenario follows.

After inserting the DB2 UDB 8.2 (8.1 FP07) for Windows CD, click Install Product




² Tivoli Security Compliance Manager Installation Guide, Pg. 7

Verify DB2 UDB Enterprise Server Edition is selected. Click Next




The following Dialog box will appear:




ITSCM Component Optimization


This Dialog box follows:



The system will be check and the following Dialog Box will appear. Click Next to continue




Accept the license agreement and click Next to continue




ITSCM Component Optimization

Select Typical as the installation type. Click Next to continue




Select Install DB2 Enterprise Server Edition on this computer. Click Next to continue



ITSCM Component Optimization


Unless required by ITCS104, accept the following location for installation. Click Next to continue



If compliance to ITSC104 is required, see the information at the following link:


<https://w3-03.ibm.com/transform/sas/as-web.nsf/ContentDocsByTitle/DB2+Universal+Database>

If not, enter a password for the default username and click Next to continue




ITSCM Component Optimization

Select Local to create a contact on this system. Click Next to continue




Click OK to continue




ITSCM Component Optimization

To instantiate the default instance, click Next to continue




Verify Prepare the DB2 tools catalog in a local database. Click Next to continue



ITSCM Component Optimization

Click Next to continue



DB2 Setup wizard - DB2 Enterprise Server Edition

Specify a local database to store the DB2 tools catalog

The DB2 tools catalog will be stored in a local database. Specify the required information for the database. If the database does not exist, it will be created for you.

Distance: DB2

Database:

☒ New: TOOLSD8

☐ Existing:

Schema:


☒ New: SYSTOOLS

☐ Existing:

InstallShield

< Back Next > Cancel Help

Verify Defer the task until after installation is complete is selected. Click Next to continue



DB2 Setup wizard - DB2 Enterprise Server Edition

Specify a contact for health monitor notification

By default, a health monitor runs on the DB2 instance you are setting up. When a health indicator threshold is breached, email or pager notification will be sent to an administration contact. Specify a new or existing contact for health monitor notification. If your contact list resides on a remote DB2 server, you must provide a user name and password for login.

Administration contact for this instance:

☐ New contact

Name:

Email address:

☐ Address is for a pager


☒ Defer the task until after installation is complete

InstallShield


< Back Next > Cancel Help

ITSCM Component Optimization

Click **Install** to continue




The following Dialog box will appear




ITSCM Component Optimization

Click Finish to complete the installation




Click Exit First Steps to leave the installation. DB2 UDB v8.2 is now installed



DB2 Installation Suggestions

The DB2 UDB database server functions as the storage area for ITSCM data. This database may be installed on the same system as the ITSCM Server, although performance would improve if it was installed on a separate, dedicated system with multiple CPUs. The following diagram illustrates a suggested configuration for a high volume setup.

Suggested Configuration



DB2 Tuning Suggestions

1. Increase Buffer Pool size
2. Place logs on a separate device
3. Increase log file size
4. Clear snapshots regularly
5. Add indexes as necessary
6. Increase maximum active applications
7. Increase maximum database files open per application
8. Enable intra-partition parallelism
9. Increase storage for lock list
10. Increase maximum number of locks

Buffer Pools

Properly defining buffer pools is one of the major keys to having a well-performing system. With 32-bit operating systems, it is important to be aware of the limit on shared memory, which restricts a database's buffers pools (that is, database global memory) from exceeding the following limits (64-bit systems do not have this limit):

- AIX - 1.75 GB
- Linux - 1.75 GB
- Sun - 3.35 GB
- HP-UX - approximately 800 MB
- Windows - 2-3 GB (use '/3GB' switch in boot.ini on NT/2000)

A buffer pool is memory storage that temporarily stores information read and changed prior to being written back to disk. The primary purpose of buffer pools is to improve database performance. Essentially, the fewer hard read/write operations made to disk, the better.

Increasing Buffer pool size

Depending on available memory, increase the JAC_BP significantly as follows (system memory $\geq$ 2GB):

```
db2 connect to jac
db2 alter bufferpool jac_bp 16384
db2 select jac_bp from syscat.bufferpools
```

Placing logs on a separate device

If placing log files on a separate device, verify the device exists prior to proceeding. Set the location for log file output:

Win32:

```
db2 update db cfg for jac using newlogpath=\\D:
```

Unix/Linux:

```
db2 update db cfg for jac using newlogpath=/newlogdevice
```

Increase log file size

A database that has a large number of update, delete, or insert transactions running against it which will cause the log file to become full very quickly, typically by the error “transaction log full”. These parameters are:

```
LOGPRIMARY  
LOGSECOND  
LOGFILSIZ
```

Note: SCM 5.1 has these properties set appropriately by default.

Clear snapshots

To delete (clear) one or more snapshots that are no long needed:

- A. Open the Storage Management View
 - a. From the Control Center window, expand the object tree until you find the database, database partition group, or table space you want to examine in the Storage Management view.
 - i. Right-click the desired database and select Manage Storage from the pop-up menu. The Storage Management view opens.
 - b. From the Storage Management view, expand the object tree until you find the database, table space, or database partition group object associated with the snapshot you want to remove.
 - c. Right-click the object and select Remove Historical Snapshots from the pop-up menu. The Snapshot Deletion window opens.
- B. Select one or more snapshots from the list of historical snapshots.
- C. Click OK. The Snapshot Deletion window closes and the specified snapshots are deleted.

Add indexes as necessary

As an example, create the index on the *cli_id* and *instance_id* columns for the JAC_IDATA.CUSTOM_COLLECTOR_V5 table using the *create index* command:

```
# db2 create index CUSTOM_COLLECTOR_INDEX on  
JAC_IDATA.CUSTOM_COLLECTOR_V5 (cli_id,instance_id)
```

To improve snapshot creation performance create indices on the snapshot table by using the following commands:

```
CREATE INDEX JAC_INDX.JAC_VIOL_SNA_ID ON JAC_SYS.VIOLATIONS
```

```
(SNA_ID)
PCTFREE 10 ALLOW REVERSE SCANS;

CREATE INDEX JAC_INDX.JAC_VIOL_COM_ID ON JAC_SYS.VIOLATIONS
(COM_ID)
PCTFREE 10 ALLOW REVERSE SCANS;
```

Increase maximum number of active applications

This parameter specifies the maximum number of concurrent applications that can be connected (both local and remote) to a database. Since each application that attaches to a database causes some private memory to be allocated, allowing a larger number of concurrent applications will potentially use more memory.

