


Deployment Guide for InfoSphere Guardium

Becoming compliance and audit ready
with InfoSphere Guardium

Understanding the InfoSphere
Guardium solution architectures

Implementing InfoSphere
Guardium solutions



Whei-Jen Chen
Boaz Barkai
Joe M DiPietro
Vladislav Langman
Daniel Perlov
Roy Riah
Yosef Rozenblit
Abdiel Santos



International Technical Support Organization

Deployment Guide for InfoSphere Guardium

March 2014

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

First Edition (March 2014)

This edition applies to InfoSphere Guardium Version 9.1.

© Copyright International Business Machines Corporation 2014. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Noticesxi
Trademarks	xii
Preface	xiii
Authors	xiii
Acknowledgement	xvi
Now you can become a published author, too!	xvi
Comments welcome	xvii
Stay connected to IBM Redbooks	xvii
Chapter 1. Solutions and architecture	1
1.1 Introduction to regulation compliance and auditing requirement	2
1.2 Database security and lifecycle	3
1.3 Architecture and functional characteristics	7
1.3.1 Product architecture components	7
1.3.2 Product functionality	8
1.3.3 Product architecture options	9
1.3.4 S-TAP architecture options	11
1.4 Integration with IT infrastructure	15
1.5 Supported platforms	17
Chapter 2. Implementation planning	19
2.1 Knowing product deployment types	21
2.2 Sizing and topology considerations	23
2.2.1 Audit level	23
2.2.2 Database to collector sizing	24
2.2.3 Collector to aggregator ratio	25
2.2.4 Central Manager ratio	27
2.3 Contingency and design considerations	27
2.3.1 Appliance location	28
2.3.2 Appliance configuration options	29
2.3.3 S-TAP agent contingency and configuration	31
2.4 Implementation approach	35
2.5 Implementation schedule	36
2.6 Roles and responsibilities	37
2.7 Installation and configuration sessions	41
2.8 Future growth considerations	42
Chapter 3. Installation and configuration	43

3.1	Schedule of Implementation activities	44
3.2	Installation and configuration planning	45
3.3	Guardium appliance overview	46
3.3.1	Hardware and software appliance modes	47
3.3.2	Appliance types	47
3.4	Database and appliance inventory	50
3.4.1	Database inventory	50
3.4.2	Collector sizing	51
3.4.3	Aggregator sizing	52
3.5	Appliance deployment considerations	53
3.5.1	Managed or stand-alone environment	53
3.5.2	Contingency	54
3.5.3	Networking	55
3.5.4	Other allocation factors	56
3.5.5	Start with a test environment	57
3.5.6	Licensing	57
3.6	Appliance installation and configuration	58
3.6.1	Rack or build appliance	58
3.6.2	Configuring the appliance	59
3.7	Agent deployment, installation, and configuration	64
3.7.1	Guardium installation manager	65
3.7.2	S-TAP	69
3.7.3	Command-line options	76
3.7.4	S-TAP and its inspection engine configuration	77
3.7.5	More information	79
3.8	Configure remaining appliances	80
3.8.1	Managed units grouping	80
3.8.2	Creating a group	81
3.9	Guardium operations	83
3.9.1	Configuration	83
3.9.2	Transferring configuration to managed units	86
3.9.3	Data management	87
3.9.4	System backup	91
3.9.5	Aggregation	92
3.9.6	Job schedules	94
3.9.7	Enterprise systems interface configuration	95
3.9.8	Self-monitoring	95
3.10	Vulnerability assessment	104
3.10.1	Creating a database account	105
3.10.2	Creating a data source	106
3.10.3	VA tests	108
3.10.4	Creating an assessment	109
3.10.5	Custom reports	114

3.10.6	Scheduling an assessment	115
3.10.7	More information	116
3.10.8	High-level steps summary	116
3.11	Where to find more help	117
3.11.1	Product documentation	117
3.11.2	Release notifications and bulletins	118
3.11.3	Product support.	119
3.11.4	User community and support	119
Chapter 4.	Monitoring and auditing	121
4.1	Regulations and compliance	123
4.2	Auditing categories	123
4.2.1	Privilege user monitoring.	125
4.2.2	Sensitive object monitoring	127
4.2.3	Comprehensive monitoring	128
4.3	Auditing requirements	129
4.4	Database activity monitoring	131
4.5	Vulnerability assessment.	132
4.6	Mapping audit requirements to the solution	136
4.6.1	Requirement 2: Do not use vendor-supplied defaults for system passwords.	136
4.6.2	Requirement 3: Protect stored cardholder data	139
4.6.3	Requirement 6: Identify systems missing patches and enforce change controls	139
4.6.4	Requirement 7: Compensating control for network segmentation	140
4.6.5	Requirement 8: Assign a unique ID for each person with a computer access.	141
4.6.6	Requirement 10: Track and monitor all access to cardholder data.	142
4.6.7	Requirement 11: Regularly test systems	143
4.6.8	SOX requirement: Prevent unauthorized changes to the financial CRM, ERP, and HR data	144
Chapter 5.	Monitoring setup.	145
5.1	Monitoring setup overview.	147
5.2	Monitoring planning.	149
5.3	Grouping	151
5.3.1	Wildcards	153
5.3.2	Hierarchical groups	153
5.3.3	Group deletion.	155
5.3.4	Public versus private groups	155
5.3.5	Groups in federated environment	156
5.3.6	Size and performance effect	156
5.4	Policy.	156

5.4.1	Policy types and policy rules	157
5.4.2	Logging granularity	160
5.4.3	Generating real-time alerts with policy	163
5.4.4	Extrusion rules	164
5.4.5	Database exceptions.	166
5.4.6	Policy for z/OS	168
5.4.7	Policy installation.	170
5.5	Reports	174
5.5.1	Reports versus queries	174
5.5.2	Domain, entities, and attributes.	175
5.5.3	Query conditions	178
5.6	Compliance workflow	179
5.7	Real-time and threshold alerting	182
5.7.1	Alert generating.	182
5.7.2	Alerter	188
5.8	Data level access control	189
5.8.1	S-TAP setup	190
5.8.2	Policy setup.	190
5.8.3	Policy violation report	190
5.9	Vulnerability assessment setup.	191
5.10	Configuration audit system setup	191
5.10.1	Prerequisites	192
5.10.2	High-level steps.	193
5.10.3	Installation and configuration	193
5.10.4	Reviewing results	201
5.10.5	Next steps	202
5.11	Entitlement reporting setup	202
5.11.1	Prerequisites	203
5.11.2	High-level steps.	203
5.11.3	Configuration steps	204
5.11.4	Review the entitlement data	207
5.12	Database auto-discovery setup.	208
5.12.1	Prerequisites	209
5.12.2	High-level steps.	209
5.12.3	Configuration steps	209
5.12.4	Viewing the results	211
5.12.5	Next steps	213
5.13	Sensitive data finder setup	214
5.13.1	Use cases and highlights	214
5.13.2	Prerequisites	216
5.13.3	High-level steps.	216
5.13.4	Configuration steps	216
5.13.5	Viewing classification process results	222

5.13.6 Next steps	224
5.14 Adding a menu tab to your portal	225
Chapter 6. Access management	227
6.1 Access management overview	228
6.1.1 Roles, portals, applications, and users	228
6.1.2 Managed configuration	229
6.1.3 Authorization versus authentication	229
6.1.4 Data-level security.	231
6.2 User accounts	231
6.2.1 Creating a user account manually.	231
6.2.2 Importing user accounts from LDAP	232
6.2.3 Modifying a user's role	233
6.3 User role browser	234
6.3.1 Custom role.	235
6.3.2 Adding a role.	236
Chapter 7. Ongoing operations	237
7.1 Performance optimization and tuning	238
7.1.1 S-TAP optimization and tuning	238
7.1.2 Inspection Core performance and unit usage	250
7.1.3 Interactive reports	266
7.1.4 Audit process reports	273
7.1.5 Back up and purge	275
7.1.6 Central management.	279
7.2 Maintenance and updates.	281
7.2.1 Appliance updates.	281
7.3 Diagnostic tools	288
7.3.1 Appliance diagnostic tools.	289
7.3.2 Agents diagnostic tests	298
7.4 Restoring audit data for forensic analysis	303
7.4.1 Restore strategies.	304
7.4.2 Restore audit process result sets	310
Chapter 8. Disaster recovery	313
8.1 Overview	314
8.1.1 System backup and restore	314
8.1.2 Data archive and restore.	316
8.1.3 Storage location	317
8.2 Appliance recovery	317
8.3 Appliance recovery steps	318
8.3.1 Stand-alone collector	318
8.3.2 Collector with Aggregator no CM	319
8.3.3 Aggregator: Not centrally managed	319

8.3.4 Centrally managed collector	319
8.3.5 Centrally managed aggregator	320
8.3.6 Dedicated central manager (no data aggregation)	321
Chapter 9. Upgrade best practices	323
9.1 Content delivery for 32-bit and 64-bit systems	324
9.2 Compatibility considerations for 32-bit versus 64-bit systems.	325
9.2.1 Central management.	325
9.2.2 Aggregation.	325
9.2.3 Central management and aggregation combo appliances	326
9.3 RAM memory considerations	327
9.4 Enterprise upgrade strategy	328
9.4.1 Change control management	328
9.4.2 Top-to-bottom rollout order	329
9.4.3 Strategies	329
9.4.4 Enterprise upgrade summary	334
9.5 Appliance upgrade methods	334
9.5.1 Resource download locator.	335
9.5.2 Upgrading appliances by using V8.2 to V9.0p50 upgrade bundle	335
9.5.3 Rebuilding appliances with V9.1p100 (64-bit) ISO image and restoring the V8.2 or V9.0 backup	338
9.5.4 Building new separate V9.1p100(64-bit) appliances and gradually retiring the old ones	339
9.6 Agents upgrade methods	339
9.6.1 GIM upgrades	339
9.6.2 UNIX S-TAP upgrade	340
9.6.3 Windows S-TAP upgrade	340
9.7 Upgrading bundle monitoring and status validation	341
9.7.1 Health check monitoring	341
9.7.2 Upgrade process monitoring	343
9.7.3 Troubleshooting upgrade issues	344
Chapter 10. Use cases	347
10.1 Creating an integrated change management report	348
10.1.1 Importing change tickets with Enterprise Integrator	350
10.1.2 Uploading the data from the external table	360
10.1.3 Defining the custom domain	363
10.1.4 Defining custom query with new external domain entities	370
10.1.5 Customizing the reports	376
10.2 Connection profiling and security best practices	386
Chapter 11. Integration with other IBM products	391
11.1 Security and audit	393
11.1.1 BigInsights integration.	393

11.2 Databases and data warehouses	397
11.2.1 Basic database activity monitoring	397
11.2.2 Advanced database activity monitoring	400
11.2.3 Basic vulnerability assessment	402
11.2.4 Advanced vulnerability assessment	403
11.3 Data lifecycle management integration	406
11.3.1 Identifying archive candidates	406
11.3.2 Auditing archive access	408
11.3.3 Test data management integration	411
11.3.4 Exchanging the sensitive information location with InfoSphere Discovery	413
11.4 Identifying the user activity	416
11.4.1 Custom identification procedures	417
11.4.2 GuardAppEvents and GuardAppUser	418
11.4.3 Setting client user ID	421
11.4.4 WebSphere application user information	422
11.4.5 CICS application user	423
11.4.6 InfoSphere MDM and Data Stage application user	426
11.4.7 Cognos application user	427
11.5 Security integration	428
11.5.1 Guardium and QRadar integration	428
11.5.2 Tivoli Netcool	434
11.5.3 IBM Security Access Manager for enterprise single sign-on	437
11.5.4 Tivoli Directory: Lightweight Directory Access Protocol	440
11.5.5 IBM Endpoint Manager and Guardium Integration	443
11.5.6 Tivoli Storage Manager and Guardium Integration	444
Related publications	449
Online resources	449
Help from IBM	449

Notices

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

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

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

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

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

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

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

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

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

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

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

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com) are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

AIX®	IMS™	Redbooks®
BigInsights™	Informix®	Redbooks (logo)  ®
CICS®	InfoSphere®	RETAIN®
Cognos®	Netcool®	S-TAP®
DB2®	Optim™	Tivoli®
developerWorks®	Passport Advantage®	WebSphere®
Guardium®	PureData™	z/OS®
IBM®	QRadar®	

The following terms are trademarks of other companies:

Netezza, and N logo are trademarks or registered trademarks of IBM International Group B.V., an IBM Company.

QRadar, and the Q1 logo are trademarks or registered trademarks of Q1 Labs, an IBM Company.

Intel, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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

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

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

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

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

Preface

IBM® InfoSphere® Guardium® provides the simplest, most robust solution for data security and data privacy by assuring the integrity of trusted information in your data center. InfoSphere Guardium helps you reduce support costs by automating the entire compliance auditing process across heterogeneous environments. InfoSphere Guardium offers a flexible and scalable solution to support varying customer architecture requirements. This IBM Redbooks® publication provides a guide for deploying the Guardium solutions.

This book also provides a roadmap process for implementing an InfoSphere Guardium solution that is based on years of experience and best practices that were collected from various Guardium experts. We describe planning, installation, configuration, monitoring, and administrating an InfoSphere Guardium environment. We also describe use cases and how InfoSphere Guardium integrates with other IBM products.

The guidance can help you successfully deploy and manage an IBM InfoSphere Guardium system.

This book is intended for the system administrators and support staff who are responsible for deploying or supporting an InfoSphere Guardium environment.

Authors

This book was produced by a team of specialists from around the world working at the IBM Littleton Massachusetts Laboratory.

Whei-Jen Chen is a Project Leader at the International Technical Support Organization, San Jose Center. She has extensive experience in application development, database design and modeling, and IBM DB2® system administration. Whei-Jen is an IBM Certified Solutions Expert in Database Administration and Application Development, and an IBM Certified IT Specialist.



Boaz Barkai is the IBM InfoSphere Data governance Services Leader with over 25 years experience in software design, implementation, and security. Before joining IBM, Boaz headed the support and services teams for Guardium. As part of his experiences with Guardium, Boaz provided valuable product implementation experiences, guidance, and support to many Guardium customers worldwide.



Joe M DiPietro is IBM InfoSphere Data Governance Center of Excellence Leader. Joe has for over 25 years experience in security and network design and implementation. Before working with IBM and Guardium, he worked at security pioneer Check Point Software for over 8 years. Previously, DiPietro was corporate systems engineer for SynOptics Communications and a member of the company's World Wide Technical Counsel (WWTC).




Vladislav Langman is the IBM InfoSphere Guardium WW L3 Engineering Leader with over 20 years of experience in software design, development, and implementation. He oversees Research and Development escalations for all aspects of Guardium solution and design of supportability and diagnostic tools. Before IBM and Guardium, Vlad lead software development teams at Amdocs, where he implemented enterprise IT solutions for telecommunication companies worldwide. Vlad holds an Engineering degree in Information Technology.




Daniel Perlov is a Senior Systems Architect with over 30 years experience in software design, implementation, and support with expertise in relational databases. Daniel joined Guardium in 2004 as a support engineer. In 2005, he became the Guardium support manager and led the support team in Guardium and IBM. Daniel moved to Research and Development in 2013 to use his customers' knowledge and experience to influence product usability and supportability.

Roy Riah is a Senior Management Consultant specializing in InfoSphere Guardium Solutions and Data Privacy for Linux, UNIX, Windows, and z/OS®. Roy leads the implementation of the InfoSphere Guardium database security solution for IBM customers in various industries. His experience includes implementing the Guardium solution for Data Activity Monitoring, Data Activity Protection, Vulnerability Assessment, Discovery, Classification, and integrating the solution with various SIEM, LDAP, and other systems and data sources. His responsibilities include requirements, analysis, design, deployment planning, configuration, testing, troubleshooting, documentation, and customer training and support.



Yosef Rozenblit is an IBM InfoSphere Guardium worldwide services leader. Before joining IBM, Yosef was the director of professional services at Guardium and was responsible for the implementation of Guardium solutions for hundreds of customers.



Abdiel Santos is a Senior Level 3 Support Engineer for InfoSphere Guardium at the Littleton, Massachusetts IBM Laboratory. He has worked with the Guardium solution since 2006, handling the most critical support issues and providing customers with custom solutions to successfully implement Guardium in their environments.

Acknowledgement

Thanks to the following people for their contributions to this project:

- ▶ David Rozenblat
- ▶ Nir Carmel
- ▶ Louis Lam
- ▶ Ron Ben-Natan
- ▶ Amy Wong
- ▶ Michael Murphy
- ▶ William Pacino

IBM USA

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

<http://www.ibm.com/redbooks/residencies.html>

Comments welcome

Your comments are important to us!

We want our books to be as helpful as possible. Send us your comments about this book or other IBM Redbooks publications in one of the following ways:

- ▶ Use the online **Contact us** review Redbooks form found at:
<http://www.ibm.com/redbooks>
- ▶ Send your comments in an email to:
redbooks@us.ibm.com
- ▶ Mail your comments to:
IBM Corporation, International Technical Support Organization
Dept. HYTD Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

- ▶ Find us on Facebook:
<http://www.facebook.com/IBMRedbooks>
- ▶ Follow us on Twitter:
<http://twitter.com/ibmredbooks>
- ▶ Look for us on LinkedIn:
<http://www.linkedin.com/groups?home=&gid=2130806>
- ▶ Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:
<https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm>
- ▶ Stay current on recent Redbooks publications with RSS Feeds:
<http://www.redbooks.ibm.com/rss.html>



Solutions and architecture

This chapter provides background information about database activity monitoring (DAM) and why it is important for organizations to include DAM in their overall compliance and security framework. We describe the architecture, capabilities, and solutions of IBM InfoSphere Guardium that satisfy the business requirements for database activity monitoring.

This chapter includes the following topics:

- ▶ Introduction to regulation compliance and auditing requirement
- ▶ Database security and lifecycle
- ▶ Architecture and functional characteristics
- ▶ Integration with IT infrastructure
- ▶ Supported platforms

1.1 Introduction to regulation compliance and auditing requirement

Corporate infrastructure evolved to allow information access over both internal and external networks: Intranet and internet. This technology provides fast speed, convenience, and flexibility in data accessing but also introduces new levels of fraud opportunities and security issues. To prevent fraud and data security breaches, proper security controls must be in place. To validate that the controls are active, regulations were developed and implemented over the years.

The Sarbanes–Oxley Act of 2002 (SOX) was passed into law that resulted from a number of corporate scandals about a decade ago. This law changes the financial reporting for public corporations and was the beginning of stricter regulatory for corporate management oversight. To provide accountability of the top management of a company, they must now individually certify the accuracy of financial information. The law also requires external auditors to verify the accuracy of certain financial information, such as balance sheet.

Now, most companies report their financial activity to be compliant with SOX electronically. If we translate what this regulation means to database activity monitoring in one sentence, it reads something similar to the following sentence:

“Monitor all changes to the financial database server to ensure that no unauthorized transaction occurred to affect the financial results of the company.”

This ensures that the financial integrity of the transactions that are stored in the database server are correct and accurate for SOX reporting.

SOX is an example of one regulation that affects database monitoring activity to ensure that the integrity of the information that is stored in the database is correct and accurate. There are numerous other regulations to which companies must adhere to be in compliance.

Figure 1-1 on page 3 outlines the basic DAM information that must be collected for some of these regulations. The following terms are used in the table:

- ▶ DDL: Data definition language is the schema or the container of the database. SQL statements, such as CREATE, ALTER, and DROP, are DDL commands.
- ▶ DML: Data manipulation language is the contents of the database or the data that is stored inside. SQL statements, such as INSERT, UPDATE, and DELETE, are DML commands.

- DCL: Data control language controls that receive access to portions of the database. SQL statements, such as GRANT and REVOKE, are part of DCL commands.

Audit Requirements	PCI DSS	COBIT (SOX)	ISO 27002	Data Privacy & Protection Laws	NIST SP 800-53 (FISMA)
1. Access to Sensitive Data (Successful/Failed SELECTs)	✓		✓	✓	✓
2. Schema Changes (DDL) (Create/Drop/Alter Tables, etc.)	✓	✓	✓	✓	✓
3. Data Changes (DML) (Insert, Update, Delete)		✓	✓		
4. Security Exceptions (Failed logins, SQL errors, etc.)	✓	✓	✓	✓	✓
5. Accounts, Roles & Permissions (DCL) (GRANT, REVOKE)	✓	✓	✓	✓	✓

Figure 1-1 Auditing requirements by regulation

1.2 Database security and lifecycle

Regulatory compliance and auditing is only one aspect of database monitoring. A complete data security solution considers security monitoring challenges, setting the monitoring goals, data security lifecycle, and infrastructure constraints as shown in Figure 1-2 on page 4.

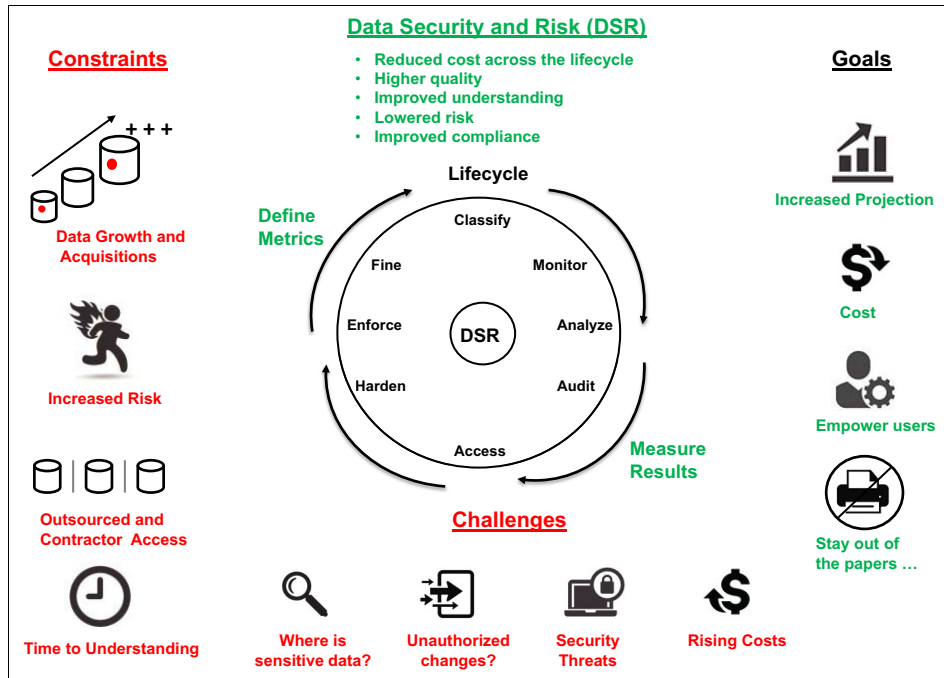


Figure 1-2 Challenges and lifecycle of database security

Many customers have the following common challenges in monitoring and auditing their environment:

- ▶ Where is my sensitive data located?
- ▶ Are there unauthorized changes to my data?
- ▶ How can I protect my information against the vast number of Security threats?
- ▶ How can I reduce my infrastructure costs that are rising dramatically?

Customers are constantly trying to balance these challenges with the following ultimate goals:

- ▶ Increasing the overall protection of information within the environment.
- ▶ Reducing the cost for compliance and security within their business.
- ▶ Empowering users with information so that they can make good decisions that positively affect the business.
- ▶ Staying away from negative publicity that can result from a data breach.

These are some of the most common goals, and the need to balance these goals within the constraints of the business includes the following factors:

- ▶ Time to understanding means how long it takes to identify that someone hacked into your environment. In many studies, hackers can infiltrate your environment for weeks, months, and sometime years before they are identified and exposed.
- ▶ Contractor and outsourced access to sensitive information is common in today's business environment. This is a constraint that every business must be able to address.
- ▶ Increased risk of information being compromised because of new and dynamic threats.
- ▶ Data growth is exploding. As data is copied and moved into data warehouses, Hadoop clusters, and other areas, it must have appropriate controls to secure this information. Another area for data growth is customers that acquire new companies. They must bring this new environment into compliance.

DAM goes through a lifecycle that includes the following tasks:

- ▶ Find your sensitive data

In many organizations, it is difficult to know where the sensitive data is stored. This is a task that must be automated because the environment is constantly changing as data moves to new databases and unstructured data stores, such as Hadoop. To know where your sensitive information is at any one time, you must continuously scan for this type of information.

- ▶ Monitor

After you know where your sensitive data is, you can set appropriate monitoring and security controls to identify who is accessing this information. For example, Payment Card Industry (PCI) requirement 10.2.1 requires you to verify that all individual access to cardholder data is logged.

- ▶ Analyze

Analyze your environment to understand where gaps are in your security policies. In today's virtualized world, resources can be quickly allocated and deployed to meet the business needs. It is critical to analyze the data that these new applications embed into your infrastructure. This analysis helps drive the audit requirements to achieve compliance and security controls for these new applications.

▶ Audit

Audit is used to validate your processes and procedures to achieve your security goals and identify gaps in the current processes. These goals should include validating the reliability of the information collected; verifying that change control processes are in effect; confirming that exception reporting is identified and working; archiving and restoring audit data for forensic events; and so on.

▶ Assess

Assess your current environment to identify missing patches or configuration issues with your servers. Assess how database infrastructure is being used. For example, is there a security risk because individuals are sharing credentials? Is there a strong password policy in place? After assessing your environment, you must remediate the found issues.

▶ Harden

Harden the environment means to close the gaps that were identified during the assessment phase. For example, we identified that there is no strong password policy in place to log in to the database. During the harden phase, this is corrected. After all of the gaps are remediated, you can improve your security policy.

▶ Enforce

Enforce means to proactively identify security policies to alert and potentially block access to unauthorized resources. For example, if the only connection to the database is coming from the application server (10.10.10.10), you can write a security policy to prevent any access from IP addresses other than 10.10.10.10. This helps secure your environment from known and unknown attacks that are started from unknown connections (everything other than 10.10.10.10).

1.3 Architecture and functional characteristics

This section provides an overview of the IBM InfoSphere Guardium architectural components and product functionality characteristics.

1.3.1 Product architecture components

From the architecture view, IBM InfoSphere Guardium components can be grouped into the following categories:

▶ Appliances

Appliances include the following subcategories:

- Collectors: The collector is the appliance that is used for real-time capture and analysis of the database activity.
- Aggregators: The aggregator appliance is used to offload reporting activity from the collectors and to provide consolidated reporting from multiple collectors.
- Central Managers: The central manager (CM) is specialized functionality that is enabled on an aggregator appliance. The CM function is used to manage and control multiple Guardium appliances.

▶ Agents

Agents include the following subcategories:

- Software TAP agent (S-TAP®): The S-TAP agent is installed on the database server and is used to monitor and relay the observed activity to the Guardium collector appliance.
- Guardium Installation Manager agent (GIM): The GIM agent is installed on the database server and is used to facilitate agent installation and the updating and configuration modification of agents.
- Change Audit System agent (CAS): The CAS agent is installed on the database server and is used to capture change audit information of configuration files and more on the database server.
- Instance Discovery agent: The instance discovery agent is installed on the database server and is used to obtain database, listener, and port information.

1.3.2 Product functionality

InfoSphere Guardium includes the following functionality to satisfy your compliance needs:

- ▶ Monitor and audit:
 - Data activity monitoring
 - Real-time alerting
 - Threshold alerting
 - Compliance reporting
 - Compliance workflow
 - User Identification
 - Security integrations
- ▶ Enforce and protect:
 - Blocking
 - Masking
 - Quarantine
- ▶ Assess and harden:
 - Vulnerability assessments
 - Configuration changes
 - Entitlement reporting
- ▶ Discover and classify:
 - Discovery of data sources
 - Classify sensitive data
 - Enterprise Integrator

1.3.3 Product architecture options

InfoSphere Guardium offers flexible and scalable solutions to support varying customer architecture requirements. In this section, we describe three architecture examples.

Basic stand-alone architecture

Figure 1-3 shows a basic architecture for monitoring several databases in one data center. This architecture example consists of one stand-alone collector appliance and several Guardium S-TAP agents that are installed on the monitored database servers. The S-TAP agents are configured to capture and send the relevant database activities to the one Guardium collector for analysis, parsing, and logging.

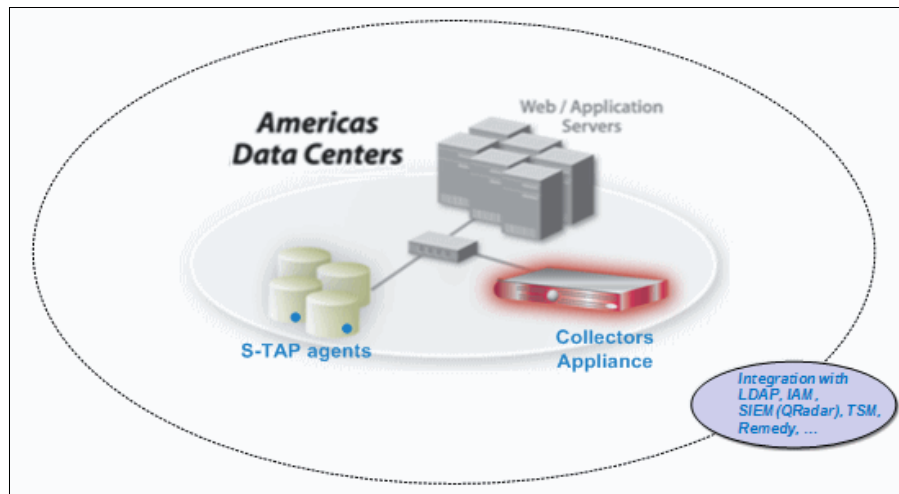


Figure 1-3 Basic architecture

Mid-size architecture

Figure 1-4 represents a mid-size architecture for monitoring numerous databases across data centers. This architecture example consists of four collector appliances and numerous S-TAP agents that are installed on the monitored database servers in each data center. The S-TAP agents are configured to capture and send the relevant database activities to the Guardium collectors for analysis, parsing, and logging. The collectors are configured to aggregate activities that are monitored to an aggregator appliance for central reporting. In this example, the aggregator appliance is also serving as the central management appliance for the solution that enables federated management capabilities, such as Access Management, patching, and metadata repository.

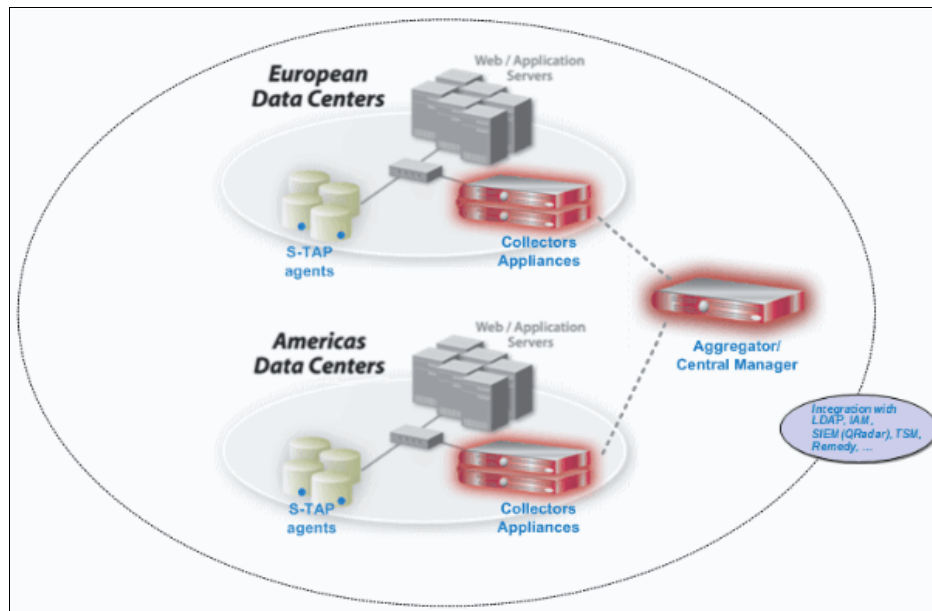


Figure 1-4 Mid-size architecture

Enterprise architecture

Figure 1-5 on page 11 represents an enterprise architecture for monitoring numerous databases across multiple data centers and continents. This architecture example consists of many collector appliances and numerous S-TAP agents that are installed on mainframe and distributed database servers across data centers. The S-TAP agents are configured to capture and send the relevant database activities to the Guardium collectors for analysis, parsing, and logging. The collectors are configured to aggregate activities that are monitored to the respective aggregator appliance for central reporting. A dedicated Central Manager appliance provides federated management capabilities, such as Access Management, patching, and metadata repository.

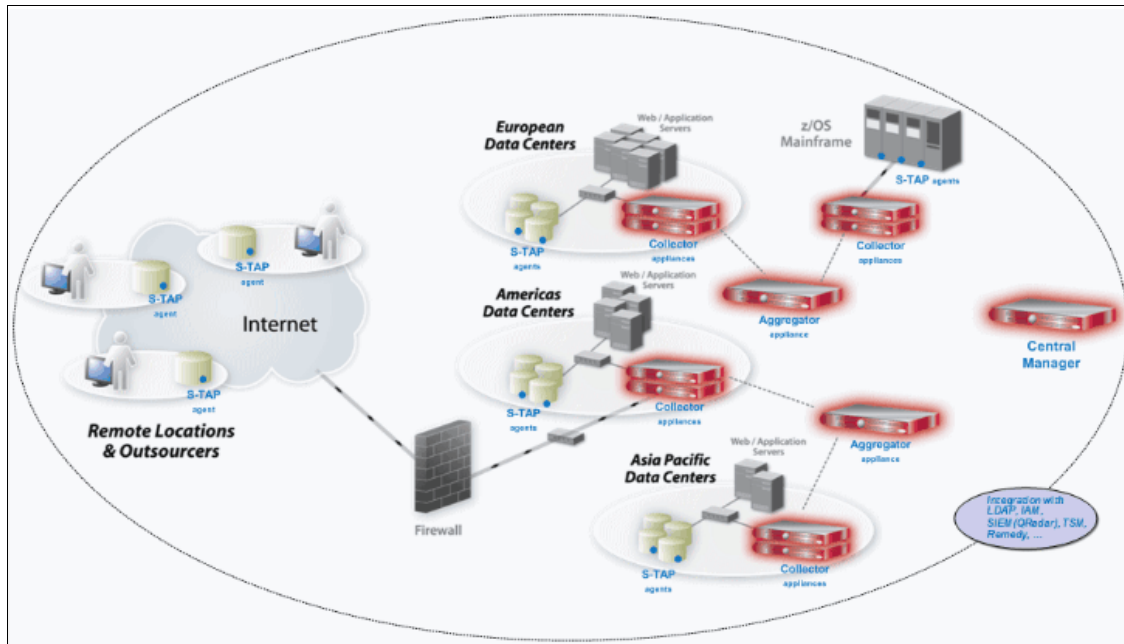


Figure 1-5 Enterprise architecture

Note: The Guardium architecture is scalable and flexible. Scaling the solution to support more monitoring capacity for existing environments or more environments can be achieved easily.

1.3.4 S-TAP architecture options

There are various Guardium S-TAP agent configuration options that can affect the overall architecture of the solution. The S-TAP can be configured to relay captured information to one collector, two or more collectors, or load balance that is captured data between multiple collectors by Guardium or through third-party load balancer.

The following common configuration options are relevant for understanding the overall architecture of the solution:

► Basic S-TAP configuration option

In this configuration (as shown in Figure 1-6), the S-TAP is configured to send traffic to one collector only. *Traffic* includes all of the relevant activity (access and results) that the S-TAP observes between the client (application, user, and so on) and the database. There is no contingency or failover that is configured with this option. Therefore, in cases of long network latency or other S-TAP collector connectivity issues, you see interruptions in monitoring. This configuration is feasible for customers that do not have more than one appliance in their environment or who determine that there is no need to have a contingency in place to support an S-TAP failover configuration.




Figure 1-6 Basic S-TAP configuration option

► Failover S-TAP configuration option

In this configuration (as shown in Figure 1-7 on page 13), the S-TAP is configured to register with multiple collectors but sends traffic only to one collector at a time. S-TAP in this configuration sends all of its traffic to one collector, unless it encounters connectivity issues to that collector that triggers a failover to a secondary collector as configured. This is the most widely used S-TAP configuration to date.




Figure 1-7 Failover S-TAP configuration option

► Load Balancing S-TAP configuration option

In this configuration (as shown in Figure 1-8 on page 14), the S-TAP is configured to send traffic to multiple collectors simultaneously while balancing the load across multiple collectors. This configuration is not widely used but might be considered in cases where there is a need to split traffic from one S-TAP across multiple collectors.

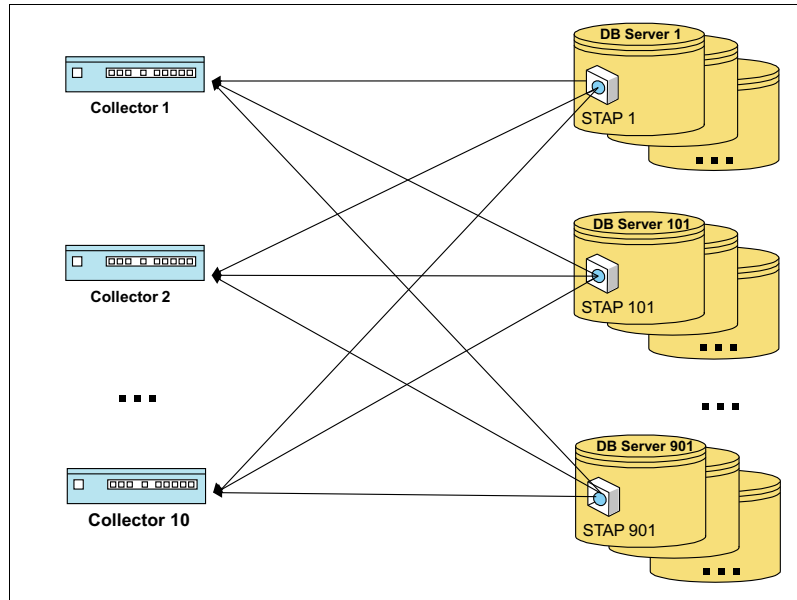


Figure 1-8 Load balancing S-TAP configuration option

► S-TAP grid configuration option

In this configuration (Figure 1-9 on page 15), the S-TAP is configured to register with multiple collectors but sends traffic only to one collector at a time. S-TAP in this configuration sends all of its traffic to one collector, unless it encounters connectivity issues that do not allow it to continue sending activity to this collector.




Figure 1-9 S-TAP Grid configuration option

Note: For more information about the S-TAP configuration, see Chapter 3, “Installation and configuration” on page 43.

1.4 Integration with IT infrastructure

The InfoSphere Guardium solution interacts with numerous applications, IT infrastructures, and products. Figure 1-10 on page 16 shows the following integration points:

- ▶ SIEM: Integration for alerting and logging purposes.
- ▶ SNMP Dashboards: Integration for polling the performance of the solution appliances.
- ▶ Change Ticketing System: Integration for consolidating change control tickets with activity that is monitored and logged on the InfoSphere Guardium solution.
- ▶ Vulnerability Standards: The solution follows industry standards for vulnerability assessments, such as Center for Internet Security (CIS) benchmarks and Security Technical Implementation Guides (STIG).

- ▶ Security Management Platforms: The solution delivers content that is identified to other security platforms.
- ▶ Application Servers: The solution delivers capabilities to identify user access for certain three-tier architectures.
- ▶ Software Deployments: The solution supports the native distribution of agents software packages.
- ▶ Long Term Storage: The solution integrates with industry long-term storage solutions.
- ▶ Data Classification and Leak Protection: The solution includes support for identifying and classifying industry-standard sensitive information patterns.
- ▶ Authentication: The solution supports integration of user access authentication with various industry standards.
- ▶ Directory Services: The solution integrates information from Directory Services, such as Active Directory and LDAP.

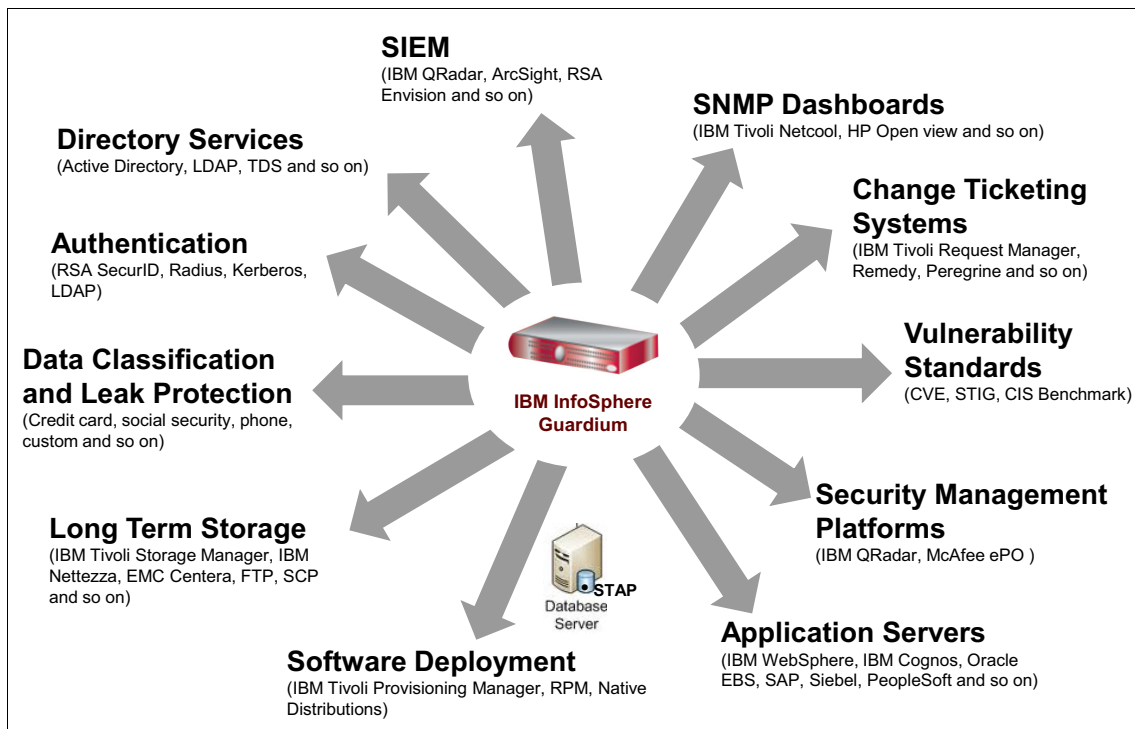


Figure 1-10 Guardium integration with IT Infrastructure

For more information about InfoSphere Guardium integration with IBM products, see Chapter 11, “Integration with other IBM products” on page 391.

1.5 Supported platforms

For the latest supported level of database management system and operating system, see the InfoSphere Guardium system requirements, which are available at this website:

<http://www.ibm.com/support/docview.wss?&uid=swg27039049>



Implementation planning

This chapter describes the preferable process for planning an InfoSphere Guardium implementation. We review important considerations that must be taken into account when you are planning and preparing for an implementation. The planning process is critical for the success of the implementation and should not be skipped or taken for granted.

This chapter provides insight and experience that is based on many successful implementations. The implementation process is described in greater detail in later chapters of this book.

When you are preparing to implement Guardium, it is important to start with a comprehensive review of the following factors:

- ▶ Guardium product components
- ▶ Your auditing goals
- ▶ Databases in scope
- ▶ Deployment time lines

This chapter includes the following topics:

- ▶ Knowing product deployment types
- ▶ Sizing and topology considerations
- ▶ Contingency and design considerations
- ▶ Implementation approach
- ▶ Implementation schedule
- ▶ Roles and responsibilities
- ▶ Installation and configuration sessions
- ▶ Future growth considerations

2.1 Knowing product deployment types

To simplify the description of various product components, we categorize the components into two main product deployment types: Database Activity Monitoring (DAM) and Vulnerability Assessment (VA). These main types are further split into Basic and Advanced, as shown in Figure 2-1.

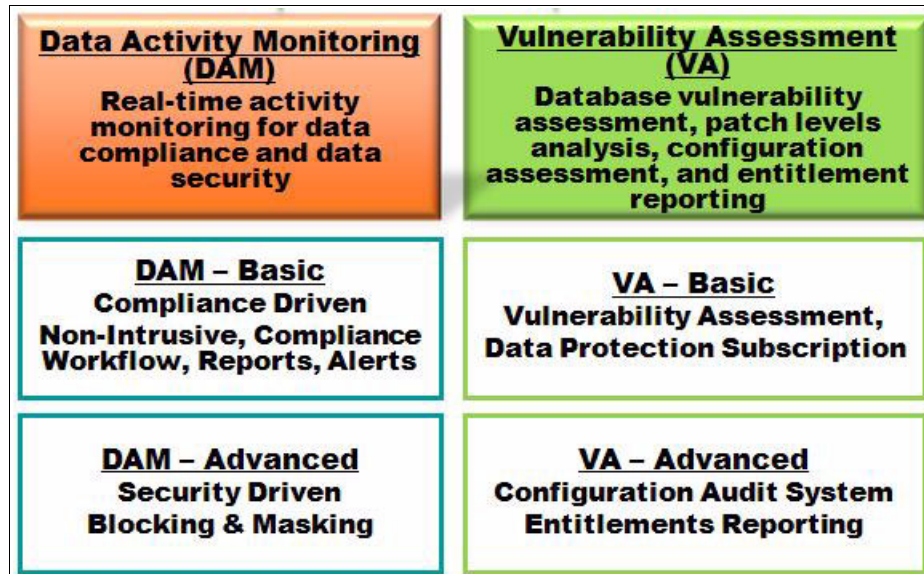


Figure 2-1 Product deployment types

The (DAM) type focuses on monitoring, reporting, and real-time alerting of all access and extrusion activities that are observed. The advanced database activity monitoring adds security-driven data level access control (DLAC) components to the mix (that is, blocking, and masking functionality).

The basic vulnerability assessment (VA) focuses on running vulnerability assessment processes against databases to report about the level of security of the databases. The advanced vulnerability assessment adds more security-driven product components, such as Configuration Audit Systems (CAS) and Entitlement Reporting. CAS probes configuration files, directories, and other database external critical components and alerts on changes that might affect the security and integrity of the databases. Entitlement reporting reports and follows changes on database user account entitlement on the various databases.

It is important to review and understand these components before you start the implementation. In the following sections of this chapter, we describe implementation considerations that are relevant to these deployment types.

Most deployments are DAM deployments and as such should follow an implementation methodology that focuses first on deploying basic Database Activity Monitoring components, followed by the various other functionality components in scope. By using this methodology, you can focus on the installation and configuration of various solution components and gain visibility of what is occurring on the database before you jump into complimentary solution components that are important but can derail the entire deployment effort if done prematurely. Getting the solution installed and operating is the most important first step.

Starting with DAM deployment first should also apply to customers that are interested in DAM and Vulnerability Assessment deployments, unless there is a compelling need to address the VA deployment first or in parallel. DAM takes longer to deploy, but it addresses the installation of the appliances and agents that are also required by the VA. After the solution-installed components are in place, you can proceed in parallel with DAM and VA, following the flow in each of the deployments that calls for the basic components first, then followed by the advanced product components.

Figure 2-2 shows the main product functionality components of the solution.

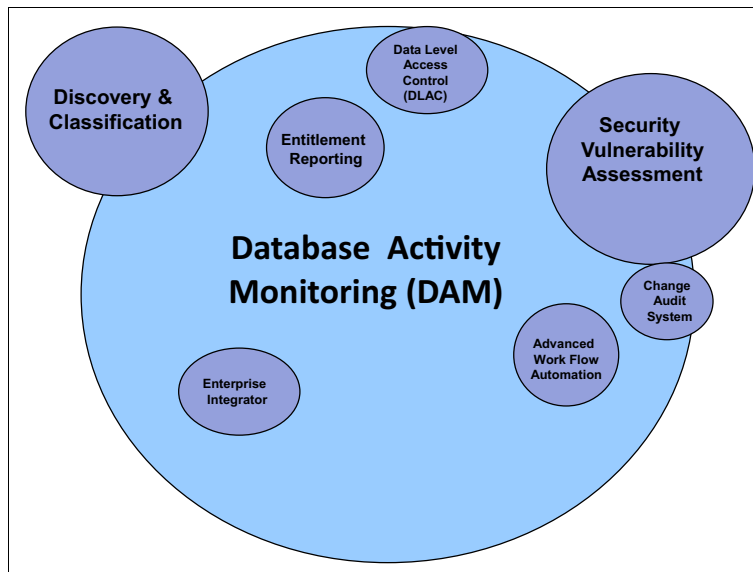


Figure 2-2 Product functionality components

2.2 Sizing and topology considerations

There is no mathematical equation to size a Guardium solution. However, there are important basic guidelines that must be considered to assess sizing needs. This section provides these guidelines. Take these guidelines into account with an understanding that there are also many unknowns in any deployment and start planning your sizing.

When you are sizing a Guardium solution, there is a need to make experienced assumptions that can translate to a deployment topology and sizing of the solution. We review the more important aspects of these assumptions.

All implementation should include recurring reviews of solution performance that might result in tweaking the solution by adjusting the various sizing aspects of the solution. There are many considerations that can be taken into account when you are sizing for the various functionality components. This section describes the relevant considerations that can most affect your sizing decisions.

2.2.1 Audit level

To determine the amount of data the solution is processing and logging, you must understand your general audit requirements that then are translated to the audit levels. Audit levels are used as the guidelines for the amounts of data that is expected to be processed.

Keep in mind that before you implement the solution, there is practical method of knowing how much information you are processing. All applications are different, provide different services to their users, and differ in size and scope. You know only the exact volumes after you implement the solution and have visibility into the volumes of information that are generated by your applications and users on the databases you are configured to monitor.

The following audit levels are important to discuss when you are planning an implementation. Audit level decisions should be taken into account when you are forming your monitoring policy:

- ▶ Privileged user audit

Audit-only specific users and ignore all other connections; the audited users should be a finite list of non-application users (meaning, real people and not application traffic). In this mode, S-TAP filters many of the sessions and only a small subset of the overall traffic is sent to the Guardium appliance (filtering is done on the session level by S-TAP).

Note: Traffic includes all of the relevant activity (access and results) that the S-TAP observes between the client (application, user, and so on) and the database.

▶ Sensitive object audit

This is also known as selective audit that audits only specific database activity, a finite list of sensitive objects, and a finite list of SQL commands (for example, only DDL commands). In this mode, S-TAP sends all of the traffic to the collector and the collector inspects all SQL statements and determines whether it is relevant.

▶ Comprehensive audit

Audit and log everything with the standard granularity (one hour). In this mode, the “Log Full Details” should be used selectively on a subset of data to avoid overloading the traffic.

Note: Comprehensive audit with values, extrusion, or both is the most comprehensive logging mode.

The following sections provide more insight into decisions on how to calculate the number of collectors that are required to monitor your environment with the consideration of the audit levels.

2.2.2 Database to collector sizing

In this section, we describe the basic parameters that are considered when you assess how many collectors are required to monitor your database inventory.

The number of collectors that are required to support monitoring of the databases in scope is roughly based on the following factors:

- ▶ The required monitoring level (audit mode)
- ▶ The type of physical or virtual collector appliance
- ▶ The capacity of the database server, as measured by its PVU (or VU for the z/OS platform)

Where there is not enough information to assess the collectors needed based on these factors, consider a ratio of ten (10) database servers per collector appliance as a good starting point. Optimization based on actual traffic can be done at a later time.

For more information, see 3.4.2, “Collector sizing” on page 51.

2.2.3 Collector to aggregator ratio

There is no fix formula that can be applied to calculate how many aggregators to deploy. The following issues should be understood and considered to help have a good start:

- ▶ Monitoring type

The type of monitoring (privileged users, sensitive objects, and comprehensive monitoring) is defined based on your auditing requirements. The selected monitor type determines the data volume to be logged on the collectors. This data must be aggregated from the collector onto the aggregator. Though you can decide the type of monitoring that you need, there is no tool to calculate how much data each monitoring level generates. After you start monitoring the environment, you can find the number of activities that your users and applications generate. You can then determine the setting granularity for logging the data.

- ▶ Online retention on aggregators

Data that is captured by the collectors and sent to aggregators is stored for a period that is consistent with the reporting needs. The retention requirements differ from customer to customer and sometimes, even between various aggregators. In general, the retention period is determined by your audit reporting requirements. For example, if you are running monthly reports, you must retain at least 30 days of online data. Keeping 30 days of data on your aggregators means that you must have enough space to store all the data.

Remember that you do not want to use up the aggregator disk space at any time. The sum of all of the data that the collectors are sending to the aggregator and the retention period that you must keep this data on the aggregator determine your ratio. The more data and the longer the retention period on the aggregator, the smaller the ration of collectors to aggregators, as shown in the following crude formula:

$$\text{Number of collectors} * \text{GByte per collector per day} * \text{days required for retention} < \text{aggregator database size}$$

When you start an implementation, you do not know the data volumes that you are logging on each of the collectors. You know this only after you start logging and establish your monitoring policy.

- ▶ Internal company needs

Other factors that are specific to a company that can affect the collector to aggregator ratio include the requirement to separate aggregation that is based on security considerations (that is, application data logged), the number of data centers, the data center locations, and network availability. Usually, the data center considerations do not factor in; however, you should discuss all aspects of your environment, including the architecture and network.

As you might conclude, it is difficult to factor in all the considerations and make an informed decision before you start your implementation because some of the factors are unknown at this stage of the implementation.

Based on our experiences, we suggest that eight collectors to one aggregator is a good starting number for planning an initial implementation. This ratio is clearly not a mathematical result but can be considered a safe ratio for addressing most of implementations of small to mid-size deployments with Sensitive Object audit levels.

When you are setting up final monitoring details, learning activity patterns, and seeing volumes being logged, you can adjust the retention periods or the collector-to-aggregator ratio. This process is expected in every implementation and must be considered when you are planning an implementation.

Note: Though calculating the exact size required during planning and initial implementation is not always possible, you can adjust sizing factors later as follows:

- ▶ Increase the aggregators database size or disk allocation (only possible when building the appliance)
- ▶ Add aggregators to an already existing architecture
- ▶ Change retention periods on aggregators anytime

For more information about Aggregator sizing, see 3.4.3, “Aggregator sizing” on page 52.

2.2.4 Central Manager ratio

The Central Manager is the unit that plays a pivotal role in managing the appliances that are registered to it. Most Guardium implementations have one Central Manager managing multiple appliances (collectors and aggregators). One Central Manager can manage hundreds of collectors across multiple data centers and geographies. However, there are various reasons that why some customers might opt to have more than one Central Manager, including following reasons:

- ▶ Territorial considerations

If you must monitor database servers across multiple countries and one of these countries has laws that require a physical separation to prevent any possible access to its appliances from out of border users, you need more than one Central Managers. You can build a separate Guardium environment on that territory that is managed by a dedicated Central Manager.

- ▶ Lack of network

If you do not have network connectivity and your appliances cannot communicate with the Central Manager, you must allocate another Central Manager to manage appliances that have no network connectivity to the primary Central Manager.

- ▶ Business considerations

If the business security policies of your company dictate that a subset of the environment must be managed and monitored as a separate entity, you might need a separate Central Manager to manage the appliances in that environment.

Note: These are only a few examples of why you might need more than one Central Manager. There are other methods to meet the separation-of-duty requirements. In a large federated environment, to provide separate access to information by users, you can use the ready-to-use data level control functionality that controls what each user of the solution is allowed to see that is based on the user and data profiling.

2.3 Contingency and design considerations

In this section, we describe s the contingency and design considerations for the deployment of the Guardium solution. The intent is to provide an understanding of the important things to consider and guidance to help make educated choices.

2.3.1 Appliance location

An understanding of the best place to locate the various appliances is paramount to building your respective Guardium architecture, including the following issues:

- ▶ Collector appliance location

The Collector appliance performs the real-time analysis of the observed database activities it receives from the S-TAP agent that is installed on the database server. These activities include parsing, analysis, logging, and real-time alerting. The Collector must be as close to the monitored database server as possible. Typically, it is the best to have the collector and the database server that the Collector monitors in the same data center to use the LAN speed. If the collector and the monitored database server cannot be in the same data center, place the Collector in the location that has the most robust network connectivity to the data center where the database server is.

Note: The ability of the S-TAP agent to relay (communicate) the activity it captures to the Collector depends on the performance of the network. When the network is not performing or the network bandwidth is not sufficient to sustain the need, this real-time communication data is lost.

- ▶ Aggregator appliance location

The Aggregator appliance does not perform any real-time analysis. The purpose of the aggregator is to run scheduled jobs mining the information it receives from the collectors for reporting. Most of the processes that are running on the aggregator are scheduled. In most cases, the aggregators run the audit jobs that provide the results that are distributed for review. Communication between the collectors and the aggregators happens once a day when the collector's daily logged content is sent to the aggregator as part of a scheduled process. The content that is sent is not network-intensive or time-sensitive. Therefore, the aggregator can be placed any location where there is network connectivity with the collectors. However, it does make sense to co-locate the aggregators with the collectors if there are data centers with multiple collectors that report to a particular aggregator. If the network between the people that must access the aggregator and review the content is best achieved with the aggregators that are in another data center closer to the users, you might consider aggregator location there as well. Most customers locate the aggregators close to the largest pool of collectors that are reporting to it.

▶ Central Manager location

A Central Manager has a pivotal role in managing the appliances that are registered to it with the following capabilities:

- Managing user access to all the appliances
- Storing all of the metadata that is used by all appliances in the solution
- Distributing policy definitions to all collectors
- Centrally managing patch distribution and installations across appliances
- Centrally distributing configuration to all appliances
- Allowing for enterprise reporting capabilities

In most cases, the Central Manager appliance is set up as a stand-alone, dedicated appliance. The Central Manager appliance can also be configured as a centrally managed aggregator appliance. The Central Manager can be placed at anywhere if there is adequate network speed. Although a robust network between the solution appliances and the Central Manager is not critical, network latency can affect the performance of the respective activities that were described in this section.

2.3.2 Appliance configuration options

In this section, we describe the following configuration options that you must consider when you are planning your appliance installation and configuration:

▶ Hardware versus virtual appliance

Guardium appliances can be built as virtual appliances or physical appliances. Most customers determine one option or a combination of options that work best in their environment that is based on their preference and standards. What is important is that regardless of the option that is chosen, the solution works. Pros and cons for the two options are in line with the general industry arguments for the two options; that is, dedicated hardware versus virtual solutions.

▶ Management port configuration

When an appliance is configured, you assign an IP and specify the Ethernet port configuration of the appliance, including the following options:

- Single port (Single IP): In this configuration, your appliance has a management Ethernet port that is configured with an IP and all access to and from the appliance is driven through one IP on an Ethernet port.

- Dual port (Dual IP): In this configuration, you dedicate two IPs to each appliance and assign each IP to a dedicated Ethernet port. This configuration often is used when there is a need to separate access to the appliance for management and data monitoring requirements. This configuration also can support agent communicate with the collector and user management access to the appliance over two separate networks. This option requires separate cabling to the appropriate Ethernet ports on the appliance.
- High availability (Port bonding)

In this option, the appliance is configured with one IP; however, two Ethernet ports are configured (bonded) to accept the traffic. This configuration often is used for hardware appliances as a contingency plan to take care of a possible Network Interface Controller (NIC) card failures.

Note: There are more port configuration items to consider than the listed items when you are determining which port configuration makes most sense for your environment. Therefore, we recommend that you discuss this issue with your network administrators. The important point here is to plan for port configuration up front so that you can prepare for things, such as cabling (hardware appliances), IP addresses, and DNS entry configurations.

- ▶ Redundant power supply (hardware appliance)

Hardware appliances have redundant power supply that requires two power sources when racking and connecting the appliance. It is important to indicate this requirement to the data center staff so that they can prepare accordingly.

- ▶ Back up and archive options (central management configuration distribution)

Each of the appliances has a repository that holds the logged monitoring data. This repository is limited in size and can vary across appliances, but eventually runs out. For this reason, you can use backup and archive on each appliance to offload data to an external backup facility.

What is important is to determine and plan up front where to back up your data. The Guardium solution supports various backup options and you can determine the option that works best for you. Neglecting backup planning can cause problems for the project and company retention policies. Plan ahead for backup and archive needs. When external backup storage is available, you can configure scheduled processes on each of the appliances to offload data to storage. In an environment with a Central Manager, you manage distribution of the backup and archive configurations onto each appliance directly from the Central Manager.

For more information, see 3.9.3, “Data management” on page 87.

- ▶ Patching appliances with latest Guardium Patch Update (GPU)

Whether you receive pre-built hardware appliances from IBM or build the physical or virtual appliances on your own, you should apply the latest patches to each appliance as part of the deployment. Patches are available on IBM Fix Central, which is available at this website:

<http://www.ibm.com/support/fixcentral/>

In an environment with a Central Manager, you can perform and schedule patching of all appliances that are registered to the Central Manager from the Central Manager directly.

2.3.3 S-TAP agent contingency and configuration

In this section, we describe the various configuration options that you must consider when you are planning your S-TAP agent deployment and configuration.

General S-TAP configuration options

This section describes the general S-TAP configuration options that should be considered when you are planning a Guardium solution. The configuration options that are described here are only a few general configuration options that you must know or apply as part of the S-TAP deployment. The following important options must be discussed during installation and configuration planning meetings:

- ▶ Data capture types

S-TAP agents can capture or exclude monitoring various types of databases traffic, including local and network traffic, such as TCP, BEQUETH, and Named Pipes. If there is any need to exclude monitoring a specific protocol or access to the database from a network segment or IP, discuss this requirement up front as it might require a different configuration that can exclude S-TAP from capturing traffic from a certain IP or protocol. You can easily modify these settings later, if needed.

- ▶ Cluster aware

S-TAP agents can be configured to support active/passive database clusters where the databases are not available or not mounted on the passive node until the failover occurs. The Guardium S-TAP agents can support this type of cluster. Configure correctly the S-TAP on the passive node so that monitoring the respective database starts automatically when the failover occurs. It is important to identify and discuss this issue in your deployment planning sessions.

► Prevention (Data-level access control)

S-TAP agents can be configured to support blocking activities or ending connections to the database, which is referred to as *data access level* control. If your implementation requires this type of capability, you must configure the S-TAP agents accordingly. It is not advised to start any implementation with blocking configured. However, it is important to discuss and identify up front whether this capability is in scope and which database servers are the target for blocking.

► Encryption (Transport Layer Security)

S-TAP agents can be configured to communicate over the network to collectors in an encrypted (TLS) manner. The default S-TAP configuration is no encryption to avoid any performance impact. When you are discussing this topic, there is a tendency to assume that the encrypted option is better. Before you determine the best choice for your environment, consider the following factors:

- Configuring the S-TAP with TLS required extra encryption time that might affect performance on the database server where the S-TAP agent is installed. The appliance (collector) also requires time to decrypt this traffic.
- If applications and database users are communicating with the database in an unencrypted manner, configuring the S-TAP agent to communicate over the network with encryption does not make your network safer.

Note: Do not assume that you need encryption because it sounds better. Instead, assess what this means to you and whether it makes sense in your environment. If you determine that you want to proceed by using this option, Guardium fully supports this option and there are configuration ramifications that you must plan.

For more information about S-TAP configuration options, see Chapter 3, “Installation and configuration” on page 43, or this website:

<http://pic.dhe.ibm.com/infocenter/igsec/v1/index.jsp>

S-TAP to collector contingency and failover configuration options

When you are planning a deployment of the agents, discuss the various failover and redundancy S-TAP options and determine which make sense for your implementation. You are not required to apply one option across all deployments. You can combine approaches that are based on what makes most sense in the particular environment, as shown in Figure 2-3 on page 33.

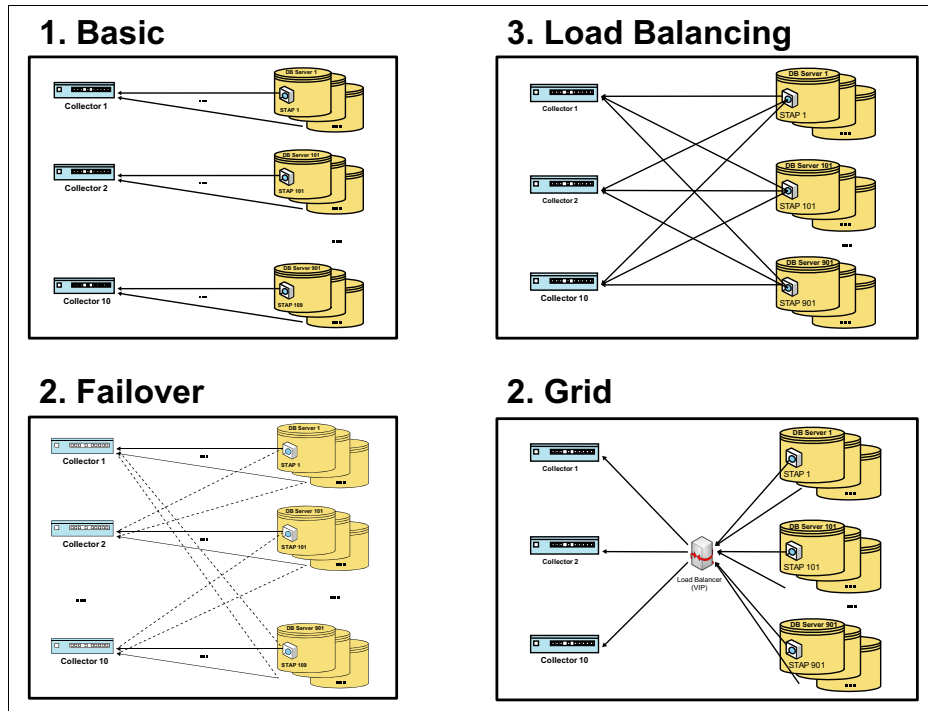


Figure 2-3 S-TAP configuration contingency options

The following failover configuration options for S-TAP are shown in Figure 2-3:

► **Basic**

This basic configuration assumes that S-TAP always sends all of the traffic to one collector and there is no failover configured. Configuring the S-TAP with only one primary collector as recipient of traffic is not suggested unless there is only one collector available and accessible.

► **Failover**

This configuration option assumes that S-TAP sends traffic to one collector (primary) and failover to one or more collectors (secondary, thirdly, and so on) as needed. This is the most common method that is used today.

In this configuration, the S-TAP agents are configured with at least a primary and one secondary collector IP. If the S-TAP agent cannot send the traffic to one collector for various reasons, the S-TAP agent automatically reverts to the other.

You can configure as many failover collectors as you want. We advise not to go beyond three collectors because there is no real reason to do so. When planning for S-TAP failover, keep all collectors in the failover group that is running at about 50% capacity. This is important in the case of failover that a collector can process the additional load.

In this failover configuration, you can use the primary collectors or have separate standby collectors as the failover collectors to multiple S-TAP agents. Besides the traditional method of the use of other active collectors as failover, our field experiences indicate that one standby for every four to five collectors is a good ratio. Choose the combination option that works best with your architecture, database, and data center layout.

- ▶ Load balancing

This configuration option assumes that there is a need to balance the traffic from one S-TAP to multiple collectors. This option might be good when you must monitor all traffic (comprehensive monitoring) of an active database. In the case where the generated traffic is large and housing the data online on a collector for an extended period is required, this method might be your best choice because it performs session-based load balancing across multiple collectors.

An S-TAP can be configured in this manner with as many collectors as needed. This is not a commonly used method because there is seldom a real need to load balance traffic from one database onto multiple collectors.

- ▶ GRID

The GRID option, as called by Guardium, was introduced in the more recent versions of Guardium. It is for customers who have many databases to be monitored and want to use the already existing load balancer technology for Guardium.

This configuration option assumes that S-TAP communicates to the collector through a load balancer, such as f5 and Cisco. With this option, the S-TAP is configured to send traffic to the load balancer. The load balancer forwards the S-TAP traffic to one of the collectors in the pool of collectors.

You also can configure failover between load balancers for continuous monitoring if a load balancer fails.

- ▶ Redundancy

This configuration option assumes that there is a need for the S-TAP to communicate its entire payload to multiple collectors. In this case, the S-TAP is configured with more than one collector (often only two) and communicates the identical content to both. This option is used when there is a need to have full redundancy of the same logged data across multiple collectors or when there is a need to log data and alert on activity at different levels of granularity for various reasons.

2.4 Implementation approach

Implementations must be planned so that each step is focused on achieving a goal consistent with moving the implementation forward. Keeping a focused approach and following your plan is important. A Guardium implementation consists of two main activities:

- ▶ Installation and configuration

The installation and configuration activities focus on getting all the solution components installed, configured, and working with each other. For example, appliances and database agents must be installed and configured first.

- ▶ Monitoring setup and verification

Monitoring setup and verification activities focus on identifying and creating the audit monitoring controls on each of the appliances so that your audit and security needs are met and your environment is protected accordingly. An example is configuring the policy and reporting components that provide your basic audit monitoring needs.

A Guardium solution consists of many applications and options that all work together. To assure that you are not losing focus on the initial goal of setting up a working Guardium solution, we advise you to focus on configuring the options that allow you to reach your basic monitoring goals first. Other functionality components that are not aligned with the basic monitoring should be phased in later stage of the implementation. A good example is Data level access control functionality. This functionality takes the monitoring and protecting your environment to another level. The data level access control allows you to block suspected activities that are inconsistent with your policy rules from performing on your database server. This means that configuring database blocking should be done after activity monitoring is done and you can determine what must be blocked. This type of functionality should be labeled as additional functionality and planned for a latter phase of the implementation. When you are planning a Guardium implementation, you must determine when each functionality component should be implemented.

Always start configuring the functionality of basic database activity monitoring or basic vulnerability assessment. Figure 2-4 on page 36 shows the concept carrying through the rollout of the solution in your respective environments (test, production, and so on).

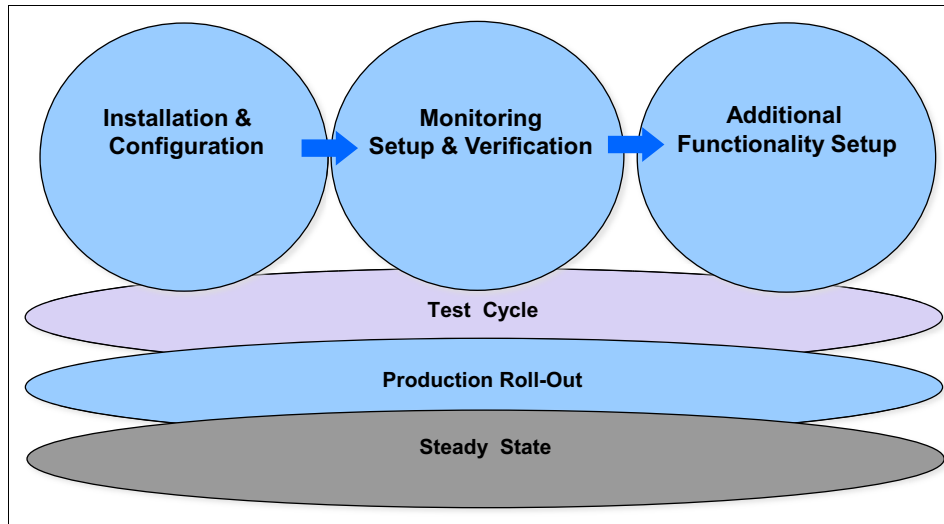


Figure 2-4 Product functionality deployment approach

2.5 Implementation schedule

Building an implementation schedule or project plan that is consistent with the implementation approach that was described in 2.4, “Implementation approach” on page 35 is advised. The two main implementation activities tracks (efforts) are installation and configuration and monitoring setup and verification.

As shown in Figure 2-5 on page 37, the activities in both tracks can run in parallel with an understanding that you cannot deploy and test any of the monitoring setup plan before the appliances and database agents are installed. Therefore, start the implementation with the installation and configuration track first. Determine the schedule of events and activities of this track first and follow up with the monitoring setup and verification track immediately afterward, taking into account the scheduling of the installation in your planning. (See the activities chart that is shown in Figure 2-5 on page 37 as a reference.) In the following section, we describe the roles and responsibilities of the personnel that you must assign to these activities.

Phase	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	
Installation & Configuration Activities																		
Installation & Configuration Planning Session & Documentation																		
Rack/Build, Install & basic configuration of systems																		
Review & Finalize systems configuration(Register units to Central Manager)																		
Guardium technical training session (Guardium Administrators)																		
Install GIM & Deploy S-tap																		
Review Traffic (Session level)																		
Self Monitoring Setup (of Guardium Systems)																		
Monitoring Setup & Verification Activities																		
Monitoring/Report building Workshop																		
Monitoring Planning & Documentation Session																		
Configure & install Monitoring Groups																		
Configure & install Security Policy Rules																		
Configure Monitoring Functionality (Reports, Alerts, Workflow)																		
Configure Additional Functionality (Protection, vulnerability, etc) - <i>IF RELEVANT</i>																		
Production Tweaking & Verification																		
Production Signoff																		
<small>Note: This is not DURATION to implement, just HL schedule of activities and when they occur. Further detailed planning & discussion will determine relevant customer schedule, testing environments etc...</small>																		

Figure 2-5 High-level Implementation schedule example

Note: The Implementation schedule that is shown in Figure 2-5 also shows the following recommended education activities:

- ▶ Technical training as one of the installation and configuration activities. This training provides your technical resources a comprehensive education of the product functionality.
- ▶ Monitoring and report building workshop as one of the monitoring and verification activities. This workshop provides the basic knowledge for the team to be working on monitoring setup to understand the basic monitoring applications.

2.6 Roles and responsibilities

When you are preparing for a typical Guardium DAM implementation, you must understand and plan your resource needs that is based on the project needs. Figure 2-6 on page 38 lists the suggested resource responsibilities aligned with the two major implementation phases: installation and configuration, and monitoring setup.

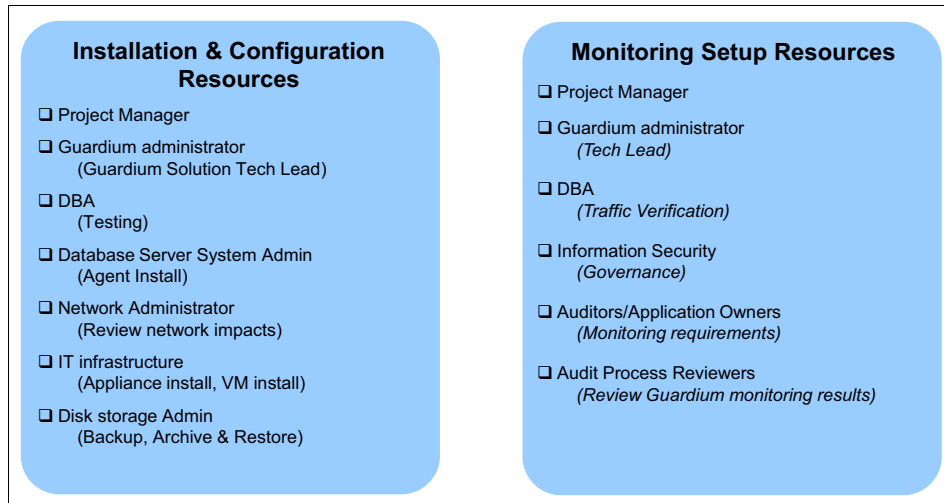


Figure 2-6 Implementation resources

The following resources are needed:

► Project manager

Similar to the most projects, there is a need to assign a project manager who is accountable for ensuring that everyone on the team understands projects goals, knows the roles of the other team members, and performs their role. The specific responsibilities of the Project Manager might vary slightly between companies and projects. However, the following responsibilities are common to all Project Managers:

- Developing the project plan
- Managing the project stakeholders
- Managing the project team
- Managing the project risk
- Managing the project schedule
- Managing the project budget
- Managing the project conflicts

- ▶ **Guardium administrator**

A Guardium administrator is the Guardium solution technical lead. This role is pivotal for the success of the project. It is expected that a person in this role is the focal point for all project-related technical activities and the lead for all installation and configuration activities during the implementation. In the course of the implementation, the Guardium administrator works with the services implementation consultants on all aspects of the deployment and learns as much as possible about the solution that is deployed. All of the technical issues should be reported directly to the Guardium administrator for initial review. If it is determined that contacting product support is required, the Guardium administrator should do this directly or instruct others to do so after reviewing the reported issue.

- ▶ **Database administrator**

The DBA role might be required at various times during the implementation. A DBA typically provides or gathers the information that is required for the database servers to be monitored, including database port and the basic configuration information. During the installation and configuration phases, the DBA might need to be apart from the team to review the database activity that is being monitored as part of the checks that the agent was configured correctly. In post S-TAP agent deployment, a DBA can help in verifying the activity that was captured by Guardium and to assist with identifying sensitive objects in the various databases.

- ▶ **Database server administrator**

The database server administrator (SA) is another important role for agent planning and installation as part of the initial installation and configuration activities. Guardium agents are installed on the database servers under root or root-like accounts that are typically reserved for SA use only. The SA of the servers where the Guardium agents are installed must be educated about the agents installation and general activities and should assess the performance effects of the Guardium agents on the servers for which they are responsible. As such, it is important that the SA representatives are involved in the installation and configuration portion of the implementation.

- ▶ Network administrator

The Guardium solution is installed in the network. Agents are installed on database servers and transfers activities over the network. Therefore, it is always important to include a network administrator role as part of the implementing team. The network administrator provides network-related information to the team, such as bandwidth, connectivity, and latency between data centers. They also provide network configuration information, such as IP, routing, resolver information, and DNS entries for the initial configuration of the Guardium appliances. Network administrators should review performance effects of the communication between the S-TAP agents and the collectors over the network.

- ▶ IT infrastructure

Guardium virtual or hardware appliances must be racked, created, and configured. The role described here assists with these activities and must be part of the initial installation and configuration planning sessions for product deployment.

- ▶ Storage administrator

The Guardium appliances must be backed up. During the initial installation and configuration planning sessions, these backup requirements are discussed. Ultimately, storage should be allocated for appliance backup, which is then configured and tested on all appliances.

- ▶ Information security

During monitoring planning sessions where the monitoring requirements are discussed, input from governance personal is critical to determine monitoring, logging, alerting, and policy needs that then are translated into a monitoring setup on the Guardium solution.

- ▶ Auditors and application owners

During monitoring planning sessions where the monitoring requirements are discussed, input from internal auditors is important to determine monitoring, logging, alerting, and policy needs that then are translated into a monitoring setup on the Guardium solution.

It might also prove necessary to request assistance from application owners to identify sensitive objects in their respective databases. It is probably not necessary to include application personal in the actual monitoring planning sessions (be prepared to involve relevant members only if needed).

- ▶ Audit process reviewers

When the monitoring solution is deployed, audit processes, alerts, and other reviewable outputs are generated by the solution. Someone must review the system monitoring output and take action to resolve incidents that are found, if any. You should also include the leaders of the review tasks in the implementation process and plan knowledge transfer sessions to train the reviewers.

2.7 Installation and configuration sessions

In preparation for deploying the Guardium solution, planning sessions or workshops are held with key stakeholders to review the solution, clarify business and technical requirements, and discuss the implementation activities and personnel.

These sessions are facilitated by IBM or IBM partner and have the following results:

- ▶ Better understanding by the customer of their tasks and responsibilities
- ▶ Deployment plan
- ▶ Project plan

In the planning stage, the following topics must be covered:

- ▶ Data center environments and networks (test, production, and so on)
- ▶ Deployment time lines, milestones, and phases
- ▶ Installation of the appliances (process and prerequisites)
- ▶ Basic configuration of the appliances
- ▶ Installation and configuration of the DAM agents (process and prerequisites)
- ▶ VA setup process and prerequisites
- ▶ Central Management functionality and setup
- ▶ Aggregation process and plan
- ▶ Backup, archiving, and purging process and plan
- ▶ Contingency plan

The following participants and their roles are included in the planning phase:

- ▶ Information security and data security compliance
- ▶ DBA and system administrators
- ▶ Network administration
- ▶ Designated Guardium administrators
- ▶ IBM or IBM Partner to facilitate the workshop

The following inputs and outcomes of this phase are included:

- ▶ Inputs:
 - Database inventory
 - Appliance and licensed modules inventory
- ▶ Outcome:
 - Deployment plan
 - Project plan

2.8 Future growth considerations

The solution is highly scalable and configurable, so adding solution components when needed is not a process that requires much planning. Methods that are used in planning the initial implementation should be used for planning future growth across newly monitored databases. However, there are simple guidelines for system growth that should be considered as part of any deployment.

When you are planning a deployment, always assume that you need more capacity than you think for two reasons: environment growth and activity growth. Because you do not know the monitoring data volumes until you determine the granularity of monitoring and have the visibility of the traffic volumes, plan more appliance capacity than you estimate and try to keep appliances at usage rates that are between 50% - 60%.

After the solution is deployed, the enterprise reports that are available on the central manager provide you the metrics that you can use to assess collectors performance, capacity, and repository available space. This is information that you need to assess whether you have the capacity to add database monitoring to a collector or must plan expansion by adding collectors.

The amount of traffic that is logged and sent from the collectors to the aggregators and the retention needs of the aggregators determine whether you need more aggregators.



Installation and configuration

In this chapter, we describe the installation and configuration process of the Guardium solution, including the components that are used for database activity monitoring (DAM) and database vulnerability assessment (VA).

This chapter includes the following topics:

- ▶ Schedule of Implementation activities
- ▶ Installation and configuration planning
- ▶ Guardium appliance overview
- ▶ Database and appliance inventory
- ▶ Appliance deployment considerations
- ▶ Appliance installation and configuration
- ▶ Agent deployment, installation, and configuration
- ▶ Configure remaining appliances
- ▶ Guardium operations
- ▶ Vulnerability assessment
- ▶ Where to find more help

3.1 Schedule of Implementation activities

The Schedule of Implementation is a high-level list of activities that is adapted to the customer's schedule, wanted solution, and environments (test, stage, and production).

Note: Because of overlap, activities for DAM and VA are listed. However, these features are often installed in separate phases; for example, first DAM, then later VA, or vice versa.

The installation and configuration of InfoSphere Guardium include the following tasks:

1. Planning session for installation and configuration:
 - a. Database and appliance inventory
 - b. Capacity planning
2. Appliance installation:
 - a. Rack or build, installation, and basic configuration of systems
 - b. Review and finalize system configuration and register units
 - c. Guardium technical training session for assigned Guardium administrators
3. DAM agent installation and verification:
 - a. Install the Guardium Installation Manager (GIM) and deploy S-TAP agents
 - b. Create inspection engines
 - c. Verify traffic
4. VA configuration activities:
 - a. Database account creation
 - b. Data source creation
5. Guardium Operations:
 - a. Appliance advanced configuration
 - b. Self-monitoring setup

Figure 3-1 on page 45 is a summary of the installation and configuration activities.

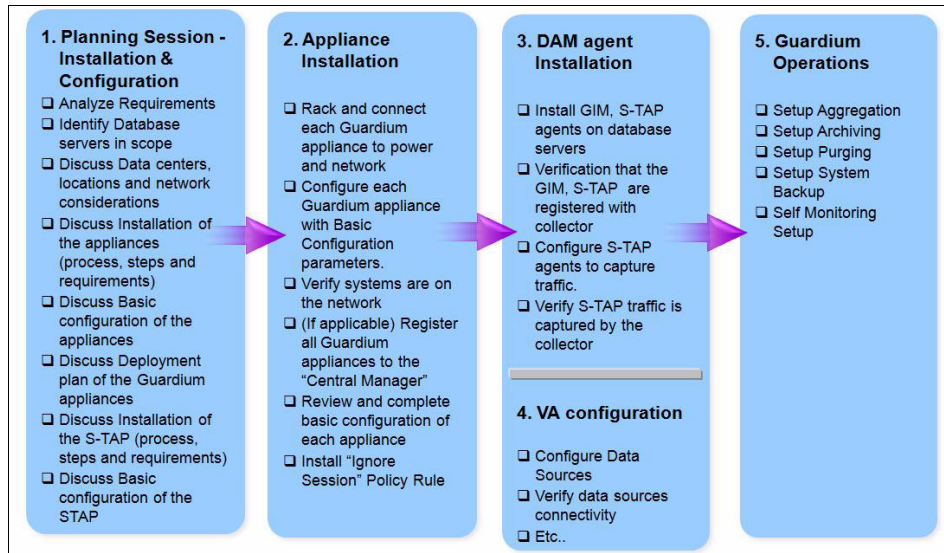


Figure 3-1 Installation and configuration activities

3.2 Installation and configuration planning

In preparation for deploying the Guardium solution, planning sessions or workshops are held with key stakeholders to review the solution, clarify business and technical requirements, and discuss the implementation activities and personnel.

These sessions are facilitated by IBM or an IBM partner, and garner the following results:

- ▶ Better understanding by the customer of their tasks and responsibilities
- ▶ Deployment plan
- ▶ Project plan

In the planning stage, the following topics are covered:

- ▶ Data center environments and networks (test, production, and so on)
- ▶ Deployment timelines, milestones, and phases
- ▶ Installation of the appliances (process and prerequisites)
- ▶ Basic configuration of the appliances
- ▶ Installation and configuration of the DAM agents (process and prerequisites)
- ▶ VA setup process and prerequisites
- ▶ Central Management functionality and setup
- ▶ Aggregation process and plan

- ▶ Backup, archiving, and purging process and plan
- ▶ Contingency plan

The following participants and their roles are part of the planning phase:

- ▶ Information security and data security compliance
- ▶ DBA and system administrators
- ▶ Network administration
- ▶ Designated Guardium administrators
- ▶ IBM or IBM Business Partner to facilitate the workshop

The following inputs and outcomes of this phase are featured:

- ▶ Inputs:
 - Database inventory
 - Appliance and licensed modules inventory
- ▶ Outcome:
 - Deployment plan
 - Project plan

3.3 Guardium appliance overview

The Guardium appliance is offered in two modes: hardware and software (as shown in Figure 3-2), and can be configured in one of two functional types: collector or aggregator (as shown in Figure 3-3 on page 47).

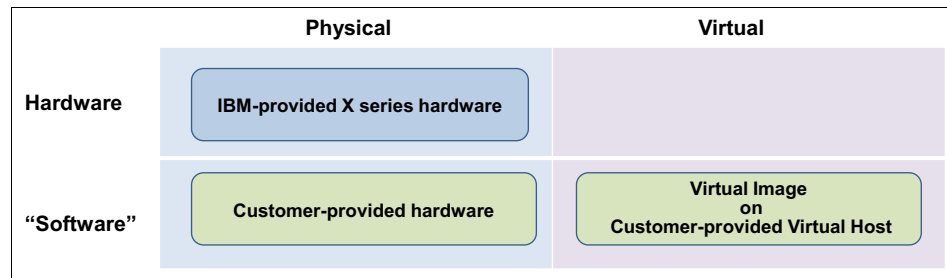


Figure 3-2 Guardium appliance mode




Figure 3-3 Guardium appliance type

3.3.1 Hardware and software appliance modes

The Guardium appliance contains the Guardium application that is integrated with a hardened operating system. An appliance is available preinstalled on IBM xSeries server (which are referred to as a *hardware appliance*) or as a software image (*software appliance*), which is installed into a virtual machine or onto pre-approved hardware that the customer provides.

A single software image is used to install the aggregator or collector. For hardware appliances, the customer requests the mode. For the software appliance, the user selects the mode during the software installation process.

For more information about the technical requirements of the installation of the IBM InfoSphere Guardium V9.1 Software Appliance, see *IBM InfoSphere Guardium V9.1 Software Appliance Technical Requirements*, 7039720. This document also includes a list of certified hardware platforms and is available at this website:

<http://www.ibm.com/support/docview.wss?&uid=swg27039720>

Note: Whether the appliance is hardware (on IBM-provided hardware) or software (on the customer's hardware or virtual host), it is not serviceable or reconfigurable by the customer.

3.3.2 Appliance types

The Guardium appliance is deployed in one of two types: collector or aggregator. An aggregator also can be designated as a Central Manger.

Collector

The collector is the workhorse appliance in the Guardium DAM solution and is used for real-time capture and analysis of the database activity.

The collector receives and processes monitored traffic in real time from the S-TAP agents that are deployed on the database servers. Think of the collector as a transaction processing system, with potentially high disk I/O, where the transactions are the database-user activity that is monitored.

Note: The collector must be network-close (that is, minimal hops) and have LAN speed connectivity to the S-TAPs to reduce network latency. If the collector is built as a software appliance, it should have low latency disk I/O.

Figure 3-4 shows multiple S-TAPs that are connected to a collector.

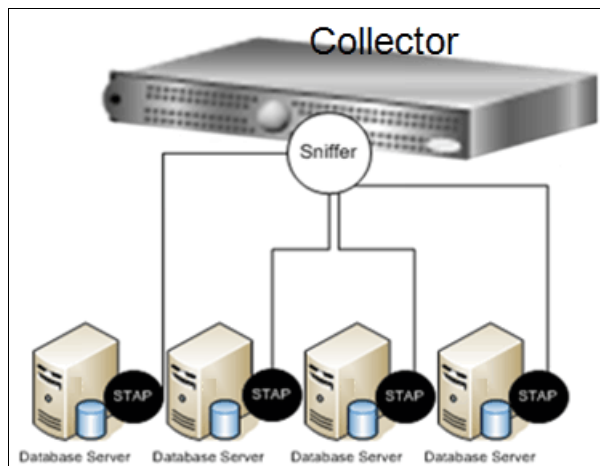


Figure 3-4 Multiple S-TAPs that are connected to a collector

Aggregator

The aggregator appliance is used to offload reporting activity from the collectors, and to provide consolidated reporting from multiple collectors. The aggregator does not collect data from S-TAPs. Instead, it receives the data from the collectors in a nightly batch file.

The aggregator is optional but is recommended if several collectors are deployed.

Although not typical, an aggregator can receive data from other aggregators to provide enterprise-wide report. This is referred to as *second-level aggregation*.

Figure 3-5 on page 49 shows an aggregator that is assigned to multiple collectors.

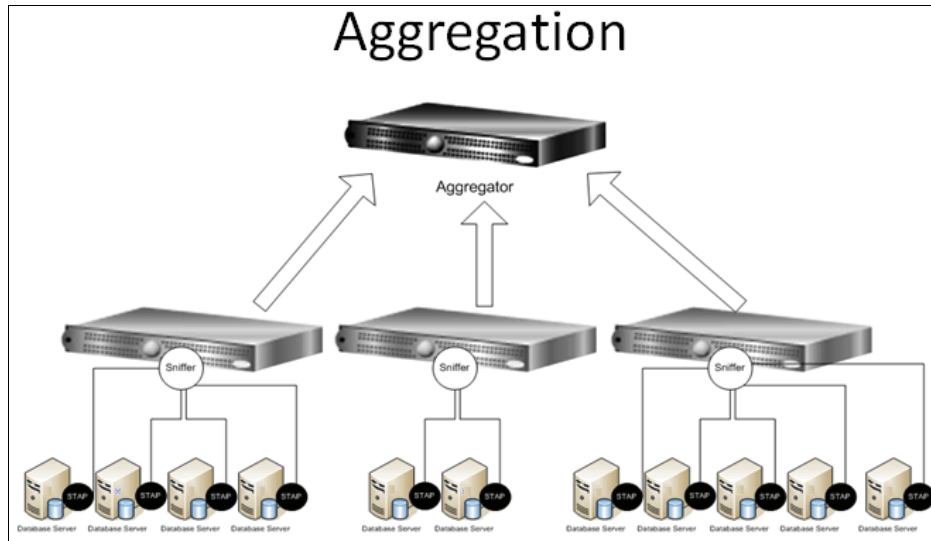


Figure 3-5 Aggregator that is assigned to multiple collectors

Central manager

The Central Manager (CM) is specialized functionality that is enabled on an aggregator appliance. The CM function is used to manage and control multiple Guardium appliances, which is referred to as a *managed environment*. This function provides patch installation, software updates, and the configuration of queries, reports, groups, users, policies, and so on.

There can be only one primary CM in a managed environment. In small environments, the aggregator appliance often serves as an aggregator and the CM. In larger environments, the CM runs on an aggregator that does not perform aggregation; that is, does not receive data from collectors, which is referred to as a dedicated CM.

The CM is optional but recommended if several appliances are deployed.

Note: Although a CM can be placed over the WAN, it should have less than 200 ms network round-trip latency to its managed units.

3.4 Database and appliance inventory

In this section, we describe the sizing and types of Guardium appliances that are required to support the in-scope database inventory.

3.4.1 Database inventory

The database inventory template is a Microsoft Excel document that is available from IBM Guardium Sales or Services. It is used for the following tasks:

- ▶ Allocating the appropriate number of collectors that are required to support a specific number of database servers on a particular subnet.
- ▶ Assign secondary or failover collectors
- ▶ Assign collectors to aggregators for consolidated reporting
- ▶ Recording database/instance configuration parameters that are required for inspection engines
- ▶ Track deployment and configuration status of the Guardium components

Although you can add or remove fields from the template, the following fields are key:

- ▶ Business application such as SAP
- ▶ Location (for example, Datacenter_1)
- ▶ DB Server host name
- ▶ DB Server IP
- ▶ InCluster
- ▶ OS Type (for example, AIX®)
- ▶ OS Version
- ▶ CPU Cores
- ▶ VU per Core*
- ▶ PVU Total (that is, CPU Cores x VU per Core)
- ▶ DB Type (for example, DB2)
- ▶ DB version
- ▶ DB Instance name
- ▶ DB listener port
- ▶ DB install dir
- ▶ Primary Collector name
- ▶ Primary Collector IP
- ▶ Collector Secondary name
- ▶ Collector Secondary IP
- ▶ Aggregator

- ▶ GIM Installed
- ▶ S-TAP Installed
- ▶ IE configured

Value Unit (VU) is an IBM metric that is used to gauge the capacity of the database server. To determine the VUs per core that is based on the database server make and model, see the processor value unit listing that is available at this website:

http://www.ibm.com/software/lotus/passportadvantage/pvu_licensing_for_customers.html

Tip: Some users find it helpful to add a separate worksheet to the database inventory workbook that contains the list and details for the appliances.

3.4.2 Collector sizing

The number of S-TAPs a collector can support is estimated on the following factors:

- ▶ The required monitoring level (that is, audit mode)
- ▶ The type of collector appliance (that is, physical or virtual)
- ▶ The capacity of the database server, as measured by its PVU (or VU for the z/OS platform)

Calculate PVU total

When the database inventory is populated, review the list of database servers and review and update the PVU Total information for each, as described in 3.4, “Database and appliance inventory” on page 50.

Tip: If there are multiple database instances or application per database server (as is often the case), there are multiple records per database server. To resolve this issue, copy the information to a separate worksheet and then remove the database instance columns and application columns so that the data can be filtered to result in a unique record for each database server.

Assign collectors

The next step is to assign a collector to each database server (or S-TAP). This assignment considers the following factors:

- ▶ Collector should be network-close to the S-TAP, specifically:
 - LAN-speed connectivity between S-TAP and collector
 - Minimal network hops to minimize network latency

- ▶ Collector should not exceed its PVU monitoring limit; sum the PVU total for each database server that is assigned to the collector. See Table 3-1 on page 52 for the current limits.

For more information about collector sizing, see *IBM InfoSphere Guardium V9.1 Software Appliance Technical Requirements*, 7039720, which is available at this website:

<http://www.ibm.com/support/docview.wss?&uid=swg27039720>

- ▶ If practical (that is, not conflicting with the prior factors), assign the same business applications to the same collector or set of collectors.

Note: Table 3-1 provides an estimate or starting point for deployment planning. However, it is necessary to periodically review the actual load on the appliances and reallocate S-TAPs to avoid overloading a collector.

Table 3-1 Collector PVU Limit

Monitoring Levels	Physical Collector	Virtual Collector
Distributed (LUW) Collector		
Privileged User and Sensitive Objects	8000	4800
Comprehensive	4000	2400
Mainframe Collector		
Privileged User and Sensitive Objects	220	110
Comprehensive	110	55

Note: The monitoring capacity of a virtual collector often is less than that for a physical collector because of the shared-everything architecture of the VM host.

3.4.3 Aggregator sizing

The number of collectors that an aggregator can support is constrained by the available space in the aggregator's database. This space is affected by the following factors:

- ▶ Online retention period (sometimes referred to as the purge period)
- ▶ Free-space that is kept in reserve for temporarily staging backup files

For example, an aggregator that is receiving files from eight collectors and keeping the data online for 60 days uses more space than if the retention is set to 30 days.

Note: The aggregator should have a minimum of 40% free disk space to provide room to create backup files, which are removed when they are successfully copied off to a designated storage server.

We recommend starting with a ratio of eight collectors to one aggregator.

3.5 Appliance deployment considerations

In this section, we describe the considerations for deploying Guardium appliances.

3.5.1 Managed or stand-alone environment

Guardium appliances often are deployed in one of two configurations: managed or stand-alone. There also is a hybrid configuration in which a managed environment might contain some stand-alone collectors.