Setting maxappls to automatic has the effect of allowing any number of connected applications. DB2 will dynamically allocate the resources it needs to support new applications.:

```
db2 update db cfg for jac using maxappls 100
```

Increase maximum number of open database files per application

This parameter specifies the maximum number of file handles that can be open for each database agent. If opening a file causes this value to be exceeded, some files in use by this agent are closed. If maxfilop is too small, the overhead of opening and closing files so as not to exceed this limit will become excessive and might degrade performance.

```
db2 update db cfg for jac using maxfilop 512
```

Enable intra-partition parallelism

This parameter specifies whether the database manager can use multiple processors. Some of the operations that can take advantage of parallel performance improvements when this parameter is "YES" include database queries and index creation.

```
db2 update dbm cfg using intra_parallel yes
```

Increase storage for lock-list

This parameter indicates the amount of storage that is allocated to the lock list. There is one lock list per database and it contains the locks held by all applications concurrently connected to the database. Locking is the mechanism that the database manager uses to control concurrent access to data in the database by multiple applications. Both rows and tables can be locked. The database manager can also acquire locks for internal use. Increasing this setting can help avoid deadlock conditions.

```
db2 update db cfg for jac using locklist 10000
```

Increase maximum number of locks

Lock escalation is the process of replacing row locks with table locks, reducing the number of locks in the list. This parameter defines a percentage of the lock list held by an application that must be filled before the database manager performs escalation. When the number of locks held by any one application reaches this percentage of the total lock list size, lock escalation will occur for the locks held by that application. Lock escalation also occurs if the lock list runs out of space.

```
db2 update db cfg for jac using maxlocks 50
```

Additional Tuning

The IBM Tivoli Security Compliance Manager is foundationally based on the database and thus efficiency of SQL statements is key not only to the operation of the IBM Tivoli Security Compliance Manager components but also to user defined queries such as those used in reports and policies and accessing tables for in-house collectors. The following describes a method that can be utilized to pinpoint sub optimal SQL execution and optimize the efficiency of the execution. As of the time of this publication the IBM Tivoli Security Compliance Manager provides support for the IBM DB2 database; the following section focuses on this database software. It is recommended that a DBA or the database manager be chosen to perform these recommended procedures.

The IBM DB2 UDB program includes provisions that allow the database administrator to enable monitors within the database program that keep track of the amount of time and work that is required for SQL statements to execute. You can view the current settings with the *get monitor switches* command.

```
# db2 get monitor switches
```

Monitor Recording Switches

Switch list for db partition number	0
Buffer Pool Activity Information	(BUFFERPOOL) = OFF
Lock Information	(LOCK) = OFF
Sorting Information	(SORT) = OFF
SQL Statement Information	(STATEMENT) = OFF
Table Activity Information	(TABLE) = OFF
Take Timestamp Information	(TIMESTAMP) = ON
Unit of Work Information	(UOW) = OFF

By default the monitor switches, save the timestamp monitor, are disabled. To enable the monitor switches use the *update monitor switches* command. Normally, it's safe to just enable all the switches; this is what you do in this example here although you are only interested in the *statements* monitor.

```
# db2 update monitor switches using bufferpool on
```

ITSCM Component Optimization

```
# db2 update monitor switches using lock on
# db2 update monitor switches using sort on
# db2 update monitor switches using statement on
# db2 update monitor switches using table on
```

Followed by:

```
# db2 get monitor switches
```

Monitor Recording Switches

```
Switch list for db partition number    0
Buffer Pool Activity Information        (BUFFERPOOL) = ON
Lock Information                        (LOCK) = ON
Sorting Information                    (SORT) = ON
SQL Statement Information               (STATEMENT) = ON
Table Activity Information              (TABLE) = ON
Take Timestamp Information              (TIMESTAMP) = ON
Unit of Work Information               (UOW) = OFF
```

After the monitor switches have been enabled, operations throughout the IBM Tivoli Security Compliance Manager that cause database activity can be performed. After these operations have been completed, take a snapshot of the monitors using the *get snapshot* command.

```
# db2 get snapshot for all on JAC > snapshot.txt
# cat snapshot.txt
```

```
...
Number of executions                    = 1
Number of compilations                  = 1
Worst preparation time (ms)             = 4
Best preparation time (ms)              = 4
Internal rows deleted                   = 0
Internal rows inserted                   = 0
Rows read                               = 2095777
Internal rows updated                   = 0
Rows written                            = 0
Statement sorts                         = 0
Total execution time (sec.ms)           = 64.718209
Total user cpu time (sec.ms)            = 7.520000
Total system cpu time (sec.ms)          = 0.350000
Statement text                          = delete from
JAC_IDATA.CUSTOM_COLLECTOR_V5 where CLI_ID=1636 and
INSTANCE_ID=37
```

The snapshot will display information about each query executed such as *rows read* and *total execution time*. You can see in this example that this particular delete query takes over a minute to execute. Using the IBM DB2 *explain* tool you can discover how DB2 is executing the query.

The explain tool output tells you that the database is using a *table scan* to locate prospective entries. Table scans are notoriously slow on larger tables (i.e. over a million rows) and are usually replaced automatically with indexes. In some instances the IBM DB2 optimizer will automatically generate indexes for large tables, in some instances it will not, and thus you may choose to manually create the indexes if the query is to be executed frequently.

In this particular instance you can create the index on the *cli_id* and *instance_id* columns for the JAC_IDATA.CUSTOM_COLLECTOR_V5 table using the *create index* command:

```
# db2 create index CUSTOM_COLLECTOR_INDEX on
JAC_IDATA.CUSTOM_COLLECTOR_V5 (cli_id,instance_id)
```

To reset the database monitors use the reset monitor command. Next you should repeat the operations you performed through the IBM Tivoli Security Compliance Manager. Finally, you can take another snapshot and/or use the *explain* tool to verify the predicted optimization.

```
# db2 reset monitor for database jac

# db2 get snapshot for all on JAC > snapshot.txt
# cat snapshot.txt

...
Number of executions                = 1
Number of compilations              = 1
Worst preparation time (ms)         = 4
Best preparation time (ms)          = 4
Internal rows deleted               = 0
Internal rows inserted              = 0
Rows read                          = 0
Internal rows updated               = 0
Rows written                       = 0
Statement sorts                    = 0
Total execution time (sec.ms)       = 0.005257
Total user cpu time (sec.ms)        = 0.010000
Total system cpu time (sec.ms)      = 0.000000
Statement text                      = delete from
JAC_IDATA.CUSTOM_COLLECTOR_V5 where CLI_ID=1636 and
INSTANCE_ID=37
```

You can see in this example that the rows read decreased from millions to 0 due to the use of the index. The corresponding decrease in time of execution was from 64 seconds to a few milliseconds. You can view the SQL execution path with the *explain* tool.

The explain tool output shows that no table scans were used, instead only indexes were referred to when executing the query. Again this method can be utilized to optimize SQL statement execution for tables created for custom collectors and for SQL statements created for reports and policies.