Managed configuration

If a CM is enabled, the deployment is referred to as a managed environment (also known as a federated environment). To complete the managed environment, the other appliances, collectors, and aggregators are registered with the CM. Before a unit (appliance) is registered, a common shared secret (password) must be stored on the CM and the to-be registered unit.

All managed (registered) units inherit the license key from the CM. Users, roles, and other definitions are synchronized from the CM across its managed units.

Note: We recommend a dedicated CM if a managed environment contains more than 10 managed units

Stand-alone configuration

The simplest deployment configuration is a single collector that is receiving data from several S-TAPs. This configuration is referred to as a stand-alone because a CM is not used. It is a typical configuration for a test environment.

A few customers deployed up to three stand-alone collectors (with no aggregator) in their production environment.

Hybrid configuration

A hybrid configuration is a combination of managed appliances with stand-alone collectors. These stand-alone collectors can, in turn, upload their data nightly to managed aggregators.

This configuration often is used in environments with a few remote or satellite locations that have low-bandwidth connectivity back to the data center where the CM is located.

3.5.2 Contingency

The appliances can be configured to minimize data and functionality loss through failover or redundancy. They also can be used for scheduled maintenance.

The failover plan is in addition to the recommended periodic backups and daily archiving.

Collector

The following contingency plans can be used for collectors:

- ▶ S-TAP failover

An S-TAP can be configured to fail over (start communicating with) to a secondary or tertiary collector if the primary collector is unreachable. When the primary collector is reachable, the S-TAP reverts to it.

The S-TAP also uses a limited memory buffer (spill file on the z/OS) to temporarily buffer data that is in transit to the collector.

- ▶ S-TAP Mirroring

If a collector fails, the data since the last daily export or archive is lost. To avoid any loss, the S-TAP can be configured to mirror its transmission to two collectors, so each collector receives the same copy of the data.

Note: For virtual appliances, avoid single-point of failures by not placing the primary and secondary collectors on the same host

Aggregator redundancy

Each collector can be configured with a secondary or standby aggregator to which automatically send its daily files should the primary aggregator be unreachable.

However, reports must run on both aggregators to view the data that is collected before and during the failover period.

CM redundancy

A failed or unavailable CM affects interactive use of the Guardium application; users cannot use most of the functions from a managed unit, such as reporting. Also, no application definitions or user changes can be made until the CM is available.

However, the collectors continue to collect and the aggregators continue to aggregate.

A managed aggregator at the same patch level as the CM can be designated as a backup CM. After it is configured, the primary CM automatically copies a backup file to the secondary every hour.

It is a manual process to make the backup the primary CM and involves the use of command-line interface (CLI) commands. After the switch occurs, the managed nodes automatically detect the switch and reconfigure themselves to communicate with the new CM.

3.5.3 Networking

By using the physical appliance as a reference, each appliance has at least four network interfaces. A typical configuration is to configure the first port (eth0) in the network. However, it is possible to have the following supplemental configuration (eth0 is always required):

- ▶ IP bonding (teaming): IP bond eth0 and eth3 for network redundancy
- ▶ Management (secondary) interface: Configure eth3 as a management interface; for example:
 - In the case of the collector, S-TAP traffic communicates over eth0, but user and CM communication occur over eth3.
 - Appliance backup and archive data can be routed over eth3 that is connected to a data network.
 - There is an option to add a static route to facilitate routing through eth3.

Virtual appliances support the configuration of a management or secondary interface. Bonding is best implemented at the VM host layer.

For more information, see the System Configuration section of the “Guardium Administration” chapter of the *Help Book Guardium*. For more information about obtaining the Help Book Guardium, see 3.11.1, “Product documentation” on page 117.

Firewalls

Depending on the Guardium feature that is used, specific network ports are required to be open across a firewall. For more information about ports, protocol, and directionality, see the *Guardium Ports Requirement* document.

3.5.4 Other allocation factors

In addition to the sizing considerations, the following factors must be considered when you are allocating S-TAPs-to-collectors and collectors-to-aggregators:

- ▶ S-TAP failover and tracking complexity
- ▶ Line-of-business reporting

The database inventory template is a useful aid for tracking these decisions.

S-TAPs-to-collectors

The following contingency plan is used for the S-TAPs-to-collectors:

- ▶ S-TAP Failover: Small environment

For a few co-located collectors, each collector can be designated as the failover for another collector. For this scheme to work, the following prerequisites must be met:

- Each collector must have approximately 50% capacity head-room; that is, not maxed to its PVU limit. Otherwise, the failover can result in the secondary collector failing because of overload.
- Each collector should be co-located or network-close.
- Plan to alert-on and correct the failover condition as soon as possible to avoid the failover collector running in a degraded mode for an extended period.

- ▶ S-TAP Failover: Medium and large environment

Although the small environment approach can be used in a medium environment or one where there are a few collectors in different data centers, it can become too complex to plan and track for larger implementations.

Instead, use one standby per a group of X collectors, where X can be four or higher; however, if more than one or two collectors from the group fail at the same time, the standby might fail because of overload.

The failover collector should export its data to the same aggregator as the group of primary collectors.

Note: If more than one or two collectors from the group fail, the standby might fail because of overload.

3.5.5 Start with a test environment

A permanent Guardium test environment is recommended. Similar to other test platforms, it is used for hands-on experience, re-creating issues, and testing new components or configurations before it is deployed to production.

This environment should consist of Guardium appliances and test database servers. For an enterprise deployment, it should include the following components:

- ▶ A test CM or Aggregator
- ▶ At least two managed collectors: one for distributed traffic and the other for mainframe, if applicable.

It should also consist of one of each type of database server (or representative sample) to be monitored in production with similar OS-and-DBMS configuration, network configuration; for example, intranet and DMZ, and clustering or zoned configuration.

Note: It is not necessary to have every version and combination of OS-DBMS represented. Instead, use the more common or business-critical platforms to comprise your test environment.

3.5.6 Licensing

Guardium product features or entitlements are enabled through specific product keys or licenses that are installed through the application interface. The following types of keys are available:

- ▶ Base key (also known as reset key)
- ▶ Append key, which is appended to base

A base key with at least one append key must be installed to enable Guardium features.

Base key

There is a base key for each type of appliance; that is, Collector and Aggregator.

Append key

An append key consists of five keys and requires a base key to be installed. Multiple append keys can be applied; for example, Central Management plus DAM Standard plus VA Advanced. The append keys are available for the following features:

- ▶ Central Management
- ▶ DAM Standard
- ▶ DAM Advanced (also includes all features in DAM Standard)
- ▶ VA Standard
- ▶ VA Advanced (also includes all features in VA Standard)

3.6 Appliance installation and configuration

In this section, we describe the steps to install and configure the Guardium appliances.

3.6.1 Rack or build appliance

The first step in deploying your Guardium solution is to rack (physical) or build (software) the appliances, including the following installation tasks:

- ▶ Unpack, inventory, and rack the appliances.
- ▶ A minimum of one network drop is required.
- ▶ Attach the power and network cables.
- ▶ Power up the appliances and verify that there are no POST errors.
- ▶ (Optional) Attach to your existing remote KVM solution (if any) for console access.

For virtual or software appliances, use the following instructions in the *IBM InfoSphere Guardium Software Appliance Installation Guide*:

- ▶ Configure the virtual guests (or hardware servers)
- ▶ Install the Guardium application

The guide is available at this website:

http://pic.dhe.ibm.com/infocenter/igsec/v1/index.jsp?topic=%2Fcom.ibm.guardium.software.app.install.doc%2FtopicsV90%2Fsoftware_appliance_installation_guide.html

Note: If you are building multiple VMs, you can clone the first appliance and use it to build the remaining appliances of the same mode.

At this stage, the appliances are powered, cabled to the network, and ready to be configured.

3.6.2 Configuring the appliance

Although the appliances are cabled to the network, they are not yet reachable because their network is not configured. For the initial configuration or basic configuration, the following console access to each appliance is required:

- ▶ For physical machines, you provide a monitor, keyboard, and mouse (some data centers have these on a “crash-cart”). You also can use the remote KVM solution if available.
- ▶ For virtual machines, the console is accessed by using the host’s hypervisor management console.

The appliance configuration steps are grouped in the following three phases to facilitate the scheduling of large deployments and to enable deploying agents in parallel with the appliance configuration:

- ▶ Basic configuration
- ▶ Advanced configuration I
- ▶ Advanced configuration II

Preparation

Before you perform the configuration collect process, prepare the following information for each appliance:

- ▶ CLI default password ¹
- ▶ CLI new password (on first login, you must change the default password)
- ▶ Appliance primary IP address
- ▶ Appliance primary network mask

¹ The default CLI password for IBM-provided hardware appliances is obtained from IBM Support or IBM Sales. For VM appliances, the CLI password is set by the person who is building the VM image.

- ▶ Default network route, or gateway, IP address
- ▶ DNS IP address (up to three can be used)
- ▶ Appliance host name
- ▶ Appliance domain
- ▶ NTP server name or IP address (up to three can be configured)
- ▶ Create a shared secret (that is, password that is used for creating a managed environment)
- ▶ The Guardium license key for the central manager, or the key for each unit in a non-managed environment

Note: Pre-configured hardware that is shipped from IBM has the license key preinstalled.

Basic configuration: Installing the appliance into the network

The basic configuration is performed on each appliance, which configures the host name, domain, and network parameters to allow each appliance to be reachable in the network.

Tip: For large implementations, the basic configuration can be done in batches or phases, in which case, always start with the Central Manager appliance followed by at least one collector appliance.

This approach allows the deployment of some Guardium agents in parallel with the appliance rollout.

Basic configuration steps

By using the information that was collected, login to each appliance and perform the following steps (the appliance is rebooted to complete this process):

1. On the console at the login prompt, enter `cli` and then enter the default password.
2. Change the default password when you are prompted to do so.

Tip: For a list of available CLI commands; enter `comm` at the CLI prompt to see a list of all available CLI commands.

You also can enter `comm <string>` to see a list of all commands that contain that string.

3. Set the primary IP address for eth0, as shown in the following example:

```
store network interface ip <ip_address>  
store network interface mask <network_mask>
```

4. Set the network gateway for the default route, as shown in the following example:

```
store network routes def <default_router_ip>
```

5. Set the DNS IP addresses (the first is required, the others are optional), as shown in the following example:

```
store network resolver 1 <dns_server_ip_1>  
store network resolver 2 <dns_server_ip_2>  
store network resolver 3 <dns_server_ip_3>
```

6. Set the host and domain names, as shown in the following example:

```
store system hostname <host_name>  
store system domain <domain_name>
```

7. Set the NTP server (use server name or IP address), as shown in the following example:

```
store system ntp server
```

At the Enter ntp server: prompt, enter the IP address or server name for the first NTP server and press Enter. Provide a second address and then a third address or press Enter to quit.

8. Enable the NTP servers, as shown in the following example:

```
store system ntp state on
```

9. Validate the settings before you complete the basic configuration, as shown in the following example:

```
show network interface all  
show network routes defaultroute  
show network resolver all  
show system hostname  
show system domain  
show system ntp all
```

10. Restart the appliance to complete the basic configuration, as shown in the following example:

```
restart system
```

The system shuts down and reboots immediately after the command is entered. Upon startup, the system should be accessible (by using CLI and GUI) through the network and uses the provided IP address or hostname.

11.(Optional) Check connectivity by completing the following steps:

- a. Verify the CLI login prompt displays `<host_name>.<domain_name>`.
- b. Log in to the CLI and ping the gateway or a server on the same subnet, as shown in the following example:

```
ping <default_router_ip>1
```

Advanced configuration I: License and registration

With the appliance configured and reachable in the network, you can now use a Secure Shell (SSH) terminal utility, such as PuTTY, to access the CLI.

The appliance's web-based graphical user interface (GUI) also can now be accessed by using a web browser and the `admin2` account, as shown in the following example:

```
https://<host_name.domain_name-or-IP-address>:8443
```

Tip: The CLI is accessed by using only the `cli` or `guardcli1` through `guardcli5` accounts.

Advanced configuration I covers the following topics:

- ▶ Applying the Guardium license key
- ▶ Setting a shared secret
- ▶ Verifying the time zone
- ▶ Registering managed units with the Central Manager
- ▶ Installing maintenance patches, if any

Advanced configuration II covers the following topics:

- ▶ Configuring interfaces to enterprise systems, such as SMTP, LDAP, and syslog-forwarding
- ▶ Configuring and scheduling system backups, daily archiving, and purges
- ▶ Configuring self-monitoring alerts

In a managed environment, start with the configuration of the Central Manager and at least one collector appliance. This configuration allows for the deployment of some Guardium agents in parallel with the appliance rollout.

² Similar to the default CLI password, the `admin` and `accessmgr` default passwords for IBM-provided hardware appliances is obtained from IBM Support or IBM Sales. For VM appliances, they are set by the person who is building the VM image. On first login, you are prompted to change the default password.

Log in to the CLI of the Central Manager appliance (or each stand-alone appliance) and complete the following steps:

1. Install specific product keys, which are based on your product entitlements, by using the following command:

store license console

(This step is only necessary for appliances you built and for the Central Manager or stand-alone units)

Copy and paste the entire license key, including any trailing equal sign (=) and press Enter, as shown in Figure 3-6 on page 63.

```
cmrb01.guard.swg.usma.ibm.com> store license console
Please paste the string received from customer services. Then press <ENTER> to continue.
Wls3 /SZF1dYGIgCEAAAAAAAAAAAAAAAAAAcW50Tk5LCUsRzUwMDAgIC
AgIC IA==

Store license successfully.
The web interface will be restarted.
Restarting gui
Changing to port 8443
Stopping.....
Safekeeping xregs
We recommend that the machine be rebooted at the earliest opportunity in order
to complete the license updating process.
ok
cmrb01.guard.swg.usma.ibm.com> █
```

Figure 3-6 Store the Guardium license key

2. Set a common shared secret on the CM and the managed units, as shown in the following example:

store system shared secret <created_shared_secret>

3. View and change the time zone, if needed (do not change the time zone and host name in the same CLI session).

Run the following command to display the current time zone:

show system clock timezone

If the time zone is incorrect, display a list of valid time zones by running the following command:

store system clock timezone list

Choose the appropriate time zone from the list and set it by using the following command:

store system clock timezone <selected time zone>

Note: When a new time zone is set up, internal services restart and any configured data monitoring is disabled during this restart.

4. Register the managed units to the CM. If this is a to-be managed unit, register it by using the following command:

```
register management <central_manager_ip> 8443
```

The default registration port on the CM is 8443.

5. Patch the appliance by using the latest Guardium patch update (GPU) downloaded from the IBM Fix Central portal for your Guardium version.

A GPU is cumulative, so it is necessary to apply only the most recently available GPU.

Patches are first uploaded to the CM (or stand-alone units) and then applied (installed) by using the CLI. After the patch is installed on the CM, it can be distributed from the CM to the managed units by using the GUI.

Note: Patches are applied first to the Central Manger, followed by the aggregators, and then the collectors.

For more information, see the “How to install patches” section of the How-to Guide Overview chapter of the *Help Book Guardium*.

Tip: During the patch installation by using the CLI, you can choose to skip the pre-patch backup.

3.7 Agent deployment, installation, and configuration

With the CM and at least one managed collector or a stand-alone collector that is configured, the agents deployment and configuration can start.

The following Guardium agents are available:

- ▶ Guardium installation manager (GIM) (recommended)
- ▶ S-TAP (includes the ATAP and S-GATE) (required for DAM)
- ▶ Instance Discovery (optional)
- ▶ Configuration Audit System (CAS) (optional)

This section describes the installation of the GIM and S-TAP agents on the Linux, UNIX, and Windows platforms.

The agent installation planning includes the following tasks:

- ▶ Download the latest version of the agent installers from IBM Fix Central.

- ▶ Determine the installation directory to use:
 - Verify that it has enough space: Approximately 400 MB - 500 MB total. Verify current space requirements by referring to the “S-TAP” chapter of the *Help Book Guardium*.
 - It should be on an internal drive. If an external drive is used, it must be mounted before the database instance is activated during server boot-up.
- ▶ Open the required ports of any firewalls (network and server) between the database server and the Guardium appliance. For more information about document for a list of ports, protocol and directionality, see *Guardium Ports Requirement*.

3.7.1 Guardium installation manager

The Guardium installation manager (GIM) allows the Guardium Administrator to install, upgrade, and configure Guardium modules, such as the S-TAP and KTAP. GIM has a client and server architecture with the clients running on the database hosts and the server running on the stand-alone Guardium appliance or CM.

As of this writing, the GIM client is available only for the Linux, UNIX, and Windows platforms.

GIM server

The GIM server is installed as part of the Guardium application on the appliance and provides the user interface (UI). By using the GIM UI, the user can install, uninstall, and upgrade Guardium bundles and modules and provide feedback about database servers, installed modules, and statuses.

An administration user can interact with GIM through the GIM CLI commands or the GUI.

GIM client

The GIM client application must be installed manually for the first time on the database server machines. The GIM client registers with the GIM server, starts requests to check for software updates, installs the new software, updates module parameters, and uninstall modules.

Installing GIM for the first time

The GIM client is OS-specific and the latest version is available for download from IBM Fix Central. It requires Perl 5.8 or higher, and the gzip file compression program to be available on the database server.

Note: In a managed environment, it is recommended the GIM client report to the CM, whereas in a stand-alone environment it should report to the collector used to monitor the database server

Installing GIM on the database server (UNIX and Linux)

Complete the following steps to install the GIM client:

1. Verify the Perl, gzip, and firewall port prerequisites.
2. Place the GIM client installer (which often is a shell script) on the database server in any directory.

Note: On a Solaris zone configuration, the GIM (and STAP) is installed in the master or global zone.

3. As root, run the installer by using the following command:

```
./<installer_name> -- --dir <install_dir> --sqlguardip <CM ip>  
--tapip <db server ip>
```

Tip: Create a simple shell script to minimize typographical errors and for reuse on other servers.

Figure 3-7 shows a GIM client installation example.

```
[root@vm-rhe54-db ~]# ./guard-bundle-GIM-guard-8.0.xx_r20992_1-rhel-5-linux-1686  
.gim.sh -- --dir /usr/local/guardium --tapip 10.10.9.14 --sqlguardip 10.10.9.8  
1 --perl /usr/bin/  
Verifying archive integrity... All good.  
Uncompressing Guard BUNDLE-GIM Installer....  
GUARDIUM, INC., AN IBM COMPANY  
SQL GUARD(TM) LOCAL TAP LICENSE AGREEMENT ("Agreement")  
  
NOTICE: CAREFULLY READ THESE TERMS AND CONDITIONS BEFORE DOWNLOADING, INSTALLING  
....  
Installing modules'....  
Installation completed successfully
```

Figure 3-7 GIM client installation example

4. Run the following command to validate that the GIM client is running:

```
ps -ef | grep gim
```

You should see two processes, as shown in Figure 3-8 on page 67.

```
[root@vm-rhe54-db ~]# ps -ef | grep gim
root      4687      1  1 13:58 ?        00:00:00 /usr/bin/perl /usr/local/guardiu
m/modules/GIM/8.0_r20992_1-1288807128/gim_client.pl
root      4725  4687  0 13:58 ?        00:00:00 ../../perl ./guard_gimd.pl
```

Figure 3-8 GIM client processes

5. After a brief wait, the client registers with the Guardium appliance.

Installing GIM on the database server (Windows)

Complete the following steps to install the GIM client:

1. Install Perl for Windows, including the IPC-Run3 and Win32-DriveInfo packages. The alternative is to wait and use the Perl that is bundled with the GIM client (you are prompted during the installation process).
2. Log in to the database server as the system administrator.
3. Transfer the GIM client (folder and contents) to the database server.
4. Run the GIM client installer, setup.exe (see Figure 3-9), and follow the prompts.

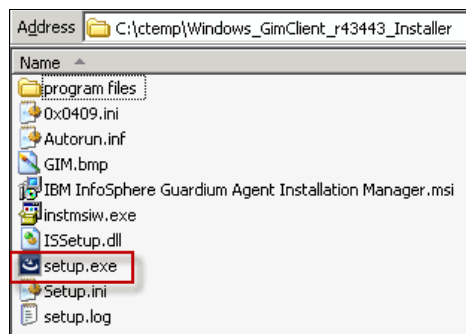


Figure 3-9 GIM Client installer on Windows

5. Accept the default complete setup-option (or choose the **Custom** option to change the default installation directory).
6. Choose which Perl distribution to use. For example, **Yes** to use the Perl that is bundled with the GIM client, as shown in Figure 3-10 on page 68.

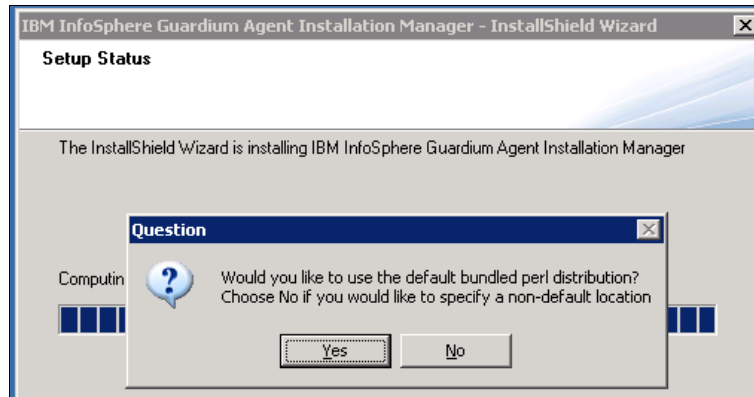


Figure 3-10 GIM client installation prompt for which Perl to use

7. Enter the following information when prompted:
 - IP Address or Fully Qualified Domain Name (FQDN) of the Central Manager or stand-alone collector appliance to which the GIM client reports.
 - IP Address of the database server (local IP address) on which the GIM client is being installed.
8. After the installation is complete, confirm that the Guardium Installation Manager service is running on the database server as Local System.

Verifying that the GIM client registered with its GIM server

To verify that the GIM client registered its GIM server, log in as the admin user to the GUI of the Guardium appliance to which the GIM client is reporting (that is, the stand-alone collector or the Central Manager) and click **Administration Console tab** → **Module Installation** → **Process Monitoring**.

An entry for the GIM client should be listed.

Note: A UNIX or Linux GIM client also has its supervisor process that is listed, as shown in Figure 3-11 on page 69.




Figure 3-11 UNIX or Linux Gim client and supervisor registered with the appliance

3.7.2 S-TAP

The Guardium S-TAP is a lightweight software agent that is installed on a database server. It monitors database traffic in real time and forwards that information to a Guardium collector appliance.

Note: The S-TAP reports only to a collector and not an aggregator.

In this section, we describe the use of the GIM to deploy the S-TAP. Although the S-TAP can be installed directly on the database server, it is recommended to use the GIM.

Preparing to install and configure an S-TAP by using the GIM

Complete the following steps:

1. Verify that the GIM client is installed on the target database server and any required firewall ports are open.
2. Download from IBM Fix Central the latest S-TAP that matches the OS and version of the target database server. As of this writing, there is only one installer for all supported Windows versions.
3. Upload the GIM version (also referred to as a *GIM bundle*) of the S-TAP installers to the GIM server on the CM or stand-alone collector, as shown in the following examples:
 - guard-winstap-v90-43443.gim is the Windows S-TAP gim bundle
 - guard-bundle-STAP-9.0.0_r48064_v90_1-rhel-5-linux-i686.gim is the Linux S-TAP gim bundle for the RHEL5 32-bit Intel x86 platform

Complete the following steps:

- a. Log in as the admin user to the GUI on the CM or stand-alone collector, and browse to **Administration Console tab** → **Module Installation** → **Process Monitoring** → **Upload**.

- b. Browse to the location where you downloaded and decompressed the S-TAP gim bundle, select the bundles that you want to upload, and then click **Upload**, as shown in Figure 3-12.

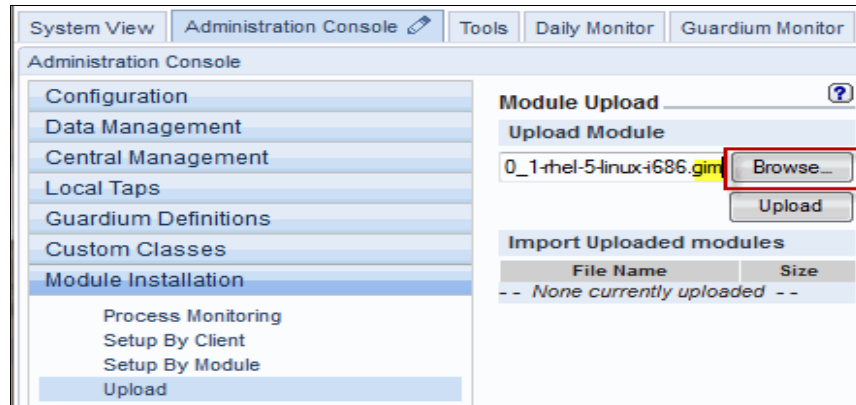


Figure 3-12 Browse and Upload GIM bundle to the GIM server

- c. Import the files by clicking the check mark icon next to each file name, as shown in Figure 3-13.

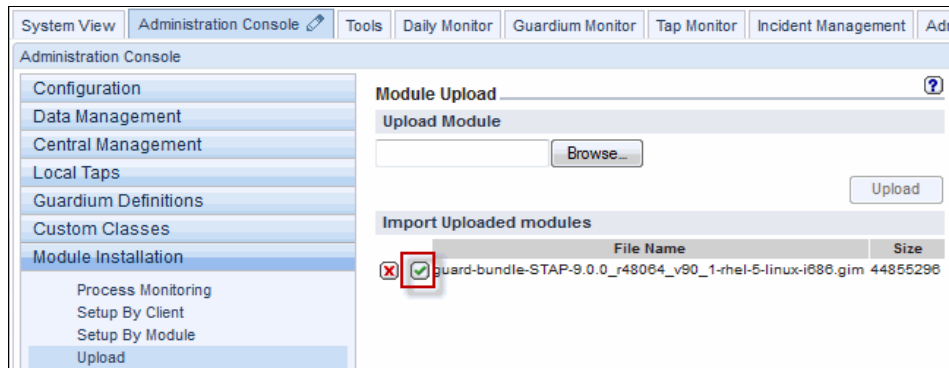


Figure 3-13 Import GIM bundle

- d. (Optional) After successfully importing the files, browse to **Administration Console tab** → **Module Installation** → **Process Monitoring** → **Setup By Module** and click **Search** to see the S-TAP bundle, as shown in Figure 3-14 on page 71.

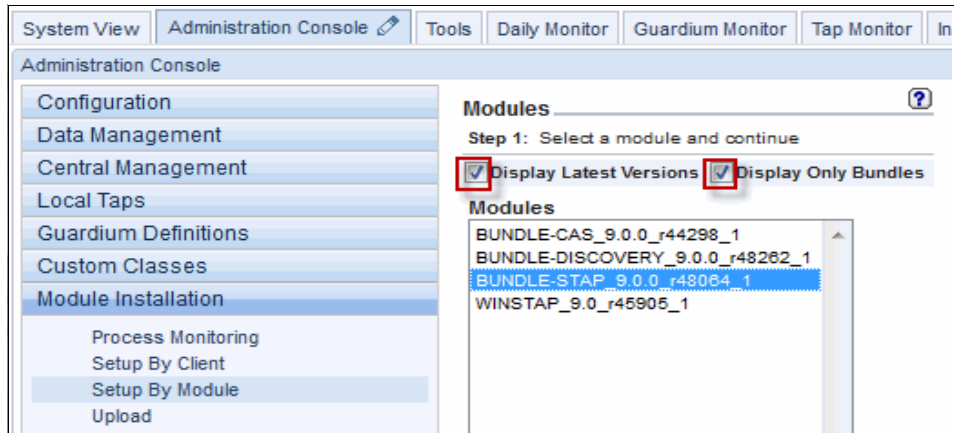


Figure 3-14 View staged modules on the GIM server

Installing and configuring an S-TAP by using the GIM (UNIX and Linux)

Complete the following steps to install and configure an S-TAP by using the GIM:

1. Log in as an admin user to the GUI on the CM or stand-alone collector and browse to **Administration Console tab** → **Module Installation** → **Process Monitoring** → **Setup By Client**. Click **Search** to start the package installation.
2. Select the UNIX database servers to install the S-TAP, as shown in Figure 3-15.

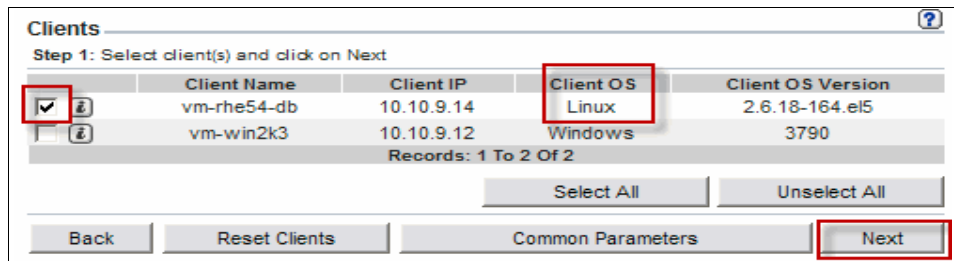


Figure 3-15 Select database client to install S-TAP

3. Select the S-TAP bundle only; for example, **BUNDLE-STAP_...**, to install the S-TAP and its components, as shown in Figure 3-16.

Common Modules

Step 2: Select a module to Config/Install/Update/Uninstall on selected client(s) and click on Next

Selected Clients

Client Name	Client IP	Client OS	Client OS Version
vm-rhe54-db	10.10.9.14	Linux	2.6.18-164.el5

Display Latest Versions Display Only Bundles

Modules	Module Status
BUNDLE-CAS_9.0.0_r44298_1	-- This module has not been set for any of the selected clients --
BUNDLE-DISCOVERY_9.0.0_r48282_1	
BUNDLE-STAP_9.0.0_r48064_1	

Figure 3-16 Select STAP bundle to install

4. Select the clients (database servers) to configure and provide the following parameters (as shown in Figure 3-17 on page 73):
 - KTAP_LIVE_UPDATE: Entering y enables the KTAP update without requiring a server reboot.
 - STAP_SQLGUARD_IP: The IP address or FQDN of the primary collector to which this STAP communicates.
 - STAP_TAP_IP: The IP address or FQDN of the database server or node on which the STAP is being installed.
 - (Optional) KTAP_ALLOW_MODULE_COMBOS: Entering Y (default is N) applies to Linux only and often is recommended.

Tips: The parameters are listed in alphabetical order. Use the horizontal scroll bar to browse the parameter list. By pointing to fields, you can see the permitted values.

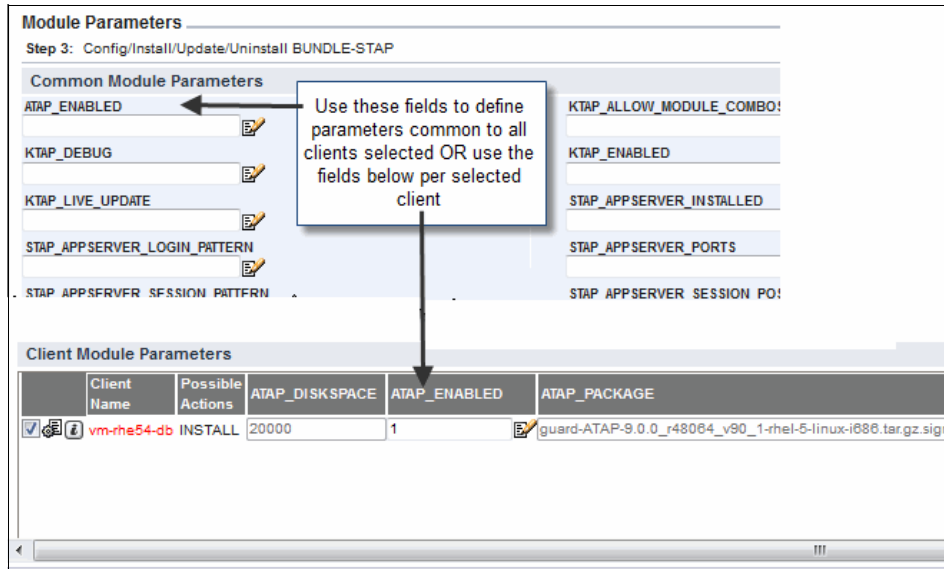


Figure 3-17 Configure STAP

5. Click **Apply to Clients**. After the confirmation message is shown, click **Install/Update**.
6. Enter now in the Schedule Date field to schedule the S-TAP installation to start now and then click **Apply**, as shown in Figure 3-18.




Figure 3-18 Enter now to schedule the S-TAP installation to start now

7. Verify the S-TAP module installation status by completing the following steps:
 - a. Click the **i** icon that is next to the client name or browse to **Administration Console** → **Module Installation** → **Setup By Client** and click the **i** icon that is next to the client.
 - b. In the status window, verify that the status is “Pending-Install”, “PI” (Pending Install), or “Installed”.
 - c. Click **Refresh** to refresh the status.
 - d. Wait until the status changes to “Installed”. You can also check the status of the GIM Events report by browsing to **Guardium Monitor tab** → **GIM Events List**.

If there is an installation error, this report includes more information, as shown in Figure 3-19.

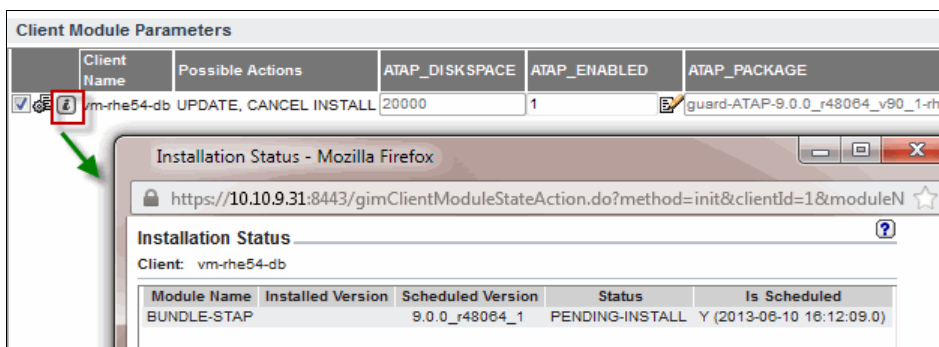


Figure 3-19 Module installation status

Note: For AIX, schedule a restart of the database instance or listener to allow the KTAP to monitor remote (TCP/IP) connections.

For DB2 and Informix® on Linux, install and configure the ATAP to allow monitoring of local (shared memory) connections. For more information, see the “S-TAP” chapter of the *Help Book Guardium*.

In both of these instances, the KTAP and S-TAP do not monitor database activity until an Inspection Engine is configured.

- Alternatively, a successfully installed S-TAP registers with the assigned collector and has an entry on the report that is found by clicking **System View tab** → **S-TAP Status Monitor**.

Installing and configuring an S-TAP by using the GIM (Windows)

The following installation process for the Windows S-TAP is similar to the UNIX and Linux S-TAP:

- Select the Windows database servers on which to install the S-TAP.
- Select the S-TAP module; for example, **WINSTAP_...**
- Select the database server from the list and provide the following parameters in the highlighted fields:
 - **WINSTAP_DBALIAS:** An alias for this STAP configuration; for example, database server host name.

- WINSTAP_INSTALL_DIR: The path to the installation directory; for example,
c: /program files (x86)/guardium/stap/. Note the following points:
 - Use / and not \ in the path.
 - You must use program files (x86) for 64-bit systems.
 - Add the “stap” subdirectory so that the STAP files are not commingled with the GIM client files.
 - e. WINSTAP_SQLGUARD_IP: The IP address or FQDN of the primary collector to which this STAP communicates.
 - f. WINSTAP_TAP_IP: The IP address or FQDN of the database server or node on which the STAP is being installed.
4. Click **Apply to Clients**. After the confirmation message is shown, click **Install/Update**.
 5. Enter now in the Schedule Date field and then click **Apply**.
 6. Verify the S-TAP module installation status by completing the following steps:
 - a. Click the **i** icon next to the client name and then click **Administration Console** → **Module Installation** → **Setup By Client**. Click the **i** icon next to the client.
 - b. In the Status window, verify that the status is “Pending-Install”, “PI” (pending install), or “Installed”.
 - c. Click **Refresh** to refresh the status.
 - d. Wait until the status changes to “Installed”. You can also check the status on the GIM Events report by clicking to **Guardium Monitor tab** → **GIM Events List**.

If there is an installation error, this report has more details.
 7. Schedule a restart of the database instances to complete the STAP installation.

Note: Until the instances are restarted, only local connections (that is, through named pipes or shared memory) are monitored (with inspection engines configured).

8. On the database server, the following three new services should be started, as shown in Figure 3-20 on page 76:
 - GUARDIUM Database Monitor
 - GUARDIUM_STAP
 - GUARDIUM DC Connector³

GUARDIUM Database Monitor	Service for monitoring DataBase Server instances activity	Manual	Local System
GUARDIUM DC Connector	GUARDIUM DC Connector	Automatic	Local System
Guardium Installation Manager	Guardium Installation Manager	Automatic	Local System
GUARDIUM_STAP	Tap DataBase activity and sends it to the Guard machine	Automatic	Local System

Figure 3-20 Guardium Windows services

3.7.3 Command-line options

The interactive method for installation and configuration can be effective for becoming familiar with the various functions and for the setup of smaller environments. However, it might not be practical for larger deployments.

Guardium provides the following methods that can be used to script various functions:

- ▶ Silent or non-interactive installers:
 - Input (arguments) is passed on the command line (default is interactive mode).
 - Use the silent mode to incorporate the agents installation into a software distribution package, such as SMS or SCCM:
 - GIM Client command-line options: For more information, see the “GIM Installation” section of the Guardium Installation Manager chapter of the *Help Book Guardium*.
 - S-TAP command-line options (only if the GIM is not used): For more information, see the UNIX and Windows S-TAP sections of the S-TAP chapter of the *Help Book Guardium*.
- ▶ GuardAPI:
 - Provides access to Guardium functionality by using the **grdapi ()** functions, which are started by using the CLI. For more information, see the “Appendices” chapter of the *Help Book Guardium* for a complete list of the **grdapi()** functions.
 - To start multiple **grdapi ()** commands, for example:
 - A user prepares several **grdapi ()** calls, then pastes these prepared statements into the CLI session.
 - Start a script that contains the **grdapi ()** statements by using an SSH client to connect to the CLI and run the statements, as shown in the following example:

³ Guardium DC service collects updates of user accounts (SIDs and user names) from the primary domain controller and then signals the changes to Guardium_S-TAP to update the S-TAP internal SID and UserName map.


```
ssh cli@myappliance.ibm.com < my_grdapi_script.txt
```

- For more information, see the “GuardAPI Reference” section of the Appendices chapter of the *Help Book Guardium*.
- CLI: CLI commands can also be scripted and started by using the same methods that are described for the `grdapi ()`.

3.7.4 S-TAP and its inspection engine configuration

The S-TAP inspection engine (also known as *inspection engine*) is a configuration that specifies the database platform and instances the S-TAP monitors on the S-TAP host (database server). An S-TAP often has many inspection engines.

Complete the following steps to configure the S-TAP and its inspection engine:

1. Log in as admin to the collector to which the recently installed S-TAP is reporting by using the following URL:
`https://<collector_ip_address-or-collector-fqdn>:8443/`
2. Browse to **Administration Console** → **Local Taps** → **S-TAP Control**.
3. Check whether the S-TAPs are listed and have a green status icon. Otherwise, click **Refresh**.
4. Click the **Edit** icon for the S-TAP (as shown in Figure 3-21 on page 78) and modify the following sections:
 - Details:
 - (Optional) alternative IPs: Add any virtual IPs that are used to connect to the database on the host. For more information, see the *Help Book Guardium*.
 - (Windows only) Shared Mem. Monitor: If the database is a 32-bit Microsoft SQL Server version, verify that the MSSQL option is checked.
 - Guardium Hosts: (Optional) Add the IP address or FQDN of the secondary (failover) collector for this S-TAP. Click **Add**.




Figure 3-21 Configure the S-TAP and inspection engine

5. Complete the following steps to configure the inspection engine:
 - a. Add Inspection Engine: Configure an inspection engine.

The fields and default values depend on the database protocol that is selected and the database server operating system.
 - b. Database protocol: Select the appropriate protocol from the drop-down list; for example, **MSSQL** (for Microsoft SQL Server).
 - c. Port Range: Enter the database listener port range; for example, 1433 - 1433 or 1521 - 1521.
 - d. KTAP DB Real Port: (UNIX only) Enter the database instance port; for example, 1521.
 - e. Client Ip/Mask: Enter the IP/MASK of the client network to monitor; for example, enter 1.1.1.1 and 0.0.0.0 to monitor all subnets.
 - f. DB Install Dir: (UNIX only) DB2, Oracle, or Informix. Enter the full path of the database installation directory; for example, /opt/oracle10/.
 - g. Process Name: (UNIX only) For DB2, Oracle, or Informix, enter the full path for the database executable; for example, /opt/oracle10/bin/oracle.
 - h. Process Name: (Windows only) For MSSQL the default is SQLSERVER.EXE.
 - i. Named Pipe: (Windows only) The named pipe that is used by Microsoft SQL Server if it is enabled for local access. The default is SQL\QUER,PIPE\SQLLOCAL.
 - j. Instance Name - MSSQL: The MSSQL instance name; for example, MSSQL2005. The default instance name is MSSQLSERVER.
 - k. Instance Name - Oracle: For Oracle that is using database encryption (Oracle ASO), enter the instance name.

- I. DB2 Shared Memory: For more information about configuring these parameters, see *Help Book Guardium*.
6. Click **Add** to save this inspection engine configuration. Click **Apply** to apply the configuration to the S-TAP.

A Windows MSSQL inspection engine example is shown in Figure 3-22.

The screenshot shows a dialog box titled "Add Inspection Engine...". It contains the following fields and values:

- Protocol: MSSQL
- Port Range: 1433 - 1433
- Client Ip/Mask: 1.1.1.1 / 0.0.0.0
- Exclude Client Ip/Mask: /
- Process Names: SQLSEVRV.EXE
- Named Pipe: SQL\QUER,PIPE\SQLLOCAL
- Instance Name: MSSQLSERVER

Buttons include "Add Pair" (two), "Add" (highlighted with a red box), "Cancel", and "Apply".

Figure 3-22 Windows MSSQL inspection engine example

Figure 3-23 shows an example of the Oracle inspection engine on Linux.

The screenshot shows a dialog box titled "Add Inspection Engine...". It contains the following fields and values:

- Protocol: Oracle
- Port Range: 1521 - 1521
- KTAP DB Real Port: 1521
- Client Ip/Mask: 1.1.1.1 / 0.0.0.0
- Exclude Client Ip/Mask: /
- Connect To Ip: 127.0.0.1
- DB Install Dir: ome/oracle10/oracle/product/10.2.0/db
- Process Name: e10/oracle/product/10.2.0/db/bin/oracle
- Encryption:
- Intercept Types:

Buttons include "Add Pair" (two) and "Add".

Figure 3-23 Linux Oracle inspection engine example

3.7.5 More information

Depending on your monitoring needs, you can explore other configuration options in the “S-TAP” chapter of the *Help Book Guardium*; for example:

- ▶ Monitoring DB2, IMS™, and VSAM on z/OS using an S-TAP
- ▶ Monitoring Oracle ASO, Oracle, or Sybase SSL-encrypted connections

- ▶ Using Auto Discovery to discover database instances to help with creating inspection engines
- ▶ Installing S-TAP without GIM

3.8 Configure remaining appliances

Complete the basic and initial advanced configuration for all remaining appliances before you continue, including registering any managed units.

3.8.1 Managed units grouping

After an appliance is successfully registered with the CM, it appears on the Central Management portal of the CM.

Figure 3-24 shows the managed unit listing on the CM, including the “Distribute...” functions. In a large managed environment, you can group managed units on the Central Manager to help with patch distribution or configuration. A managed unit or appliance can belong to more than one group.

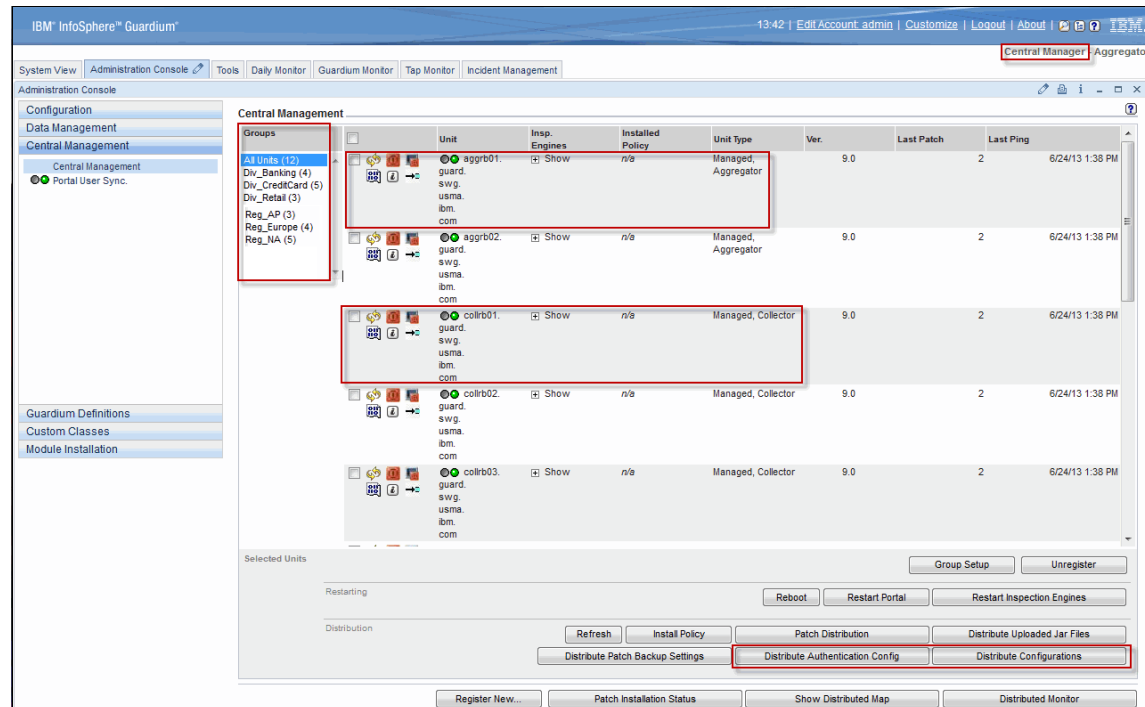


Figure 3-24 Central Management view showing managed units

The Central Manager view that is shown in Figure 3-25 shows 4 of 12 managed units in the group “Div_Banking”.

Central Management					
Groups		Unit	Insp. Engines	Installed Policy	Unit Type
All Units (12)		aggrb01.	Show	n/a	Managed, Aggregator
Div_Banking (4)		guard.			
Div_CreditCard (5)		swg.			
Div_Retail (3)		usma.			
		ibm.			
		com			
		collrb01.	Show	n/a	Managed, Collector
		guard.			
		swg.			
		usma.			
		ibm.			
		com			
		collrb02.	Show	n/a	Managed, Collector
		guard.			
		swg.			
		usma.			
		ibm.			
		com			
		collrb03.	Show	n/a	Managed, Collector
		guard.			
		swg.			
		usma.			
		ibm.			
		com			

Figure 3-25 Central Manager view showing 4 of 12 managed units in a group

3.8.2 Creating a group

Complete the following steps to create a group of managed units that is called All_Aggregators, which contains two aggregators:

1. On the CM, browse to **Administration Console** → **Central Management** → **Central Management**.
2. Select both aggregators, then click **Group Setup** to open the Group Setup window.
3. Enter the group name All_Aggregators and click **Add** to add the group, as shown in Figure 3-26 on page 82.