Hints and Tips

General Notes

Before installing any of the SCM components it is important to check the latest information from the README.TXT that is shipped with the code, and the latest product Release notes. This is particularly important as the supported set of platforms for each SCM component, i.e. Server, Client, Admin GUI and Admin CLI are different.

DB2 Installation

On a Windows platform, ensure that you reboot the system after completing the DB2 installation, as you will probably get a jdbc error when trying to access the SCM Admin GUI for the first time.

During the installation of DB2 selecting the default installation options will be sufficient for an SCM installation. You do not need to create a DB2 instance as the installation of the SCM server will do this automatically. It is suggested to use the default user names (and groups in a UNIX installation) for all DB2 services. It would probably be a good idea to use the same password for all these services also.

SCM Server Installation

Before installing the SCM server component you need the following information:

- Name of an email server, and an email address to which server notices will be sent
- Server and client ports, the defaults are 1955 and 1950 respectively, it is recommended to use these defaults, unless they are already in use
- Passwords to be used for the master and server keystores. Note: During installation you will have the opportunity to stashing the server keystore password for automatic startup by checking a selection box. It is recommended that you check this box, as this will instruct the installation program to register the SCM server process as a Windows service, otherwise you will need to run the jacserv command to start and stop the SCM server service.
- The user id and password (with administration access) for communicating with the DB2 database
- For remote databases or Windows installations, you will need the location of the jar or zip containing the JDBC driver and the URL for database connectivity, for Windows this is c:\program files\sqllib\java\db2java.zip
 - i. By default SCM will install to /opt/IBM/SCM on AIX, and should not be installed on the root file system, therefore if /opt is on the root file system you will need to do one of the following:
 - ii. Create a separate file system for /opt
 - iii. Create a separate file system for /opt/IBM or /opt/IBM/SCM
 - iv. Create a symbolic link from /opt to /usr/opt
 - v. Create a symbolic link from /opt/IBM or /opt/IBM/SCM to some other file system

ITSCM Component Optimization

Execute one of the following to start the server installation wizard:

- For windows execute scm_win32.exe from the code source
- For AIX execute ./scm_aix.bin as the root user.
- For Solaris execute ./scm_solarisSparc.bin as the root user

When installing the SCM server on Solaris ensure that the file system used for temporary disk space, typically /var, is large enough for the installation, otherwise the installation program will not be able to run. If the default file system for temporary disk space is less than 45 MB, you should specify another location on a file system with adequate disk space for temporary installation data to be stored, you can do this by specifying the temp directory, i.e.

./scm_solarisSparc.bin

- For Linux execute ./scm_linux.bin as the root user
- When presented with the installation directory, C:\Tivoli\SCM, either accept this default or change. If more than one SCM component is to be installed on the same machine, you must use the *same* base install directory for each one.
- Ensure that the IBM SCM Server is selected to be installed, by default the SCM client is selected
- On the Server E-mail configuration page enter the Email server name and an email address to send server messages to. You cannot leave these fields blank, however, as long as the email address is in the correct format, i.e. name@domain, you can proceed.
- Enter the port numbers for the Server and client, or accept the defaults.
- Enter passwords for the master keystore password and the server keystore password
- Specify whether the database is local or remote
- In the Database configuration enter the database user id and password that SCM will use to connect to DB2
- Enter the location of the JDBC driver, usually c:\program files\sqllib\java\db2java.zip
- If you are using a local database you do not need to make any changes to the field labeled "URL to use for database connectivity". If the database is remote add the URL to this field
- If you are using a remote database you will be prompted on the next page to check that the JAC database has been created and that DB2 has been enabled to use the JDBC interface.
- Enter a user name and password for the SCM administrator
- Before installation begins you will be presented with some information about the configuration, i.e. the location SCM will be installed to and the components to be installed – i.e. the SCM server, and the SCM Database Configuration which is automatically selected for installation

Installing the SCM Client

SCM client – this component can be installed on AIX, Solaris, Windows 2000 Server, Windows Advanced Server, Windows 200 Professional, Windows XP Professional, Windows NT 4.0 Server and Workstation, Windows 2003 (Standard and Enterprise), Redhat Linux, Redhat Enterprise Linux, SUSE Linux, SUSE Enterprise Linux server.

Before installing, you should have the following information:

- The port numbers the client and server will be communicating over, this should be the same port numbers selected during the server installation.
- The communication mode – Push or pull. In push mode the client initiates communication with the server, in push mode the Server initiates communication with the client
- As with the Server installation, if the platform you are installing on is HP-UX (or Solaris?), you will need obtain the Java Runtime environment for that platform as it is not supplied with the SCM code

To begin the installation of the SCM client execute one of the following:

- For windows execute scm_win32.exe from the code source
- For AIX execute ./scm_aix.bin.
- For Solaris execute ./scm_solarisSparc.
- For Linux execute ./scm_linux.bin.
- For HP-UX execute ./scm_hp11x.bin.

Here is the procedure for installing the SCM client:

- i. When presented with the installation directory, C:\Tivoli\SCM, either accept this default or change. If this client is on the same machine as the Server, it must be installed in the same location as the server was.
- ii. Ensure the SCM client has been selected for installation, it should as it is the default
- iii. Enter the port number the client will be listening on, this should be the same port value specified for the client during the server installation
- iv. Select whether the Client will be running in Push or Pull mode. In Push mode communication can be initiated by either the client or server, in Pull mode the server will initiate communication. Why?
- v. If you select the Push mode, you will be prompted to enter the server name (or server IP address) and port number the server will be listening on. The port number should be the same as that specified for the server during the server installation. Also you will have the option of specifying whether the client uses DHCP to obtain its IP address, if this check box is selected you will need to enter a name that the server can use as a client for your DHCP client.

Note: If you select Push mode, then check the box to specify that the client uses DHCP, and then go back a couple of pages to change your option from Push to Pull mode, you will be prompted to enter a DHCP alias name, this is a bug in the beta code build, the workaround to this is to go back to the page where you select Push or Pull mode and select Push mode, click next for the next page and *unselect* the DHCP option, then go back one page and change your selection to Pull mode and continue.

Installing the SCM Administration Utilities

This component provides you with the SCM command line interface and if installed on a Windows platform it will also install the SCM Admin GUI. The admin GUI does not need to be installed on a machine with any other SCM components, the machine it is installed on will need TCP/IP connectivity to the SCM server machine.

Post Installation

After you have installed the server, the client, and the administration console or command line interface or both, be sure to install the Tivoli Security Compliance Manager collectors and policies from the product CD into a directory that the server can read and write to. See the *IBM Tivoli Security Compliance Manager Administration Guide* for more information about using the collectors.

After you have installed Tivoli Security Compliance Manager, it is important to review log space in DB2 to determine if you have adequate space for Tivoli Security Compliance Manager logging requirements. You must have administrator authority for DB2 to issue the following commands, which will enable you to review the logging requirements and make changes, if necessary.