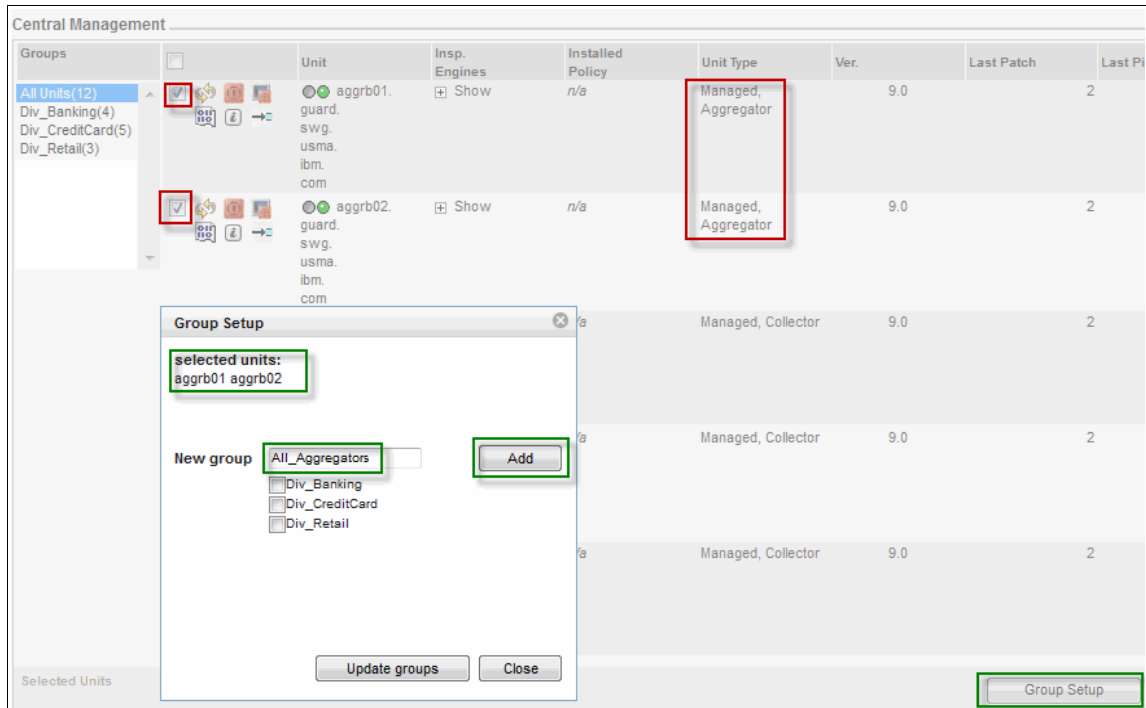


Figure 3-26 Adding a managed unit group

4. Select the new group **All_Aggregators** and click **Update groups** to add the selected aggregators to this new group, as shown in Figure 3-27.

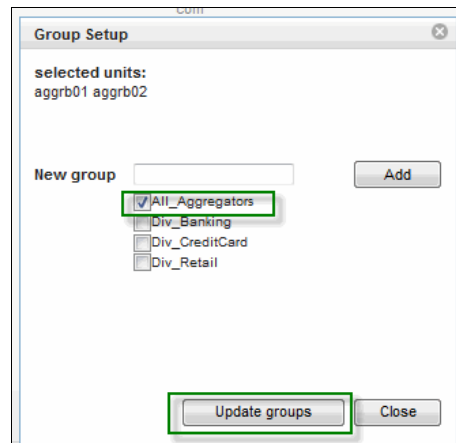
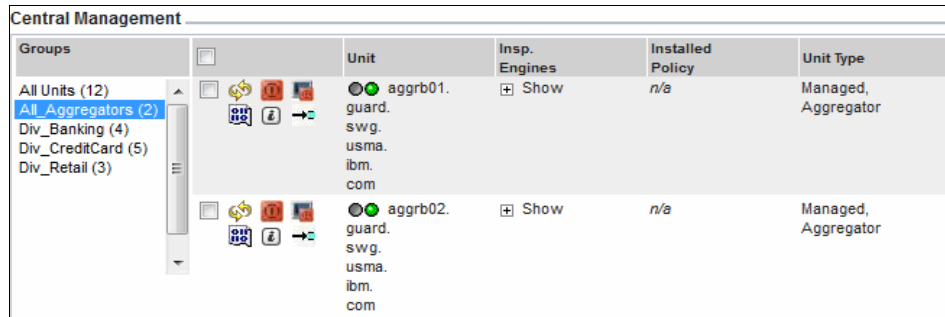


Figure 3-27 Updating a managed unit group

Figure 3-28 shows the newly created group with two assigned members.



The screenshot shows the 'Central Management' interface. On the left, a tree view under 'Groups' shows 'All Units (12)' expanded, with 'All_Aggregators (2)' selected. The main table displays two units:

Groups	Unit	Insp. Engines	Installed Policy	Unit Type
All_Aggregators (2)	aggrb01. guard. swg. usma. ibm. com	Show	n/a	Managed, Aggregator
	aggrb02. guard. swg. usma. ibm. com	Show	n/a	Managed, Aggregator

Figure 3-28 Newly created group with two assigned members

3.9 Guardium operations

In this section, we describe the recommended steps that are used to maintain your Guardium infrastructure, including the remaining advanced configuration items.

In a managed environment, start with the Central Manager.

Note: In a managed environment, these configurations can be repeated on each managed unit or selectively distributed from the CM by using the distribute functions.

3.9.1 Configuration

In this section, we describe the following common configuration options:

- ▶ Alerter
- ▶ IP-to-hostname aliasing
- ▶ Global profile

Log in as admin to the GUI on the CM or stand-alone collector and browse to the **Administration Console tab** → **Configuration** section and configure alerter, IP-to-hostname aliasing, and global profile.

Alerter: Configure SMTP

SMTP and SNMP are used to notify external users and systems of local events on the appliance through email (SMTP) or traps (SNMP). Complete the following steps:

1. For SMTP, enter the IP Address/Host Name, Port, and Return Email Address information.
2. Click **Test Connection** to check whether this appliance can access the configured SMTP host and port. The test connection function does not check whether emails can be forwarded.
3. If the SMTP server uses authentication, enter a valid user name and password.
4. For SNMP, if required, see the “Alerter Configuration” section in the Guardium Administration chapter of the *Help Book Guardium*.
5. Select **Active on startup** and then click **Restart**.
6. To send a test email, use the CLI diag utility. For more information, see the *Help Book Guardium*.

Setting up and scheduling IP-to-Hostname aliasing

The IP-to-Hostname aliasing process builds an alias list for the IP addresses by using the hostnames that are retrieved from the DNS. Alias (hostnames) can then be displayed in reports, if needed.

Complete the following steps:

1. Select **Generate Hostname Aliases for Client...**
2. Select **Update existing...**

This selection updates previously defined aliases that do not match the current DNS hostnames (which often indicate that the hostname for an IP address changed). You might not want to perform this step if you assigned some aliases manually.
3. Click **Apply** to save the configuration.
4. Click **Define Schedule** to define a schedule for running this task (as shown in Figure 3-29 on page 85) and enter the following information:
 - Set a Start Time of **6 a.m.** (DNS updates often are done overnight)
 - Schedule by **Day/Week**
 - Select **Every Day** so that the schedule runs daily
 - Save and activate the configured schedule

Figure 3-29 Schedule IP-to-Hostname aliasing process to run daily at 6 am

Global profile

This is an optional configuration step. The global profile section contains configuration settings for various features, but only the following settings are described:

- ▶ Check **Use aliases in reports....** to show aliases; for example, IP-hostname aliases, by default.
- ▶ Add a suitable PDF footer text; for example, “Copyright <your company name>”, which is printed on the footer of the output PDF reports.
- ▶ (Optional) Clear the **Disable accordion menus** option.
- ▶ Add a suitable login message for your organization; for example, a message that notifies the user that unauthorized access of this system is prohibited.
- ▶ Select **Show login message** to enable the previously entered login message.
- ▶ Enable the **Concurrent login from...** option to prevent concurrent logins by the same account from different IP addresses.
- ▶ (Optional) Use the Upload logo image option to upload a file that contains your corporate logo that is to be displayed in the upper right corner of the Guardium window (which replaces the default IBM logo).

Figure 3-30 on page 86 shows examples of the options for configuration items that include associated schedules.




Figure 3-30 Functions that can be scheduled

3.9.2 Transferring configuration to managed units

In a managed environment, configurations can be repeated on each managed unit or selectively distributed from the CM.

Complete the following steps to use the Distribute Configuration function:

1. Log in as admin to the GUI on the CM and browse to the **Administration Console** → **Central Management** → **Central Management**.
2. Select the managed units or group to which the configuration is distributed and click **Distribute Configuration** at the bottom of the window.
3. Select the options to distribute and then click **Distribute**.

Figure 3-31 on page 87 shows the Distribute Configuration menu with the configuration and schedule options to distribute selected configurations.

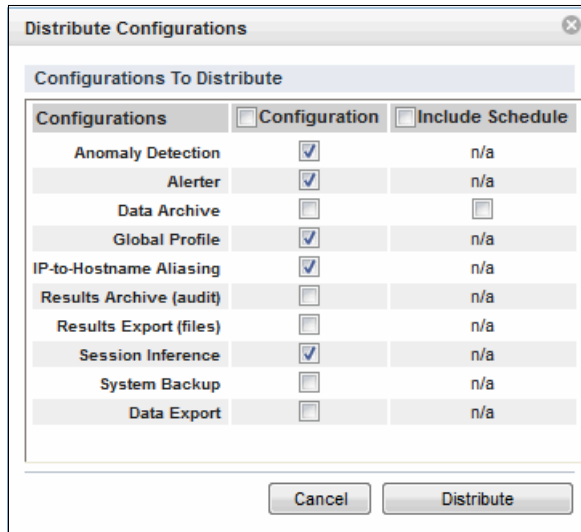


Figure 3-31 Distribute Configuration menu

Note: The distribute configuration option distributes a copy of the settings from the CM for the items that are selected to the selected managed units.

For example, if a group of managed units should use `smtp_server_A` while another uses `smtp_server_B`, you complete the following steps:

1. Set the Alerter configuration on the CM to **smtp_server_A**, then select the managed units and distribute the Alerter configuration only.
2. Change the Alerter configuration on the CM to **smtp_server_B**.
Select the second group of managed units and distribute the Alerter configuration.
3. Set the Alerter configuration on the CM to the SMTP server the CM communicates with (if it is not `smtp_server_B`).

3.9.3 Data management

In this section, we describe the following common data management options:

- ▶ Data archive
- ▶ Data export
- ▶ Data import
- ▶ System backup

Prerequisites

The data archive and system backup functions require access (network and user accounts) to a third-party storage (with allocated capacity) on which to copy and store these daily archive files and the periodic backup files.

The following interfaces (protocols) are available for transferring these files from the appliances:

- ▶ FTP
- ▶ Secure Copy (SCP)
- ▶ Tivoli Storage Manager
- ▶ EMC/Centera

Note: It is not possible to directly attach external storage, such as external disks or a SAN, or install different backup clients for backup or archiving purposes on the appliance.

Data archive and purge

More disk space cannot be dynamically added without rebuilding and reconfiguring the appliance. Therefore, it is necessary to maintain the disk space by periodically archiving and purging the data.

The appliance often has the most recent number of days of data online, where the number of days is determined by your configuration. As a suggested starting point (which can be adjusted later) set the number of days to one of the following values:

- ▶ Less than or equal to 15 days on a managed collector
- ▶ 30 - 60 days on an aggregator

The archive, followed by the purge, is run daily on the appliance and the archived files include the following characteristics:

- ▶ Compressed and encrypted before they are moved to the external storage
- ▶ Available for use in recovering the appliance
- ▶ Can be retained offline, depending on your corporate offline data retention policy, and later restored for forensic investigation

Aggregator configuration

Complete the following steps to configure an aggregator to archive data by using SCP or FTP:

1. Log in as admin to the GUI on the CM (or stand-alone collector) and browse to **Administration Console tab** → **Data Management** → **Data Archive**.

Note: Tivoli Storage Manager or EMC Centera requires other configuration by using the CLI, including the upload of specific configuration files `dsm.sys` (Tivoli Storage Manager) and `pea` file (Centera). For more information, see the *Help Book Guardium*.

2. Select **Archive** and complete the following steps:
 - a. Enter 1 Archive data older... field and enter 2 in the Ignore data older... field to collect the previous day's data only.
 - b. Verify that Archive Values is selected (it is selected by default).
 - c. For Protocols, select **SCP** or **FTP**.
 - d. Enter the fully qualified host name or IP address of the SCP or FTP host, destination directory, and credentials to connect to the archive destination.
 - e. For Port, leave the default of 0 if the default SCP (or FTP) ports are used; otherwise, enter the non-default listener port that is used for the SCP or FTP server.
3. Verify that Purge is selected (it is selected by default). Complete the following steps:
 - a. Update the Purge data older... field to the suggested initial values (the default value is 60 days).
 - b. Clear the **Allow purge without...** option (this option is selected by default to allow purging until the archive is configured).
4. Select **Apply**. An attempt is made to send a test file to the archive destination, and, if successful, the configuration is saved. The scheduling configuration also is enabled.

Figure 3-32 on page 90 shows an aggregator Data Archive configuration example.

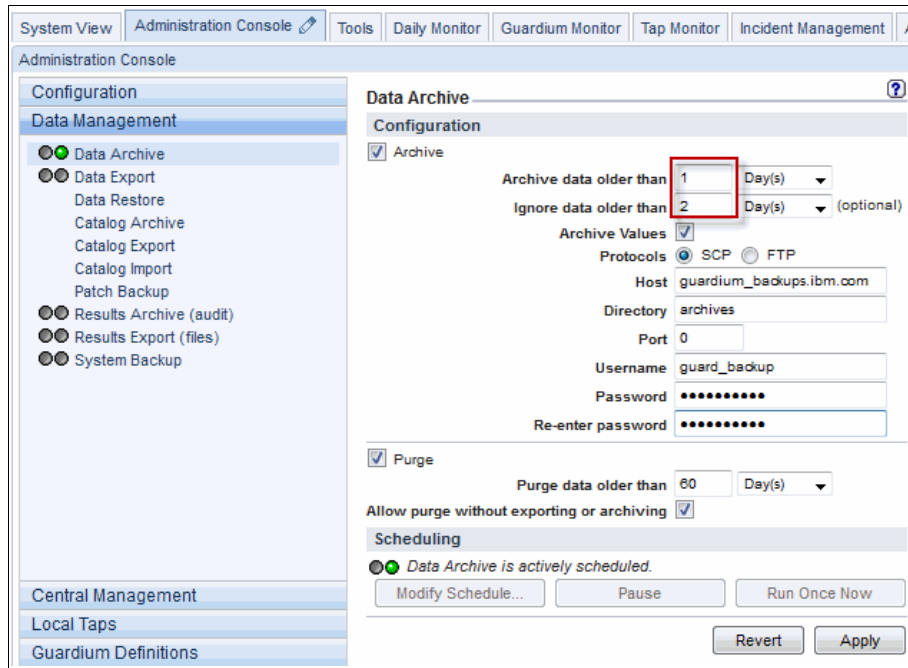


Figure 3-32 An aggregator Data Archive configuration example

5. For scheduling, click **Modify Schedule**. Complete the following steps:
 - a. For Start Time, enter the following information:
 - For aggregator appliances, including the CM, enter 7:00 p.m.
The aggregator archives are scheduled after business hours after the daily import of data from the collectors so that data can be included in the archive, and any audit processes are run.
 - For collector appliances, enter 3:00 a.m.
Archives are scheduled after data is exported to the aggregator.
 - b. Schedule daily and then click **Save** (leave all other fields with defaults).

Collector configuration

The steps to configure daily archives on the collector is similar to those for the aggregator, except that the archive is scheduled to occur after the data export. Therefore, schedule the archive to start at 3:00 a.m.

Note: Some customers choose not to archive from managed collectors. Instead, they rely on the archives of the aggregator to which data is exported.

Background purge

In addition to the purge process that is configured by using the GUI, there is a separate purge process for other objects. This background purge uses different limits, which are configurable by using the CLI and should not be adjusted unless directed to do so by IBM Support.

3.9.4 System backup

A system backup is a full backup of the Guardium database and selected configuration files from the appliance. It is used to restore the appliance in case of hardware failure, and as such, it is not necessary to keep more than three rolling copies.

The backup is written to a single file that is compressed and encrypted and sent to the specified destination by using the transfer method that is configured for backups on the appliance.

It is important to back up the aggregators. A weekly backup is recommended, especially for the CM; however, some users might opt for a slightly longer cycle.

Tip: In a managed environment with aggregation, you might choose not to back up managed collectors. However, stand-alone collectors should be backed up.

Complete the following steps to back up SCP or FTP:

1. Log in as admin to the GUI on the CM or stand-alone collector and browse to **Administration Console tab** → **Data Management** → **System Backup**.
2. For Protocols, select **SCP** or **FTP**.
3. Enter the fully qualified host name or IP address of the SCP or FTP host, destination directory, and credentials to connect to the backup destination.
4. For Port, leave the default of 0 if you are using the default SCP (or FTP) ports; otherwise, enter the non-default listener port that is used for the SCP or FTP server.
5. For Backup, check **Configuration** and **Data**.
6. Click **Apply**. An attempt is made to send a test file to the backup destination, and, if successful, the configuration is saved. Also, the Scheduling configuration is saved.
7. For scheduling, click **Modify Schedule** and schedule the System Backup to occur weekly on Sunday during a quiet period; for example, 7 a.m.

3.9.5 Aggregation

Aggregation refers to the importing and merging of data on an aggregator. The data source often is the collector, but it can be other aggregators (which are referred to as *second-level aggregation*).

The primary purpose of aggregation is to offload the reporting and analysis function from the collectors, while also providing a consolidated view of the data from multiple collectors. It is this need for a consolidated view that might justify second-level aggregation.

Note: Aggregation does not summarize or roll-up the data. Instead, it merges the records.

The data is transferred through daily batch files by using SCP. A daily data export is scheduled on the source and a corresponding data import is scheduled on the aggregator. There is an option to use a secondary aggregator in case the primary aggregator is unreachable.

Note: The export and import are manually synchronized. That is, import is scheduled to occur after all sources are expected on the aggregator.

Aggregator configuration

Complete the following steps to configure an aggregator to import and aggregate data:

1. Log in as admin to the GUI on the aggregator and browse to **Administration Console tab** → **Data Management** → **Data Import**.
 - ▶ Select **Import data from** and click **Apply**.
 - ▶ Click **Modify Schedule...** and schedule the import to occur daily at 2:30 a.m. EST.

This schedule allows up to two hours for the export files from the collectors to be ready on the aggregator.

Tip: The import start time can be adjusted later in production to more closely match the actual wait time. Click **Guardium Monitor** → **Aggregation/Archive Log** to view the actual duration over a period of several days.

Collector configuration

Complete the following steps to configure and schedule the data export on each collector:

Note: Although you can selectively distribute the configuration from the CM, you must schedule the data export individually on each collector by using the GUI.

1. Log in as admin to the GUI on the collector and from the **Administration Console** tab, click **Data Management** → **Data Export**.
2. Select **Export** and complete the following steps:
 - a. Enter 1 in the Archive data older... field and 2 in the Ignore data older... field to collect the previous day's data only.
 - b. Verify that the Export Values option is selected (it is selected by default).
 - c. For Host, enter the fully qualified host name or IP address of the aggregator.
 - d. (Optional) For Secondary Host, enter the fully qualified host name or IP address of the secondary aggregator.
3. The purge option already should reflect the settings that were selected during the archive configuration. If they are not set, complete the following steps:
 - a. Verify that the Purge option is selected (it is selected by default).
 - b. Update the Purge data older... field to the suggested initial values.
 - c. Clear the Allow purge without... option (by default, this option is selected to allow purging until export or archive is configured).
4. Click **Apply**. An attempt is made to send a test file to the specified aggregator, and, if successful, the configuration is saved. Also, the Scheduling configuration is enabled.
5. For scheduling, click **Modify Schedule** and schedule the data export to occur daily at 12:30 a.m.

This setting allows up to two hours for the export files from the collector to be prepared and sent to the aggregator before the aggregator starts its data import.

Note: Avoid scheduling any jobs between midnight and 12:30 a.m. to allow the appliance to complete its start-of-day processing.

3.9.6 Job schedules

The proper maintenance of the Guardium solution requires the setup, scheduling, and tracking of various jobs. Table 3-2, Table 3-3, and Table 3-4 summarize the schedules that are described in this chapter.

Table 3-2 *Aggregator schedules summary*

Aggregator schedules summary		
Data Archive and Purge	Daily	7:00 p.m. Purge is initially set to 60 days. Archives are scheduled during a quiet period after the following tasks are complete: <ul style="list-style-type: none"> ▶ The daily import of data from the collectors so that data can be included in the archive ▶ Any audit processes are run
Data Import	Daily	2:30 a.m.
System Backup	Sunday	7:00 a.m.

Table 3-3 *Collector schedules summary*

Collector schedules summary		
Data Export and Purge	Daily	12:30 a.m. Purge is initially set to 15 days
Data Archive (stand-alone)	Daily	3:00 a.m.
System Backup (stand-alone)	Sunday	7:00 a.m.

Table 3-4 *Common schedule*

Common schedule		
IP-to-Hostname Aliasing	Daily	6:00 a.m.

As the admin user at the **Guardium Monitor** tab, click **Aggregation/Archive Log** report to review the start times and duration of data management-related jobs.

Tip: Copy this information to a worksheet in the deployment plan spreadsheet and use it as a quick reference to track and coordinate these schedules.

3.9.7 Enterprise systems interface configuration

In addition to the interfaces discussed, Guardium provides other enterprise systems interfaces. For more information, see the following relevant sections of the *Help Book Guardium*:

- ▶ Active Directory and LDAP for user authentication.
- ▶ Active Directory and LDAP for importing user accounts.
- ▶ SIEM integration using syslog forwarding; for example, QRadar® and Arcsight.
- ▶ SCP and FTP server for delivering results data in a CSV format.

3.9.8 Self-monitoring

By using Guardium, you can set appliances to monitor themselves (self-monitoring) to ensure that the Guardium solution is available, functioning properly, and to alert users of problems.

The following approaches for self-monitoring are available and supplement each other:

- ▶ Threshold alerts (which are also known as *correlation alerts*):
 - Use queries to check key measures (for example, processor usage) against a specified threshold and alert if the threshold is breached.
 - The alert can be emailed, sent to syslog for forwarding, sent as an SNMP trap, or sent to a custom alerting, user-provided Java class.
 - Are automatically distributed to all managed units from the CM where they are created and activated.
 - Guardium provides several predefined threshold alerts and the user can also create their own threshold alerts. The predefined alerts must be configured and activated, as needed.
- ▶ SNMP polling:
 - Poll the appliance by using a combination of standard and custom metrics that use the object identifiers (OIDs) that are published by the UCD-SNMP-MIB and HOSTRESOURCES-MIB.
 - Each appliance must be polled individually.

In this section, we describe configuring threshold alerts with email delivery. For more information about configuring SNMP polling, see the *Help Book Guardium*.

Prerequisites

To configure alert delivery through email, the following prerequisites must be met:

- ▶ The Alerter must be configured with the SMTP option.
- ▶ Guardium user accounts with email addresses must be created. Alerts are emailed to these addresses.

Note: Alert receivers can be individual accounts or (preferably) a group email address of personnel who should respond to the alert that is defined at your mail server and associated with a pseudo or non-functional user in Guardium.

Recommended self-monitoring threshold alerts

The following self-monitoring threshold alerts are recommended:

- ▶ Aggregation/Archive Errors (predefined)
Alert once a day on any aggregation or archive-related errors, including data import/export and backups jobs.
- ▶ Scheduled Jobs Exceptions (predefined)
Alert on any scheduled jobs-related errors; for example, an error with the daily IP-to-Hostname aliasing job.
- ▶ Inactive S-TAPs Since (predefined):
 - Alert once an hour on all S-TAPs that are not heard from for a specified period.
 - For S-TAPs that are configured with a primary and secondary collector, if the S-TAP cannot communicate with the primary for any reason (such as network issues), it fails over to the secondary. Unless the former-primary collector can ping the S-TAP, it generates an inactive S-TAP alert.
 - S-TAPs in a failover-configuration might generate false alerts if configured incorrectly.
- ▶ No Traffic (predefined)
Alert when there is no traffic that is collected from a server from which the Guardium system was collecting traffic at some point during the last 48 hours. For more information, see the *Help Book Guardium*.
- ▶ Sniffer restarts (user created):
 - Alerts if the sniffer on a collector restarted at least three times in an hour.
 - The sniffer is a major component on the collector and is responsible for receiving and processing data from the S-TAPs; therefore, frequent restarts are one indication of overload.

- ▶ Disk Space: One for the Aggregator and another for the collector (user created):
 - Alert when the database available free space falls below a configured threshold.
 - The database storage engine currently is different on an aggregator that is compared to the collector; therefore, two different alerts are used.

Configuring and activating predefined alerts

Complete the following steps to configure and activate the predefined threshold alerts:

1. Log in as admin to the GUI on the CM or stand-alone collector and browse to **Tools tab** → **Config & Control** → **Alert Builder**.
2. From the list, select the **Aggregation/Archive Errors** alert and click **Modify**, as shown in Figure 3-33.

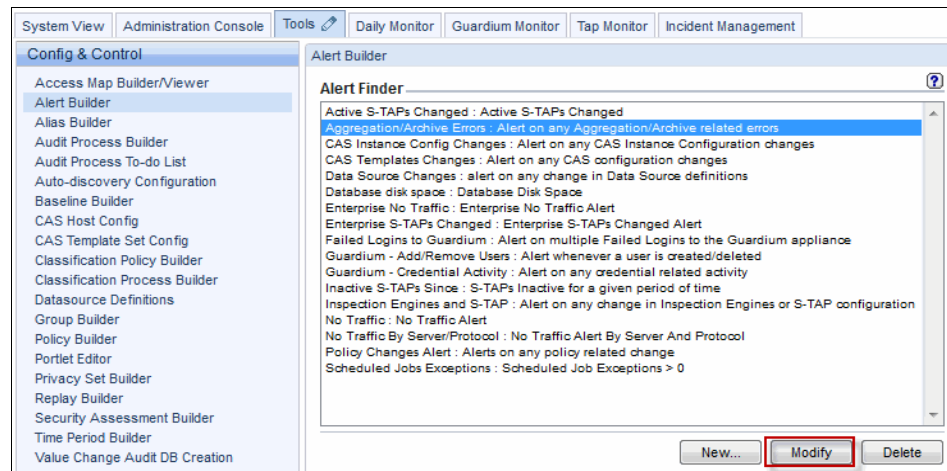


Figure 3-33 Selecting a predefined threshold alert to modify

3. Select **Active** to activate the alert.
4. Scroll down to Alert Receivers section and click **Add Receiver** to add a new receiver (the default receiver is SYSLOG).
5. In the Alert Receiver Selection window, select a Notification Type of MAIL and use the Search User feature or click **Alert Receiver** to select the account (with the email address) to receive this alert, as shown in Figure 3-34 on page 98.

The screenshot shows the 'Alert Builder' interface with the following configuration:

- Settings:**
 - Name: Aggregation/Archive Errors
 - Description: Alert on any Aggregation/Archive related errors
 - Category: Self Monitoring
 - Classification: Export
 - Severity: LOW
 - Run frequency: 1440 (minutes)
 - Active:
 - Log policy violation:
- Alert Definition:**
 - Query: Aggregation Errors
 - Query Parameters: [Empty field]
- Notification:**
 - Notification frequency: 1400 (minutes)
- Alert Receivers:**
 - SYSLOG [Delete](#)
 - [Add Receiver..](#) (highlighted with a red box)
- Buttons:** Add Comments, Roles..., Apply, Back

Figure 3-34 Activate a predefined alert and add a new receiver

6. After all of the receivers are added, click **Apply** to save the changes, including the activation of this alert.
7. When an alert is active, it is listed in the Anomaly Detection. To view the alert, click **Administration Console** → **Anomaly Detection**.
8. Repeat these steps for each of the other predefined alerts that are described in “Recommended self-monitoring threshold alerts” on page 96. For the Scheduled Jobs Exception alert, change the run frequency to 60 minutes.

Figure 3-35 on page 99 shows the user accounts with email addresses to add as email receivers.




Figure 3-35 The Alert Receiver Selection dialog

Figure 3-36 shows the list of active threshold alerts that are managed by the Anomaly Detection process.




Figure 3-36 List of active threshold alerts managed by the Anomaly Detection process

Configuring and activating user-created (custom) alerts

A threshold alert uses a query to test the required metric that is collected by the Guardium solution. In this section, we assume that you understand the use of the Guardium query builder to build the queries for the user-created (custom) alerts as described in “Recommended self-monitoring threshold alerts” on page 96.

In this section, we show the following user-create alerts examples:

- ▶ Sniffer restarts alert
- ▶ Disk space alerts

Sniffer restarts alert

Complete the following steps to create an alert for sniffer restarts if the sniffer on a collector restarts at least three times an hour:

1. Create a query by using the Sniffer Buffer Usage domain with the columns and fields as shown in Figure 3-37. There are no query conditions.

The screenshot displays the 'Sniffer Buffer Usage Tracking' interface. On the left, the 'Entity List' shows 'Sniffer Buffer Usage'. The main area is titled '-MySnifferRestart' and shows 'Main Entity: Sniffer Buffer Usage'. Below this, a table lists query fields:

Seq.	Entity	Attribute	Field Mode
1	Sniffer Buffer Usage	Timestamp	Count
2	Sniffer Buffer Usage	Sniffer Process ID	Count

A green arrow points from the 'Query Fields' table to the 'Result of running the above query' section. This section shows the query results for '-MySnifferRestart' with the following details:

- Start Date: 2013-06-28 13:19:58
- End Date: 2013-06-28 16:19:58
- Aliases: OFF
- Summary table:

Count of Timestamp	Count of Sniffer Process ID
176	1

The interface also includes a 'Records' control showing '1 to 1 of 1' and various utility icons at the bottom.

Figure 3-37 The sniffer restart query and sample output

2. Create and activate an alert for this query by clicking **New** on the **Tools** → **Config & Control** → **Alert Builder**.

Figure 3-38 on page 101 shows the custom Sniffer restart alert with key fields highlighted.

Figure 3-38 The custom Sniffer restart alert with key fields highlighted

Disk space alert

Complete the following steps to create two alerts, one for the Aggregator and another for the collector (user-created), that alert when the database available free space falls below a configured threshold:

1. Create two similar but separate queries by using the Sniffer Buffer Usage domain; one for the collector and the other for the aggregator because the database size is fixed on the collector but dynamic on the aggregator (up to the size of the Var partition).

Figure 3-39 on page 102 shows the query for the aggregator disk space alert.

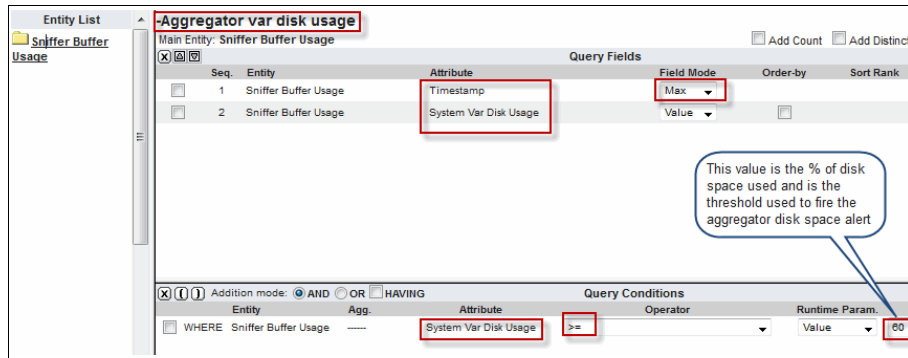


Figure 3-39 Query for the aggregator disk space alert

2. Create a similar query that is named -Collector MySQL disk usage. However, instead of using System Var Disk Usage, use the Mysql Disk Usage column.
3. Create and activate an alert for each query by clicking **New** on the **Tools** → **Config & Control** → **Alert Builder** and with the appropriate values for the aggregator.

Figure 3-40 on page 103 shows the aggregator disk space usage alert.

Alert Builder

Settings

Name:

Description:

Category:

Classification:

Severity:

Run frequency: (minutes)

Active:

Log policy violation:

Alert Definition

Query:

Accumulation interval: (minutes)

Note: Alerts run on aggregators will be based only on data within the defined merge period

Log full query results:

Column: (optional)

Alert Threshold

Threshold:

Alert when value is: threshold

Threshold Evaluated: per report
 per line

Threshold Used: As absolute limit
 As percentage change within period:
From: To:

As percentage change for the same "Accumulation Period" on a relative time:
Ending at:

Notification

Notification frequency: (minutes)

Figure 3-40 Aggregator disk space usage alert

Monitoring the monitor

Most customers choose to implement threshold alerts because of the following factors:

- ▶ The broader spectrum of available metrics to monitor.
- ▶ The automatic distribution to managed units by the CM in a managed environment.
- ▶ The finer control over the alerting; for example, setting thresholds, how often and when the alert fires.
- ▶ Email delivery of the alerts to assigned personnel.

If the appliance is down or incapacitated, however, the alerts might not run or be delivered. In this case, it is important to also monitor the appliance externally by using one of the following methods:

- ▶ SNMP polling; for example, periodically check processor or memory metric. If the poll is unsuccessful, assume that the appliance might be in trouble and take further steps.
- ▶ (Recommended) Use a URL monitoring tool to check the URL of the GUI login.
- ▶ A ping test, while not preferred, is better than no external monitoring because a ping response indicates only that the network is OK.

3.10 Vulnerability assessment

InfoSphere Guardium Vulnerability Assessment (VA) runs scheduled scans against selected database instances looking for vulnerabilities and platform-specific issues as reported by various database security organizations and the database vendors. Any vulnerabilities that are identified should be prioritized and remediated to help harden the database against threats.

VA can be deployed with another Guardium module that is called Configuration Change Audit (CAS). CAS is an agent-based utility that periodically scans the database configuration files at the operating system level for mis-configurations or changes.

The VA scan, which is referred to as an *assessment*, is a user-selected list of database-platform specific tests. These tests can be a combination of predefined (provided by IBM) or custom (user-created).

For VA to scan a database, the appliance from which the assessment is run connects to the database by using a minimum-necessary privileges database account through a JDBC connection. This connection is referred to as a data source.

VA does not require an S-TAP and if it is deployed without CAS, it sometimes referred to as *agent-less*.

Tests are added or updated by using a downloaded file that is available quarterly on IBM Fix Central. This file is referred to as the Database Protection Knowledge base Subscription (DPS).

Use the assessment-related `grdapi ()` functions, if available, to configure and manage VA.

For more information about VA and CAS usage, see the “Assess and Harden” chapter of the *Help Book Guardium*.

Figure 3-41 shows the high-level VA process flow (without CAS).




Figure 3-41 High-level VA process flow (without CAS)

3.10.1 Creating a database account

For VA to scan a database, the appliance from which the assessment is run connects to the database by using a minimum-necessary privilege database account.

Guardium provides the scripts per database platform, which has the minimum-necessary privileges to create the database role that is needed to successfully run the scan.

The scripts are available for download from IBM Fix Central (first, check Fix Central for an updated version of the scripts, if any) or from IBM Passport Advantage® portal. For example, search for the following information

“InfoSphere_Guardium_Database_User_Role_Definitions”

The DBA should review the script and be comfortable with the privileges granted, which are primarily read-only against selected catalog objects.

After the script is run by the DBA in the database instance to be scanned, the DBA provides the user name and password with other database configuration settings to create the Guardium data source.

Note: The script should be used to create the account to minimize errors during the scan because of missing privileges.

3.10.2 Creating a data source

By using the account information that is provided by the DBA (as described in 3.10.1, “Creating a database account” on page 105), create the data source (JDBC connection) on the Guardium appliance.

In a managed environment, the data sources are created and stored on the CM but available to the managed units. Complete the following steps:

1. Log in as admin to the GUI on the CM or stand-alone collector and browse to **Tools tab → Config & Control → Datasource Definitions**.
2. In the Application Selection window, select **Security Assessment** and click **Next**.
3. In the Datasource Finder window, click **New** to open the Datasource Definition window.

The inputs vary slightly based on the database type that is selected, but most of the information is provided by the DBA.

Figure 3-42 on page 107 shows a data source definition for an Oracle instance that uses the custom (not open source) JDBC driver.

Datasource Builder

Datasource Definition ?

Name: dev_ora_rhe54db_orcl102_sqlguard ← meaningful name

Database Type: Oracle Change to: Oracle (DataDirect) ← DataDirect is the opensource driver

Severity classification: NONE

Description:

Share Datasource: ← recommended

Authentication

Save Password:

Login Name: sqlguard

Password:

Location

Host Name/IP: 10.10.9.14

Port: 1521

Service Name: orcl102

Informix Server:

Schema:

Connection Property:

Custom URL:

CAS

Database Instance Account:

Database Instance Directory:

Roles

No roles have been assigned to this datasource Roles...

Buttons: Add Comments, Test Connection, Apply, Back

After 'Apply', click 'Test Connection' to initiate a test connection to the database instance. You should see a 'The datasource can be successfully connected' message

Figure 3-42 Data source definition

Tip: Use a naming standard for the data source names to easily identify from the name which environment, database platform, server, instance, and user, the data source is for; for example, prod_ora_dbsrv1_orainst_grduser.

Note: Use the `grdapi` data source functions to help create many data sources.

Custom JDBC drivers

JDBC drivers are included with InfoSphere Guardium. However, for certain database-specific features for Oracle or Microsoft SQL Server, you might have to download and install a custom (or vendor-provided JDBC driver); for example, to use Windows domain authentication for Microsoft SQL Server.

For more information, see the *Help Book Guardium*.

3.10.3 VA tests

An assessment is a selected list of database platform-specific tests that can be a combination of the following tests:

- ▶ IBM-provided (predefined) tests:
 - Database platform-specific SQL-based tests
 - Database-specific Common Vulnerabilities Exposures (CVE) tests, which are is maintained by the MITRE organization
 - Non-DBMS-specific Observed tests (rarely used)
- ▶ User-created (custom) tests that use SQL or procedural SQL; for example, TSQL.

The following test characteristics are available:

Some tests have modifiable thresholds; for example, the Oracle test DBA Profile FAILED_LOGIN_ATTEMPTS Are Limited can be set to check that this setting is less than or equal to 5.

- ▶ Some tests allow exceptions by using an exception group; for example: MS SQL Server, No Non-Exempt Public Privileges, can use an exception group to exclude specified accounts from this restriction.
- ▶ Tests have external references, for mixable, STIG DO0286.
- ▶ Test results include pass or fail recommendations.
- ▶ Tests can be deferred for a specific period by creating a test exception.
- ▶ Tests are assigned to the following categories and severity classifications:
 - Test categories:
 - Privileges: Checks users and system level rights.
 - Authentication: Password policies, profile limits, and so on.
 - Configuration: CVE tests and database-specific settings, such as Oracle listener.ora settings
 - Version: Version numbers and patch levels

- Other: Permissions on configuration files such as sqlnet.ora and directories (often requires CAS)
- Tests severity classification (can be changed by the user):
 - Critical
 - Major
 - Minor
 - Cautionary
 - Informational

Table 3-5 shows a few examples of predefined VA tests.

Table 3-5 Examples of predefined VA tests

DB platform	Test name	Description
Oracle	No Roles with the Admin Option	This test checks whether Oracle privileges were granted with the ADMIN option to users with no DBA role, which allows the grantee to make grants to other users. The ADMIN option reduces administrative control and creates an unwarranted vulnerability.
DB2	No Individual User Objects Privileges	This test checks for object privileges that are granted to individual users. Privileges that are granted to individual users are difficult to maintain and create risk of misuse. Roles of user groups should be created by the DB administrator and privilege grants made instead to these roles or groups.
MSSQLSERVER	No Select Privileges On System Tables/Views in Application Databases	This test checks for grants of the SELECT privilege on system tables in application databases. Users with these privileges can access sensitive information about other users' objects or data.

3.10.4 Creating an assessment

Because there are as many as hundreds of tests for the supported database platforms, it is useful to have requirements to help determine which tests to select for a database platform.

Tip: Create a report of all the available tests by database platform by using the Tools tab and clicking **Report Building** → **VA Test Tracking** domain in a similar format as shown in Figure 3-5 on page 109.

Complete the following steps to create an assessment:

- ▶ Log in as admin to the GUI on the CM or stand-alone collector and browse to the Tools tab and click **Config & Control** → **Security Assessment Finder**.
- ▶ Click **New** to open the Security Assessment Builder and complete the following steps:
 - a. Enter a meaningful description for this assessment (use a similar naming standard as for the data source name).
 - b. Click **Add Datasource** and select the data sources for this assessment. Select data sources for the same database type.
 - c. Click **Apply**, then click **Configure Tests** to start Assessment Test Selections.
 - d. Select the database tab (for example, Oracle) and click the **Tests available for addition** option to filter the tests that are displayed for selection.
 - e. By using your requirements for this platform, select the tests and click **Add Selections** to add them to the assessment.
 - f. Click **Return** when you are finished with the test selection to return to the Security Assessment Finder.

Note: As of this writing, assessments have a limit on the number of tests, which should not exceed 10,000 tests, as shown in the following example:

```
Assessment-A has 1 test and 2 datasources.  
Assessment-A therefore has 2 tests i.e. # of tests X # of datasources
```

To work with this limitation, clone the assessment and add a portion of the data sources to one assessment and the remainder to the clone.

Figure 3-43 on page 111 shows the creation of an assessment and assigning a data source (or multiple data sources).

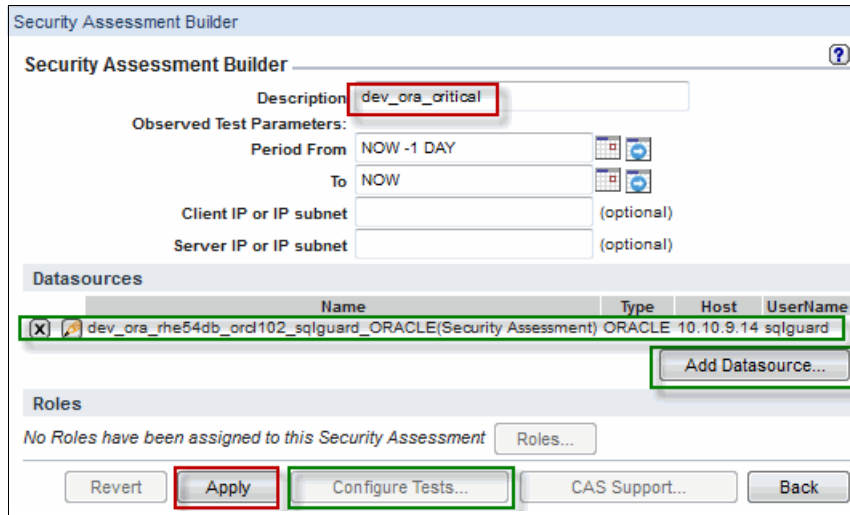


Figure 3-43 Creation of an assessment and assigning data sources

Figure 3-44 on page 112 shows the test selection process with other key test characteristics.




Figure 3-44 VA test selection process along with other test characteristics

Running an assessment

Complete the following steps to immediately run an assessment and view the results by using the predefined VA report:

1. Select the assessment in the Security Assessment Finder and click **Run Once Now** to queue the test for run.

As shown in Figure 3-41 on page 105, the listener process schedules the assessment to start within a minute or so if no other assessment is running.

2. After a few minutes (to allow sufficient time for the assessment to be scheduled and to complete), click **View Results** to start the predefined Security Assessment Results report.

Tip: To view the queue and run status, including how long it took the assessment to complete, browse to the Guardium Monitor tab and click **Guardium Job Queue**. However, you must return to the Security Assessment Builder to view the results

Figure 3-45 shows the assessment results by using the predefined report.

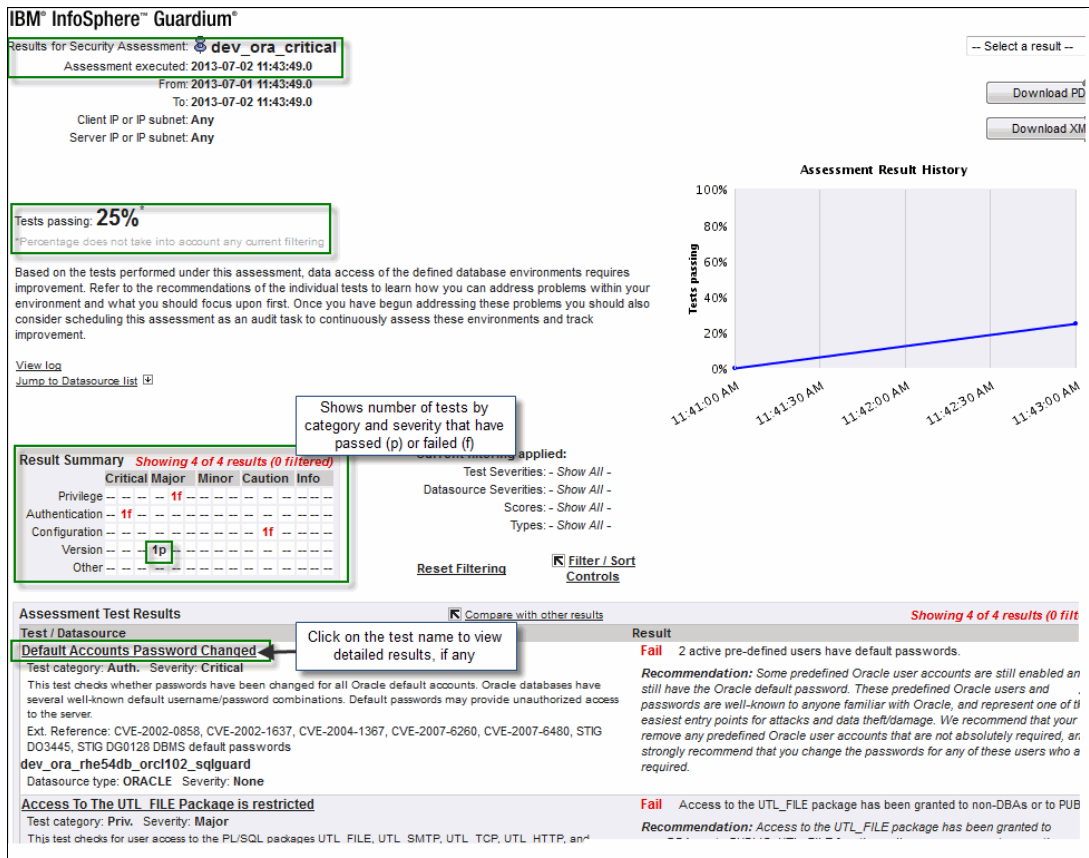


Figure 3-45 Assessment results using the predefined report

Figure 3-46 on page 114 shows the result details (if any) for a selected test from the predefined assessment report.

IBM InfoSphere Guardium

Results for Security Assessment: **dev_ora_critical**
 Assessment executed: 2013-07-02 11:43:49.0
 From: 2013-07-01 11:43:49.0
 To: 2013-07-02 11:43:49.0
 Client IP or IP subnet: Any
 Server IP or IP subnet: Any

Default Accounts Password Changed
 Test category: Auth. Test severity: Critical
 dev_ora_rhe54db_orcl102_sqlguard
 Datasource type: ORACLE Datasource severity: None

Fail

2 active pre-defined users have default passwords.

Short Description: This test checks whether passwords have been changed for all Oracle default accounts. Oracle databases have several well-known default username/password combinations. Default passwords may provide unauthorized access to the server.

External Reference: CVE-2002-0858, CVE-2002-1637, CVE-2004-1367, CVE-2007-6260, CVE-2007-6480, STIG DO3445, STIG DG0128 DBMS default passwords

Recommendation: Some predefined Oracle user accounts are still enabled and still have the Oracle default passwords to anyone familiar with Oracle, and represent one of the easiest entry points for attacks and data theft/damage. We are not absolutely required, and we strongly recommend that you change the passwords for any of these users who

Details:
 DIP
 SCOTT

Detailed results for selected test

PASS

FAIL

11:41:00

Figure 3-46 Result details for a selected test from the predefined assessment report

3.10.5 Custom reports

The predefined assessment report is useful for testing or refining assessments. However, for remediation or reporting purposes in large distributed environments, it is typical to generate and distribute custom reports by using the following different criteria:

- ▶ Failed tests by database platform by environment; for example, DB2 failed tests in production
- ▶ Failed tests for a particular category by database platform by environment; for example, critical DB2 failed tests in production
- ▶ Errored tests by platform and environment
- ▶ Tests details for failed tests by environment and platform
- ▶ Version and patch-related tests only, and so on

Complete the following steps to create custom tests by using the results from a prior scan. (By using this option, it is assumed that you understand creating Guardium queries and reports):

1. Create queries by using the Security Assessment Result Tracking domain and generate the reports.
2. After the scan completes, run the reports and adjust as needed.

Figure 3-47 shows an example of a custom query for a VA report.




Figure 3-47 Custom query for a VA report that is selecting specific columns for tests with a PASS status

3.10.6 Scheduling an assessment

For periodic scans, such as monthly or quarterly, schedule the assessments by using the Guardium Audit Process builder. For more information, see 5.5, “Reports” on page 174.

Remember the following points as you plan the schedule:

- For many assessments and tests (that is, # of tests x # of data sources), use a staggered schedule to minimize the effect on the appliance by spreading the assessments over a few days.

- ▶ As of this writing, there is a limitation on the number of tests in a single audit process. The limit is a total of 35,000 tests across all assessments. Do not exceed more than 10,000 tests per assessment; that is, # of tests X # of data sources.
- ▶ Avoid scheduling the assessments during the period when the databases are often scheduled for maintenance.

Scheduling custom reports

On the day following the last scheduled assessment, use the Audit Process Builder to schedule the run and delivery of the custom reports. Remember the following points:

- ▶ Assessments (scans) must be complete before reporting is done.
- ▶ Custom reporting might take as long as the actual scans, so plan on staggering the schedule, if necessary.

3.10.7 More information

See the “Assess and Harden” chapter of the *Help Book Guardium* for more information about the following topics:

- ▶ Creating custom VA tests (referred to as *custom queries*)
- ▶ Applying the DPS
- ▶ Creating exemptions (that is, modify the results) for certain tests by using the following methods:
 - The **grdapi ()** test exception process
 - Populating and using an exception group

3.10.8 High-level steps summary

The following high-level steps are involved in configuring and managing VA:

1. Create account on each database instance.
2. Obtain platform-specific, user-created scripts from Guardium.
3. Have DBAs review and run scripts.
4. Select the Guardium appliance from which to run VA; for example, aggregator.
5. Create a data source for each database instance.
6. Apply DPS (subscription).

7. Review your requirements and map to predefined tests (create custom tests, if necessary) for each DBMS type.
8. Create an assessment for each database platform by using the mapping.
9. Schedule the assessments by using the Audit process builder.
10. Create, schedule, and deliver custom assessment reports.
11. Remediate findings.
12. Maintain environment and assessment tests.

3.11 Where to find more help

In this section, we describe the various channels for obtaining product documentation and support, including peer support.

3.11.1 Product documentation

You can download the product manual, *Help Book Guardium*, from the IBM Passport Advantage portal, or downloaded from within the Guardium application, as described in this section.

Product documentation is available online by using the application GUI for context-sensitive help, searching, or for downloading to a PDF format. Figure 3-48 shows the Application Context Help icon.



Figure 3-48 Application Context help icon to get help on the current topic

Figure 3-49 shows how to access the Help System to search or save the product help in PDF format.

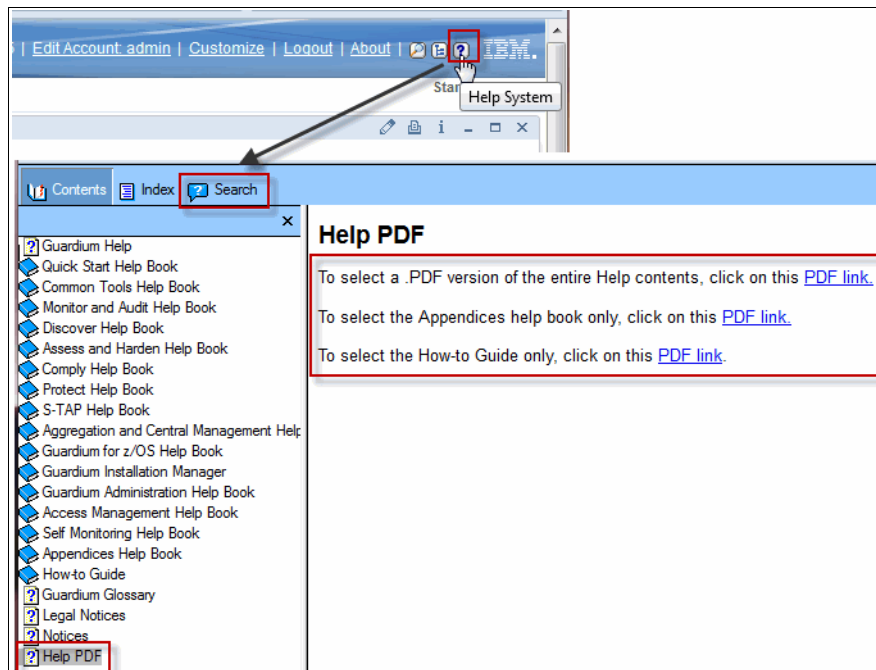


Figure 3-49 Accessing the Help System to search or save the product help in PDF format

The CLI also provides command syntax and usage help:

- ▶ To see a list of available commands, enter `comm` or `command`.
- ▶ To see commands with the string “interface”, enter `comm inter`.

3.11.2 Release notifications and bulletins

Sign up at the IBM Passport Advantage portal to be notified through email of new releases and product bulletins. The portal is available at this website:

http://www-01.ibm.com/software/lotus/passportadvantage/pao_customer.htm
1

3.11.3 Product support

For InfoSphere Guardium product support, contact IBM Guardium Support by using one of the following methods:

- ▶ Online: Passport Advantage (account required)
- ▶ Phone: 800-426-7378 (US/Canada Only)

3.11.4 User community and support

The following product and usage information sources are available:

- ▶ Guardium Developer Works Forum:
<https://www.ibm.com/developerworks/forums/forum.jspa?forumID=2648&start=45>
- ▶ Guardium DAM User Group on LinkedIn:
<http://www.linkedin.com/groups/IBM-Guardium-Database-DB-Activity-3746392>
- ▶ Guardium YouTube Channel:
<http://www.youtube.com/user/InfoSphereGuardium>
- ▶ Guardium Information Center:
<http://publib.boulder.ibm.com/infocenter/igsec/v1/index.jsp>
- ▶ Guardium Data Security Library:
<http://www-01.ibm.com/software/data/guardium/library.html>



Monitoring and auditing

This chapter introduces regulatory requirements and how Guardium functionality can be mapped into these requirements. There are many different regulations that customers must be concerned about to protect sensitive and personal information. This is not an inclusive list by any means, and the following regulations are concerned with accessing financial or sensitive data:

- ▶ Sarbanes-Oxley (SOX)
- ▶ Payment Card Industry (PCI) Data Security Standard
- ▶ Gramm–Leach–Bliley
- ▶ European Union Directive 95/46/EC
- ▶ Health Insurance Portability and Accountability Act (HIPAA)
- ▶ ISO 27002

Many of these regulations provide a framework for protecting information or validating the integrity of information as it is used within a corporation. It is important to understand the strategy to be in compliance with these regulations from a Database Activity Monitoring (DAM) perspective, and adopt industry best practices in this approach.

This chapter includes the following topics:

- ▶ Regulations and compliance
- ▶ Auditing categories
- ▶ Auditing requirements
- ▶ Database activity monitoring
- ▶ Vulnerability assessment
- ▶ Mapping audit requirements to the solution

4.1 Regulations and compliance

The regulations were developed to ensure that corporations comply with these regulations because most were overlooking these best practices in their day-to-day activity.

Figure 4-1 lists the most common regulations that we describe in this chapter to provide you a framework of understanding how the Guardium functionality can help you achieve regulatory compliance for DAM.

Audit Requirements	PCI DSS	COBIT (SOX)	ISO 27002	Data Privacy & Protection Laws	NIST SP 800-53 (FISMA)
1. Access to Sensitive Data (Successful/Failed SELECTs)	✓		✓	✓	✓
2. Schema Changes (DDL) (Create/Drop/Alter Tables, etc.)	✓	✓	✓	✓	✓
3. Data Changes (DML) (Insert, Update, Delete)		✓	✓		
4. Security Exceptions (Failed logins, SQL errors, etc.)	✓	✓	✓	✓	✓
5. Accounts, Roles & Permissions (DCL) (GRANT, REVOKE)	✓	✓	✓	✓	✓

Figure 4-1 Audit requirements versus regulation

4.2 Auditing categories

When you are monitoring databases, Guardium typically defines the following general auditing categories to help articulate what must be monitored from the flow of information into and out of the database:

- ▶ Privilege user
- ▶ Sensitive object
- ▶ Comprehensive
- ▶ Comprehensive with extrusion and values

Each of these categories should be understood in context with the way that databases communicate to the users or applications by using the database. For example, when someone enters the following information:

```
SQL> select CARDID,FIRSTNAME, LASTNAME, CARDNUMBER from creditcard where
CARDID=3;
```

CARDID	FIRSTNAME	LASTNAME	CARDNUMBER
3	Joe	Jones	1234567890123453

SQL>

The input information flows from client to the database, is processed, then the result is returned to the client, as shown in Figure 4-2.




Figure 4-2 Database information flow to clients

In this example, the SELECT statement is considered the SQL query or SQL access, as shown in the following example:

```
select CARDID,FIRSTNAME, LASTNAME, CARDNUMBER from creditcard where CARDID=3;
```

This is how the user or application interacts with the database to get the information that wanted.

The information that is returned from the database is called the “result set”, as shown in the following example:

CARDID	FIRSTNAME	LASTNAME	CARDNUMBER
3	Joe	Jones	1234567890123453

In some cases, you want to audit this information to see who accessed sensitive information, such as credit card data. Remember that this can generate much information and you must plan accordingly for your storage requirements.

If someone runs the following SQL statement, the error message (ORA-01031) is returned:

```
SQL> grant select on joe.creditcard to john;  
grant select on joe.creditcard to john
```

*

```
ERROR at line 1:  
ORA-01031: insufficient privileges  
SQL>
```

The error message “ORA-01031: insufficient privileges” is an exception message that is returned to the user from the database server. These exception messages should always be included in the auditing policy, if possible.

The following types of activities can be monitored:

- ▶ SQL Query (or SQL Access)
- ▶ Result Sets
- ▶ Exceptions

4.2.1 Privilege user monitoring

Privilege user monitoring is defined as monitoring database users that have elevated privileges to access sensitive information within the database. Table 4-1 lists the privilege users for some of the regulations.

Table 4-1 Regulation versus privilege users

Regulation	Privilege user	Database users
Sarbanes-Oxley (SOX)	Anyone that can modify the integrity of the financial records inside the database. This includes any SQL changes, such as insert, update, delete (directly or indirectly) through stored procedures, functions, and views.	<ul style="list-style-type: none">▶ Oracle: system, sys▶ DB2: db2admin, db2inst1▶ SQL Server: sa, administrator▶ informix: informix▶ sybase: sa▶ Other: All other database users that have DML capability on financial objects that are defined in the database

Regulation	Privilege user	Database users
PCI	<p>Anyone that can view credit card or cardholder data inside the database.</p> <p>This includes any SQL select statements or indirect access through procedures, functions, views, and so on.</p>	<ul style="list-style-type: none"> ▶ Oracle: System, sys ▶ DB2: db2admin, db2inst1 ▶ SQL Server: sa, administrator ▶ informix: informix ▶ sybase: sa ▶ Other: All other database users that have select capability on creditcard objects that are defined in the database
HIPAA	<p>Anyone that can view patient information inside the database.</p> <p>This includes any SQL select statements or indirect access through procedures, functions, views, and so on.</p>	<ul style="list-style-type: none"> ▶ Oracle: System, sys ▶ DB2: db2admin, db2inst1 ▶ SQL Server: sa, administrator ▶ informix: informix ▶ sybase: sa ▶ Other: All other database users that have select capability on HIPAA objects that are defined in the database
European Union Directive	<p>Anyone that can view personal information inside the database.</p> <p>This includes any SQL select statements or indirect access through procedures, functions, views, and so on.</p>	<ul style="list-style-type: none"> ▶ Oracle: System, sys ▶ DB2: db2admin, db2inst1 ▶ SQL Server: sa, administrator ▶ informix: informix ▶ sybase: sa ▶ Other: All other database users that have select capability on objects that contain personal information that is defined in the database

Privilege user monitoring is a key requirement for any database security policy because it is these accounts that can be misused or hacked. When you have access to these accounts, you want to ensure (through a verifiable audit trail) that no activity from these database accounts violates your security policies.

The typical privilege user monitoring includes SQL queries and exceptions, but not result set.

4.2.2 Sensitive object monitoring

Sensitive object monitoring is defined as monitoring access to database objects (tables, views, and so on) that contain sensitive information. This includes who accessed these objects, at what time, from what IP address, and the specific SQL statement that was used, as shown in Figure 4-3.

OS User	Server Type	Client IP	Server IP	Network Protocol	Session Id	DB User Name	Source Program	Full Sql	Timestamp
ROOT	ORACLE	10.10.9.56	10.10.9.56	BEQUEATH	20141	SYSTEM	SQLPLUS@OSPREY	select * from creditcard	2013-03-18 13:15:25.0

Figure 4-3 Access to CreditCard object

Almost every enterprise application that is written these days has a database infrastructure with the application. It is critical to understand what tables and views hold the sensitive information for these applications. For example, consider an SAP application that stores credit card information. In this configuration, the table `but0cc` should be included in the group of objects to be monitored because it contains credit card information, as shown in Figure 4-4.

```
30
39 select MANDT, PARTNER, CCARD_ID, CCINS, CCNUM, CCDEF, CCACCNAME, CARD_GUID from but0cc
40
41
42
```

49:1	INS							
1	MANDT	PARTNER	CCARD_ID	CCINS	CCNUM	CCDEF	CCACCNAME	CARD_GUID
2	-----	-----	-----	-----	-----	-----	-----	-----
3	800	10001	632	Discover	6011261437467568	1	Joe DiPietro	10001

Figure 4-4 SAP object with credit card information

When these sensitive objects are monitored, this information can be accessed by using one of the following methods:

- ▶ Through the application (such as SAP)
- ▶ Direct access from the database

It is important that you monitor these database objects from both access methods to ensure that access to sensitive information has appropriate security controls.

The good news is that there are automated tools that can help you identify where your sensitive information is located, as shown in Figure 4-5 on page 128.

Catalog	Schema	Table Name	Column Name	Rule Description	Comments
<input type="checkbox"/>	JOE	CC1	CARDNUMBER	V9 PoT PCI Classification Rule	Date: Thursday, November 21, 2013 11:08:46 AM EST Datasource: DB2 10.10.9.57: 50000 sample Object: JOE.CC1 CARDNUMBER Category: 'PCI' Classification: 'CreditCard' Comprehensive: true Rule: Search For Data: V9 PoT PCI Classification Rule TABLE_TYPE='TABLE', DATA_TYPE='TEXT', COLUMN_NAME_LIKE='%CARD%', SEARCH_VALUE_PATTERN='[0-9]{16}', SHOW_UNIQUE_VALUES='false' Action: Add to Group of Objects: add_to_sensitive_objects Object Group='Sensitive Objects', Replace Group Content='false', Add Member Type='FULLNAME'
<input type="checkbox"/>	JOE	PATIENT	CARDNUMBER	V9 PoT PCI Classification Rule	Date: Thursday, November 21, 2013 11:08:46 AM EST Datasource: DB2 10.10.9.57: 50000 sample Object: JOE.PATIENT CARDNUMBER Category: 'PCI' Classification: 'CreditCard' Comprehensive: true Rule: Search For Data: V9 PoT PCI Classification Rule TABLE_TYPE='TABLE', DATA_TYPE='TEXT', COLUMN_NAME_LIKE='%CARD%', SEARCH_VALUE_PATTERN='[0-9]{16}', SHOW_UNIQUE_VALUES='false' Action: Add to Group of Objects: add_to_sensitive_objects Object Group='Sensitive Objects', Replace Group Content='false', Add Member Type='FULLNAME'

Figure 4-5 Classifier results of sensitive data

The typical sensitive objective monitoring includes SQL queries and exceptions, but not result set.

4.2.3 Comprehensive monitoring

Comprehensive monitoring is defined as monitoring all access into the database. In many audit policies, there are different connections that you might ignore such as nightly batch or extract, transform, and load (ETL) jobs. You want to record when these sessions started, but not necessarily record the contents of these jobs because they might contain millions and millions of SQL statements.

4.3 Auditing requirements

Some industries are more formalized than others in their security controls. For example, in the financial services sector, they audit many systems and the granular details can include all three types of activities SQL query, result sets, and exceptions.

Database Activity Monitoring Maturity Model helps customers identify how to improve their security posture for accessing and protecting sensitive information. Figure 4-6 shows different maturity levels that companies try to achieve.

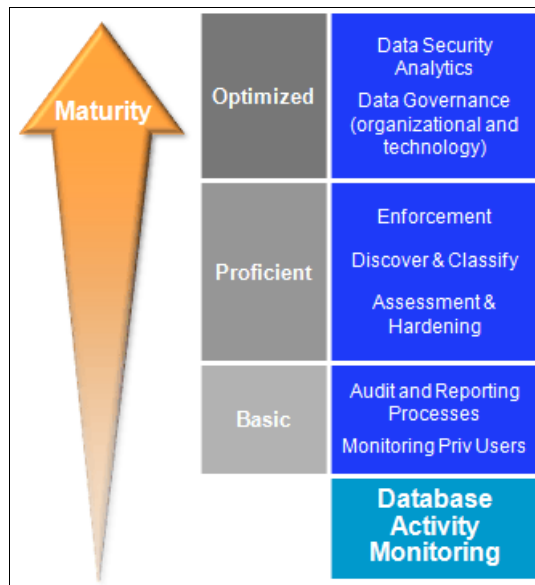


Figure 4-6 Database activity monitoring maturity model

Figure 4-6 also shows the following different levels of maturity as it relates to database activity monitoring:

- ▶ Basic
- ▶ Proficient
- ▶ Optimized

All customers should strive to be in the optimized maturity level. However, a realistic self evaluation helps you understand what area you currently evaluate yourself, and where you must close the gaps. We define the following levels to put some context around the following terms:

▶ Basic

Basic maturity includes the following content:

- Separation of duties
- Well-defined audit processes with cross departmental teams (information security, database administration and architecture, audit, risk and compliance, and so on)
- Basic database monitoring to include privilege user activity
- Real time alerting for security policy violations
- Database entitlements review for understanding what privileges the users that are defined in the system have to access sensitive information

▶ Proficient

Proficient maturity includes all of the items that are listed in the Basic level and the following content:

- Integration with external security systems, such as Security Information Event Managers (SIEM)
- Organizational procedures that are defined for reviewing audit data to enforce separation of duties
- A policy that proactively blocks unauthorized connections to the database
- Understanding of where your sensitive data is within the database
- Regular vulnerability assessments to ensure that the database is configured to the appropriate security standard for the organization
- Regular scanning for sensitive data and classification of sensitive data

▶ Optimized

Optimized maturity includes all of the items that are listed in the Proficient level and the following content:

- Data governance areas that identify processes and procedures to help automate the governance of data access and handling of audit reports
- Automatic update of security policy to include the newly identified sensitive tables and database objects that are identified in the proficient stage to reduce the risk of unauthorized access

- Updated list of connection profile, which includes authorized connections to the database with Client IP, Server IP, Source Program, and other connection information. For more information about this process, see this website:

<http://youtu.be/yRoRkAExVz0>

- Workflow with the business owners to include their approval and understanding of the risk that is associated with new connections to the database that might access their information. For more information about this process, see this website:

<http://youtu.be/NwndWdCmAic>

- Forensic analysis for outliers in the audit data to understand what is abnormal behavior

This list is a small subset of items that help organizations understand where they are in the maturity model to help them improve their security posture.

4.4 Database activity monitoring

Database activity monitoring is the ability to monitor and audit activities into and out of the database server. These events often are SQL events, such as Select, Insert, Update, Delete, Drop, Create, and Alter. However, Database Activity Monitoring evolved over to include more items than only the activity into the database. As described in 4.3, “Auditing requirements” on page 129, the following other areas are important in database activity monitoring as well:

- ▶ Blocking capabilities to prevent unauthorized access
- ▶ Data discovery and classification to help you locate sensitive data
- ▶ Importing external information to enrich the audit reports
- ▶ Identifying the application user from a pooled database user connection
- ▶ Monitoring entitlement reports to identify privilege users and high risk accounts that elevated privileges to sensitive data
- ▶ Vulnerability assessment to identify gaps in the configuration of the database server

4.5 Vulnerability assessment

Vulnerability assessment (VA) is a critical process in the security of database servers. VA is the process of identifying the following types of potential issues:

- ▶ Database tier issues
- ▶ Operating system (OS) tier issues
- ▶ Database user activity in how they are using the database issues

Figure 4-7 shows three types of VA categories. Each of these categories helps identify potential issues that should be resolved through your audit processes.




Figure 4-7 Three types of vulnerability assessment categories

The VA process identifies the following potential issues:

- ▶ Database tier issues

To assess the database tier, you include the following typical tests:

- Identifying patches
- Who has permissions and entitlements to database objects
- Strong password policies to log in to the database
- Configuration information

Automated tests to verify correct patch levels for database servers is a critical part of the security process. Figure 4-8 on page 133 shows a sample result of the Guardium patch VA test.

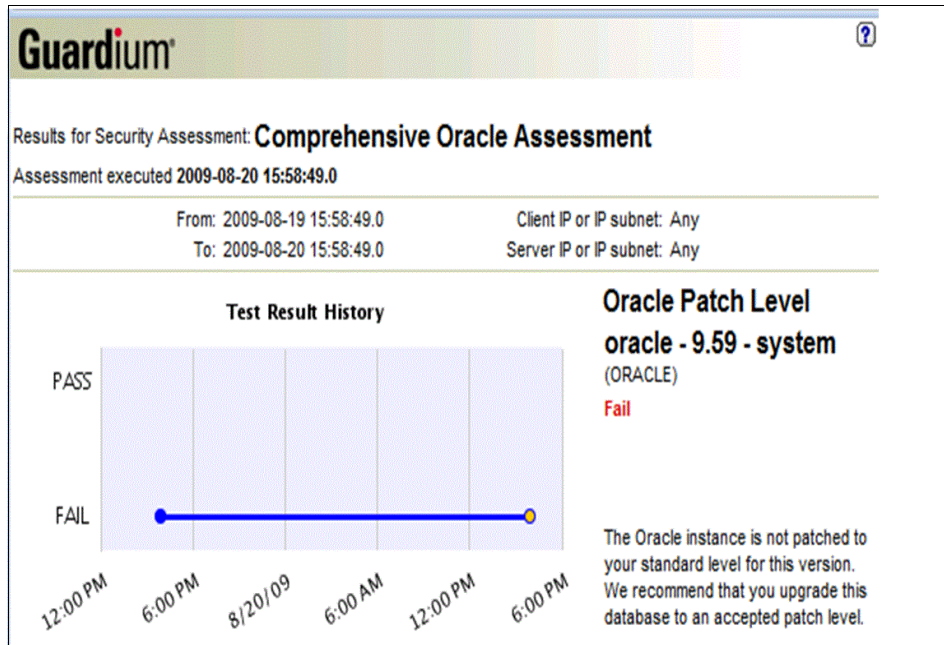


Figure 4-8 Patch vulnerability assessment test

► OS tier issues

OS issues can include permissions on files, environment variables, and registry settings.

Figure 4-9 shows an example of someone changing permissions for a critical file. Permission change from `rwxr-xr-w`, which means only the owner can change the file, to `rwxrwxrwx` so that anyone can change the file.

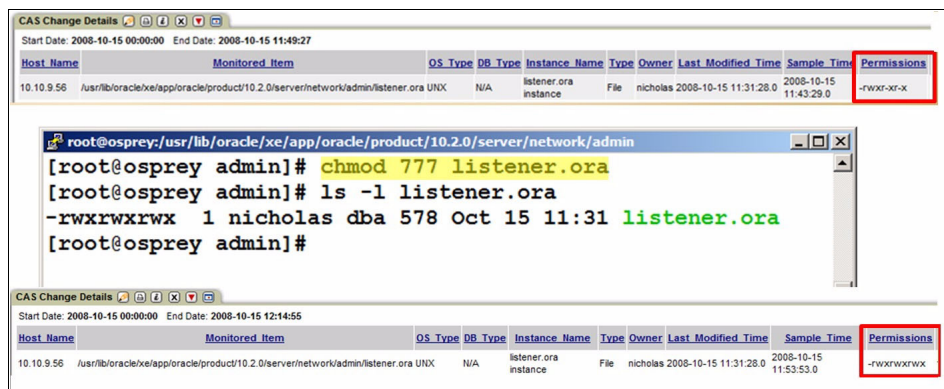


Figure 4-9 Permission change for listener.ora file

In this case, the risk is that someone can modify the configuration file and change to which Transmission Control Protocol (TCP) port the database server is listening. You want to ensure that you are monitoring this type of activity so that you can identify the differences if a modification to the file is made.

Figure 4-10 shows an example of monitoring the changes to the Operating System files. In this case, the listener port was changed from port 1525 to port 1529.

Guardium		Selected Record Differences	
New		Previous	
Line #13		Line #13	
013: (DESCRIPTION_LIST =	013: (DESCRIPTION_LIST =	013: (DESCRIPTION_LIST =	013: (DESCRIPTION_LIST =
014: (DESCRIPTION =	014: (DESCRIPTION =	014: (DESCRIPTION =	014: (DESCRIPTION =
015: (ADDRESS = (PROTOCOL = IPC)(KEY = EXTPROC_FOR_XE))	015: (ADDRESS = (PROTOCOL = IPC)(KEY = EXTPROC_FOR_XE))	015: (ADDRESS = (PROTOCOL = IPC)(KEY = EXTPROC_FOR_XE))	015: (ADDRESS = (PROTOCOL = IPC)(KEY = EXTPROC_FOR_XE))
016: (ADDRESS = (PROTOCOL = TCP)(HOST = osprey)(PORT = 1521))	016: (ADDRESS = (PROTOCOL = TCP)(HOST = osprey)(PORT = 1521))	016: (ADDRESS = (PROTOCOL = TCP)(HOST = osprey)(PORT = 1521))	016: (ADDRESS = (PROTOCOL = TCP)(HOST = osprey)(PORT = 1521))
017:)	017:)	017:)	017:)
018:)	018:)	018:)	018:)
019:)	019:)	019:)	019:)
020: LISTENER =	020: LISTENER =	020: LISTENER =	020: LISTENER =
021: (DESCRIPTION_LIST =	021: (DESCRIPTION_LIST =	021: (DESCRIPTION_LIST =	021: (DESCRIPTION_LIST =
022: (DESCRIPTION =	022: (DESCRIPTION =	022: (DESCRIPTION =	022: (DESCRIPTION =
023: (ADDRESS = (PROTOCOL = TCP)(HOST = osprey)(PORT = 1529))	023: (ADDRESS = (PROTOCOL = TCP)(HOST = osprey)(PORT = 1525))	023: (ADDRESS = (PROTOCOL = TCP)(HOST = osprey)(PORT = 1525))	023: (ADDRESS = (PROTOCOL = TCP)(HOST = osprey)(PORT = 1525))
024:)	024:)	024:)	024:)
025:)	025:)	025:)	025:)
026:)	026:)	026:)	026:)
027: DEFAULT_SERVICE_LISTENER = (XE)	027: DEFAULT_SERVICE_LISTENER = (XE)	027: DEFAULT_SERVICE_LISTENER = (XE)	027: DEFAULT_SERVICE_LISTENER = (XE)
028:)	028:)	028:)	028:)

Figure 4-10 Changes to the Listener.ora file

Best practices for VA include operating system components to ensure that you have a complete view of the database environment.

- ▶ Database user activity in how they are using the database issues

There are some common tests that can help you identify how users interact with the database and see whether the appropriate controls are actively enforced.

One such test is excessive administrative log ins. If the total amount of traffic that is observed shows that administrators are logging in to the database server as opposed to regular, non-privileged users, there is a separation of duties issues with your configuration. In this particular case, the amount of logins can help you identify certain users that have job functions with elevated privileges.

Why is this a problem? These administrative users have access to, or the ability within the database to get access to, sensitive information that is outside of their job responsibility.

From a security best practices perspective, you should use the concept of least privileges, which provide only enough privileges for the user to perform their job. There is no reason why a database administrator should be able to see someone's salary information.

Another such test is the amount of after hours logins or sharing of database credentials. The following examples of vulnerability assessment include database user activity:

- Excessive SQL errors
- After hours logins
- Excessive administrator logins
- Checks for calls to extended stored procedures
- Checks that user IDs are not accessed from multiple IP addresses
- Default users access
- Access rule violations
- Execution of Admin, DDL, and DBCC commands directly from the database clients
- Excessive login failures

Figure 4-11 shows some typical types of observed database activity behavior to understand if the database is used correctly.

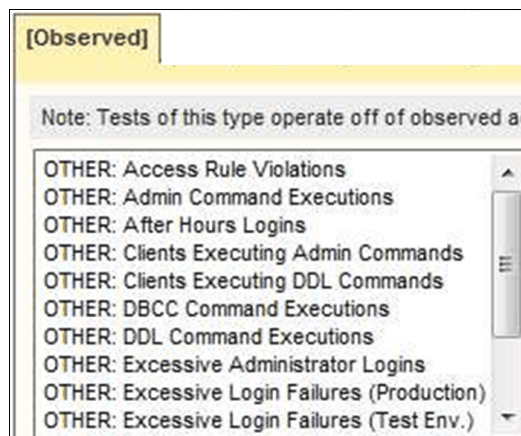


Figure 4-11 Observed database activity that should be monitored

4.6 Mapping audit requirements to the solution

In this section, we describe a few key compliance drivers and how the Guardium solution can help satisfying the following compliance requirements:

- ▶ PCI DSS
 - Requirement 2: Do not use vendor-supplied defaults for system passwords.
 - Requirement 3:
 - Protect stored cardholder data.
 - Compensating control for column-level encryption.
 - Requirement 6: Identify systems missing patches and enforce change controls.
 - Requirement 7: Compensating control for network segmentation.
 - Requirement 8: Assign a unique ID to each person with computer access.
 - Requirement 10: Track and monitor all access to cardholder data.
 - (Requirement 11: Regularly test systems.
- ▶ SOX, MAR (NAIC), and COBIT:
 - Prevent unauthorized changes to the financial CRM, ERP, and HR data.
 - Includes changes to data (DML) and schemas (DDL).

4.6.1 Requirement 2: Do not use vendor-supplied defaults for system passwords

The PCI requirement about the system password states:

“2.2 - Develop configuration standards for all system components. Assure that these standards address all known security vulnerabilities and consistent with industry accepted system hardening standards”.

Guardium VA module provides the capability to address this requirement and that the module has predefined test for weak passwords for all supported platforms, as shown in Figure 4-12.

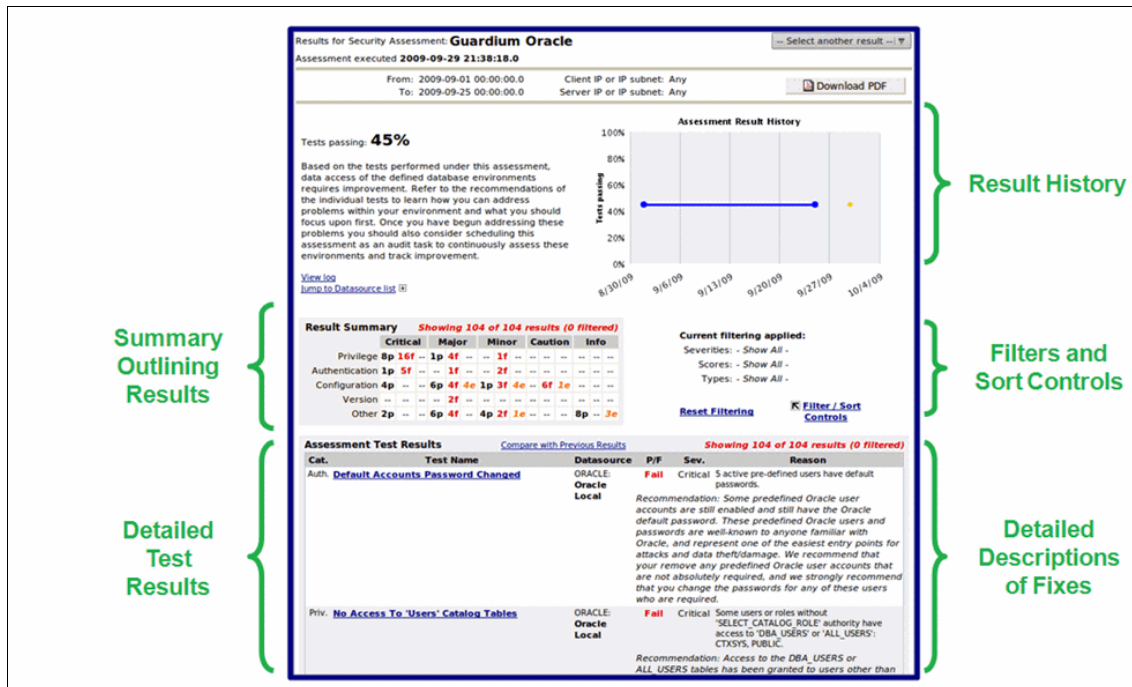


Figure 4-12 VA module default password test

For more information about setting the VA module, see 5.9, “Vulnerability assessment setup” on page 191.

The Guardium Entitlements reports provide the functionality to check privileges that are assigned to users. Figure 4-13 shows an example of an Oracle report.

Grantee	User Or Role	System Privilege	Admin Option
ALEX	User	CREATE SESSION	NO
ALEXSA	User	CREATE SESSION	NO
ANONYMOUS	User	CREATE SESSION	NO
APEX_PUBLIC_USER	User	CREATE SESSION	NO
AQ_ADMINISTRATOR_ROLE	Role	CREATE EVALUATION CONTEXT	YES
AQ_ADMINISTRATOR_ROLE	Role	CREATE RULE	YES
AQ_ADMINISTRATOR_ROLE	Role	CREATE RULE SET	YES
AQ_ADMINISTRATOR_ROLE	Role	DEQUEUE ANY QUEUE	YES
AQ_ADMINISTRATOR_ROLE	Role	ENQUEUE ANY QUEUE	YES
AQ_ADMINISTRATOR_ROLE	Role	MANAGE ANY QUEUE	YES
BI	User	ALTER SESSION	NO
BI	User	CREATE CLUSTER	NO
BI	User	CREATE DATABASE LINK	NO
BI	User	CREATE INDEXTYPE	NO
BI	User	CREATE OPERATOR	NO
BI	User	CREATE PROCEDURE	NO
BI	User	CREATE SEQUENCE	NO
BI	User	CREATE SESSION	NO
BI	User	CREATE SYNONYM	NO
BI	User	CREATE TABLE	NO
BI	User	CREATE TRIGGER	NO
BI	User	CREATE TYPE	NO
BI	User	CREATE VIEW	NO
BI	User	UNLIMITED TABLESPACE	NO
CONNECT	Role	CREATE SESSION	NO
CTXSYS	User	ALTER SESSION	NO
CTXSYS	User	CREATE CLUSTER	NO
CTXSYS	User	CREATE INDEXTYPE	NO
CTXSYS	User	CREATE OPERATOR	NO
CTXSYS	User	CREATE PROCEDURE	NO
CTXSYS	User	CREATE PUBLIC SYNONYM	NO
CTXSYS	User	CREATE SEQUENCE	NO
CTXSYS	User	CREATE SESSION	NO
CTXSYS	User	CREATE SYNONYM	NO
CTXSYS	User	CREATE TABLE	NO
CTXSYS	User	CREATE TRIGGER	NO
CTXSYS	User	CREATE TYPE	NO
CTXSYS	User	CREATE VIEW	NO
CTXSYS	User	DROP PUBLIC SYNONYM	NO
CTXSYS	User	UNLIMITED TABLESPACE	NO
DATAPUMP_EXP_FULL_DATABASE	Role	ADMINISTER RESOURCE MANAGER	NO
DATAPUMP_EXP_FULL_DATABASE	Role	ADMINISTER SQL MANAGEMENT OBJECT	NO
DATAPUMP_EXP_FULL_DATABASE	Role	BACKUP ANY TABLE	NO
DATAPUMP_EXP_FULL_DATABASE	Role	CREATE SESSION	NO
DATAPUMP_EXP_FULL_DATABASE	Role	CREATE TABLE	NO
DATAPUMP_EXP_FULL_DATABASE	Role	EXECUTE ANY PROCEDURE	NO
DATAPUMP_EXP_FULL_DATABASE	Role	EXECUTE ANY TYPE	NO
DATAPUMP_EXP_FULL_DATABASE	Role	READ ANY FILE GROUP	NO
DATAPUMP_EXP_FULL_DATABASE	Role	RESUMABLE	NO
DATAPUMP_EXP_FULL_DATABASE	Role	SELECT ANY SEQUENCE	NO

Figure 4-13 Oracle entitlement report

For more information about setting Entitlement Reports module, see 5.11, “Entitlement reporting setup” on page 202.

4.6.2 Requirement 3: Protect stored cardholder data

Guardium features predefined policies and rules to protect data. Figure 4-14 shows an example of policy rule that uses predefined PCI track data patterns for evaluation.

3.2.1 - Do not store the full contents of any track from the magnetic stripe (located on the back of a card, equivalent data contained on a chip, or elsewhere). This data is alternatively called full track, track, track 1, track 2, and magnetic-stripe data.

Access Rule Definition

Rule #14 of policy - PCI, SAP Production

Description: guardium://PCI_TRACK_DATA Unauthorized Users - Violation

Category: PCI, SAP Classification: Violation Severity: LOW

Not Server IP and/or Group

Not Client IP and/or Group

Not Client MAC and/or Group

Not Net Prtcl. and/or Group

Not DB Type

Not Svc. Name and/or Group

Not DB Name and/or Group

Not DB User and/or Group (Public) Authorized Users

Not Client IP/Src App./DB User/Server IP/Svc. Name

Not App. User and/or Group

Not OS User and/or Group

Not Src App. and/or Group

Not Field and/or Group

Not Object and/or Group

Not Command and/or Group

Object/Cmd. Group

Object/Field Group

Pattern

XML Pattern

App Event Exists Event Type Event User Name

App Event Values Text and/or Group

Numeric Date Replacement Character

Time Period

Minimum Count 0 Reset Interval 0 minutes Message Template Default

Quarantine for 0 minutes Records Affected Threshold 0 Rec. Vals. Cont. to next rule

Actions

LOG ONLY

guardium://PCI_TRACK_DATA

Detects two patterns of magnetic stripe data. The first pattern consists of a semi-colon (;), 16 digits, an equal sign (=), 20 digits, and a question mark (?). For example:

;1111222233334444=11112222333344445555?

The second pattern consists of a percent sign (%), the character B, 16 digits, a caret (^), a variable-length character string terminated by a forward slash (/), a second variable-length character string terminated by a caret (^), 31 digits, and a question mark (?). For example:

%B1111222233334444^xxx/xxxx x^111122223333444455566667777888?

- **DO NOT store the full contents of any track from the magnetic stripe**
- **DO NOT store the card-validation code (three-digit or four-digit value printed on the front or back of a payment card (e.g., CVV2 and CVC2 data))**
- **DO NOT store the PIN Verification Value (PVV)**

Figure 4-14 PCI track data pattern in policy rule

4.6.3 Requirement 6: Identify systems missing patches and enforce change controls

The Guardium VA module includes predefined test for patch levels and missing patches for all supported platforms. The Guardium monitoring supports change control process.

Chapter 4. Monitoring and auditing 139

Figure 4-15 shows an example of the change control reconciliation report.

6.4.5 - Change control procedures for the implementation of security patches and software modifications. Procedures must include the following:

6.4.5.2 - Documented change approval by authorized parties.

- **Automatic change control reconciliation using bracketing / delineating beginning and end of specific activity for specific DBA**

- **Integrating Remedy, Peregrine or other change control system by importing change control instructions and reporting back when DB change occurs**

Figure 4-15 Change Control reconciliation

4.6.4 Requirement 7: Compensating control for network segmentation

The PCI requirements states:

“7.1 - Limit access to system components and cardholder data to only those individuals whose job requires such access.”

The Guardium Data Level Access Control module provides the capability to address this requirement and protects sensitive cardholder data by ending or quarantining suspicious or malicious database activity while allowing valid access to the same sensitive PCI data (which is available in distributed supported platforms in version 9).

For more information about setting the Data Level Access Control module, see 5.8, “Data level access control” on page 189.

4.6.5 Requirement 8: Assign a unique ID for each person with a computer access

Guardium monitoring provides the capability to log necessary information and reports that can be used to satisfy the PCI requirement 8.

In many environments, users log in with their OS account and then switch to a generic shell account that has the required privilege to access the database. You can use the User Identification (UID) chain functionality to identify the privileged users who use the generic OS accounts.

Figure 4-16 shows an example of database access through a generic OS account (Oracle) and user identification through UID chain information.

8.5.16 - Authenticate all access to any database containing cardholder data. This includes access by applications, administrators, and all other users. Restrict user direct access or queries to databases to database administrators.

Timestamp	Client IP	Server IP	Network Protocol	Uid Chain Compressed	OS User	DB User Name	Source Program	Full Sql	Uid Chain
2011.07.29 14.02.16.0	10.10.9.56	10.10.9.56	BEQUEATH	joe	ORACLE	SYSTEM	SQLPLUS@OSPREY	select * from creditcard;	(1,root,init [3])-->(2572,root,usr/sbin/sshd)-->(1971,root,sshd: joe [priv])-->(1979, joe,sshd: joe@pts0)-->(1980,joe,bash)-->(2061,joe,su - oracle)-->(2068,oracle,bash)-->(2100,oracle,sqlplus)-->(2107,oracle,oracleXE (DESCRIPTION=(LOCAL=YES)ADDRESS=(PROTOCOL=beq)))

```

joe@osprey~$ su - oracle
Password:
-bash-3.00$ sqlplus system

SQL*Plus: Release 10.2.0.1.0 - Production on Fri Jul 29 14:02:16 2011

Copyright (c) 1982, 2005, Oracle. All rights reserved.

Enter password:

Connected to:
Oracle Database 10g Express Edition Release 10.2.0.1.0 - Production

SQL> select * from creditcard;

NAME                                CARDNUMBER          CARDID
-----
Joe D                                1234567890123456    1
Harry S                              2345678901234567    2

SQL> quit
Disconnected from Oracle Database 10g Express Edition Release 10.2.0.1.0 - Production
-bash-3.00$

```

What's wrong with the access to this production server?

- DB User = System account with access to production
- Table accessed Patient = All patient accounts
- Source Program = SQLPLUS, not the patient app
- Generic Accounts – Who did it? (Joe)

Figure 4-16 UID chain

A similar issue is that the generic IDs are used from application server to access the database. Guardium has support to identify users for major enterprise applications and for custom applications.

Figure 4-17 shows the list of supported enterprise applications.

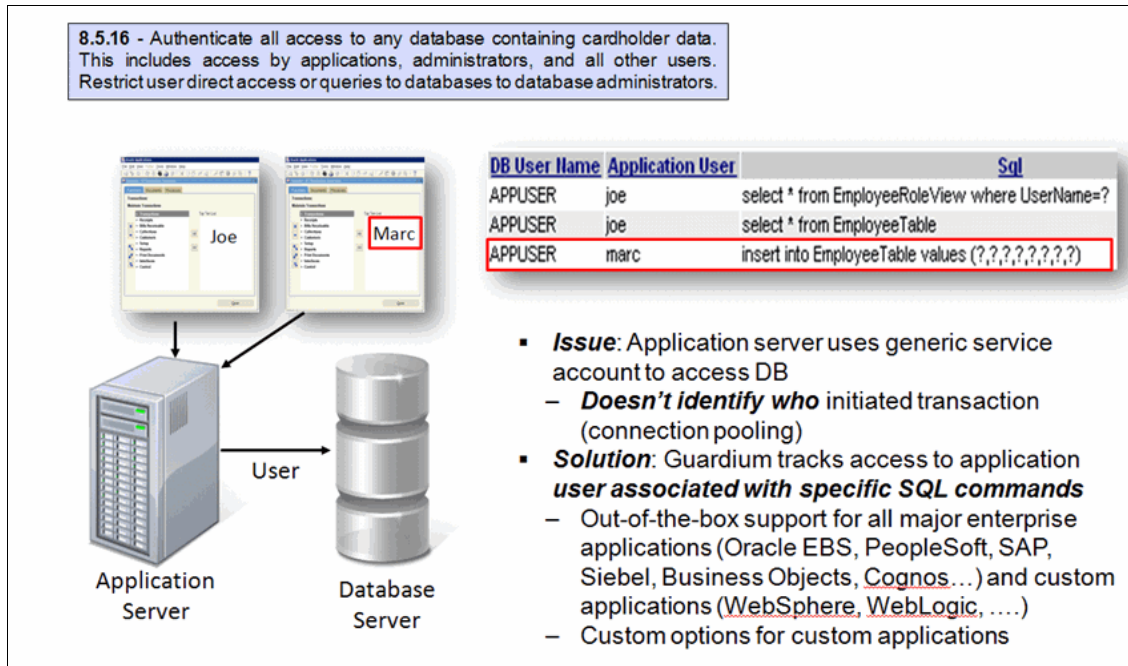


Figure 4-17 Application user identification

4.6.6 Requirement 10: Track and monitor all access to cardholder data

Guardium monitoring provides the capability to monitor access to objects that have sensitive data and use Compliance Workflow Automation to distribute the monitoring reports for review and sign off.

Figure 4-18 shows an example of compliance workflow.

10.2 - Implement automated audit trails for all system components to reconstruct the following events:
10.2.1 - All individual accesses to cardholder data

Receiver	Status	Action Required
pod(POC IBM)	Viewed not Signed	Review and Sign
admin(admin admin)	Not Viewed	Review and Sign
pot(pot training)	Not Viewed	Review and Sign
DataPrivacy(Data Privacy)	Not Viewed	Review and Sign

Comments:

- Report: Failed Login Reports (Failed User Login Attempts) _Overall Value: 0
- Report: Creates to the database I: SQL Trace _Overall Value: 7
- Security Assessment: Sybase Security Assessment (Sybase Assessment) _Overall Value: 71
- Security Assessment: Db2 Assessment (DB2 Assessment) _Overall Value: 46
- Security Assessment: SQL Server (SQL Server Assessment) _Overall Value: 78
- Security Assessment: Oracle Assessment (Oracle Product Assessment) _Overall Value: 71

- Automates entire compliance workflow
 - Report distribution to oversight team
 - Electronic sign-offs
 - Escalations, comments & exception handling
- Addresses auditors' requirements to document oversight processes
- Results of audit process stored with audit data in secure audit repository

Figure 4-18 Compliance workflow

4.6.7 Requirement 11: Regularly test systems

Guardium Change Audit System (CAS) ensures database configuration auditing and can satisfy the PCI requirement about test systems, as shown in the following statement:

"11.5 - deploy file-integrity monitoring tools to alert personnel to unauthorized modification of critical system files, configuration files, or content files; and configure the software to perform critical file comparisons at least weekly"

A CAS agent can monitor files, the output from OS and SQL scripts, environment variables, and Windows registry entries. Built-in templates are included for all supported (distributed) platforms.

Figure 4-19 shows a list of predefined CAS templates for different supported platforms.

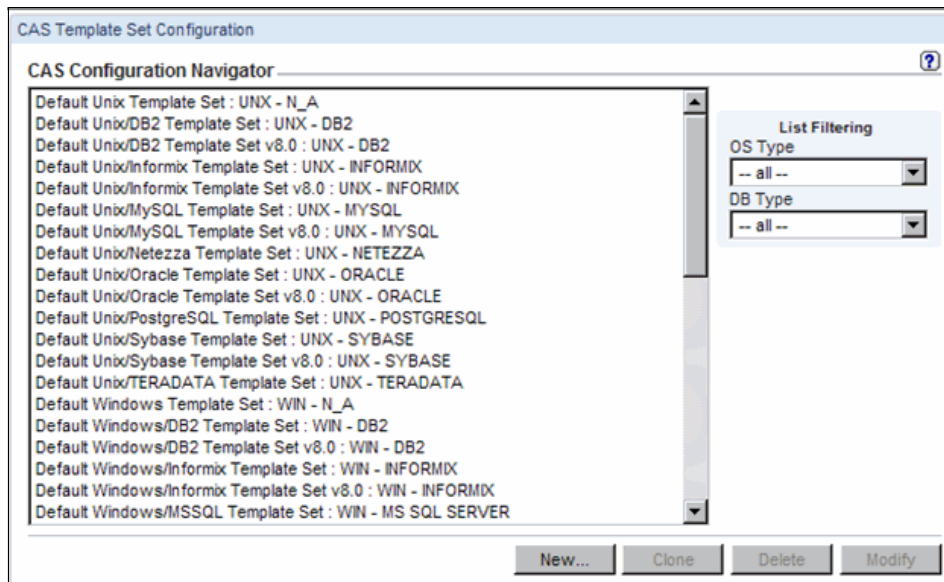


Figure 4-19 CAS templates

For more information about setting the CAS module, see 5.10, “Configuration audit system setup” on page 191.