The values in the following commands are the minimum size needed for a reasonably loaded system. If you are using a more heavily loaded system, you might need to increase the log file size and or the number of log files.

- i. To see the current DB2 settings, issue the command:

```
db2 get db cfg for JAC
```

- ii. To set the log file to 1000 4K pages, issue the command:

```
db2 update db config for JAC using LOGFILSIZ 1000
```

- iii. To set DB2 to use 10 circular log files, issue the command:

```
db2 update db config for JAC using LOGPRIMARY 10
```

- iv. After you make these changes, stop and then restart the server.

Installing an Email server

You would only need to install an email server if you are working on a demo/POC environment where you do not have access to a corporate email system. The ArGoSoft SMTP/POP3 mail server is a freeware product that is quick to install and configure, however it runs only on Windows platforms.

Configuration

There are two main configuration files which you can edit in order to modify SCM:

- client.pref (default path /opt/IBM/SCM/client/client.pref or C:\Tivoli\SCMclient\client.pref)

- **server.ini** (default path /opt/IBM/SCM/server/server.ini)

```
[JAC]
# Hostname or IP address of the server you wish to contact
jac_host=9.173.178.14

# Port number on the JAC server
jac_port=1955

# Port number the JAC server uses to contact the client with to execute commands
# Use -1 to allow RMI to choose
export_port=1950

# Type of client to act as
# PULL: Server contacts client first - Use for firewalls
# PUSH: Client contacts server first - More efficient
client_type=pull

# Print debug messages to STDOUT
debug=false

# Base directory of the client
base_directory=C:\Tivoli\SCM\client

# Once the number of messages in the queue > MAX_MESSAGES, the flush the queue
immediately
max_messages=100

# Don't allow more than MAX_QUEUE_SIZE messages in the queue
max_queue_size=200

# Time in milliseconds messages in the client queue are sent to the server
flusher_sleep_time=6000

# Time in milliseconds the client pings the server
pinger_sleep_time=60000
dhcp=false
#alias=dhcp9207.pic.uk.ibm.com
```

Figure 4, client.pref

The main settings you may need to modify are:

- **jac_host** this can either be the SCM server's name, or IP address. If the client is acting only as a proxy relay, the jac_host will be set to localhost, i.e. jac-host=localhost
- **export_port** this is the port that this client will be listening on, the default is 1950, but may be changed, need to be consistent with the definition for the client in the GUI
- **client_type** this is either pull or push. If you want to change the type of the client you will need to edit this value as well as the client setting from the GUI or CLI.
- **flusher_sleep_time** the SCM client stores data to be sent to the SCM server as messages in a queue, this value specifies the interval in milliseconds between sending each message. Typically in test scenarios you may want to decrease this value.
- **pinger_sleep_time** this is the time interval between client "pings" to the SCM server. The default is 60000 ms, i.e 10 minutes, again in a test environment you may want to decrease this value.

ITSCM Component Optimization

- `dhcp=false` this value can be either true or false. The omission of this value or commenting it out with a preceding `#` character is the same as setting the value to false.

Note: If a client is configured as a DHCP client in the `client.prefs` file, the IP address of the client machine can be changed without the client losing its registration to the SCM server. Note this feature is intended for use by client machines picking up an IP address from a DHCP server, however, it is functionally independent of DHCP, i.e. a machine with a static IP address can be registered as a DHCP client without any functional problems.

```
# #####
# General JAC parameters
# #####
# jac.debug - shows debug information
# jac.server.resolve.clients - if false JAC assumes UnicastRemoteObject.getClientHost returns an IP address
#                               Set this to true if it returns a hostname instead
# jac.server.process_threads - Number of processing threads in the JAC server
# jac.server.no_activity_timeout - milliseconds before the server will consider a client inactive
# jac.server.dump_file - File used to store messages queued when the server shutdown
# jac.adminserver.separate_db_connection - if true each admin server object has it's own database connection
# jac.run_cleaners - if true, start the table cleaner thread
# jac.run_scheduled_queries - if true, start the scheduled queries thread
# jac.cleaner.sleep_time - milliseconds between table cleanings

bootstrap.services.filename=./services
jac.debug=false
jac.base_directory=C:\Tivoli\SCM\server\
jac.server.resolve_clients=false
jac.server.max_process_threads=20
jac.server.process_threads=5
jac.server.no_activity_timeout=300000
jac.server.dump_file=./server.dump
jac.cleaner.sleep_time=3600000
jac.data.tablespace=JAC_DTS
jac.index.tablespace=JAC_DTS
jac.data.noindex=false
jac.server.queue.high_wm=10000
jac.server.queue.low_wm=9500
jac.server.queue.scratch_directory=scratch/
jac.adminserver.max_token_age=30000
jac.server.exceptions_directory=exceptions/
client.pull.secret=secret
server.inactive_timeout=900000
db.serverpoolsize=5
db.adminpoolsize=5
jac.security.authenticator=com.ibm.jac.server.JACAuthenticator
#jac.security.authenticator=com.ibm.jac.server.CWAAuthenticator
#authenticator.cwa.ldaphost=
#authenticator.cwa.defaultdomain=

# DHCP Client Auto Registration - Uncomment to enable
#jac.dhcp_auto_register=true

# #####
# Mail parameters used by JAC when sending mail
# #####
# mail.smtp.host - Hostname of the SMTP server to send mail. Can be localhost if you have sendmail running
# mail.mailid - Specifies who to send mail as. This will be the "From" for JAC mailings

mail.smtp.host=secureway6.pic.uk.ibm.com
mail.mailid=user1@secureway6.pic.uk.ibm.com
```

```
# #####
# RMI parameters
# #####
# rmi.port - port the server listens for remote calls on
# rmi.debug - turns on RMI debugging.  LOTS OF OUTPUT!

rmi.port=1955
rmi.debug=false
client.pull.listener.port=1950
export_port=1951

# #####
# Database parameters
# #####
# db.is_remote - if false connect to the database using the CLI.  Otherwise use the network driver
# db.userid - userid to log onto the database
# db.password - database password for db.userid
# db.local.url - JDBC url of your local database
# db.remote.url - JDBC url of your remote database
# db.local.jdbc.driver - fully qualified class of the local JDBC driver. Make sure this is in your classp
# db.remote.jdbc.driver - fully qualified class of the remote JDBC driver. Make sure this is in your clas

db.is_remote=false
db.userid=db2admin
db.password=IyMjcGFzc3dvcmQ=
db.local.url=jdbc:db2:JAC
db.remote.url=jdbc:db2://localhost/JAC_REM
db.local.jdbc.driver=COM.ibm.db2.jdbc.app.DB2Driver
db.remote.jdbc.driver=COM.ibm.db2.jdbc.net.DB2Driver
db.enable.batch=false
db.default.max_age=720

# #####
# Security
# #####

keystore.directory=C:\Tivoli\SCM\server\keystores
provider.keystores=ibm.jks
provider.certs=ibm
provider.keystore.ibm.password=IyMjMTIzNDU2
server.keystore=server.jks
server.keystore.password=IyMjcGFzc3dvcmQ=
server.sslkey=sslcert
server.authkey=authcert
authorizer.max.keys=4
licensecert.password=IyMjbGljZW5zZW51cnQ=
auditcert.password=IyMjYXVkaXRjZXJ0

# SSL session cache size for push clients
server.session.cache.size=5000

# SSL session cache size for pull clients
client.session.cache.size=1000

# #####
# ACO Security Manager settings
# #####

aco.objectList=com.ibm.jac.server.JACClientImpl,com.ibm.jac.server.JACGroupImpl,com.ibm.jac.server.JACCol
aco.disableAccessChecks=false
bootstrap.services.filename=C:\Tivoli\SCM\server\services
jac.server_directory=C:\Tivoli\SCM\server\
db.jdbc_jar=D:\Program Files\IBM\SQLLIB\java\db2java.zip
```