4.6.8 SOX requirement: Prevent unauthorized changes to the financial CRM, ERP, and HR data

Guardium database activity monitoring and data access level control support monitoring and prevention of specified activities, such as DML type commands and DDL type commands on the defined set of objects for all or specific users.

For more information about monitoring and data access level control setting, see 5.8, “Data level access control” on page 189.



Monitoring setup

In this chapter, we describe the steps that are required to successfully implement database monitoring. In this context, we define monitoring as the review of database activities that can pose compliance or security violations. The result of these steps is a process that automatically delivers your required reports to the appropriate staff members on a scheduled basis by using workflow automation and delivers required real-time alerts to defined destination.

This process includes the following steps:

- ▶ Monitoring approach
- ▶ Monitoring planning
- ▶ Groups
- ▶ Policy
- ▶ Reports
- ▶ Compliance workflow
- ▶ Real time and threshold alerting
- ▶ Data level access control
- ▶ Vulnerability assessment
- ▶ Configuration audit system
- ▶ Entitlement reports
- ▶ Discovery and classification

This chapter includes the following topics:

- ▶ Monitoring setup overview
- ▶ Monitoring planning
- ▶ Grouping
- ▶ Policy
- ▶ Reports
- ▶ Compliance workflow
- ▶ Real-time and threshold alerting
- ▶ Data level access control
- ▶ Vulnerability assessment setup
- ▶ Configuration audit system setup
- ▶ Entitlement reporting setup
- ▶ Database auto-discovery setup
- ▶ Sensitive data finder setup
- ▶ Adding a menu tab to your portal

5.1 Monitoring setup overview

The monitoring deployment is an important aspect of a successful Guardium implementation. An efficient monitoring setup not only is part of the solution that fulfills the business requirements but also contributes to more efficient capacity usage of the solution and overall reduction of cost.

Figure 5-1 shows a generic monitoring setup order that is applicable for most of the database activity monitoring (DAM) deployments.

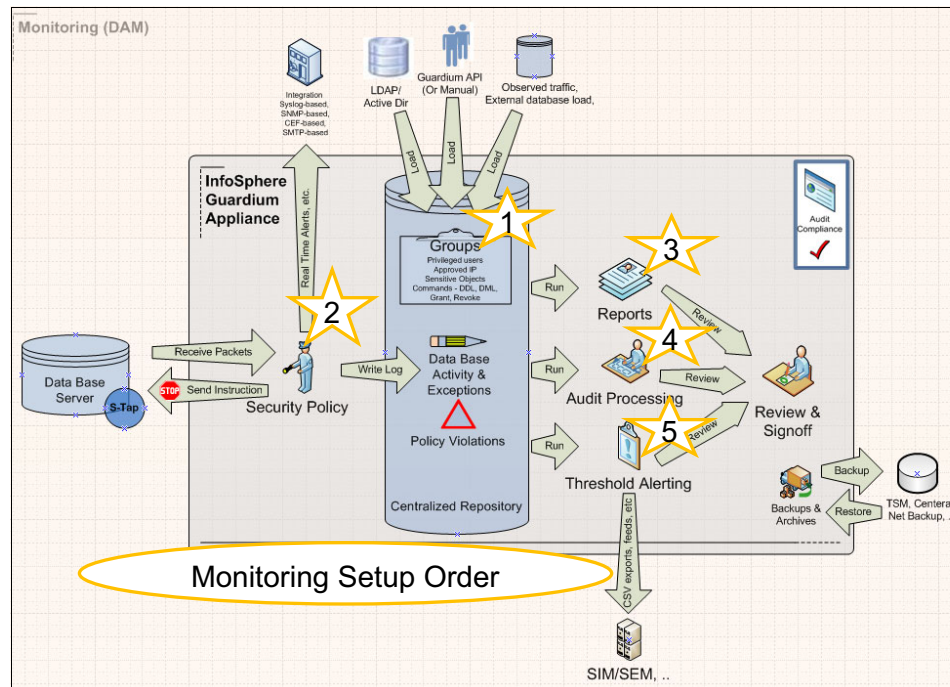


Figure 5-1 Monitoring setup

The following steps are shown in Figure 5-1:

1. Uploading available metadata from external sources into the Guardium environment.
2. Defining security policies that enable logging of database activity, exceptions, and policy violations into internal repository, integrating real-time alerts with external targets, and blocking of the activities that are based on defined conditions.
3. Identifying existing system reports and defining the custom reports to present the data that is available in internal repository.

4. Defining the compliance workflow automation to support business process of review and sign off from audit data.
5. Defining threshold alerts to be delivered to external targets.

The volume of the activity that is logged in to internal repository often depends on the volume of database activities and audit level requirements.

The following main audit levels are available:

- ▶ Privileged user audit
- ▶ Sensitive object audit
- ▶ Comprehensive audit

For more information about these audit levels, see 2.2.1, “Audit level” on page 23.

A Guardium policy defines the required audit level. It is best to start with restrictive policy and then open the policy gradually to include more audit activities. For example, start from auditing session level activities, such as login and logout information, then add access level rules (IPs, users, source programs, and so on) to support privileged users audit, and then continue to command level rules, object level, result sets, patterns, and so on. Other applications (such as vulnerability assessment, configuration audit system, entitlement reports) can be configured after the database activity monitoring is deployed (at least the initial phase of it).

The team that is involved in database activity monitoring deployment usually consists of the following members:

- ▶ Project manager
- ▶ Guardium administrator: Technical lead responsible for setup of relevant solution components
- ▶ Information security, auditors, and application owners to define monitoring requirements
- ▶ Audit process reviewers: Responsible for reviewing and signing off on audit results
- ▶ DBA and application team members: Perform necessary functional tests

The monitoring planning session is an opportunity to get the team together and start the project.

5.2 Monitoring planning

A monitoring planning session is held to discuss and capture in detail all monitoring requirements. This session also reviews the relevant Guardium applications (that is, policy, reporting, compliance workflow automation, and so on) and the best practices for monitoring. The outcome of the session is the monitoring plan document that includes all captured requirements that are to be used during the setup of the monitoring.

The following topics are reviewed in this session:

- ▶ Customer audit requirements (and governance controls)
- ▶ Policy rules and data filtering needs
- ▶ Reporting requirements (access and exceptions)
- ▶ Compliance workflow automation requirements
- ▶ Best practices around monitoring
- ▶ Alert methods (threshold and real time)
- ▶ User access, accounts, roles, and privileges on Guardium appliance
- ▶ Implementation of the monitoring plan, timelines, and milestones

The outcome of the sessions is a monitoring document.

The monitoring requirements are often determined by the auditors of the organization, especially in SOX, PCI, or HIPAA implementations, but also can be determined by the internal security rules.

The answers to following discussion points can assist in requirements definition and preparation for the monitoring planning session (especially if DAM is the main focus):

- ▶ Logging and real-time alerting:
 - Who must be monitored: Privileged users, specific applications, or everyone?
 - What type of activities must be monitored:
 - DDL, DML, DCL, and activities on specific objects?
 - Exceptions (failed logins, SQL errors)?
 - Information returned within result sets?
 - What type of activities should prompt real-time alerts, what mechanism should be used to send these alerts, and who should be the recipients of these alerts?
 - What type of activities should prompt more logging granularity (that is, logging of values and masked values within SQL statements)?

- Who are your privileged users? Can they be logically grouped (such as production DBAs, non-production DBAs, production system administrators)? Can this information be loaded into Guardium from external sources and be maintained ongoing from these sources?
- What are your sensitive objects? Can they be logically grouped and can this information be loaded into Guardium from external sources?
- ▶ Reporting:
 - What reports are required?
 - What information should appear on the reports (layout)?
 - Under what conditions should the information appear in the reports?
- ▶ Audit processes (compliance workflow automation):
 - Who receives the monitoring reports?
 - How frequently should the reports be delivered?
 - What action should the recipients of the reports take: review or sign off?
 - Should the delivery of the reports to specific users be contingent on other users to complete their action in a workflow process?
- ▶ Threshold (correlation) alerting:
 - What type of activities should prompt threshold (correlation) alert over a period?
 - What mechanism should be used to send these alerts?
 - Who should be the recipients of these alerts?
- ▶ Data level access control (blocking):
 - Who must be blocked?
 - What type of activities should be blocked?
- ▶ Vulnerability assessment:
 - What tests are part of security assessment?
 - How often should the scans run?
- ▶ Audit system configuration:
 - What files and folders must be monitored?
 - What type of changes should be reported?
- ▶ Entitlements reports:
 - What reports and entitlements are part of the scope?
 - How often should the reports run?

- ▶ Discovery and classification
 - What are the patterns to search for and in what objects?
 - What is the sampling approach?
 - Should identified objects be added to any of Guardium groups?

5.3 Grouping

Grouping is one of the basic tools in the Guardium solutions. By using this feature, you can combine similar elements of the same type to use later in queries, policies, reports, and so on. A list of financial servers, list of privileged users, or list of sensitive objects are typical examples of groups that are used throughout the Guardium solutions. The same set of groups is used by all Guardium users and all Guardium applications. Also, in a federated environment, all appliances that are managed by the same central manager share the set of groups.

Each group has its own group type and might contain only members from the same data type. For example, a group of weekdays contains weekdays, and a group of server IPs contains data that is formed as an IP address. Data types are predefined.

There is a special category of group types called *tuples*. A tuple allows multiple attributes to be combined to form a single group member. Tuple groups include the following examples:

- ▶ Object/Command
- ▶ Server IP/Server Port
- ▶ Client IP/Source Program/DB User
- ▶ Server IP/Instance Name/Port

Figure 5-2 shows the Group Builder tool of Guardium.

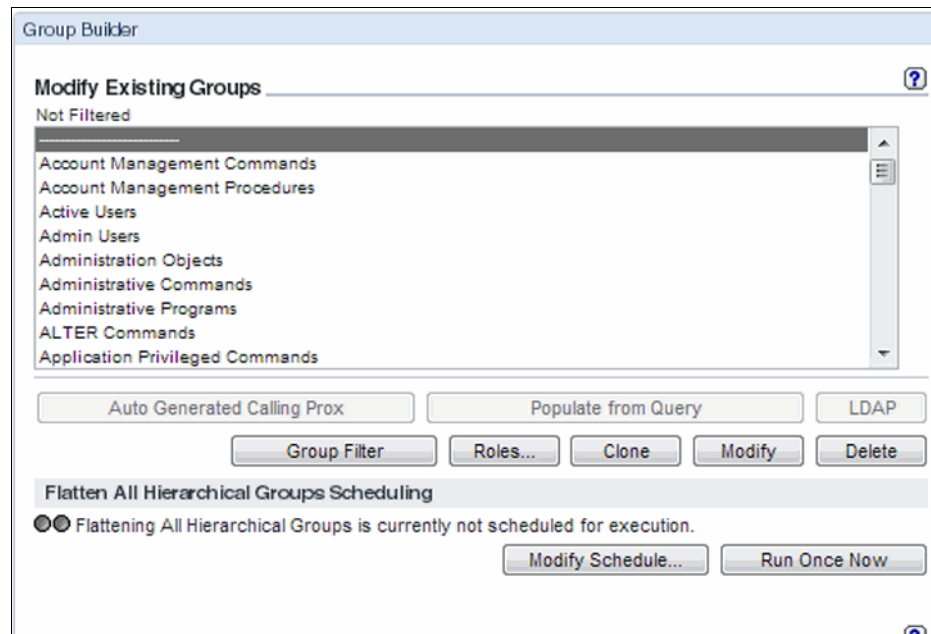


Figure 5-2 Group Builder tool

Use the Group Builder tool to modify the content of the existing group or to create a group. By using Group Builder, you can populate a group with members by entering them manually or through other Guardium tools.

The following methods are available to populate groups:

- ▶ Populating group users directly from LDAP or Active Directory servers.
- ▶ Populating from query: You can use the reports on Guardium appliance to populate various groups.
- ▶ Auto-generated calling process functionality is based on the process that is running on clients database server. This process scans stored procedures in client's database server and adds members to group. For example, process might scan stored procedures for sensitive objects and add them to a group of sensitive objects.
- ▶ Creating and updating groups and group members by using the GuardAPI interface. You can use GuardAPI to script groups updates from external applications or within Guardium.

- ▶ Creating or populating groups that are based on the results of classification process. This process searches through customer data for the predefined patterns and adds its findings to the groups as group members; for example, search for credit card number pattern in a set of tables and add tables that contain credit cards to a group of sensitive objects.
- ▶ Importing group members from a flat file by using Secure Shell (SSH).
- ▶ Entering group members manually.

5.3.1 Wildcards

Group function supports the use of wildcards with group members. Table 5-1 shows examples of wildcards usage.

Table 5-1 Group wildcard

Group member	Matches	Does not match
abc%	abc123	zyxabc123
%abc	zyxabc	zyxabc123
%abc%	zyxabc123	aabbcc

5.3.2 Hierarchical groups

A hierarchical group is a group of another groups from the same type. You can add other groups to a hierarchical group but not to a regular group, at least not explicitly. To create hierarchical group, select **Hierarchical** when you create a group, as shown in Figure 5-3.

The screenshot shows a 'Create New Group' dialog box with the following fields and values:

- Application Type: [Dropdown menu]
- Group Description: MyCommands
- Group Type Description: COMMANDS
- Group Sub Type Description: [Empty text box]
- Category: [Empty text box]
- Classification: [Empty text box]
- Hierarchical: (highlighted with a red arrow)

An 'Add' button is located at the bottom right of the dialog.

Figure 5-3 Creating a hierarchical group

After you define the new hierarchical-type group, you can add other groups of same type to that group. For example, you can add other groups from the COMMANDS type to MyCommands that are defined in Figure 5-3 on page 153 as shown in Figure 5-4.

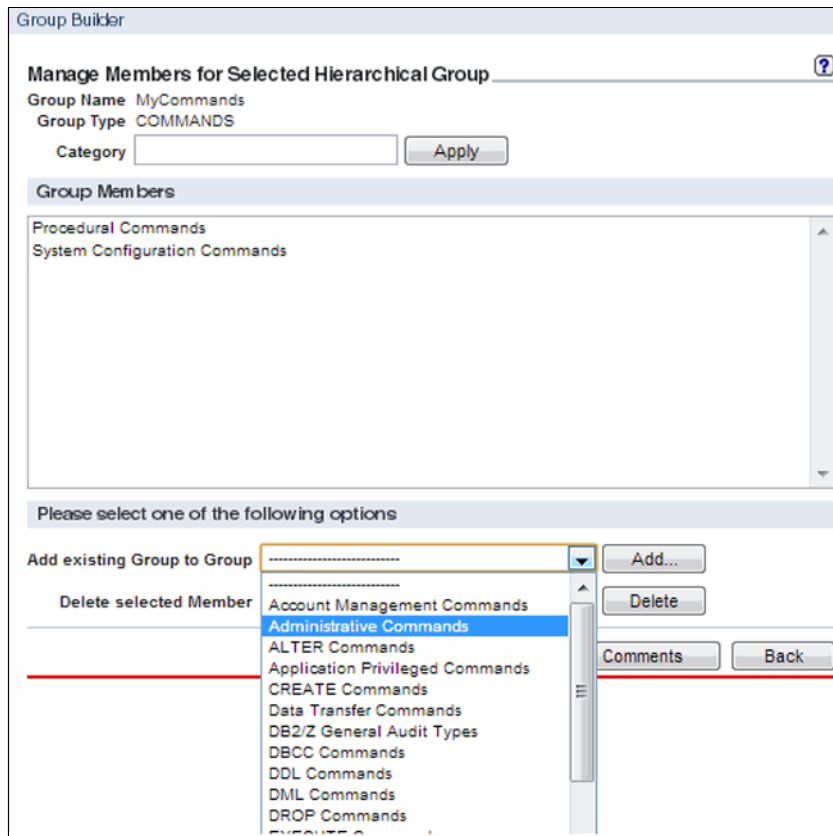


Figure 5-4 Adding groups to a hierarchical group

Multiple levels of hierarchy are allowed. You can add a hierarchical group to a hierarchical group. Figure 5-4 shows that you can add child groups to parent groups.

To make a hierarchical group usable by queries and policies, the group must be *flattened*. Flattening populates parent groups with members of child groups. If the child group is changed later, the parent group must be flattened again to pick up the changes. Flattening is a process that can be run on-demand once or run periodically by schedule.

The flattening process is not group-specific. When the flattening process runs, it flattens all hierarchical groups that are defined in the system. Figure 5-5 shows the part of group builder to run the flattening process.




Figure 5-5 Flattering process

5.3.3 Group deletion

The group builder tool includes a delete button to remove the unwanted groups. The groups that are used in queries and policies are not deleted and an appropriate message is displayed. The warning message includes a list of objects that are using that group.

5.3.4 Public versus private groups

When you are creating a group, you can decide whether this new group is to be available to all applications on the appliance (public) or is limited to only particular applications (private).

You specify the public or private group type through the application type when you are creating the group, as shown in Figure 5-6. For example, you might want to have private groups that are created and used only in policies. If you select Policy Builder in the drop-down menu for application type when you are creating a group, this new group is visible only in Policy Builder and not in other applications.

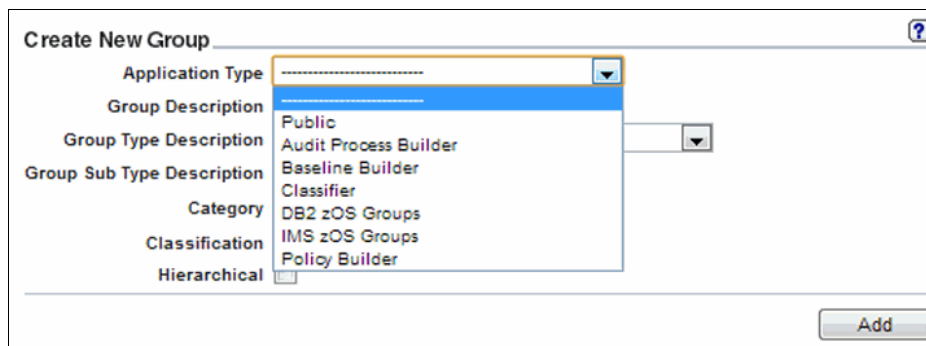


Figure 5-6 Public versus private groups

5.3.5 Groups in federated environment

In a federated environment, all appliances share the set of groups. The change that is made to the group on one appliance is propagated to all other appliances through a process that is known as *portal synchronization*.

When a group is created anywhere in the federated environment, it is recorded directly on the Central Manager. Because all application builders, such as query builder and policy builder, receive the list of existing groups directly from the Central Manager, you obtain the same group list regardless of from which appliance you retrieve the list of the existing groups. However, at run time, applications use the local copy of the group. For example, when you run a report on a managed unit that uses a group, the group must first be copied locally to the appliance where the report runs. The definition of the group is available immediately on all of the appliances in a federated environment. However, the content of the new group is available locally only after it is propagated from the Central Manager to the managed unit. This physical group copy process to all appliances is done by the portal synchronization process. The portal synchronization process usually runs every 30 minutes (which is configurable).

5.3.6 Size and performance effect

There is no physical limit to the size of a group. However, unreasonably large groups (tens of thousands of members in a single group) most likely have negative effect on overall system performance.

5.4 Policy

A policy is a set of rules and actions that are applied in real time to the traffic as it is captured by Guardium collector. The policy defines what traffic is ignored or logged, what activities require more granular logging, and what activities should trigger alert or block access to the database.

The traffic is evaluated against policy rules sequentially until the rule fires (meets criteria) and multiple actions can be taken upon rule firing. More rules of the same type (access and exception) are not evaluated unless the fired rule has the Continue condition set. All extrusion rules are evaluated whether one fires. If the criteria that is set in the rules is not met for any of the rules, the access traffic (SQLs) is logged in non-selective policy and not logged in selective policy.

5.4.1 Policy types and policy rules

Guardium features the following policy types:

- ▶ Non-selective: A non-selective policy logs the following information:
 - All the client-to-server access traffic is logged in to the collector with minimum granularity.
 - All the database exceptions are logged.
 - Server-to-client result sets that have explicit policy rules that are specified to log.
- ▶ Selective: A selective policy logs the following information:
 - Database exceptions.
 - Client-to-server access that has explicit policy rules that are specified to log.
 - Server-to-client result sets that have explicit policy rules that are specified to log.

Figure 5-7 shows the configuration panel in which the selective or the non-selective policy type is defined.

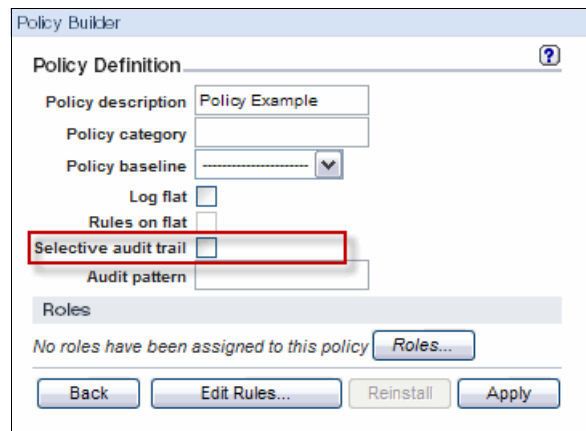


Figure 5-7 Selective audit trail configuration

The following types of policy rules are available (as shown in Figure 5-8 on page 158):

- ▶ Access rules are applied to the database traffic that comes from client to server (accessing the database).
- ▶ Exception rules are applied to database exceptions, such as failed logins and SQL errors.

- ▶ Extrusion rules are applied to database traffic from the database to client (results sets). The database exceptions are excluded.

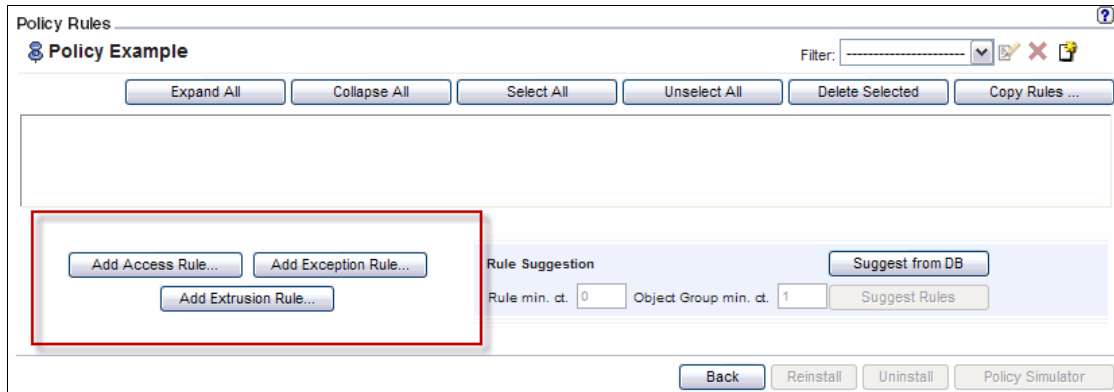


Figure 5-8 Policy rule types

When you are setting up the policy rules, start from a policy that ignores all traffic while you are working on the group population and policy rules definition. Such policy logs all database logins and logouts information and allows the use of session level reports to evaluate connection profiles. The policy has only one access rule with no filters and features an action of “Ignore S-TAP session”.

Use selective policy if it is applicable for your business requirements. If the business requirements are to log small subset of activities and there is a well-established process to keep the groups that are used in the policy rules updated, the selective audit policy can be a good choice.

You should evaluate what sessions are considered trusted and their traffic can be ignored and not logged in to the collector. Use “Ignore S-TAP session” rule for trusted traffic to significantly improve performance of the Guardium solution and reduce network usage. The collector instructs S-TAP to stop sending traffic for specific sessions if the rule with such action fires.

The following example shows setting up a policy rule that ignores traffic for connections profiles that are identified as trusted applications. This rule is applicable for non-selective and selective policies. In the Access Rule Definition, select **Trusted Connections** for database access, as shown in Figure 5-9 on page 159.

Access Rule Definition

Rule #1 of policy Policy Example

Description: Ignore Trusted Connections

Category: [] Classification: [] Severity: INFO [v]

Not Server IP [] / [] and/or Group [v] [] []

Not Client IP [] / [] and/or Group [v] [] []

Not Client MAC []

Net Prtcl. [] and/or Group [v] [] []

DB Type [v]

Not Svc. Name [] and/or Group [v] [] []

Not DB Name [] and/or Group [v] [] []

Not DB User [] and/or Group [v] [] []

Client IP/Src App./DB User/Server IP/Svc. Name Trusted Connections [v] [] []

Not App. User [] and/or Group [v] [] []

Figure 5-9 Policy rule: Trusted Connections

For action, use the Ignore S-TAP session option (as shown in Figure 5-10) instead of the Ignore Session or Skip logging options. The Ignore Session option causes the sniffer to ignore traffic on the collector but the S-TAP still sends the traffic. The Skip logging option causes the sniffer to ignore traffic at SQL level. The Ignore S-TAP session or Ignore session options often are used when the filter criteria is at the session level. The Skip logging option is used when the filter criteria is at the SQL level.

Quarantine for [0] minutes Records Affected Threshold [0] Rec. Vals. [v] Cont. to next rule []

Actions

IGNORE S-TAP SESSION

Add Action

Back Save

Figure 5-10 Policy rule: IGNORE S-TAP SESSION

If evaluating result sets from the database or logging exceptions (except failed logins) is not required, add the Ignore responses per session action to the policy rule, as shown in Figure 5-11. If this action is applicable for a subset of sessions, S-TAP does not send server-to-client traffic to the collector, which often is the majority of the overall traffic.

Quarantine for [0] minutes Records Affected Threshold [0] Rec. Vals. [v] Co

Actions

IGNORE RESPONSES PER SESSION

Figure 5-11 Policy rule: IGNORE RESPONSES PER Session

5.4.2 Logging granularity

Now that we described the policy actions that help filter unnecessary traffic that is logged to the collector, we now describe the actions that log the data to different domains with different granularity.

By default, all of the SQLs are logged in to the Access Period entity with 60 minutes logging granularity setting for non-selective audit policy. Although you can configure the logging granularity parameter, 60 minutes is the suggested logging time. The logging granularity parameter is available in Inspection Engine Configuration window, as shown in Figure 5-12.

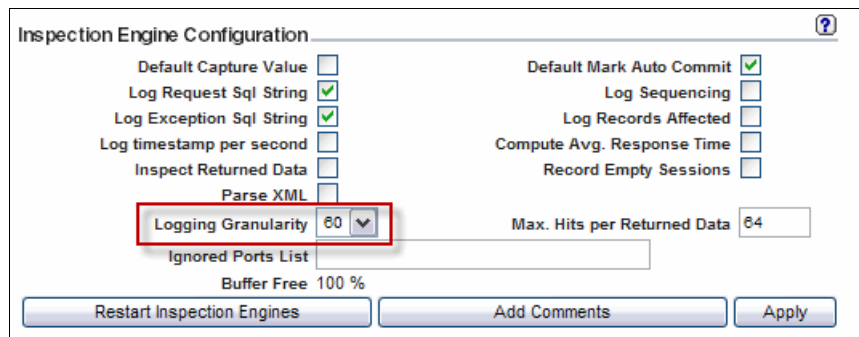


Figure 5-12 Logging granularity for non-selective audit policy

To have “logging into the Access Period entity with granularity of 60 minutes”, use the “Audit only” action for the selective policy, as shown in Figure 5-13.




Figure 5-13 Logging granularity for selective audit policy

Each record in the Access Period entity represents the number of times a specific construct (SQL) ran within a specific session during the 60-minute time frame. The parameters of the SQL statement are replaced with “?” in record.

For example, if session A is started at 14:30 and the `select * from employee where employee_id=10` SQL statement is run at 14:30, one record is written into the Access Period entity that represents the `select * from employee where employee_id = ?` construct for session A. This record is for the time frame between 14:00:00 and 14:59:59 and the counter is 1.

If `select * from employee where employee_id=20` is run in session A at 14:32, only the count of the existing record is updated to 2.

If same SQL statement is run as part of session B, a record is created to represent the entry for session B.

If same SQL statement is run at 15:30 as part of session A, then a record is created to represent time frame between 15:00:00 and 15:59:59.

If the SQL statement is run another 500 times before 16:00 in session A, only the counter of the existing entry is updated to 501 (regardless of what parameter values were used in the query).

With such logging (default and audit only action), the logged data is presented in the Access Tracking domain and the Access Period entity.

Figure 5-14 shows the Access Tracking panel of the Report Builder.




Figure 5-14 Access Tracking domain

Figure 5-15 shows the Access Period entity setting.




Figure 5-15 Access Period entity

If the exact time stamp of a query that was run is required or storing the values that are used by query is required, you can use the access policy rule with the “Log full details” action or “Log masked details” action. Each SQL statement is logged individually into the Full SQL domain and can be presented through the Access Tracking domain and Full SQL entity, as shown in Figure 5-16.

Figure 5-16 Full SQL entity

Alternatively, you can log the individual SQL statements into the Policy Violation domain by using the Log only action and present the record through Policy Violation Tracking domain and the Policy Violation Rule entity, as shown in Figure 5-17 and Figure 5-18 on page 163.

Figure 5-17 Policy Violation Tracking domain




Figure 5-18 Policy Rule Violation entity

5.4.3 Generating real-time alerts with policy

By using the Guardium solution, you can generate real-time alerts through the policy. If alerting in real time that is based on specific conditions is required, you can use the policy rule with the following actions:

- ▶ Alert per match
- ▶ Alert once per session
- ▶ Alert daily
- ▶ Alert per time granularity

Figure 5-19 shows the ALERT PER MATCH action.



Figure 5-19 ALERT PER MATCH action

The Individual SQL statement is logged into the Policy Violation Tracking domain and the Alert Tracking domain, as shown in Figure 5-20.

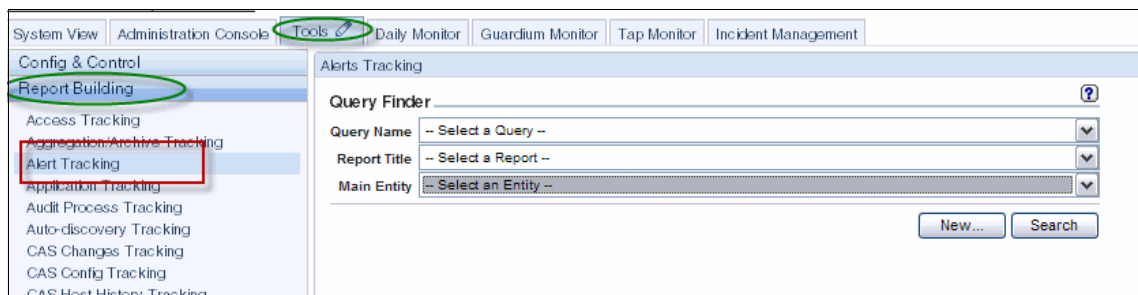
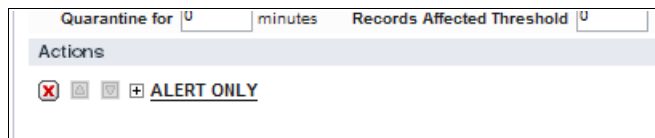


Figure 5-20 Alert Tracking domain

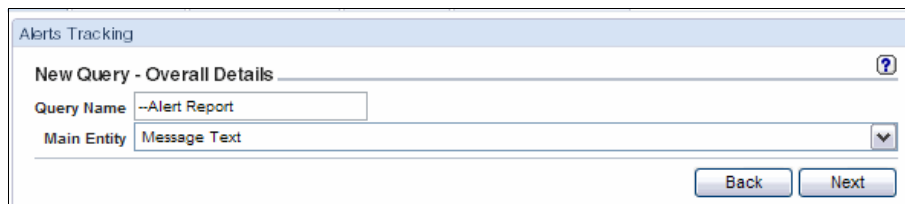
If alert only is sufficient and recording the policy violation is not required, you can use the “Alert only” action (see Figure 5-21) in the policy rule instead of the “Alert per match”.



The screenshot shows a configuration window for Alerts Tracking. At the top, there are two input fields: "Quarantine for" followed by a text box containing "0" and the word "minutes", and "Records Affected Threshold" followed by a text box containing "0". Below this is a section titled "Actions" with a light blue background. Inside this section, there are several icons: a red 'X' in a square, a square with a magnifying glass, a square with a checkmark, and a square with a plus sign. To the right of these icons is the text "ALERT ONLY" in blue, underlined.

Figure 5-21 ALERT ONLY action

You can generate the report about the alerts from the Alert Tracking domain and Message Text entity, as shown in Figure 5-22.



The screenshot shows a configuration window titled "Alerts Tracking". Below the title bar is a section "New Query - Overall Details" with a help icon (question mark in a circle). Under this section, there are two input fields: "Query Name" with the text "--Alert Report" and "Main Entity" with a dropdown menu showing "Message Text". At the bottom right of the window, there are two buttons: "Back" and "Next".

Figure 5-22 Message Text entity

5.4.4 Extrusion rules

The extrusion rules log data into the following entities based the actions that are used:

- ▶ The Full SQL entity when the “Log full details” or “Log extrusion counter” actions are used
- ▶ The Policy Violation Rule entity when the “Log only” or one of “Alert...” actions are used
- ▶ The Message Text entities when one of the “Alert ...” actions is used

Configure the inspection engine setting to enforce the evaluation of the result sets by the sniffer. If pattern match is configured in the policy rule, the sniffer evaluated the result sets until the value of the “Max. Hits per Returned Data” parameter is reached. The default value is 64 per packet. If the default value does not satisfy the requirements and the sniffer should continue to evaluate packet to identify all matching patterns, “Max. Hits per Returned Data” should be increased from the default value of 64 (see Figure 5-23). Increasing the value can negatively affect performance.




Figure 5-23 Inspection Engine Configuration

Figure 5-24 on page 166 shows an example of an extrusion rule for the results set that is returned from the database to the privileged users. The “Log Full Details” action is used to record detail information if the results set has data that matches the credit card pattern and passes the Luhn algorithm validation. The Luhn algorithm validation is activated by the name template of the Description field in the policy rule. You also can use the “Log Extrusion Counter” action or the “Log masked Extrusion Counter” action if the result sets do not have to be logged or must be masked.

Extrusion Rule Definition

Rule #1 of policy **Deployment Best Practice**

Description: **guardium://CREDIT_CARD** Log Extrusions for priv users

Category: _____ Classification: _____ Severity: INFO

Not Server IP _____ and/or Group _____

Not Client IP _____ and/or Group _____

Not Client MAC _____

Net Prtol. _____ and/or Group _____

DB Type _____

Not Svc. Name _____ and/or Group _____

Not DB Name _____ and/or Group _____

Not DB User _____ and/or Group **Production Users**

Client IP/Src App./DB User/Server IP/Svc. Name _____

Not App. User _____ and/or Group _____

Not OS User _____ and/or Group _____

Not Src App. _____ and/or Group _____

Data Pattern: **[0-9][4]-.[0-9][4]-.[0-9][4]-.** RE Replacement Character _____

Sql Pattern _____ RE

Time Period _____

Minimum Count 0 Reset Interval 0 minutes

Quarantine for 0 minutes Matched Returned Data Threshold 0 Rec. Vals. Revoke

Actions

LOG FULL DETAILS

Add Action

Figure 5-24 Extrusion rule example

5.4.5 Database exceptions

Database exceptions are logged into the Exception domain by default in non-selective and selective policies. The exceptions can also trigger alerts by using one of the “Alert...” actions and log data into the Policy Violation Rule entity and the Message text entity.

Figure 5-25 shows an exception rule to generate an alert for 10 failed logins within a 5-minute period.

The screenshot displays the 'Exception Rule Definition' window for 'Rule #1 of policy Deployment Best Practice'. The rule is titled 'Failed Logins - Alert' with a severity of 'HIGH'. The configuration includes various fields for filtering exceptions, such as Server IP, Client IP, Client MAC, Net Prtol, DB Type, Svc. Name, DB Name, DB User, App. User, OS User, Src App., and Error Code. The 'Exopt. Type' is set to 'LOGIN_FAILED'. The 'Minimum Count' is 10 and the 'Reset Interval' is 5 minutes. The 'Actions' section shows 'ALERT PER MATCH' selected. The 'Add Action', 'Back', and 'Save' buttons are visible at the bottom.

Figure 5-25 Exception rule

If logging all or part of the exceptions is not required, use the “Skip Logging” action for the policy rule. Figure 5-26 on page 168 shows an exception rule to skip logging all Oracle SQL errors.

Exception Rule Definition ?

Rule #1 of policy **Deployment Best Practice**

Description: Severity: **INFO** ▼

Category: Classification:

Not Server IP / and/or Group: ▼

Not Client IP / and/or Group: ▼

Not Client MAC

Net Prtcl.: and/or Group: ▼

DB Type: **ORACLE** ▼

Not Svc. Name and/or Group: ▼

Not DB Name and/or Group: ▼

Not DB User and/or Group: ▼

Client IP/Src App./DB User/Server IP/Svc. Name: ▼

Not App. User and/or Group: ▼

Not OS User and/or Group: ▼

Not Src App. and/or Group: ▼

Not Error Code and/or Group: ▼

Not Excpt. Type: **SQL_ERROR** ▼

Masking Pattern: RE Replacement Character:

Time Period: ▼

Minimum Count: Reset Interval: minutes

Quarantine for: minutes Rec. Vals. Cont. to next rule

Actions

SKIP LOGGING

Figure 5-26 Skipping logging

5.4.6 Policy for z/OS

Because the S-TAP behaves differently between the z/OS and the distributed systems, the policy concept is also different in two environments. In a distributed system, the S-TAP sends all traffic for the monitored database by default. However, S-TAP does not send traffic on a z/OS system by default. On z/OS, the policy must be configured on the appliance and to be pushed down to the z/OS system to instruct what traffic the S-TAP should collect and send to the Guardium collector. The policy rules that are pushed down to the z/OS system are of “DB2 connection profile” type (if monitored database on z/OS is DB2) and the only action that is available for the rules is “z/OS Audit” (see Figure 5-27 on page 169). Other fields in the rules must be used to define the actual filters; that is, what S-TAP should send to the collector.

The screenshot shows a configuration form with several fields and a list of actions. The 'DB Type' dropdown is highlighted with a red box and set to 'DB2 COLLECTION PROFILE'. Below it are fields for 'Svc. Name', 'DB User', 'App. User', 'OS User', 'Object', 'Command', and 'Client Info', each with an 'and/or Group' dropdown. A 'Time Period' dropdown is also present. The 'Actions' section at the bottom has a red box around the 'Z/OS AUDIT' action, which is preceded by a minus sign icon. An 'Add Action' button is located at the bottom right.

Figure 5-27 z/OS AUDIT action

The rest of the rules in the policy (or in another policy if multiple policies installed) instruct the Guardium collector what actions should be taken on the traffic that was sent by S-TAP. So, the actions that are used for distributed environment (such as, Allow, Audit only, Log Full, or Masked Details) can and should be used in the z/OS environment to support business requirements. Actions such as, Ignore S-TAP session or Ignore Responses per session, are irrelevant to the z/OS environment because the collector does not send the verdict back to S-TAP in a z/OS environment.

Different from the distributed S-TAP, the z/OS S-TAP has other SQL-related information already parsed (such as, Verb and Object) and sends this information to the collector. The use of the parsed information directly saves the parsing process of the collector or sniffer. If filtering or reporting at the Field level is not required, you can use the “QUICK PARSE NATIVE” action in the policy (see Figure 5-28) to instruct the collector to use the parsed information that is provided by S-TAP directly. This can provide significant performance improvement (mostly in analyzer component of the sniffer).

The screenshot shows a configuration form with a 'Quarantine for' field set to 0 minutes and a 'Records Affected Threshold' field set to 0. There are checkboxes for 'Rec. Vals.' (checked) and 'Cont. to next rule' (unchecked). The 'Actions' section has a red box around the 'QUICK PARSE NATIVE' action, which is preceded by a minus sign icon. 'Add Action', 'Back', and 'Save' buttons are visible at the bottom.

Figure 5-28 QUICK PARSE NATIVE action

5.4.7 Policy installation

For the policy to take effect, you must install the policy on the collector. The policy is not applicable to the aggregator or central manager appliances. If policy rules are changed or the members of the user groups are changed, the policy must be reinstalled for the change to take effect.

Internally, the policy rules, groups, and their members are copied into a set of installed policy tables and this set of tables is used by the sniffer process. All subsequent changes to the rules or groups do not affect the installed policy tables until the policy is reinstalled. Because the policy rules use groups and the group members are updated frequently in most deployments, it is the best to schedule the policy installation daily at the end of the day after all manual updates to the group members are complete and all processes to update groups are run.

Figure 5-29 on page 171 shows the panel that is used to schedule the daily policy installation.




Figure 5-29 Scheduling policy installation

There is an option to install multiple policies on the collector. When this option is used, the rules are processed sequentially in a similar manner as though only one policy is installed. Rules of the first policy are processed first then the rules of second policy, and so on. If one of the rules fires and the Continue condition is not set on that rule, the evaluation of the rules stops and no other rules are evaluated regardless if these rules are in the same policy or in the next policies.

Figure 5-30 shows a policy installation with an option to define a sequence within the installed policies. In this example, policy Example 2 is installed after policy Example; that is, the first rules from policy Example are evaluated and then the rules of the Example 2.

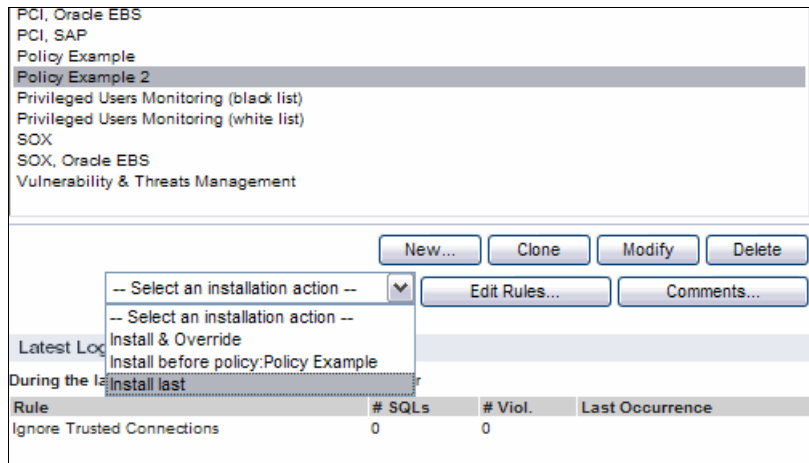


Figure 5-30 Policy installation on a collector

In a federated environment, the policy can be installed and reinstalled from the Central Manager on multiple managed units, as shown in Figure 5-31 and Figure 5-32 on page 174.

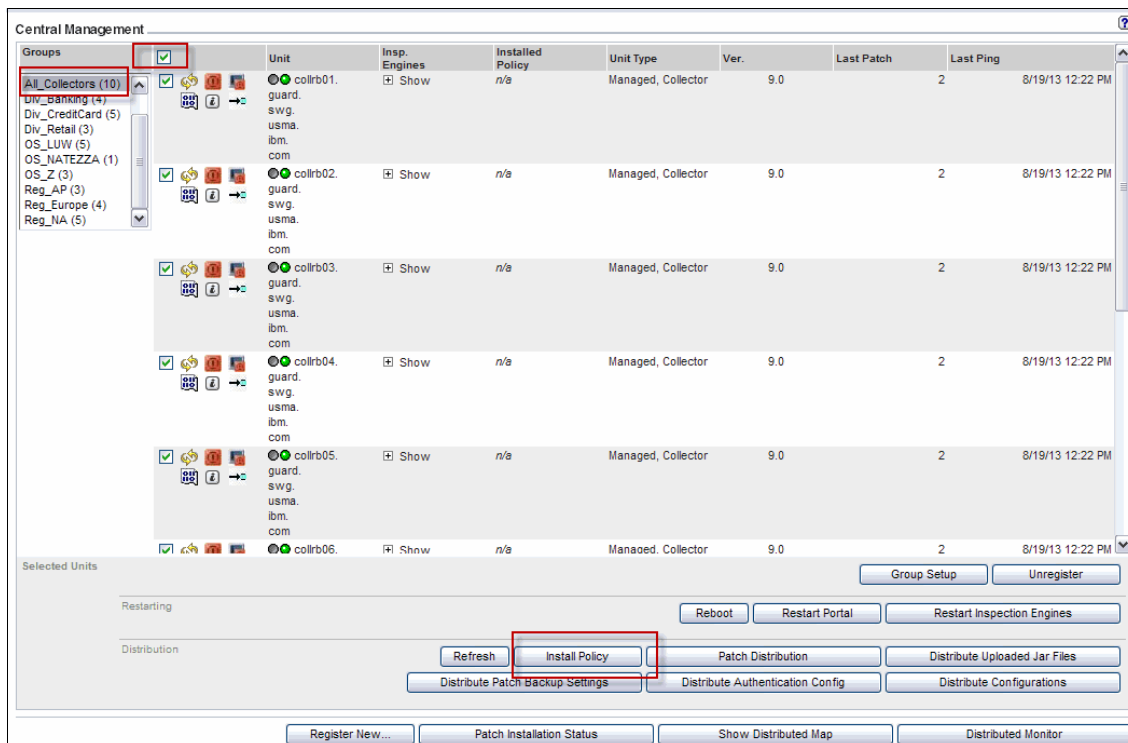


Figure 5-31 Installing policy in a federated environment 1




Figure 5-32 Installing policy in a federated environment 2

5.5 Reports

InfoSphere Guardium provides a significant number of pre-built reports that are available for customers and are ready to use. If you have specific business needs that call for specific reports, Guardium also provides a report generating tool that you can use to customize the existing reports or to build new reports.

In this section, we describe the Guardium reporting mechanism and how the reported data is stored internally on the appliance.

5.5.1 Reports versus queries

The reporting mechanism has two major components: queries and reports. Queries define what data you want to see and reports define how the retrieved data is presented. For example, you might have a tabular report and a separate graphical report that are based on the same query.

Frequently, the term *report* is used ambiguously to refer to queries and the report components. To create an efficient report really means to create an efficient query. In the following sections, we focus on query building.

The definitions of the queries, reports, and other metadata are stored in an internal database of the appliance. In a federated environment, all queries and reports definitions are stored on Central Manager. When a query is defined on one of the managed units, it is immediately visible and ready for use on all the appliances in that environment.

5.5.2 Domain, entities, and attributes

To create a query, you first identify the domain tracking tool that should be used to create a query. Each query is associated with a particular predefined set of data that is called *data domain*. For example, the Access domain contains captured traffic data, the Exception domain contains captured error messages from the database server, and the Guardium Activity domain contains data that is used for self monitoring on activities that are performed by Guardium users.

Each domain comprises a number of entities that contain specific data of this domain. For example, the Group domain is comprised of the Group entity, the Group Type entity, and the Group Members entity; the Alert domain is comprised of the Alerts entity, the Message Header entity, and the Message Text entity; and the Policy Violation domain comprises the entities, such as Session, SQL, Policy Rule, and more than 10 other entities that contain data that is important for policy violations tracking.

Tip: You might think of the entities of tables and domains as groups of the logically related tables.

There are approximately 40 predefined domains that are available in InfoSphere Guardium V9. You also can create custom domains.

Figure 5-33 on page 176 shows a list of available domains on the left side. The entity list in the Access domain is shown in the middle of the figure.