Figure 5, server.ini

This configuration file is provided with comments for each setting, note that it has important information such as the location of password keystores, databases and the SMTP mail server, DB2 access information and debug settings.

Note the `jac_dhcp_auto_register` setting, this needs to be set to true if you want DHCP clients to get automatically registered at the SCM server.

Implementing a Proxy Agent

The basic steps to installing a proxy agent are as follows

- i. Install the SCM client on the machine where you want to run the proxy agent
- ii. During installation ensure that the client type is defined as PULL
- iii. Use the SCM GUI or CLI to Register this client on the SCM server
- iv. Add the `com.ibm.jac.server.JACProxy.jar` collector to the client. If this collector has not been installed into SCM yet, you would have to install, authorise and register it before this step
- v. Edit the proxy collector – do this from the Clients view, by right clicking on the collector field and selecting Edit Collector from the drop down menu. You will need to add INBOUND and OUTBOUND IP and port values for the proxy, e.g.
INBOUND 192.100.120.54/32:1960
OUTBOUND 192.100.120.2/16:1950
- vi. where in this case 32 and 16 specify net masks to apply to the IP address to define the range of address to accept/broadcast to, i.e. 32 = 4 x 8 (4 IP octets) therefore the mask applies to the whole IP address so that only traffic from the single IP address is accepted, so in this case only traffic from the SCM server, which has an address of 192.100.120.24, the proxy will be listening on port 1960. For the outbound traffic we have specified a netmask of 16, i.e. two octets, which means the traffic from the proxy agent, will be broadcast to IPs in the 192.100.X.X address range, the clients will all be listening on port 1950.
- vii. Select the proxy tab on the Admin GUI and right click in the lower panel and select Create Proxy from the menu.
- viii. Enter a meaningful name for the Proxy
- ix. Enter the IP address of the machine the proxy agent resides on
- x. Next, install and register the clients that will connect to the SCM server via the proxy. These clients need to be PULL clients also; Push clients are not supported for proxy operation. Existing clients may need their `client.pref` files modifying as well as configuration from the GUI
- xi. From the Client panel, click on the client name you wish to connect to the proxy. From the client configuration section click on the arrow by the Proxy field open a list with the names of available proxies, select the proxy and then click on the Save button.
- xii. Right click on the `com.ibm.jac.server.JACProxy` collector field (whilst still in the Clients panel), and select Run Collector now from the menu.

The client should connect to the SCM server via the proxy at the next “heartbeat”, if this does not occur try restarting the clients and server

Miscellaneous

Starting and stopping the SCM client

To start the SCM client in Windows you can use the jac.bat command from the c:\Tivoli\SCM\client directory, however, it is recommended to start the SCM client using the Windows Services utility, either from the GUI or with the following command:

```
net start "SCM Client"
```

this is because some collectors may not have the authority to access data if the client was not started as the SYSTEM account. To stop the SCM client use the Services utility or the net stop "SCM Client" commands.

To start or stop the SCM client on a UNIX platform use the jacclient command, e.g.

```
./jacclient start  
./jacclient stop
```

You will probably find that you cannot stop the client properly using the kill -9 command.

Using the Build SQL function

When using the Build SQL function, you should always carefully check the syntax of the automatically generated query, as on occasions you may find minor errors which will prevent the query from successfully running, e.g. you may find that spaces have not been inserted between words where they should have.

Also, care should be taken when building a query using this utility, especially when the query is anything more than trivial, as each time you run the BUILD SQL function the last query is lost. It may be better to edit the query manually if changes are necessary.

Troubleshooting

As SCM is quite a "light" product, with few pre-requisite products, troubleshooting tends to quite straightforward. At the time of writing, the problems that tend to be difficult to troubleshoot are more likely to be issues with the beta product code rather than configuration or communication problems, as these tend to be relatively simple to remedy.

CLIENT

Client flusher doesn't restart after reset

In rare occasions during client soft resetting, a client could mistakenly think the flusher threading is still running and therefore omit to start a new flusher thread. If the flusher thread of a client is not running, the collector messages will not be sent to the server in a timely manner and instead of the flusher flushing the client's message queue every minute, the client will only send all messages in its message queue after the number of messages in the queue reaches 101.

This problem will be resolved in the first service refresh and the current work-around is to hard-reset the clients if they are not sending collector messages to the server in a timely manner.

Performing a client register, unregister and reregister does not complete successfully

After a client is active and registered, the unregister process is performed. This process works successfully on the server, but the client doesn't clean up all of the information (client.id) from the previous register process. An attempt to reregister the client will be appear successful from the server, but the client never becomes active. A circumvention to this reregistration problem is to shut down the client after it has been unregistered (jaccient stop command on UNIX or stopping the service in windows), then removing the client.id file that exists in the JACclient directory (where the client code was installed). Restart the client (jaccient start command or start the service). Then register the client as normal.

A client IP address sometimes shows on the unregistered clients list even though the client is currently registered

If a client machine is installed as a push client before it is registered to the JAC server then that client's IP address is added to the unregistered clients list. If you then use the register button from the show unregistered client popup to register the client, the client will be successfully registered and the IP address will be removed from the list. However, if you instead used the client menu bar action Register client to register the new client, the IP address remains on the unregistered clients list.

If you encounter a situation where the client is validly registered and the IP address still shows on the unregistered clients list, you can simply remove it from the unregistered clients list by selecting the IP address and pressing the remove button. If you select the IP address and attempt to register it when it is already registered, a warning will be displayed that you are attempting to register a duplicate IP address.