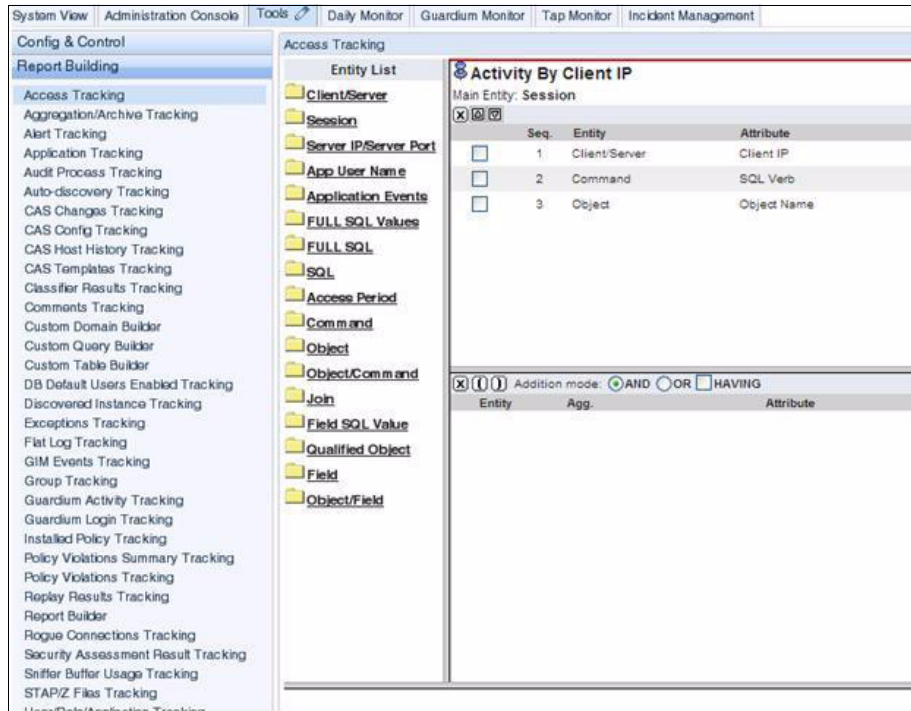


Figure 5-33 Available domains

Complete the following steps to create a typical query:

1. After you select your domain, click the corresponding tracking tool of the domain from the list in the left pane of the Tools tab (see Figure 5-33) to start the Query Builder.
2. Select main entity. The main entity defines the data on which that report focuses. For example, when you are creating query in Access domain, one of the following entities can be selected as a main entity (these are only a few of the frequently used choices, many more are available):
 - Session: Used when the session level information is the focus of the query; for example, logins to the database servers.
 - SQL (or Full SQL, if logging full details): Used when the SQL level information must be provided in the report and one line per SQL statement is required.
 - Command: Used when commands are the main focus of the report. Each individual command within each SQL appears on its own line in the report.

- Object: Used when the actual object name that is accessed by an SQL statement is required (an SQL statement must be parsed). Because each object appears on a separate line in the report, some SQL statements appear on multiple lines in the report if the SQL statement has multiple objects.

Figure 5-34 shows the window for selecting the main entity for a new query.

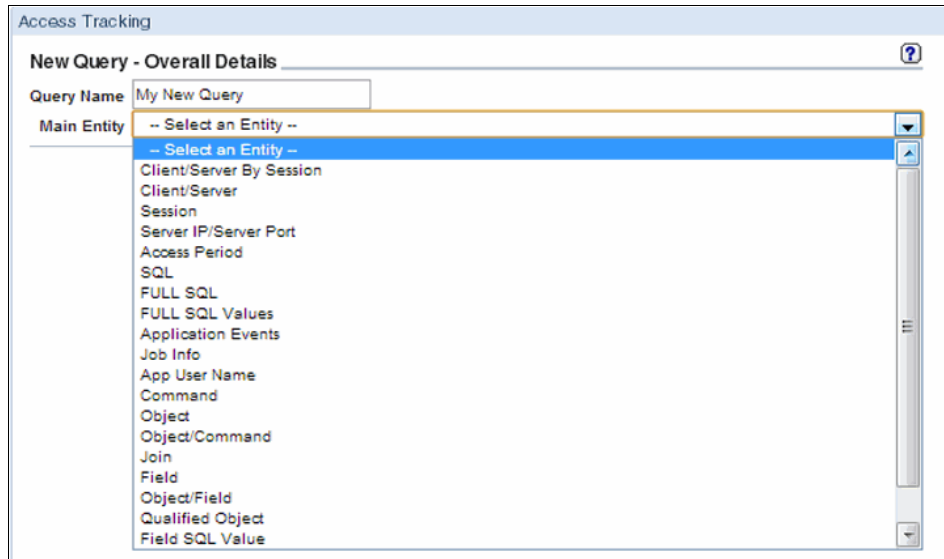


Figure 5-34 Selecting main entity for a new query

3. Select attributes: Select the attributes (fields) from the list of available attributes. You can drag the attributes into the upper portion of the pane to add the attributes to the list of fields that the query retrieves.

The following optional flags at the top of the Query Fields pane (see Figure 5-35 on page 178) can be applied to the query:

- Add Count: This option automatically adds another counter column to the query. The counter shows how many times each unique combination of the fields in the row occurred for an observed period.
- Add Distinct: This option displays only the unique combinations of field values in the row. Each unique combination is displayed once only. This produces a shorter report and might boost performance.
- Sort By Count: Use this option if you want the results sorted by the counter values.




Figure 5-35 Optional flags: Add Count, Add Distinct, and Sort By Count

5.5.3 Query conditions

You can drag any attribute into the lower portion of the Query Builder tool to be used in the query conditions. A Query condition can be composed of any number of single-line conditions that are connected with AND or OR operators. Usage of brackets is also supported.

Figure 5-36 shows the Query Conditions pane.




Figure 5-36 Using AND and OR operators in query condition

The Query builder supports a list of operators for query conditions. When you are creating query conditions, you have a choice of selecting a fixed value for the attribute in the condition or a runtime parameter. The use of a runtime parameter gives an opportunity to run a report for different attribute values without changing a query.

The following general considerations improve query performance when you are building a query with conditions:

- ▶ Use the In Group operator instead of several conditions that are connected by the OR operator.
- ▶ Use = instead of Like where possible.
- ▶ Use Like instead of Like Group.
- ▶ Use In Group instead of Like Group if possible.
- ▶ Use Not In... Group with caution.

For more information about the available operators, see the Guardium online Help.

For more information about creating custom reports (including performance considerations), see Chapter 7, “Ongoing operations” on page 237.

5.6 Compliance workflow

By using the Compliance workflow function, you can schedule the reports and deliver the result to users for review and sign off. The workflow results can be delivered to individual users, group of users, or roles. It is best to deliver the workflow results to roles because more than one user can review and sign off a result. It is also easier to manage employee's absence and turnover.

To use this function, you must define users and roles in Guardium.

When you are defining roles in your organization, consider the answers to the following questions:

- ▶ Who should receive reports and what is the job function of each receiver (for example, DBA, manager, and internal audit)?
- ▶ Which users have the same job function and can provide an equivalent review and sign off?

You can define roles by using the Access Management UI or the Guardium predefined roles.

To create users, use the Access Management UI and assign an appropriate role to the user.

Note: Integration with an external system can be used to manage users in Guardium; for example, LDAP.

To develop the workflow, define the audit process that includes the following information:

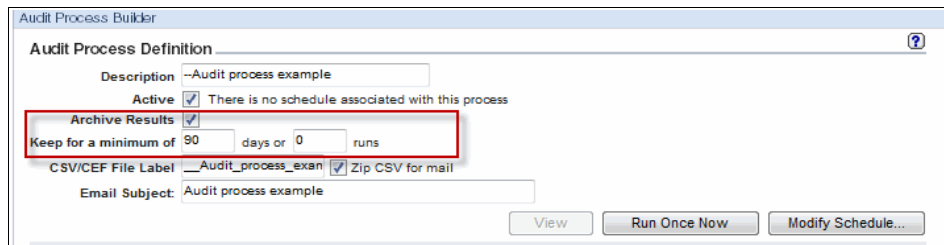
- ▶ Who receives the reports.
- ▶ Which reports are delivered and how often.
- ▶ The order in which the users or groups receive the reports.
- ▶ If review or sign off is required.
- ▶ If the delivery should stop at any user or role until they complete the required action (review and sign off).

The following deployment considerations are applicable to the operations aspect of the audit processes:

- ▶ Audit process results have an individual purge schedule that is based on the definitions of specific audit process and are independent from data purge definition. Therefore, if the audit requirement is to keep audit results online for 90 days, it does not necessarily mean that you must keep underlying data for 90 days on the appliance. Keeping audit results and their workflow (review, sign off, and comments information) might be sufficient. Also, audit process results can be archived and then restored (without restoring the underlying data). The results can be restored into the investigation center of aggregator. For more information about the investigation center configuration, see Guardium Information Center, which is available at this website:

http://pic.dhe.ibm.com/infocenter/igsec/v1/topic/com.ibm.guardium91.doc/aggregation_cm/topics/investigation_center.html

Figure 5-37 shows an example of an audit process with 90 days retention of audit process results and archive results option checked.



The screenshot shows the 'Audit Process Builder' window. The 'Audit Process Definition' section includes the following fields and options:

- Description: --Audit process example
- Active: There is no schedule associated with this process
- Archive Results:
- Keep for a minimum of: 90 days or 0 runs
- CSV/CEF File Label: __Audit_process_exan Zip CSV for mail
- Email Subject: Audit process example

Buttons at the bottom: View, Run Once Now, Modify Schedule...

Figure 5-37 Audit process archive definition

Figure 5-38 shows where the Results Archive can be configured.

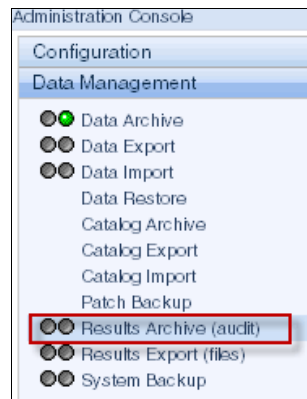


Figure 5-38 Results Archive configuration

- ▶ Audit process results are not purged from the appliance if the workflow is not complete, even if the retention criteria is met. For example, if the configuration is to keep audit process results for 90 days and there are three receivers who should sign on the results but only two of them signed on a specific run, the audit results are not purged, even after 90 days passed until the third receiver signs. A list of the outstanding audit process workflow is available on the Outstanding Audit Process Reviews report, as shown in Figure 5-39.

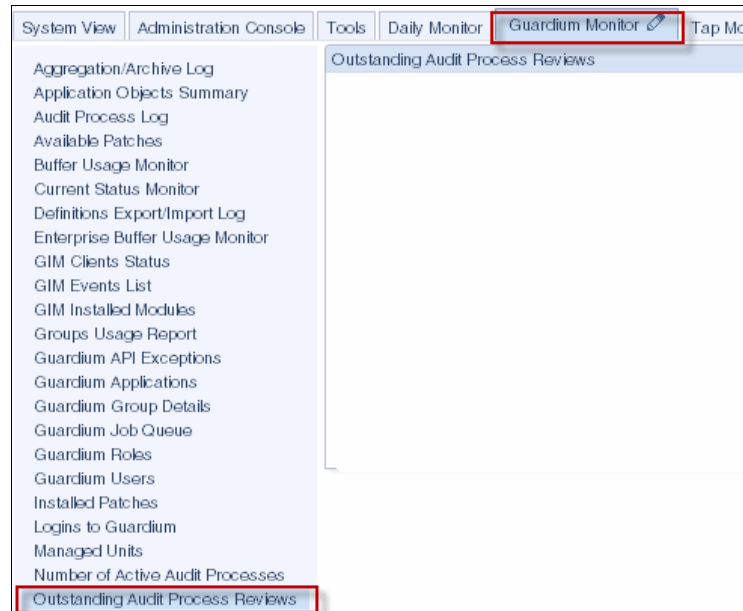


Figure 5-39 Outstanding Audit process Reviews

- ▶ The audit workflow is maintained outside Guardium and the audit process results are exported through CSV to external system, as shown in Figure 5-40 on page 182. The recommendation is not to have any receivers assigned to the audit process and to ensure that the export CSV process is scheduled. Such configuration ensures prompt purge of audit results from the appliance (based on retention definition of each audit process).

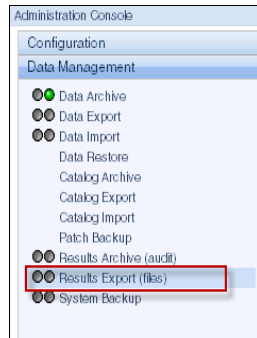


Figure 5-40 Results Export configuration

5.7 Real-time and threshold alerting

Guardium features the following alert types:

- ▶ Real-time alert: This alert is generated by using policy rules.
- ▶ Threshold alert: This alert is generated based on query. This alert is also known as a *correlation alert*.

Each alert type has its advantages and drawbacks that are important to understand to use these alerts efficiently.

The alerting mechanism consists of the following parts:

- ▶ The first part checks the predefined condition. If the condition matches, the alert message is generated and inserted into a queue.
- ▶ The second part is an asynchronous process that checks the queue periodically and sends out the newly generated alerts that are waiting in the queue.

5.7.1 Alert generating

Guardium alerts are implemented through policy (real-time alert) or queries (threshold alert).

Real-time alerts

A policy rule associates with various types of actions. One group of actions is used to generate alerts that are based on the rule conditions.

Figure 5-41 shows the Policy Builder window for selecting the wanted alert notification frequency from a list of rule actions.

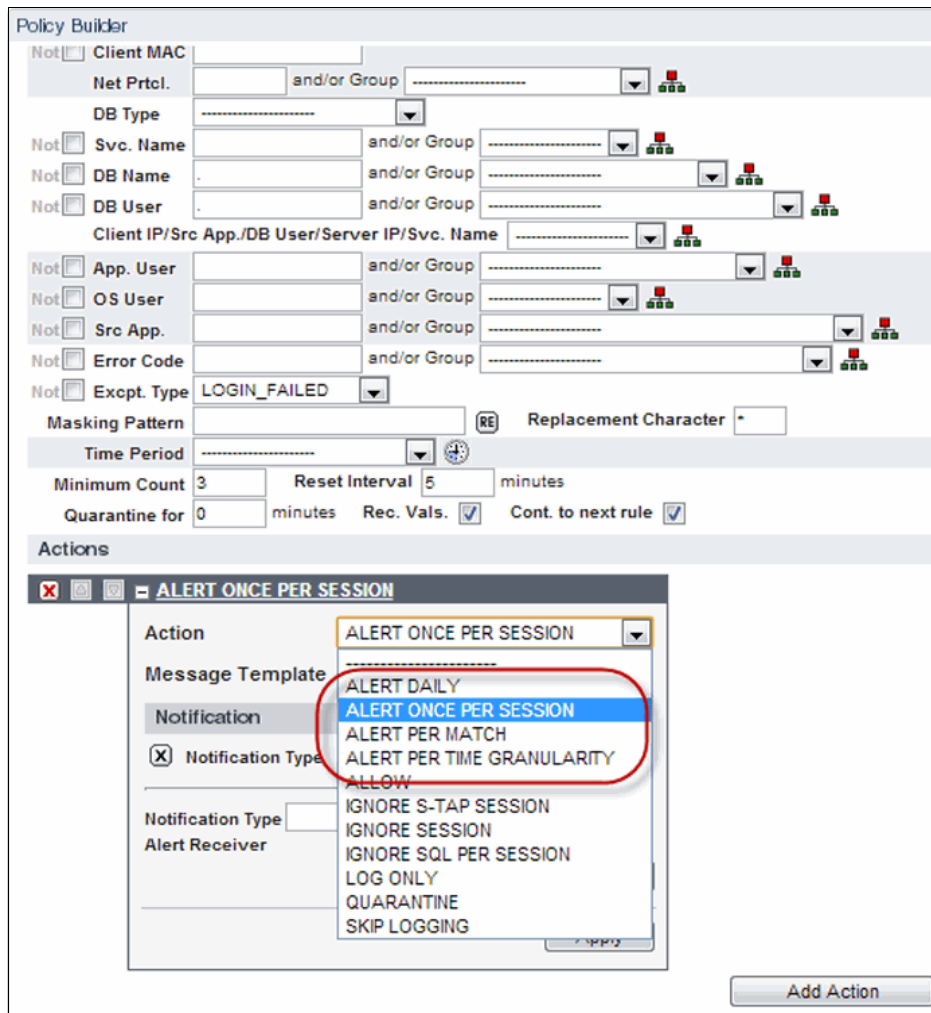


Figure 5-41 Selecting the wanted alert notification frequency

One group of actions is dedicated to alerting. All alert options work the same way. The main difference among the actions is the notification frequency. For instance, generate an alert on every occurrence or only once a day.

The real-time alerts are ideal for notifying about the events that are related to the observed traffic, as shown in the following examples:

- ▶ Repetitive failed login attempts to database
- ▶ Unauthorized access to sensitive data
- ▶ Attempt to extract massive amount of sensitive data in one query

Real-time alerts often use less system resources; therefore, they are more efficient. However, real-time alerts can fire only on the events that are related to the observed traffic. If you must create an alert on multiple failed logins to the Guardium appliance or on changes that are made in the Guardium configuration, you must use the threshold alerts because a policy does not see events that are related to the Guardium appliance.

Therefore, the rule is to use real-time alerts on all of the events that are related to the audit data monitoring and to use threshold alerts for everything else.

Threshold alerts

You can use the Alert Builder (see Figure 5-42) to set up threshold alert.

The screenshot displays the 'Alert Builder' interface. On the left is a 'Config & Control' sidebar with a list of tools including 'Alert Builder'. The main area is titled 'Alert Builder' and contains a 'Modify Alert' section. This section is divided into several panels: 'Settings', 'Alert Definition', 'Alert Threshold', 'Notification', and 'Alert Receivers'. The 'Settings' panel includes fields for Name, Description, Category, Classification, Severity (set to INFO), and Run frequency (1440 minutes). The 'Alert Definition' panel shows the Query as 'Guardium - add/remove users' and an Accumulation interval of 1450 minutes. The 'Alert Threshold' panel is configured with a Threshold of 1.0, an operator of '>=' threshold, and 'per report' evaluation. The 'Notification' panel shows a frequency of 1440 minutes. The 'Alert Receivers' panel lists 'SYSLOG' with a 'Delete' link and an 'Add Receiver...' button.

Section	Field	Value
Settings	Name	Guardium - Add/Remove Users
	Description	Alert whenever a user is created/deleted
	Category	Self Monitoring
	Classification	
	Severity	INFO
	Run frequency	1440 (minutes)
	Active	<input type="checkbox"/>
Alert Definition	Query	Guardium - add/remove users
	Accumulation interval	1450 (minutes)
	Note	Alerts run on aggregators will be based only on data within the defined merge period
	Log full query results	<input checked="" type="checkbox"/>
Alert Threshold	Threshold	1.0
	Alert when value is	>= threshold
	Threshold Evaluated:	<input checked="" type="radio"/> per report
	Threshold Used:	<input checked="" type="radio"/> As absolute limit
	Ending at	
Notification	Notification frequency	1440 (minutes)
Alert Receivers	Receivers	SYSLOG Delete

Figure 5-42 Alert builder

Because the threshold alert is based on query, you must select the available query from the drop-down menu to set up the alert, as shown in Figure 5-43.




Figure 5-43 Select existing query for threshold alert

You also can create a query of your own to use in an alert. To be included in alert builder drop-down menu, a query must include a time stamp field and a numeric field. If you cannot select a numeric field for your query when you are building your query with the Query Builder, you can use the aggregation functions (MIN, MAX, COUNT, and so on) with one of the fields or you use the report counter.

Figure 5-44 on page 187 shows a query example that satisfies both of the time stamp and numeric fields requirements. This query features a time stamp field and counter; therefore, it can be used to generate a threshold alert.

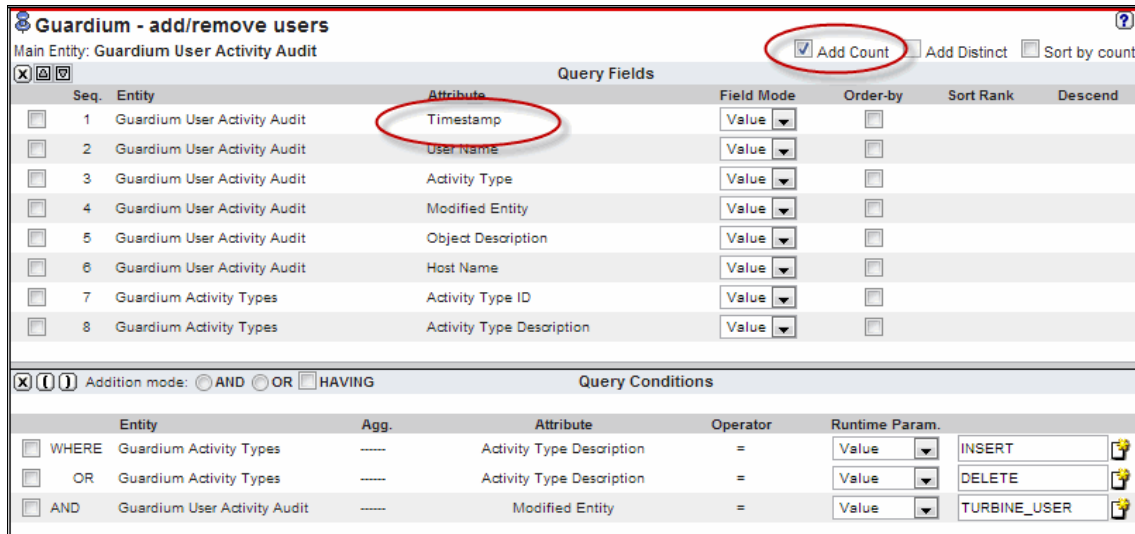


Figure 5-44 Query for threshold alerts

The following time-related fields are in Alert Builder, as shown in Figure 5-45 on page 188:

- ▶ Run frequency field: Indicates how frequently the query that is used in the alert should run.
- ▶ Accumulation interval field: Specifies a period (in minutes) for the query to analyze. The time is counted from now backwards. For instance, 1440 in this interval indicates that query should review data for the last 24 hours from the moment the query runs.
- ▶ Notification frequency field: Indicates how often the notification should be sent to the receiver.

The screenshot shows the 'Alert Builder' window with the following configuration details:

- Settings:**
 - Name: Guardium - Add/Remove Users
 - Description: Alert whenever a user is created/deleted
 - Category: Self Monitoring
 - Classification: (empty)
 - Severity: INFO
 - Run frequency: 1440 (minutes)
 - Active:
 - Log policy violation:
- Alert Definition:**
 - Query: Guardium - add/remove users
 - Accumulation interval: 1450 (minutes)
 - Note: Alerts run on aggregators will be based only on data within the defined merge period
 - Log full query results:
 - Column: (empty) (optional)
- Alert Threshold:**
 - Threshold: 1.0
 - Alert when value is: >= threshold
 - Threshold Evaluated:
 - per report
 - per line
 - Threshold Used:
 - As absolute limit
 - As percentage change within period:
 - From: (empty)
 - To: (empty)
 - As percentage change for the same "Accumulation Period" on a relative time:
 - Ending at: (empty)
- Notification:**
 - Notification frequency: 1440 (minutes)
- Alert Receivers:**
 - SYSLOG [Delete](#)
 - [Add Receiver..](#)

Figure 5-45 Three time intervals

5.7.2 Alerter

When an alert message is generated by the threshold alert or by the real-time alert, it is added to the message queue. The alerter then delivers the message to a recipient.

You can use the alerter configuration window in the Administration Console (see Figure 5-46 on page 189) to configure and activate the alerter process. The window includes the following information:

- ▶ Polling Interval defines how frequently the Alerter should check the queue for new waiting messages. Use the default value of 60 seconds.

- ▶ Configure at least one protocol, SMTP or SNMP (you can configure both).
- ▶ In the SMTP section, the User Name and Password are required only if you select **Auth** in the Authentication method field; otherwise, you can leave them blank.
- ▶ Check **Active on startup** for alerter to start run automatically when the appliance is restarted.

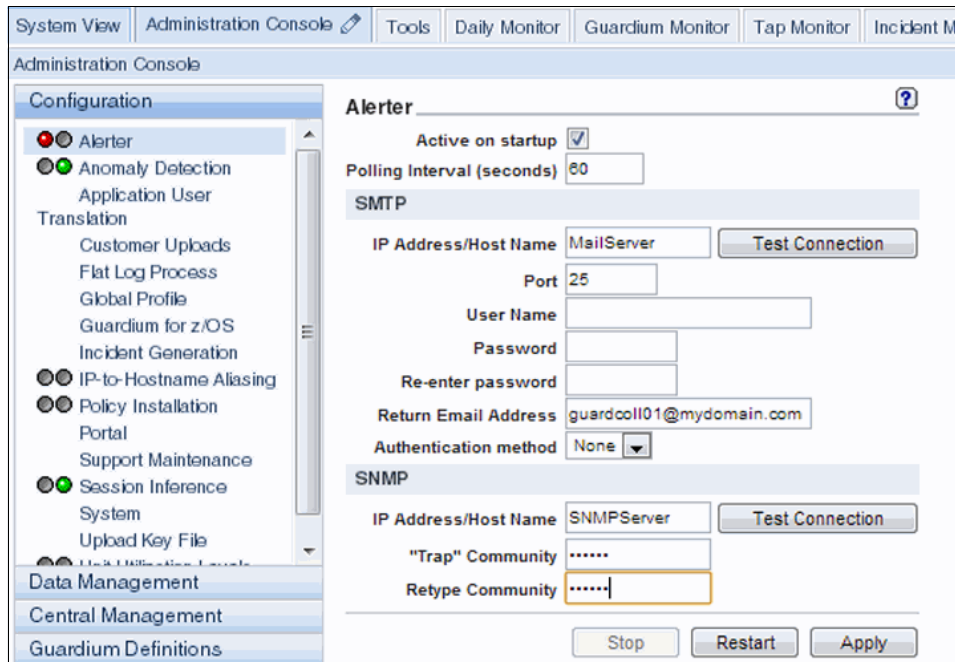


Figure 5-46 Alerter configuration

5.8 Data level access control

Data level access control requires the following tasks:

- ▶ Setting up the access control at the S-TAP level to enable the firewall.
- ▶ Configuring the policy rules to define which activities should cause a session to end. The stop can be performed only if Guardium kernel module or driver is enabled. K-TAP must be enabled for the UNIX platform and the Lhmon driver must be enabled for the Windows platform.

5.8.1 S-TAP setup

Configure the following settings to set up access control at S-TAP:

1. Enable blocking (S-GATE functionality): Set `firewall_installed = 1`.
2. Decide whether blocking is to be configured in Open or Close mode: Use `firewall_default_state = 0` or `firewall_default_state=1`.
3. Decide which action should be taken if S-TAP does not receive a verdict from the collector within the number of seconds that are set in the `firewall_timeout` parameter. End the connection or let the statements pass through by using `firewall_fail_close = 1` or `firewall_fail_close = 0`.
4. Configure `firewall_timeout` to define how long S-TAP can wait for a verdict from the collector. Use the `firewall_fail_close` parameter to control whether a connection must be ended or the statement can be let through.

5.8.2 Policy setup

If blocking is configured in the Open mode, define an access rule with action as S-GATE attach. Select **Continue** for the connections that should be watched for potential blocking.

If blocking is configured in Close mode, define an access rule with action as S-GATE detach for the connections that should not be watched for potential blocking. If any other rules are in the policy for this traffic, you might need to select **Continue**.

Define an access rule with the S-GATE terminate action to specify under what condition the connection must be ended.

Note: During deployment of the S-GATE functionality initially, replace the S-GATE terminate action in the rule with the Alert per match or Alert once per session action to evaluate the correctness of the policy rules that are defined and to review through alerts whether the connections should be ended.

5.8.3 Policy violation report

The policy violation is created when the session is ended. The security team can generate a report of all policy violations that were created by termination rules or generate real-time alerts upon the session ending.

5.9 Vulnerability assessment setup

InfoSphere Guardium Vulnerability Assessment (VA) runs scheduled scans against selected database instances and looks for vulnerabilities and platform-specific issues when they are reported by various database security organizations and the database vendors. The identified vulnerabilities must be prioritized and remediated to help harden the database against threats.

For more information about configuring VA, see 3.10, “Vulnerability assessment” on page 104.

5.10 Configuration audit system setup

CAS is a light-weight, Java based agent that is installed on the database server. It is used to track changes to the database operating environment that can affect the database security, including the following components:

- ▶ Configuration files; for example, the `sqlnet.ora` or `tnsnames.ora` files of Oracle
- ▶ Windows registry variables; for example, `LoginMode` for Microsoft SQL Server
- ▶ File permissions or changes to key executable files; for example, `oracle` or `sqlplus`
- ▶ Environment variables and other database or operating system components

CAS also complements the Guardium VA module, that is, there are some VA tests that use the data that is collected by CAS.

CAS highlights

CAS includes the following highlights:

- ▶ Uses a lightweight agent on the database server to periodically run predefined and custom tests. There is one agent per database server.
- ▶ Requires Java 1.4.2_13 or higher on the database server.
- ▶ Polls randomly (that is, not in real time) within a user-defined polling period for changes to the tracked objects.
- ▶ Can track changes to files or variables that are based on content (what changed) or modification time stamp (when).
- ▶ Can use MD5sum to detect content change, and if specified, report the new and updated values.

- ▶ Is often installed by using the Guardium Installation Manager (GIM) and runs as root on UNIX and Linux, and the administrator on Windows.
- ▶ Requires the S-TAP on the Windows platform, but not on the UNIX and Linux platform.
- ▶ Uses a template (a collection of related tests) to define the objects to monitor.
- ▶ After a template is deployed, the CAS agent expands the test to the actual instance elements; for example, resolve wildcards in file names.
- ▶ Tests can be run as a specific operating system user; for example, `db2inst` or `oracle`.
- ▶ Tests that are used by the VA module are grouped in the special Assessment templates.
- ▶ Reports to a collector appliance, but data can be aggregated to an aggregator for reporting or for use by VA.
- ▶ In a managed environment, CAS templates can be managed on the Central Manager.

5.10.1 Prerequisites

The following prerequisites must be satisfied before CAS is installed:

- ▶ Java 1.4.2_13 or higher is installed on the database servers (CAS host)
- ▶ CAS-module license is applied to the Guardium appliance
- ▶ Install S-TAP installer or GIM bundle for the Windows platform, and CAS installer or GIM bundle for UNIX and Linux (download from IBM Fix Central or Passport Advantage)
- ▶ Firewall ports are open between the database server and collectors; 16017 for all platforms (if TLS is used to encrypt the communication, open port 16019 instead)

Note: For more information, see the *Configuration Auditing System* and *CAS* topics in the *Assess and Harden* chapter of the online *Help Book Guardium*.

5.10.2 High-level steps

The process for installing and configuring CAS includes the following high-level steps:

1. Install and enable the CAS agent on the database server.
 - Windows: CAS is installed automatically during the S-TAP installation by using the GIM.
 - UNIX and Linux: CAS is a separate installer and can be installed without the S-TAP (or after the S-TAP) by using the GIM.
2. Configure CAS templates:
 - a. Determine which predefined CAS templates to deploy to the agents.

A template (template set) is a collection of tests the agent uses to monitor a specific operating system (OS) or a database that is running on that OS; for example, Windows, or Microsoft SQL Server running on Windows.
 - b. Clone and modify the templates to remove or update tests to suit the environment or requirements, if needed.
3. Configure CAS hosts to deploy templates from each collector to the hosts:
 - a. Create database-specific data sources for each host, or modify any existing VA data source for that same database instance.
 - b. Use the CAS Host Configuration navigator to assign one or more templates and a data source to each CAS host (agent). The combination of templates and data source is referred as a CAS *instance*.
4. View the results.

Verify the results of the CAS tests by using the predefined CAS Change Details and CAS Saved Data reports, and the create custom reports.
5. Complete the following tasks:
 - a. If any VA CAS-based tests are used, complete the following:
 - i. Verify that the appropriate assessment template set is deployed.
 - ii. Run and review the output of the VA assessments, which references the CAS tests.
 - b. Adjust the templates and tests as needed.

5.10.3 Installation and configuration

In this section, we describe the general installation and configuration steps.

Installing and enabling the CAS agent on the database server

When CAS is installed and enabled, ensure that the CAS service or processes are started and the agent is reporting to the appliance.

Windows platform

CAS is automatically installed and enabled during the Windows S-TAP installation by using the GIM. Therefore, if the database server already has an S-TAP installed, it already has CAS. CAS is installed in a subdirectory of the S-TAP installation directory.

However, the CAS agent runs as a separate service, Change Audit System. Similar to the S-TAP, CAS agent uses the same `guard_tap.ini` file for some of its configuration settings. As a result, CAS reports to the same collector appliance as the S-TAP.

Note: CAS configuration parameters are not available from the GIM interface; however, it is not typical to change the default settings.

Verifying that the CAS agent is enabled on Windows

If Java is not found during the CAS installation, the S-TAP installation still completes and the CAS Change Audit System service is started, but not enabled. You can verify whether the CAS agent is enabled through the GUI or log. The following conditions can indicate that the CAS agent is not enabled:

- ▶ Agent is not registered with collector, as shown in Figure 5-48 on page 197.
- ▶ Errors exist in the `CasService.log` file on the database server in `<install directory>\stap\cas\Logs`.

To enable the CAS agent, complete the following steps:

1. Install the required Java version.
2. Edit the `cas.cfg` file (in `<install directory>\stap\cas\conf\`). Update and uncomment the JVM argument and the `[RuntimeLib]` tag.
3. Restart the CAS service.

UNIX and Linux platform

Different from the Windows platform where CAS agent is included in the S-TAP installer, the CAS agent is a separate installer for the UNIX and Linux platform. It is available as a GIM bundle (that is, can be installed by using the GIM) and as a shell installer (can be installed directly on the database server if the GIM is not used).

On UNIX and Linux platform, CAS can be installed with or without the S-TAP, and can be installed before or after the S-TAP. When you are installing CAS with GIM, provide the following parameters:

- ▶ **CAS_JAVA_DIR**: Java installation directory; for example, `/usr/java/jre1.6.0_18/`
- ▶ **CAS_SQLGUARD_IP**: The IP or FQDN of the collector appliance
- ▶ **CAS_TAP_IP**: The IP or FQDN of the database server

Tip: To find the Java installation directory, use the `find / -name java_vm -print` command.

Location of some key CAS files on the database server

On database server, the following directories are key CAS file locations:

- ▶ CAS log files:
`<install directory>/modules/CAS/current/files/logs`
- ▶ CAS configuration files:
 - `/usr/local/guardium/modules/CAS/current/files/etc/guard_tap.ini`
 - `/usr/local/guardium/modules/CAS/current/conf/`

After the CAS agent installation, verify the following CAS agent status:

- ▶ The agent is started by use the `ps -ef | grep cas` command.
- ▶ The agent is registered with the collector, as shown in Figure 5-48 on page 197. If the CAS agent is not started, review the log files for more information.

Configuring CAS templates

A template or template set is a collection of tests. There are many predefined templates with hundreds of tests. Although the predefined templates can be used as they are, you should clone and modify the template or update or remove unnecessary tests to meet your requirements and reduce system load.

Tip: The predefined templates and tests are a good source of usage information. Review them to get ideas about what is possible and the syntax.

To view or work with the templates, complete the following steps:

1. Click **Tools** → **Config & Control** → **CAS Template Set Config**, as shown in Figure 5-47 on page 197.
2. Use the List Filtering options to filter by the OS or database types.

3. From the filtered list, select the template to work with and click **Modify** to review the tests.
4. To change the test, use the Clone action, then modify the cloned template. For more information about the attributes and parameters of the test, see the “CAS Templates” topic in the Assess and Harden chapter of the online *Help Book Guardium*.

The following types of templates are available, as shown in Figure 5-48 on page 197:

- ▶ **System:** Contains tests of OS files or variables but does not contain any database-specific tests.
- ▶ **OS/DB:** Contains database-specific tests for a particular OS; for example, Windows or Microsoft SQL server.
- ▶ **Assessment:** Contains tests that are linked to the VA module for the specified database platform and OS. This type of template must first be deployed to collect data for the related VA tests to work.

Each template can contain one or more of the following types of tests:

- ▶ **OS script:** Used to run a script, executable file, or OS-specific commands.
- ▶ **SQL query:** Used to run SQL commands after the agent connects to the database. Although available, use the VA query-based tests instead.
- ▶ **Environment or registry variable:** Used to check changes to an environment or registry variable. You can also use an OS script test to check these variables.

Note: CAS is a 32-bit Java application, so it does not have access to the 64-bit Windows registry keys. On a 64-bit server, it accesses only the 32-bit application keys that are stored in the WOW6432Node area; for example, HKLM\Software\WOW6432Node. This is important to note when you are creating your own custom test. The predefined tests already take this issue into account.

- ▶ **File:** Used to monitor file changes and file permissions. For predefined VA, CAS-based tests allow the VA module to compare the reported file ownership and permissions with what is expected.
- ▶ **File pattern:** Used to specify file tests against a collection of files that match a pattern; for example, monitor changes to any `init.*ora` files, where the pattern is defined using regex.

Figure 5-47 shows some predefined templates.

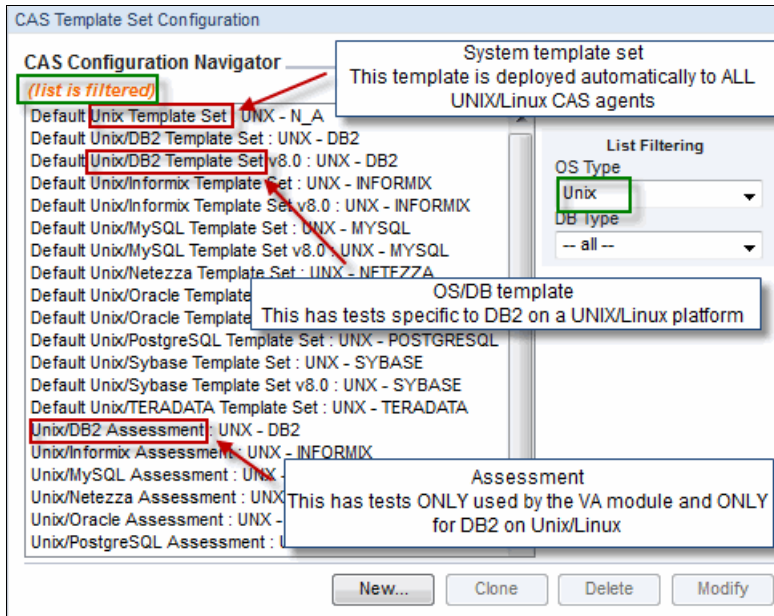


Figure 5-47 Predefined templates

Figure 5-48 shows the CAS status window on the collector appliance.

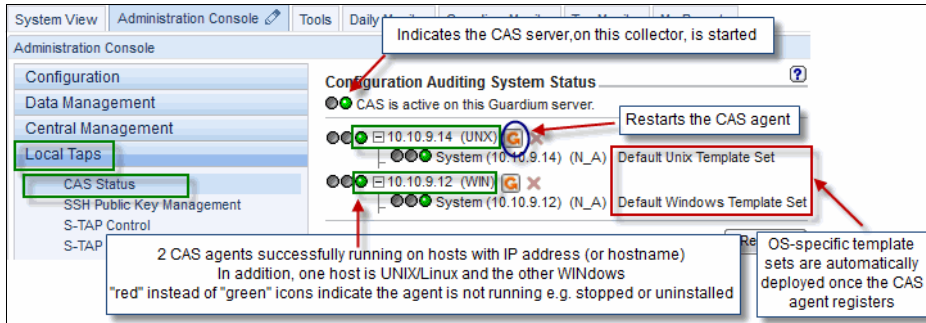


Figure 5-48 The CAS status window on the collector appliance

Configuring CAS hosts

The default system template is automatically deployed to the host when the agent registers. It is automatically redeployed if it is inadvertently removed from the host.

A CAS instance specifies a template set and any parameters that are required to run the tests; for example, OS user or directory path. It is typical to deploy a mix of system and database-specific templates to each host; for example, one CAS instance to monitor OS items (added by default) and another CAS instance to monitor databases.

Before you assign a database-specific template, a data source is required (the system templates use a system-generated data source).

You can use the CAS section of the data source form to specify the database instance account and the installation directory to use for running the OS tests for that specific database instance. For more information, see the *Help Book Guardium*.

Figure 5-49 shows the CAS specifications for an Oracle database on UNIX data source.

Datasource Definition

Name: dev_ora_rhe54db_orcl102_sqlguard

Database Type: Oracle (Change to: Oracle (DataDirect))

Severity classification: NONE

Description:

Share Datasource:

Authentication

Save Password:

Login Name: sqlguard

Password: ●●●●●●

Location

Host Name/IP: 10.10.9.14

Port: 1521

Service Name: orcl102

Informix Server:

Schema:

Conn Pr: An Oracle database is running on this host the Oracle OS user information is provided here

Custom Uri:

CAS

Database Instance Account: oracle10

Database Instance Directory: /home/oracle10/oracle/product/10.2.0/db

Figure 5-49 CAS specifications for Oracle on UNIX data source

To assign a template to a CAS host, complete the following steps:

1. On the collector appliance, click **Tools** → **Config & Control** → **CAS Host Config**.
2. The CAS Configuration Navigator lists all of the hosts (agents) that are registered with this collector appliance. Select the host, or use the List Filtering options to filter the list, and then click **Modify**.

Figure 5-50 shows the CAS Configuration Navigator.




Figure 5-50 CAS Configuration Navigator

3. On Host Instance Definitions, complete the following steps (as shown in Figure 5-51):
 - a. Select the template set from the pick list.
 - b. Click **Add Datasource** and select the data source for this database on this host.
 - c. After the instance is added, the Monitored Items shows a count of 0 items (0 enabled). However, after the agent processes and expands the test (for example, directory items or wildcards), the enabled item count increases. Click **Refresh** to see the updates.

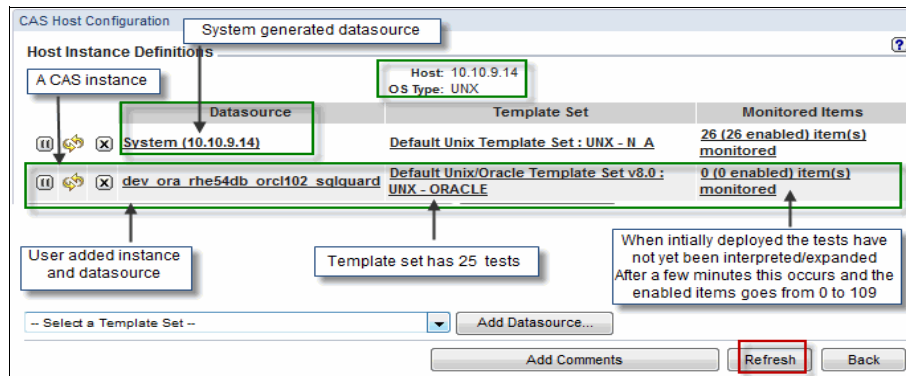


Figure 5-51 CAS host instances

Figure 5-52 shows a partial example of expanded tests.

The screenshot displays two panes. The left pane, 'Monitored Item Template Definitions', shows a list of templates under 'Template Set Name: Default Unix/Oracle Template Set v8.0'. The right pane, 'Monitored Item Definitions', shows the expanded items for 'Datasource dev_ora_rhe54db_orcl102_sqguard' under 'Template Set Default Unix/Oracle Template Set v8.0'. A callout box states: 'Template tests are expanded to a large number of actual tests i.e. wildcards are interpreted as is \$ORACLE_HOME'. Red arrows indicate the expansion from the 'Apache' template in the left pane to the corresponding file paths in the right pane.

Item	Type	Period	Use	MD5	Ki
\$ORACLE_HOME/dbs/*	File Pattern	1h			
\$ORACLE_HOME/inventory/Templates/bin/*	File Pattern	1h			
\$ORACLE_HOME/jdk/jre/bin/*	File Pattern	1h			
\$ORACLE_HOME/jre/*bin/*	File Pattern	1h			
\$ORACLE_HOME/network/admin/*	File Pattern	1h			
\$ORACLE_HOME/oui/bin/./so	File Pattern	1h			
\$ORACLE_HOME/Apache/Apache/bin/*	File Pattern	1h			
\$ORACLE_HOME/Apache/Apache/open_ssr/bin/open_ssr...	File Pattern	1h			
\$ORACLE_HOME/assistants/dbma/mep.cfg	File	1h			
\$ORACLE_HOME/ds/bin/*	File Pattern	1h			
\$ORACLE_HOME/javavm/admin/.so	File Pattern	1h			
\$ORACLE_HOME/jdk/bin/*	File Pattern	1h			
\$ORACLE_HOME/ocommon/ntsbuilder/builder	File	1h			
\$ORACLE_HOME/olap/ol*	File Pattern	1h			
\$ORACLE_HOME/soap/bin/*	File Pattern	1h			
\$ORACLE_HOME/syndication/bin/*	File Pattern	1h			
\$ORACLE_HOME/sysman/admin/OMSRepositoryConstraints.properties	File	1h			
\$ORACLE_HOME/sysman/config/.properties	File Pattern	1h			
\$ORACLE_HOME/xdk/admin/xml.properties	File	1h			
ORACLE_BASE	Environment Variable	10m			
ORACLE_HOME	Environment Variable	10m			

Figure 5-52 Partial example of expanded tests

- To remove an instance, click the **x** icon next to the instance.

5.10.4 Reviewing results

There are several predefined CAS-related reports and two of them are the main reports for viewing the test results.

To view these reports, click **Tap Monitor** → **CAS** → **Changes** on the collector. This pane lists the following reports:

- ▶ **CAS Change Details:** Lists the changes that are observed for each monitored item (test).

Note: The results (baseline) for all tests are reported from the first run. However, on subsequent runs (determined by the period setting), only changes from the baseline are reported.

- ▶ **CAS Saved Data:** Lists the actual detected changed data value for those monitored items (tests) with the Keep Data selected.

In addition to the predefined reports, custom queries and reports can be created by using the CAS Changes Tracking domain and query builder.

5.10.5 Next steps

After the CAS agent is installed and configured, complete the following tasks:

- ▶ Review the results and adjust the tests as needed.
- ▶ Deploy other CAS host instances.
- ▶ If using VA, consider the following tasks:
 - Deploying the appropriate CAS assessment templates to the target hosts.
 - Adding the VA CAS-based tests to the VA assessments, then running and reviewing the output of the VA assessments.

A No CAS Data - No data available for this test error is an indication that the CAS test was not deployed or has not yet run.

- ▶ Create and schedule any required custom CAS reports.

5.11 Entitlement reporting setup

Entitlement (privilege) reports allow for the review of database accounts (users) to validate these users have only the privileges required to perform their duties. Auditors usually review database accounts and their privileges, and some of the following reports are requested:

- ▶ All accounts that are defined on a database instance system
- ▶ All accounts with certain object and or system privileges

The designated Guardium appliance connects to each database instance and uploads accounts and privileges information into Guardium from the database catalogs and the data-dictionary tables. This function does not use an S-TAP.

After the data is in Guardium, this information can be reported or used to populate Guardium groups that are used in database activity monitoring.

For more information about the entitlement reports, see the “Database Entitlement Reports” topic of the Appendices chapter of the online *Help Book Guardium*.

5.11.1 Prerequisites

The following prerequisites must be satisfied to use the entitlement reports:

- ▶ Determine which platform and what entitlement information is required. For more information, see the “Database Entitlement Reports” topic of the Appendices chapter of the online *Help Book Guardium* and make a list of the entities to use, as shown in the following examples:
 - ORA Accts of ALTER SYSTEM: Accounts with the ALTER SYSTEM and ALTER SESSION privileges.
 - DB2 Column-level Privileges (SELECT, UPDATE, and so on)
- ▶ The Entitlement Reports module license is applied to the Guardium appliance.
- ▶ The Entitlement database account-creation scripts is downloaded. The scripts are available for download from IBM Fix Central or from IBM Passport Advantage.

Note: If you are using, or plan to use the VA module, you can use the VA creation scripts and or the same database accounts to minimize the setup effort.

5.11.2 High-level steps

The process that is used to configure Entitlement reporting includes the following high-level steps:

1. DBAs create a database account (by using the Guardium-provided script) on each database instance from which to collect entitlement information.
2. In Guardium, create a data source for each of the database instances by using the account information that is provided by the DBAs.
3. Configure the Guardium appliance to periodically retrieve the required entitlement information from the source databases.
4. Prepare and review the predefined Entitlement Reports.
5. (Optional) Schedule reports for delivery.
6. (Optional) Create and schedule customized entitlement reports.

5.11.3 Configuration steps

The configuration process includes the following steps:

1. DBA uses the database platform-specific scripts to create the database accounts; for example, `gdent-0racle.sql` or `gdmmonitor-ora.sql`.

Note: The `gdent` prefix is used for the scripts that are used for entitlement reporting; the `gdmmonitor` prefix is used for VA scripts. The VA scripts includes the privileges that are required for entitlement reporting.

2. Create the data source.

For more information about creating the data sources and the database accounts, see 3.10.2, “Creating a data source” on page 106.

3. Identify entities.

As the Guardium Administrator, click **Tools** → **Report Building** → **Custom Table Builder**. Scroll through the list to identify the entities that were selected during the prerequisite phase; for example, `ORA Acnts` of `ALTER SYSTEM`.