Unregistering of an Active Pull Client Times Out

When unregistering an active pull client, the server can take a long time (up to several minutes) to quiesce the client. On occasion, this delay results in a read timeout and an unmarshaling exception on the Administration GUI. After a timeout, close the exception

screen and continue with normal processing as the client has been successfully unregister.

Adding a Client to a Group

When adding a client to a group, you can only select one group at a time.

Debugging Collectors

Some of the collectors will do a `printStackTrace()` when they get an exception. In this case, the exception trace will show up in the `log.txt` file on the JAC client.

If you want to look at error return data from collectors, you need to register the `com.ibm.jac.basecollector.ERROR_LOG` collector and add it to the clients. The error data will be put in the `JAC_DATA.ERROR_LOG` table.

Collector Error : Collector doesn't run here

To correct, remove the collectors from the client and then repush via `softreset`. To remove collectors delete the `/classes/*` directory from the client server. From the GUI console, right click on the agent and choose `softreset`. This could take about 10 minutes for the collectors to actually be pushed to the client.

USERINFO_NG.EXE Windows file causing JAC errors in WIN_USER Table

Error: 'COM.ibm.db2.jdbc.DB2Exception:[IBM][CLI Driver][DB2/6000]SQL0302N The value of a host variable in the EXECUTE or OPEN statement is too large for its corresponding use.SQLSTATE=22001' appeared in Message Exceptions of Server table in my JAC admin GUI. To correct this problem, replace the existing `userinfo_ng.exe` file on the client with the patched version below

Column Width(s) Too Small

COM.ibm.db2.jdbc.DB2Exception: [IBM][CLI Driver][DB2/NT] SQL0302N The value of a host variable in the EXECUTE or OPEN statement is too large for its corresponding use. SQLSTATE=22001

This message indicates that the data returned was too large for the corresponding column width. Unfortunately, SQL does not return which column width needs to be adjusted. Examine the returned data against the actual column width to determine what needs to be fixed.

To adjust column widths, either correct the problem in the collector and reregister or use the following SQL command (requires SADMIN):

```
ALTER TABLE tablename  
ALTER COLUMN columnname  
SET DATA TYPE VARCHAR(new#)
```

SERVER

Installation Passwords

Record the installation passwords used. Remember the server keystore password for use when importing and authorizing policies.

JAC Server Install completes but JAC Server will not start

Refer to the jacserver.log file for error details and suggestions.

Make sure permission to /home/db2inst1 is 755.

Make sure jac account owns all files where JAC server installed. Default:
/opt/IBM/JACserver

Make sure you are using the correct db2inst1 password when installing JAC server and that the db2inst1 password has not expired.

Make sure DNS and/or host files are correct. JAC Server will not function properly with incorrect DNS or Host file entries.

Try stopping all JAC processes and DB2. Then restart DB2, JAC server, then the JAC client.

To get more details for problems, start the JAC server in DEBUG mode:

```
IBM_JAVA_OPTIONS=-Dis.debug=1 <executable bin file>-log '!<output filename>'  
'@ALL'
```

On UNIX:

```
export JAC_DEBUG=-Dlog.stdout_echo=1
```

start as usual

On Windows:

edit the jac.bat to pass the -Dlog.stdout_echo=1 parameter

start by executing jac.bat (make sure you stop the service first)

What this will do is echo all the debug info to the screen, and keeps the client or server from forking into the background. This is nice, because not only do you get the log messages, but you also can just ctrl-C to stop the program. (And it won't keep trying to restart....)

Unexpected server behavior when the client port number is incorrect

When a wrong client primary port number is defined to the server, the server may still be able to detect the client is active and the client is also able to send messages to the server. But if you try to send commands such as run collector, immediate collector execution, ping, soft or hard reset, you will see a read time out exception. If such a problem occurs, select the client on the client tab of the Admin GUI and select client information tab in the lower right portion of the screen. Check to see if the client port matches the client port defined when the client was installed.

Policy Bundle Schedules Supersede Specific Client Schedule

If a policy bundle containing one or more collectors is assigned to a client or a group, then the schedules and parameters in the policy bundle supersede any definitions that may be specifically assigned to the client. If the schedules or collector parameters in the policy bundle are not correct for that client, then you cannot assign the policy bundle to that client. Instead, you should assign each of the individual collectors (from the collector panel or the add collector menu option) to the client and then set the parameters and schedules accordingly.

Importing a Policy Twice Can Result in Duplicate Entries

When a policy is imported, the code does not check if a duplicate name already exists. If you import an already existing policy, a duplicate policy (with duplicate compliance queries) will be created.

Importing a Policy Twice Can Result in Duplicate Entries

When a policy is imported, the code does not check if a duplicate name already exists. If you import an already existing policy, a duplicate policy (with duplicate compliance queries) will be created.

If you need to reimport an existing policy for some reason, it is recommended that you first delete the existing policy, then do the import step.

When deleting policies

Remove policy from all groups prior to deleting policies from the JAC server. If this is not done, you may start receiving errors in the GUI when clicking on the Groups 'Null Pointer Exception'.

'Null Pointer Exception' errors after deleting or importing policies.

Steps to correct:

stop the JAC server

```
su - db2inst1
```

```
db2 connect to jac
```

From the DB2 prompt:

```
db2 "delete from jac_sys.pol_gro where pol_id not in(select id from jac_sys.policies) "
```

```
db2 "delete from jac_sys.pol_cli where pol_id not in(select id from jac_sys.policies) "
```

```
db2 terminate
```

restart the JAC server

COLLECTORS

Collector data typed into parameter field is not saved.

When entering a collector parameter, you **MUST** press enter before pressing save. If you do not press enter, the data typed into the parameter field is not saved

An RMI RemoteException: Collector does not run here is received if a collector is assigned to a client and then a run collector now or immediate collector execute is done for that collector

You need to wait until the client heartbeat occurs or you must run a soft reset on the client, and then wait until the client becomes active. Collectors are not sent to the client and cannot be run until one of those actions occurs.

A collector immediate execute or collector run now does not work from the popup screen where you can select the collector to run

When you do a collector immediate or collector run action, it should be possible to run the action by only selecting client you wish to run the collector on, and then selecting the menu action. A popup window is then displayed where you can select the collector to run. This function is currently broken.

To circumvent this problem, always select both the client to run and the collector to run before selecting the client immediate exec or run collector now option.

JAC Proxy ACL Restrictions Fail to Work for Some IP Addresses

When adding an Inbound or Outbound ACL definition to a JAC proxy machine, any value in the IP address that has a value above 128 will cause the IP mask to be incorrectly interpreted. To bypass this problem, do not make JAC Proxy ACL definitions that limit the IP addresses.

For example, the definition of the following in the Outbound ACL field for a JAC Proxy should limit the outbound connection to the defined IP address and port, but it will fail because one segment of the IP address contains a value of 133.

```
9.27.133.99:1950
```

To bypass the problem, remove the IP address check by coding the following definition instead:

```
9.27.133.99/0:1950
```

The /0 in the above definition stops the code from doing any IP address check and will only limit the outbound connections to port 1950 at any IP address.

Admin GUI

Error: "syscollect_admin_factory_server" when logging on using JAC Admin GUI

When a user attempts to login to a server while the server is starting up the error message "syscollect_admin_factory_server" will be displayed.

Problem: The JAC Server service/daemon is starting up Resolution: Wait a couple of minutes and try again

Error: "java.lang.ClassCastException: com.ibm.jac.HostInfo"

Symptoms:

When attempting to expand the a group in the clients panel, user receives the following exception:

```
java.lang.ClassCastException: com.ibm.jac.HostInfo
```

```
at java.lang.String.compareTo(String.java(Compiled Code))
```

at java.util.Arrays.mergeSort(Arrays.java:1065)

Cause:

Group_members tables contain host ips that are not listed in the JAC_SYS.HOSTS table.

Solution:

Check database tables to make sure that all hosts in the group_members are listed in the jac_sys.host tables, if not delete them and restart the JAC server.

Replace the syscollect.jar file in the install dir (eg: "C:\Program Files\IBM\JAC Administration GUI") the with copy attached to resolve GUI error and restart the GUI.

Run the following SQL from the SQL query panel in the JAC GUI to find the group members that are not in the JAC_SYS.HOSTS table:

List:

```
"select hostname , group from jac_sys.group_members where hostname not in (select  
hostname from jac_sys.hosts)"
```

Cleanup:

```
"delete from jac_sys.group_members where hostname not in (select hostname from  
jac_sys.hosts)"
```

restart the JAC Server to refresh server cache.

PS: Apply only to Admin Clients connecting to JAC Server version 3.2

Refresh Button on Administration GUI

The Refresh button is often necessary after GUI actions to update display information.

To refresh the display, press the refresh icon in the upper left corner of the Administration GUI (right next to JAC on the menu bar).

Examples of cases where the refresh button should be pressed:

If you add a collector and then switch to the table tag, the newly added table is not displayed unless you press the refresh button.

If you add or remove groups to other groups, a refresh is necessary to properly display the group layout in the client window.

Snapshot creation

When creating snapshots you may receive Warning/Error when creating a snapshot: 'An error occurred while generating snapshot for <> processing will continue. Exception: Warning: Violation Message will be truncated.' Process will continue, however, some of the violation details have been truncated. Execute the following SQL to increase the length of the message field for snapshot reports.

```
db2 alter table jac_sys.violations alter column message set  
data type varchar(1024)
```

Exporting Policies

When exporting policies, once you receive the window that the export has completed, you will need to manually cancel the progress window.

Collector Error : Collector does run here

To correct, remove the collectors from the client and then repush via softreset. To remove collectors delete the /classes/* directory from the client server. From the GUI console, right click on the agent and choose softreset. Could take about 10 minutes for the collectors to actually be pushed to the client.

Known Errors

Error

Connection Refused to Host: <server name>; nested exception is:
java.net.ConnectException: Connection refused: Connect.

Solution

This error can be caused by a personal firewall installed on the workstation from which you're attempting to access JAC Admin GUI. Check your firewall to see if it has been configured to prevent this communication. This error has been seen numerous times on jacsprat. The cause is currently unknown and investigation is continuing to find a resolution.

Error

'Null Pointer Exception' errors after deleting or importing policies.

Solution

stop the JAC server

```
su - db2inst1
```

```
db2 connect to jac
```

From the DB2 prompt:

```
db2 "delete from jac_sys.pol_gro where pol_id not in(select id from jac_sys.policies) "
```

```
db2 "delete from jac_sys.pol_cli where pol_id not in(select id from jac_sys.policies) "
```

```
db2 terminate
```

restart the JAC server

Error

SQL Exception: [IBM][CLI Driver][DB2/6000]

SQL0302N The value of a host variable in the EXECUTE or OPEN statement is too large for its corresponding use. SQLSTATE=22001

Solution

This is likely a problem with the win.any.win_user. Run the following query to increase the table size:

```
ALTER TABLE jac_idata.win_user ALTER COLUMN last_logon SET DATA TYPE VARCHAR(60)
```

FAQs (Frequently Asked Questions)

Q: Will ITSCM run on another database (i.e. not DB2)?

A: *The current plan is to support DB2 as the data store, the setting up and maintenance of the database is generally managed from within SCM.*

Q: Has IBM provided canned reports

A: *Yes, reports are available with SCM 5.1*

Q: Is there a direct link to IBM Tivoli Data Warehouse or is it always going to be via Risk Manager?

A: No plans for this release

Q: Will the IBM Tivoli Risk Manager adaptor be on every client or the IBM Tivoli Security Compliance Manager server?

A: The only Risk Manager adapter required will be on the SCM Server/Database and send non-compliant events to the RM/TEC Server.

Q: Is the 10,000 client limit a result of DB schema design?

A: The object ID is maintained in a 32-bit integer, so there is room for a VERY large number of client object IDs (billions). The 10K client limit is not a fixed limit, but a target for scalability, in terms of clients per server. In one sense, it relates to "active" clients that are pushing messages to the server. It also relates to the GUIs ability to manage and present information for a large number of objects.

Q: Is the SCM client web based or java client?

A: All components are Java, including the SCM Client

Q: Is the Java GUI on Win32 only?

A: Windows only. Linux support is future item. However, if a customer has a specific requirement in using the GUI on Linux let us know. The primary issue preventing the GUI from being used on Linux is test/support.

Q: Does client run as root/administrator and can it be configured?

A: This is feasible, but the Lab hasn't tested the scenario. Assuming none of the collectors require root access to do their job, necessary changes would include updating the start-up scripts to su to a non-root account and modifying file ownership/permissions as needed to run under the selected account. However, at this point we must discourage this, since ownership/permissions could affect more than just collector execution (collector update, for example), and we haven't really tested this scenario.

Q: Are there plans to merge the IBM Tivoli Security Compliance Manager Client and the Tivoli client?

A: IBM Tivoli is working on general plans to reduce the number of clients/agents it uses to one or two. This will be done as a migration and not a rip and replace, using standards based technologies, over a period of time.

Q: Can you use Tivoli to load clients (i.e., SWD from Configuration Manager)?

A: SCM 5.1 has the SCM Client pre packaged to be distributed by IBM Tivoli Configuration Manager as well as manual install. The Client can be installed using silent install and therefore could be packaged to use any other Software distribution program.

Q: Do you have a collector per platform or a collector per info set?

A: *A collector is required for each group of information collected. Therefore each client will generally have many collectors.*

Q: Can I create my own collectors?

A: *SCM 5.1 has a SDK that will allow for the development and signing on collector code by customers, Business Partners or a Services organisation.*

Q: What exactly is a collector?

A: *Java code that can read a file, execute a program to collect data, etc.*

Q: Is the policy schedule separate from the Collector schedule?

A: Yes

Q: Can you archive data for future reference?

A: *Yes, either by use of the SCM delta table or through Database back ups depending on the exact requirements*

Q: Does multiple use of a collector on one client result in multiple data tables

A: *SCM 5.1 will create multiple data tables with "Instance ID", each having different collector schedules or parameters for example.*

Q: What is the Collector Reuse function?

A: *"Collector Reuse" is the ability to associate the same collector "code" with multiple configurations (schedules and parameters) as it gets assigned to different clients, groups, policies, etc.*

Q: Are JVMs included for all platforms?

A: *All except the SCM Client for HP. The JVM needs to be downloaded from HP web site for this.*

Q: Is GSKIT used?

A: *NO, this is all Java using native JCE interfaces. The underlying JCE provider is the IBM JSSE v1.4. Java savvy users may notice we are using a 1.4 JSSE with a 1.3 JVM. And it works very nicely.*

Q: Is SCM supported on AIX 4.3.3?

A: SCM 5.1 does not support this version as AIX 4.3.3 is out of support.

Q: Do the collectors or server write to DB2?

A: A configured number of threads are deployed to handle messages as they are received at the Server. The messages are queued up as they are received and wait processing on one of the message threads (for insertion into the DB). The number of dedicated threads can be increased or decreased as necessary.

Q: Is the server to DB write single or multi threaded?

A: Multi- Threaded

Q: Can you change the IP address of a fixed IP client without unregistering it or losing its SCM Client ID?

A: Yes. You can change the IP address from the client panel directly (or make a database update). Simply enter the new value and hit "save".

Q: What is the Master Key store and Server key store?

A: The "Master Key" provides possible separation between the SCM Admin and another party. (Typically the admin of the actual server machine). The master key is used mainly to allow future growth and scalability. (Redundant servers, key rollovers, etc...) The Server Key store is used for the SSL certificate

Q: Does the CLI have a subset of the GUI that does SSL?

A: The CLI commands are a subset of the current GUI function and a userid and pw must be supplied, plus the flow is SSL encrypted.

Q: What are the IBM Policies based upon?

A: They were derived from IBM internal policies which were interpretations or commercial standards

Q: Can the DB2 instance be created manually to an organisations DBA policy default?

A: Yes, the DB2 instance may be manually created. The SCM Install Shield contains the DDL necessary to setup the JAC database.

Q: Can you change server.ini and then restart server?

A: Yes

Q: Does the install set autostart = true for service and process?

A: The server will only be started at reboot if the option to stash the password is selected during install. If you add the password to server.ini, you will have to

issue the mkitab command to ensure that the server is started on reboot. Note: you can add the password plain text to server.ini and when the server is started, it will be overwritten with the obfuscated value. The client is automatically set to start on reboot

Q: Can we install multiple components at one time?

A: No

Q: What is the impact of duplicate Alias Names - DHCP?

A: *In a DHCP setup, two systems with the same alias but different "fingerprints" will appear as two separate systems (both in the GUI and the database). They will have different client IDs and different fingerprints however. The "fingerprints" are the digital fingerprint of the client's SSL certificate, so even an exact duplicate machine has virtually zero chance of a collision.*

Q: What is a "message" in the client queue as referred to in the .pref file - max number of messages 100 etc?

A: *The array of information a collector returns to the server is called a "message". It might represent data going into a single row in the database, or multiple rows. More specifically, a message is generally the output a single collector execution. The "max messages" preference refers to the maximum number of messages the client will hold in local memory before it begins to discard the oldest ones. This is used in cases where the server cannot be reached for an extended period of time to prevent the client from consuming too much memory and/or disk space.*

Q: What is the server hardware spec?

A: See SCM HW Requirements DOC

Q: Does the CLI create SSL connections?

A: Yes. A userid and pw must be supplied, plus the flow is SSL encrypted.

APPENDIX

Additional tools

db2look

To generate all the DB2 CLP commands and Data Definition Language (DDL) statements required to make a copy of a database, use the following:

db2look syntax:

db2look Version 8.2

Syntax:

```
db2look -d DBname [-e] [-u Creator] [-z Schema] [-t Tname1 Tname2...TnameN] [-tw Tname] [-h] [-o Fname]
[-a] [-m] [-c] [-r] [-l] [-x] [-xd] [-f] [-fd] [-td x] [-noview] [-i userID] [-w password] [-v Vname1 Vname2 ...
VnameN] [-wrapper WrapperName] [-server ServerName] [-nofed]
```

```
db2look -d DBname [-u Creator] [-s] [-g] [-a] [-t Tname1 Tname2...TnameN] [-p] [-o Fname] [-i
userID] [-w password]
```

```
db2look [-h]
```

-d: Database Name: This must be specified

-e: Extract DDL file needed to duplicate database

-u: Creator ID: If -u and -a are both not specified then \$USER will be used

-z: Schema name: If -z and -a are both specified then -z will be ignored

-t: Generate statistics for the specified tables

-tw: Generate DDLs for tables whose names match the pattern criteria (wildcard characters) of the table name

-h: More detailed help message

-o: Redirects the output to the given file name

-a: Generate statistics for all creators

-m: Run the db2look utility in mimic mode

-c: Do not generate COMMIT statements for mimic

-r: Do not generate RUNSTATS statements for mimic

-l: Generate Database Layout: Database partition groups, Bufferpools and Tablespace

-x: Generate Authorization statements DDL excluding the original definer of the object

-xd: Generate Authorization statements DDL including the original definer of the object

-f: Extract configuration parameters and environment variables

-td: Specifies x to be statement delimiter (default is semicolon(;))

-p: Use plain text format

-s: Generate a postscript file

-g: Use graph to show page fetch pairs for indices

-i: User ID to log on to the server where the database resides

-w: Password to log on to the server where the database resides

-noview: Do not generate CREATE VIEW ddl statements

-wrapper: Generates DDLs for federated objects that apply to this wrapper

-server: Generates DDLs for federated objects that apply to this server

-nofed: Do not generate Federated DDL

-fd: Generates db2fopt statements for opt_buffpage and opt_sortheap along with other cfg and env parameters.

-v: Generate DDL for view only, this option is ignored when -t is specified

```
db2look -d <dbname> -e -m -a -o db2look.out
```

Verify the correct instance if necessary:

```
echo %DB2INSTANCE%
```

To set the instance:

```
set DB2INSTANCE=<instance name>
```

```
db2look -d <db> -u <username> -e -o <file>
```

example:

```
db2look -d sample -u db2admin -e -o db2look.sql
```

To recreate:

```
db2 -tf db2look.sql
```