Note: Click the help icon **?** to get a description of each custom table, or see the “Database Entitlement Reports” topic of the Appendices chapter of the online *Help Book Guardium*.

4. Select the entity and click **Upload Data** to configure the following settings:
 - Overwrite or purge options.
 - Assign data sources.
 - Assign a schedule to periodically upload the entitlement data.
 - Fetch records from source database catalog into Guardium.

Figure 5-53 on page 205 shows the selection of an entitlement reporting entity for configuration.

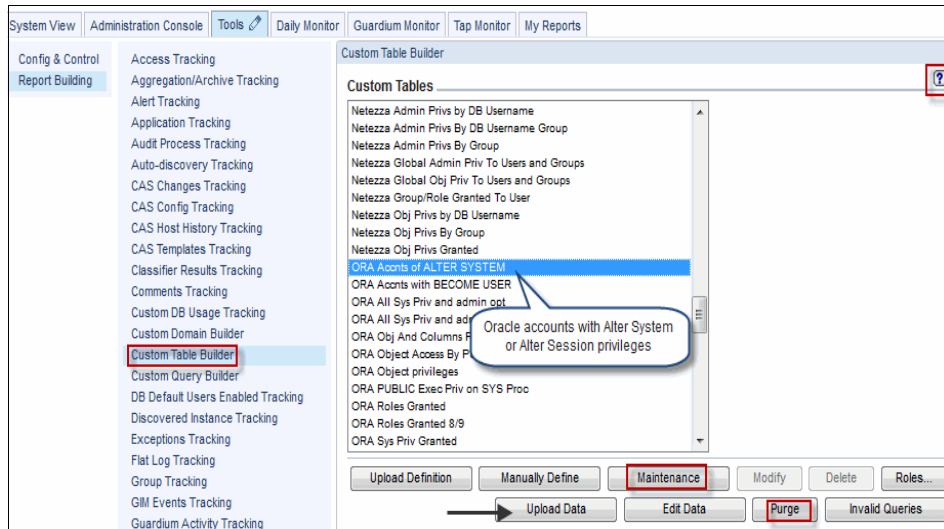


Figure 5-53 Selecting an entitlement reporting entity for configuration

5. (Optional) After these items are configured, click **Run Once Now** to load data from the data sources so you can review the data.

Figure 5-54 on page 206 shows the configuration options for uploading entitlement data.

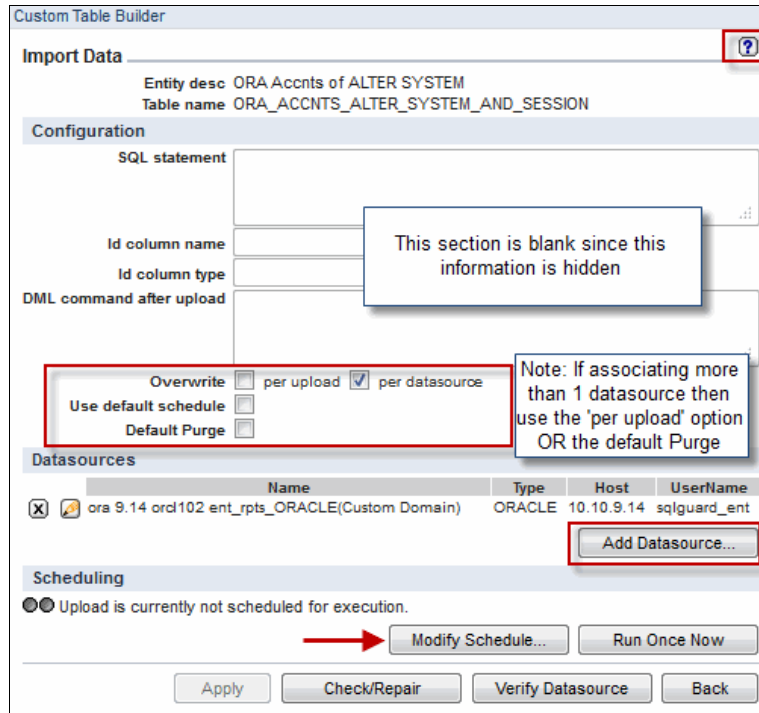


Figure 5-54 Configuration options for uploading entitlement data

Figure 5-55 shows the options for setting and scheduling the purge of a custom table invoked from **Purge** on the Custom Table Builder window.

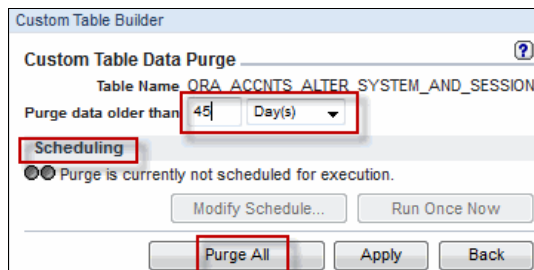


Figure 5-55 Options for setting and scheduling the purge of a custom table

6. Repeat these steps for each identified entity.

Note: The entitlement data uses the Guardium custom table feature. As a result, it has functionality that is associated with custom tables; for example, add indexes, change database storage, and customized purge.

5.11.4 Review the entitlement data

To review the entitlement information that is uploaded to the Guardium appliance, use the predefined reports or create custom reports that are based on the predefined reports.

As the Guardium Administrator, create a menu tab on your portal to stage the reports or use the My New Reports tab, if it is available (that is, your account has the admin and user roles). For more information, see 5.14, “Adding a menu tab to your portal” on page 225.

Accessing predefined reports

To access the predefined reports, complete the following steps:

1. Click **Tool** → **Report Building** → **Report Builder**.
2. Click **Search** (do not select anything on this window).
3. In the Report Search Results listing, scroll down and select the report that has the same name as the entitlement report entity or custom table that was configured; for example, ORA Accnts of ALTER SYSTEM.
4. Click **Add to Pane...** and select the menu tab on which to place the report; for example, the newly created menu tab My Reports.
5. After the confirmation is shown that the report was added to the tab, click the tab to view and run the report.

Custom reports

Although you can create a custom query and report from scratch, it is easier to clone the predefined query, make changes (for example, add columns or query conditions) and generate a new report.

Figure 5-56 on page 208 shows a cloned, customized query.

The screenshot shows a query builder interface. On the left is an 'Entity List' with items like 'ORA Accnts of ALTER SYSTEM', 'Grantee', 'Privilege', 'Admin Option', 'SqlGuard Timestamp', 'Datasource Details', 'DataSource Id', 'DataSource Name', 'DataSource Host', 'DataSource Port', 'DataSource', 'DataSource DB', and 'DataSource Type'. The main area is titled '-My ORA Accnts of ALTER SYSTEM' and contains a 'Main Entity: ORA Accnts of ALTER SYSTEM'. Below this are two tables: 'Query Fields' and 'Query Conditions'. The 'Query Fields' table has columns for Seq., Entity, Attribute, Field Mode, Order-by, Sort Rank, and Descend. The 'Query Conditions' table has columns for Entity, Agg., Attribute, Operator, and Runtime Param. The 'Query Conditions' table shows three conditions: 'WHERE ORA Accnts of ALTER SYSTEM Grantee <> Value DBA', 'AND ORA Accnts of ALTER SYSTEM Grantee <> Value SYS', and 'OR ORA Accnts of ALTER SYSTEM Admin Option = Value YES'. The 'Addition mode' is set to 'AND'.

Seq.	Entity	Attribute	Field Mode	Order-by	Sort Rank	Descend
1	ORA Accnts of ALTER SYSTEM	Grantee	Value	<input checked="" type="checkbox"/>	1	<input type="checkbox"/>
2	ORA Accnts of ALTER SYSTEM	Privilege	Value	<input checked="" type="checkbox"/>	2	<input type="checkbox"/>
3	ORA Accnts of ALTER SYSTEM	Admin Option	Value	<input checked="" type="checkbox"/>	3	<input type="checkbox"/>
4	ORA Accnts of ALTER SYSTEM	Datasource Details	Value	<input type="checkbox"/>		
5	ORA Accnts of ALTER SYSTEM	SqlGuard Timestamp	Value	<input type="checkbox"/>		

	Entity	Agg.	Attribute	Operator	Runtime Param.
<input type="checkbox"/>	WHERE ORA Accnts of ALTER SYSTEM	-----	Grantee	<>	Value DBA
<input type="checkbox"/>	AND ORA Accnts of ALTER SYSTEM	-----	Grantee	<>	Value SYS
<input type="checkbox"/>	OR ORA Accnts of ALTER SYSTEM	-----	Admin Option	=	Value YES

Figure 5-56 A cloned, customized query

5.12 Database auto-discovery setup

The database auto-discovery module scans hosts (servers) in the network to discover new or unknown database instances. It does not require an agent. This process is useful for discovering databases instances that might have been created outside the normal provisioning process and might include sensitive data.

Auto-discovery starts the following actions:

- ▶ A scan of hosts for open ports that are defined by the user
- ▶ A probe for any database services that are listening on those open ports

After it is configured, the process can be run on an ad hoc or scheduled basis.

For more information, see the “Database Auto-discovery” topic in the Discover chapter of the online *Help Book Guardium*.

Note: The database auto-discovery is not the same as the instance discovery. The instance discovery uses a lightweight Java agent that is deployed with the GIM client on the database servers to check for new database instances.

5.12.1 Prerequisites

The following prerequisites must be met for database auto-discovery:

- ▶ Database and Sensitive Data Finder module license are applied to the Guardium appliance.
- ▶ An appliance from which to run the network scans. The appliance must have network access to the network segments to be scanned. In a managed environment, this can be the Central Manager.
- ▶ The scan process is disabled by default, but it is enabled by applying a special patch to the appliance. From the IBM Fix Central portal, download InfoSphere_Guardium_DB_Discovery_Activation_Patch. IBM Fix Central is available at this website:

<http://www.ibm.com/support/fixcentral/>

5.12.2 High-level steps

The process that is used to configure database auto-discovery includes the following high-level steps:

1. Apply the activation patch to the appliance.
2. Configure one or more auto-discovery processes to search for hosts with database services.
3. Run the auto-discovery process on-demand or create a schedule.
4. View auto-discovery reports or create custom reports.

5.12.3 Configuration steps

After the activation patch is installed, complete the following configuration steps:

1. As the Guardium Administrator, click **Tools** → **Config & Control** → **Auto-discovery Configuration** and then click **New**.
2. Enter a description for this process; select **Run probe after...** and click **Apply**, as shown in Figure 5-57 on page 210.

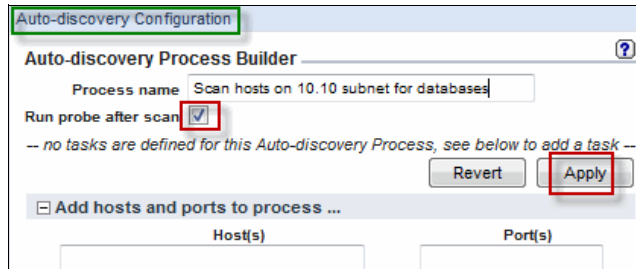


Figure 5-57 Creating a database auto-discovery process

3. Add specific IP addresses or subnets (by using wildcards), and ports, including port ranges, to scan. As each combination is added, the list of potential hosts and ports to scan are calculated and displayed.
4. When you are finished entering the combinations, click **Run Once Now** for the Scheduling - Scan for open ports task to start the scan and probe.

Figure 5-58 shows the configuring of hosts and ports.

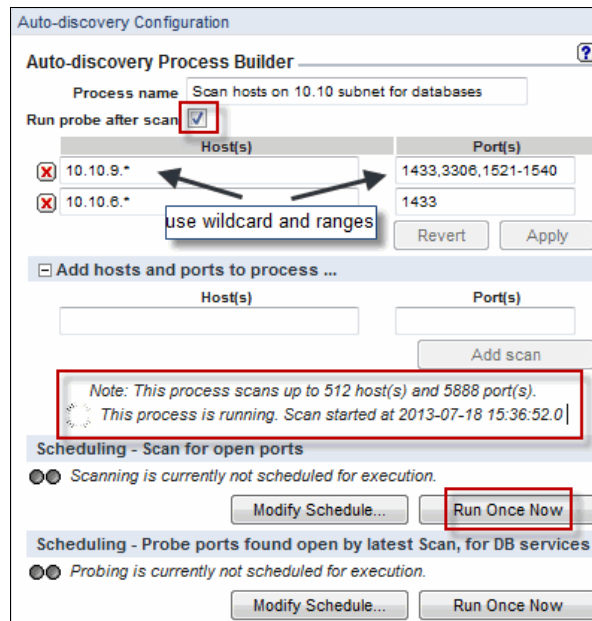


Figure 5-58 Configuring hosts and ports

5. To check on the progress of a scan and probe, select the process and click **Progress/Summary**, as shown in Figure 5-60 on page 212.

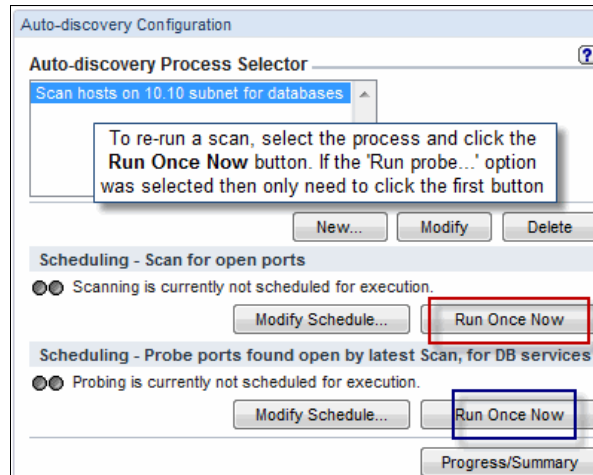


Figure 5-59 Steps to rerun a process

6. To re-run a process, select the process in the Auto-discovery Process Selector list and click **Run Once Now** for the Scheduling - Scan for open ports task. If the Run probe after scan option was not selected, you must also click **Run Once Now** for the Scheduling - Probe ports.
7. Click **Modify Schedule** in the Auto-discovery Configuration window to schedule the selected process.

5.12.4 Viewing the results

The following reports are associated with the auto-discovery process:

- ▶ A high-level or summary report that can be viewed directly from the Auto-discovery Configuration
- ▶ A detailed predefined report of discovered databases

Figure 5-60 on page 212 shows the Progress/summary report.




Figure 5-60 Viewing the summary report for a selected auto-discovery process

Figure 5-61 shows the auto-discovery predefined detailed report.

Time Probed	Server IP	Server Host Name	DB Type	Port	Port Type	#
2013-07-18 15:45:24.0	10.10.9.14	10.10.9.14	Oracle	1521	tcp	1

Figure 5-61 Example of the auto-discovery predefined detailed report

In addition to the predefined report, a custom query and report can be created by using the Auto-discovery Tracking domain and query builder. For example, create a custom query with conditions to exclude known databases from the results by using Guardium groups.

Figure 5-62 shows the domain to use for creating an auto-discovery custom query.

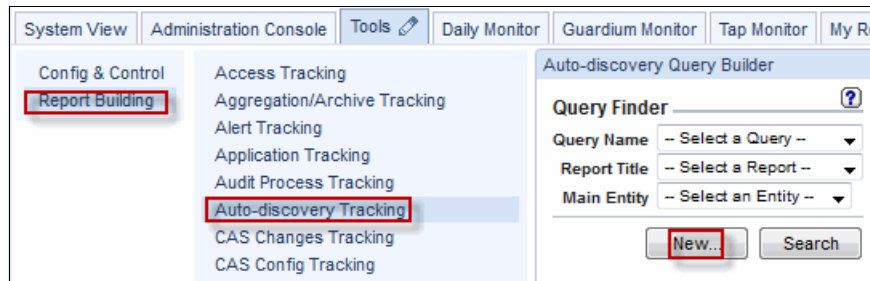


Figure 5-62 Domain to use for creating an auto-discovery custom query

5.12.5 Next steps

The following next steps are suggested:

- ▶ Cross check the results from the reports with known databases, identify the owners of the unknown databases, and determine if the instances must be monitored by Guardium.
- ▶ Update the known database groups to filter the reports and update the auto-discovery processes to exclude specific hosts, if possible.

5.13 Sensitive data finder setup

The classifier assists with the discovery and classification of sensitive data to help inform access policy decisions. Although this function does not require a Guardium agent, it does require a connection to the database instance that is checked for sensitive data.

5.13.1 Use cases and highlights

The classifier features the following use-cases and highlights:

- ▶ The classifier uses regular expressions (regex), pattern matching, or SQL command syntax to search within a database for the following components:
 - Data search:
 - Sensitive data; for example, credit card numbers, social security numbers, or financial data
 - Can be used to identify whether more than one type of data is present in a single table; for example, social security number and date-of-birth
 - Catalog search: Objects and or columns with names having a certain pattern; for example, SAP_
 - Permission search:
 - Objects and their users and or roles having certain types of grants; for example, Alter, and whether those grants were granted with the admin option
 - Who has access to privileged commands; for example, CREATE USER or BACKUP DATABASE
 - Unstructured data: Search a text file by using a regex pattern; for example, search a CSV file for credit card numbers
- ▶ Regex and pattern matching expressions are not precise and can result in false-positives. Therefore, test and tweak these types of search patterns to minimize false positives.
- ▶ When the specified data is found, the classifier can start specified the following actions:
 - Add to Group of Objects or Group of Objects and Fields
Add the object name and the column name to a Guardium groups. For example, add the table name that contains credit card numbers to a sensitive object group. That sensitive object group is used in the DAM policy.

- Send Alert

Send an alert by using the Guardium alerting mechanism. For example, alert DBA management who has access to highly privileged commands.

Figure 5-63 shows a full list of rule actions.

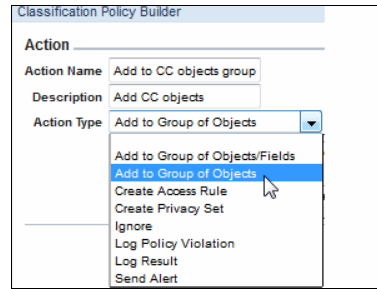


Figure 5-63 List of rule actions

Figure 5-64 shows the Add to Group of Objects action.




Figure 5-64 Add to Group of Objects action

- ▶ The classifier searches can be run on-demand or scheduled
 - Searches use default sample sizes to minimize the time and effect on the database.
 - For some searches, the processing time is dependent on the amount of objects in the database, but there are techniques for tuning these searches.

Using the Luhn algorithm

The Luhn algorithm is a widely-used algorithm for validating credit card numbers.

When a classification rule name begins with `guardium://CREDIT_CARD` and there is a valid credit card number pattern (for example, `[0-9]{16}`), the classification policy uses the Luhn algorithm to check the value that is found.

Note: For more information, see the Classification topics in the Discover chapter of the online *Help Book Guardium*.

5.13.2 Prerequisites

The following prerequisites should be met for using classifier:

- ▶ Database and Sensitive Data Finder module license is applied to the Guardium appliance.
- ▶ DBA creates a database account for the databases to be scanned. This account must have the SELECT access on the database objects you plan on scanning; for example, tables, views, catalogs.

Note: Guardium does not provide an account creation script. This account should have a strong password and, if possible, should be disabled after it is used.

- ▶ List of criteria or type of data to search for and what patterns to use. It is preferable to be familiar with regex.

5.13.3 High-level steps

The process that is used to create and run the classifier (that is, a classification process) includes the following high-level steps:

1. Define a Classification Policy: Collection of search rules.
2. Define the Classification Process: To associate a policy with data sources.
3. Run the Classification Process.
4. View Results of the Process.
5. Complete Next Steps.

5.13.4 Configuration steps

In this section, we use a simple example of one feature of the classifier (data search for credit card numbers and the Luhn algorithm) to show the configuration. For more information about other features, see the Classification topics in the Discover chapter of the online *Help Book Guardium*.

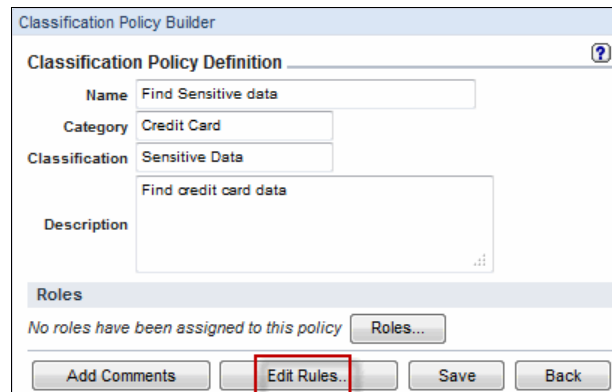
Defining a classification policy

By using the Classification Policy Builder, define a policy that is composed of one or more search rules and the actions to take; for example, add the results to a Guardium group.

This policy can be independent of a database type and often targets a certain data domain; for example, sensitive data, privileged commands, or object permissions.

To create a classification policy, complete the following steps:

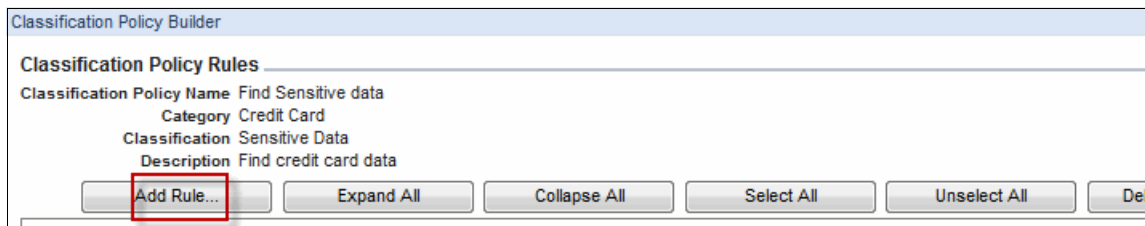
1. As the Guardium Administrator, click **Tools** → **Config & Control** → **Classification Policy Builder** and then click **New** to define a new policy.
Enter a policy name, category, and classification (a description is optional). Click **Save** to save the definition and enable the Edit Rules option, as shown in Figure 5-65.



The screenshot shows the 'Classification Policy Builder' window. The title bar reads 'Classification Policy Builder'. Below the title bar is a section titled 'Classification Policy Definition' with a help icon. It contains four input fields: 'Name' with the text 'Find Sensitive data', 'Category' with 'Credit Card', 'Classification' with 'Sensitive Data', and 'Description' with 'Find credit card data'. Below these fields is a 'Roles' section with the text 'No roles have been assigned to this policy' and a 'Roles...' button. At the bottom of the window are four buttons: 'Add Comments', 'Edit Rules..', 'Save', and 'Back'. The 'Edit Rules..' button is highlighted with a red rectangle.

Figure 5-65 Initial step to create a classification policy

2. Click **Edit Rule** and then **Add Rule** to add a search rule, as shown in Figure 5-66.



The screenshot shows the 'Classification Policy Builder' window. The title bar reads 'Classification Policy Builder'. Below the title bar is a section titled 'Classification Policy Rules'. It displays the policy details: 'Classification Policy Name Find Sensitive data', 'Category Credit Card', 'Classification Sensitive Data', and 'Description Find credit card data'. Below these details are several buttons: 'Add Rule...', 'Expand All', 'Collapse All', 'Select All', 'Unselect All', and 'De'. The 'Add Rule...' button is highlighted with a red rectangle.

Figure 5-66 Adding a rule to classification policy

The following types of rules or searches are available that require different inputs:

- Search for data: Search one or more columns.
 - Catalog Search: Search database catalog for table or column.
 - Search by Permissions: Search the catalog for objects with certain permissions.
 - Search for Unstructured Data: Searches non-database files. However, the files, such as the databases, must be accessible from the appliance.
3. Complete the following steps to create a Search for Data rule that uses a Search Expression to search for credit card numbers and then uses the Luhn algorithm to verify the match:
- a. Enter the Rule Name (the Category and Classification are pre-populated from the policy).
 - b. Check Table only. In this example, it is not necessary to check the System Table because it is unlikely that the system tables include credit card information. Similarly, skip Synonym or View, because they reference a table.
 - c. Select the data type of Text and specify a Maximum length of 16.
 - d. The Search Expression field requires a regex. Use the regex builder to create and test the expression.

This process is shown in Figure 5-67 on page 219 and Figure 5-68 on page 219.

Classification Policy Builder

Classification Rule #1 For Classification Policy "Find Sensitive data"

Rule Name:

Category: Credit Card

Classification: Sensitive Data

Description: Find Tables with Credit card data and use the Luhn algorithm to verify. Rule name MUST start with guardium://CREDIT_CARD to trigger the use of the Luhn to check the value

Continue on Match:

Rule Type: Search For Data

Table Type: Synonym System Table Table View

Table Name Like:

Data Type: Date Number Text

Column Name Like:

Minimum Length:

Maximum Length:

Search Like:

Search Expression:

Evaluation Name:

"Fire only with" Marker:

Hit Percentage:

Compare to Values in SQL:

Compare to Values in Group: -- none --

Show Unique Values:

Unique Value Mask:

Classification Rule Actions

Annotations: These lengths are checked against the column definition e.g. a column with a length of 20 will NOT be checked. to launch the regex builder and tester

Figure 5-67 Classification rule example

Build Regular Expression

Regular expression:

Select element:

(Case sensitive)

Text to match against:

Figure 5-68 Regex builder and tester

- After the rule is created, click **Save**.

5. In the Classification Policy Rule listing (see Figure 5-69), click the pencil icon to re-open the rule form and then click **Add Action**.




Figure 5-69 Edit the rule

6. Add the actions to take if the rule is satisfied; that is, data is found that matches the search pattern. Click **New Action**, which opens the Action window.

This example uses the Add to Group of Objects action which includes the following tasks:

- Add the names of the tables to the group that is specified in the Object Group field.
 - The Actual Member Content field specifies how the table name should be added. In this example, it is prefixed with a wildcard; for example, %cc_card, where cc_card is the table.
7. You can add more rules to this classification if necessary; for example, rules to look for credit card numbers that might include the embedded dashes.

Defining the classification process

With a classification policy created and populated with rules, complete the following steps to use the Classification Process Builder to create a process and associate data sources (that is, database connections) to run the policy:

1. Select **Tools** → **Config & Control** → **Classification Process Builder** and click **New** to define a new process.
2. Enter a Process Description. From the Classification Policy pick list, select the previously defined policy and add one or more data sources that the process uses to connect to the target databases and run the classification policy.

For more information about the Comprehensive Search and Sample size settings, see to the “Classification Process” topic in the Discover chapter of the online *Help Book Guardium*.

Figure 5-70 on page 221 shows the classification process with an Oracle data source.

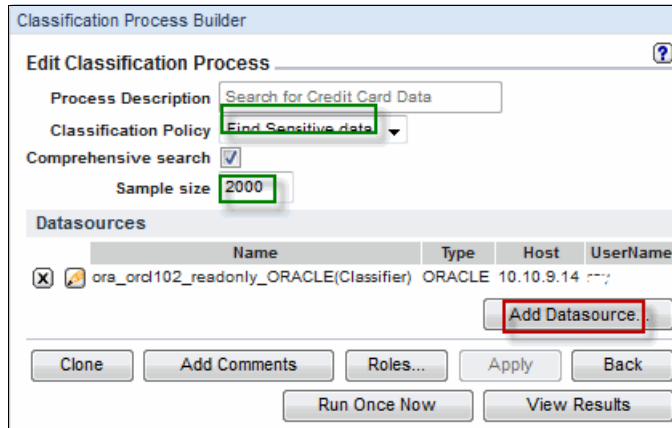


Figure 5-70 Classification process with an Oracle data source

Running the classification process

By using the Classification Policy Builder, run the classification process immediately (that is, Run Once Now) to review the results and adjust the rules, if needed, to minimize any false-positives.

The processing time depends in large part on the following factors:

- ▶ Number and type of search rules
- ▶ Number of target databases (data sources) to search
- ▶ Sample size and Comprehensive search settings

On-demand running

To run a classification process now, complete the following steps:

1. Select **Tools** → **Config & Control** → **Classification Process Builder**. In the Classification Process Finder, select the process and click **Run Once Now**, as shown in Figure 5-71.

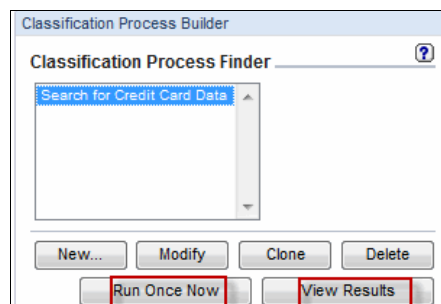


Figure 5-71 Run or view the results of the selected process

2. To check the process status, click **Guardium Monitor tab** → **Guardium Job Queue**.
3. After the job completes, review the results. For more information, see 5.13.5, “Viewing classification process results” on page 222.

Figure 5-72 shows the classification process in the job queue

Process Run Id	Process Type	Status	Process Id	Report Result Id	Guardium Job Description	Task Description	Queue Time	Start Time	End Time	Datasources
32	CLASSIFICATION COMPLETED	20000	1	1	Search for Credit Card Data		2013-07-22 11:26:00.0	2013-07-22 11:26:38.0	2013-07-22 11:28:16.0	ORACLE ora_orcl102_readonly

Figure 5-72 Classification process in the job queue

Scheduled running

To schedule the running of a classification process, use the Audit Process Builder and add a Classification Process task. For more information, see 5.6, “Compliance workflow” on page 179.

5.13.5 Viewing classification process results

The following types of reports are available to view the classification process results:

- ▶ Predefined reports
- ▶ Custom reports

Predefined reports

After the job completes as per the Guardium Job Queue report, click **View Results** on the Classification Process Finder window to view the output, as shown in Figure 5-71 on page 221.

The predefined classification report shows the process log and details of the search, as shown in Figure 5-73 on page 223.

IBM InfoSphere™ Guardium®

Results for Classification Process / Policy: **Search for Credit Card Data** -- Select another result of this Process --

From: 7/22/13 11:26 AM To: 7/22/13 11:28 AM

Datasources: ORACLE ora_orcl102_readonly

Process Run Log:

2013-07-22 11:26:43	Process Start	Started processing of Classification Process: 'Search for Credit Card Data'...
2013-07-22 11:26:43	Info	Connecting to Datasource: 'ORACLE ora_orcl102_readonly'
2013-07-22 11:26:59	Info	Successfully connected to Datasource: 'ORACLE ora_orcl102_readonly' at 'ORACLE 10.10.9.14: 1521 orcl102' in: 15.93 seconds
2013-07-22 11:27:32	Error	SQLException: ORA-00600: internal error code, arguments: [8003], [], [], [], [], [], [] SqlState: 60000 Error Code: 600 Unable to retrieve results for: 'SYS.CDEFS' columns: 'RRULES.CONDITION'.
2013-07-22 11:27:43	Error	SQLException: Protocol violation: [3] SqlState: 99999 Error Code: 17401 Unable to retrieve results for: 'SYS.IDL_CHARS' columns: 'PIECE'.
2013-07-22 11:28:02	Error	SQLException: ORA-00600: internal error code, arguments: [8015], [], [], [], [], [], [] SqlState: 60000 Error Code: 600 Unable to retrieve results for: 'SYS.VIEWS' columns: 'TEXT'.
2013-07-22 11:28:15	Datasource Statistics	"Statistics for Datasource: 'ORACLE ora_orcl102_readonly' at 'ORACLE 10.10.9.14: 1521 orcl102'"
2013-07-22 11:28:16	Process Complete	...Completed processing of Classification Process: 'Find Sensitive data'

Report details: horizontal vertical Show original values Use Aliases

Catalog	Schema	Table Name	Column Name	Rule Description	Comments	Classification Name	Category	Datasource Description
<input checked="" type="checkbox"/>	SCOTT	CC_NUMB_16	CC	guardium://CREDIT_CARD	Date: Monday, July 22, 2013 11:27:18 AM EDT Datasource: ORACLE 10.10.9.14: 1521 orcl102 Object: ROY.CC_NUMB_16 CC Category: 'Credit Card' Classification: 'Sensitive Data' Comprehensive: true Rule: Search For Data: guardium://CREDIT_CARD TABLE_TYPE='TABLE', DATA_TYPE='TEXT', MAXIMUM_LENGTH='16', SEARCH_VALUE_PATTERN='[0-9]{16}', SHOW_UNIQUE_VALUES='false' Action: Add to Group of Objects: Add to CC objects group Object Group='Sensitive Objects', Replace Group Content='false', Add Member Type='%NAMELIKE'	Sensitive Data	Credit Card	ora_orcl102_readonly : ORACLE : 10.10.9.14 : orcl102 :: 1521 :

Select All Unselect All **Adhoc Action**

Figure 5-73 Predefined process results

When the results are viewed, in addition to the action that is defined in the policy, you have the option of starting an ad hoc action; for example, send an Alert.

In addition to the process report, verify the results of the actions if any were specified. In this example, objects (tables) that are found to contain credit card numbers were to be added to the Sensitive Objects group.

Click **Tools** → **Config & Control** → **Group Builder**, then select the group and verify that the tables reported were added, as shown in Figure 5-74 on page 224.

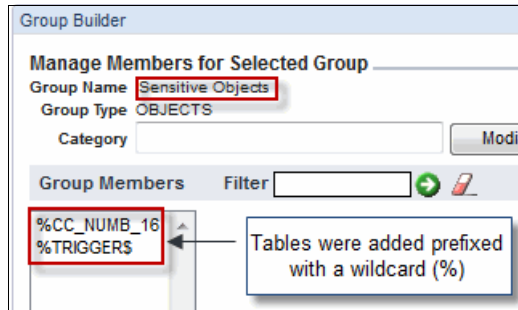


Figure 5-74 Tables added to the Sensitive Objects group

Custom reports

In addition to the predefined report, custom queries and reports can be created by using the Classifier Results Tracking domain and query builder.

5.13.6 Next steps

The following next steps are suggested:

- ▶ The example in this section shows one capability of the Classifier module. For more information about the other types of rules, rules configuration options, and types of actions, see the Classification topics in the *Discover* chapter of the online *Help Book Guardium*.
- ▶ Based on the classification process results, the following next steps are possible:
 - Tweak and test the policy search rules to minimize false positives and reduce processing time, if necessary.
 - Review and adjust the rule actions.
 - Add other rules or data sources, if needed.

5.14 Adding a menu tab to your portal

Users with the admin role only do not have a convenient place on their portal to add custom reports. To address this issue, the following options are available:

- ▶ Add the user role to the admin account
- ▶ Create a menu tab on the portal

A menu tab is a tab or page that is divided into two vertical sections: a menu section down the left side and a display section on the right. Examples of a menu tab are the Guardium Monitor tab or the My New Reports tab. A menu tab is a convenient layout for a list of reports.

For more information, see the “Customize the Portal” topic in the Common Tools chapter of the online *Help Book Guardium*.

To create a menu tab, complete the following steps:

1. In the Guardium application GUI, click **Customize** in the upper right corner of the window.
2. Click **Add Pane** and enter a name; for example, My Reports. Click **Apply**. (A *pane* is the basic building block of the portal.)

The new pane, which is named My Reports, should now be listed on your portal. The next step is to change the pane layout from the default to a “menu” layout

3. Click the link for the new pane. In the Layout field, select **Menu pane**. Click **Save** then click **Save** again to return to the main portal.

The new tab, My Reports, is displayed and ready to receive reports or other applications.

Figure 5-75 on page 226 shows the steps for adding a menu tab to a portal.




Figure 5-75 Steps for adding a menu tab to a portal



Access management

In this chapter, we describe the functions for creating and managing user access to the Guardium solution and data.

This chapter includes the following topics:

- ▶ Access management overview
- ▶ User accounts
- ▶ User role browser

6.1 Access management overview

Access management is the function of managing user access to the Guardium application. This function is separated from the system administration duties to allow for separation-of-duties.

Access management is performed by a user with the `accessmgr` role (privileges).

6.1.1 Roles, portals, applications, and users

A *role* is a set of privileges and application (function) access; for example, the user role has different privileges and access to functionality than the admin role.

Some (that is, not all) roles have a unique panel layout or *portal*; for example, the portal for the `accessmgr` role is different from that for the admin role, as shown in Figure 6-1 on page 229.

Note: On the first login, the user's portal is generated based on their assigned role and associated with the account. Users that are assigned multiple roles have a portal that combines the portals for each role (if there is a portal that is associated with each role) for example, user and admin.

An *application* is a high-level Guardium feature or function; for example, Audit Process Builder or Group Builder. Access to applications are controlled by updating which roles have access to that application.

Note: To avoid crippling access to the Guardium solution, do not remove applications from the admin or `accessmgr` roles.

A user (account) gets access to privileges and applications by being assigned to a role. A user can have multiple roles, with a few exceptions.

The following key predefined users (account) are available on each Guardium appliance:

- ▶ `admin`: Used for system administration by using the Guardium GUI
- ▶ `accessmgr`: Used to manage user accounts by using the Guardium GUI
- ▶ `cli`: Used for system administration (command-line interface only)

Note: The `admin`, `accessmgr`, and `cli` can refer to a user or a role, depending on the context.

Accountability

To provide accountability or traceability (that is, associate a change with a specific user), limit the use of the default admin and accessmgr accounts, and instead create named accounts and assign the appropriate roles.

Note: Do not disable these default accounts.

6.1.2 Managed configuration

In a managed configuration, access management information can be viewed on any managed unit, but changes can be done only on the Central Manager. Any changes are automatically synchronized with the managed units on an hourly basis by using the User Portal Sync process.

For more information, see the Access Management chapter of the *Help Book Guardium*.

Figure 6-1 shows the portals for the three key Guardium roles.




Figure 6-1 Portals for the three key roles

6.1.3 Authorization versus authentication

When a user attempts to log in to the Guardium application through the GUI, the application first checks if the user is authorized, that is, has an account. If the user is authorized, the second step is to authenticate the user, that is, verify the password that the user provided.

A user is authorized by creating a Guardium user account in the Guardium application manually or by LDAP import.

User authentication depends on the following Guardium configuration that is used by the administrator:

► Internal

Local: The user's password is stored (encrypted) in the Guardium database and is used for authentication. This is the default.

► External: The following mechanisms are available:

- Lightweight directory access protocol (LDAP), including Active Directory (AD): The user's password is stored in LDAP and not in Guardium. The account and password that is provided on the login window are sent to LDAP for validation.
- Remote Authentication Dial-In User Service (RADIUS): Similar to LDAP.

Note: External and internal authentications are mutually exclusive, except for the three default accounts: admin, accessmgr, and cli, which are locally authenticated. This allows access to the appliance if the LDAP or RADIUS servers are unreachable.

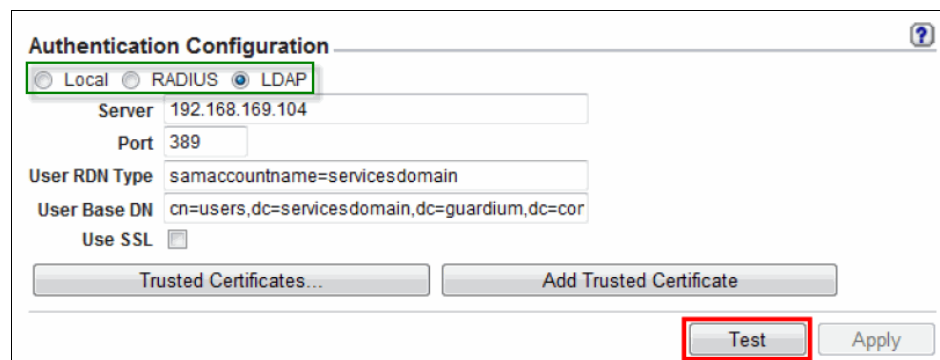
Configuring authentication

The authentication configuration is performed by the Guardium administrator by using the Administration Console tab and clicking **Configuration** → **Portal** → **Authentication Configuration**.

The external configuration is done on each appliance. In a managed environment, it can be distributed from the Central manager to all managed units.

For more information, see the “Portal Configuration” topic of the Guardium Administration chapter of the *Help Book Guardium*.

Figure 6-2 shows an example of the authentication configuration for LDAP.



The screenshot displays the 'Authentication Configuration' window. At the top, there are three radio buttons: 'Local', 'RADIUS', and 'LDAP'. The 'LDAP' radio button is selected and highlighted with a green box. Below the radio buttons, there are several input fields: 'Server' with the value '192.168.169.104', 'Port' with the value '389', 'User RDN Type' with the value 'samaccountname=servicesdomain', and 'User Base DN' with the value 'cn=users,dc=servicesdomain,dc=guardium,dc=cor'. There is also a 'Use SSL' checkbox which is unchecked. At the bottom of the window, there are two buttons: 'Trusted Certificates...' and 'Add Trusted Certificate'. In the bottom right corner, there are two more buttons: 'Test' and 'Apply'. The 'Test' button is highlighted with a red box.

Figure 6-2 Example of the authentication configuration for LDAP

6.1.4 Data-level security

Data-level security is used to control user access to data that is logged by Guardium. This control is achieved by mapping user accounts to the database servers from which the data was collected.

A typical use case allows users to view only data that is collected from databases for which they are responsible.

For more information, see the following resources:

- ▶ The “Data Security - User Hierarchy and Database Associations” topic in the Access Management chapter of the *Help Book Guardium*.
- ▶ The “How to define User Hierarchies” topic in the How-to Guide Overview chapter of the *Help Book Guardium*.

6.2 User accounts

Guardium accounts can be added manually or by importing from LDAP. Users are expected to have differing privileges that are based on their Guardium responsibilities.

6.2.1 Creating a user account manually

Complete the following steps to manually add a user account:

1. Log in as the `accessmgr` user to the GUI of the Central Manager or stand-alone appliance, browse to the Access Management tab and click **User Browser**.
2. Click **Add User** and enter the user name, password, first and last names, and email address, if any. Remember the following points:
 - Password complexity is enforced.
 - On first login, the user is prompted to change their password.
3. Clear the **Disabled** option and click **Add User** to save the account.

All accounts are defaulted to the user role.

Figure 6-3 shows the Add User functionality.




Figure 6-3 Adding a user

6.2.2 Importing user accounts from LDAP

Guardium provides an interface to LDAP, including AD, which allows the import and creation of user accounts in Guardium.

Tip: It is important to enlist the help of your LDAP administrator to provide the inputs and syntax for the import process.

For more information, see the “Import Users from LDAP” topic in the Access Management chapter of the *Help Book Guardium*.

Accounts do not have to be imported from LDAP to use LDAP authentication. For example, manually create the user account in Guardium (for authorization), matching their LDAP account but with a bogus password. When that user then attempts to log in, they provide their LDAP password (and not the bogus password) and are authenticated against LDAP.

In organizations with tens or thousands of employees, it is typical that only a handful of users directly use Guardium and therefore need Guardium accounts. In these cases, some customers choose to manually create the accounts but still use LDAP authentication.

6.2.3 Modifying a user's role

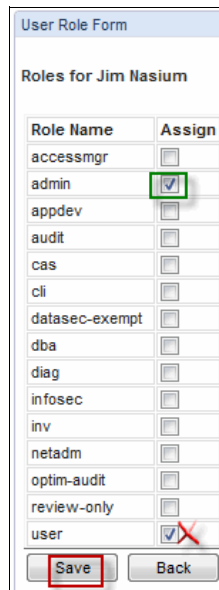
All new accounts default to the user role. To change an account's role (for example, from user to admin), complete the following steps as the accessmgr user:

1. Browse to the Access Management tab, click **User Browser**, and then click **Roles** next to the account that you want to modify.
2. In the User Role Form, select **admin** and clear the **user** option, then click **Save**, as shown in Figure 6-4. If the user logged in before the role change, they might have to reset their portal to reflect the changes, as shown in Figure 6-5 on page 234.

Remember following points about accounts:

- ▶ An account can have multiple roles; for example, user and admin.
- ▶ An account cannot have admin and accessmgr roles because of separation-of-duties.
- ▶ An account must have one of the following roles: user, admin, accessmgr, or cli.

Note: Because of separation-of-duties, the admin and accessmgr roles cannot be granted to the same user.



Role Name	Assign
accessmgr	<input type="checkbox"/>
admin	<input checked="" type="checkbox"/>
appdev	<input type="checkbox"/>
audit	<input type="checkbox"/>
cas	<input type="checkbox"/>
cli	<input type="checkbox"/>
datasec-exempt	<input type="checkbox"/>
dba	<input type="checkbox"/>
diag	<input type="checkbox"/>
infosec	<input type="checkbox"/>
inv	<input type="checkbox"/>
netadm	<input type="checkbox"/>
optim-audit	<input type="checkbox"/>
review-only	<input type="checkbox"/>
user	<input checked="" type="checkbox"/>

Figure 6-4 Changing a user's role from user to admin

Figure 6-5 shows the steps for a user to reset their portal.

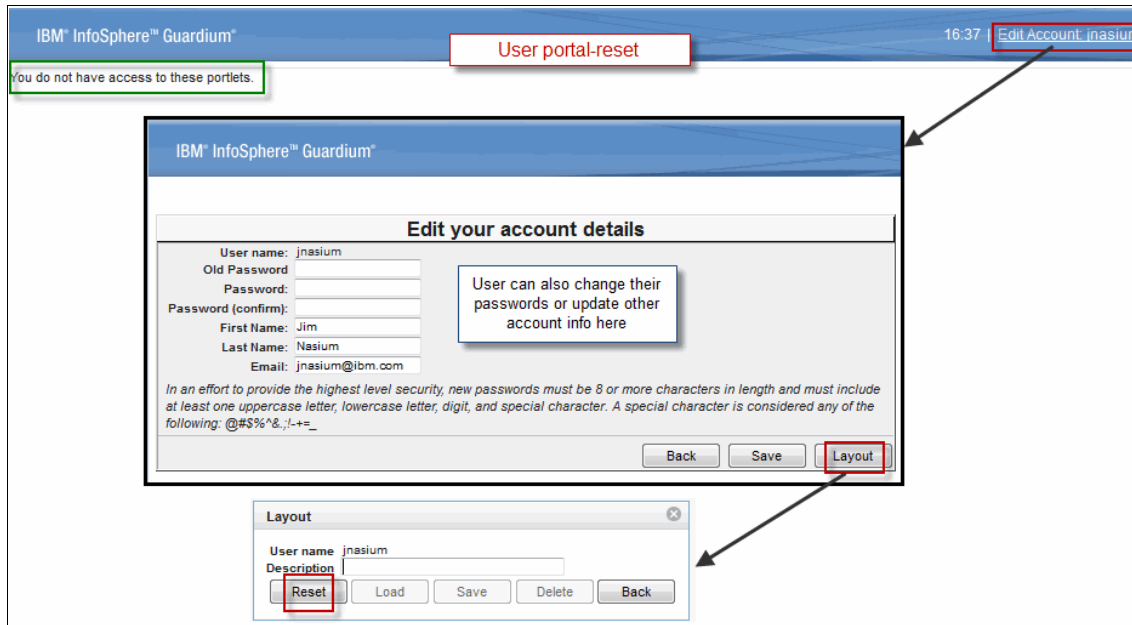


Figure 6-5 Steps for a user to reset their portal

6.3 User role browser

Guardium provides a set of default roles and sample roles. There are also add-on product roles that are used to control access to add on products; for example, sox for the Sarbanes-Oxley (SOX) accelerator.

Unlike the sample roles, the default roles cannot be deleted.

Figure 6-6 on page 235 shows the default and sample roles the use the User Role Browser.


Access Management 		Data Security																																				
<ul style="list-style-type: none"> User Browser User Role Browser User Role Permissions User LDAP Import User & Role Reports 		<table border="1"> <thead> <tr> <th colspan="2">Role Browser</th> </tr> <tr> <th>Role Name</th> <th>Actions</th> </tr> </thead> <tbody> <tr><td>accessmgr</td><td></td></tr> <tr><td>admin</td><td></td></tr> <tr><td>appdev</td><td>Delete</td></tr> <tr><td>audit</td><td>Delete</td></tr> <tr><td>cas</td><td>Delete</td></tr> <tr><td>cli</td><td></td></tr> <tr><td>datasec-exempt</td><td></td></tr> <tr><td>dba</td><td>Delete</td></tr> <tr><td>diag</td><td></td></tr> <tr><td>infosec</td><td>Delete</td></tr> <tr><td>inv</td><td></td></tr> <tr><td>netadm</td><td>Delete</td></tr> <tr><td>optim-audit</td><td>Delete</td></tr> <tr><td>review-only</td><td></td></tr> <tr><td>user</td><td></td></tr> <tr> <td colspan="2" style="text-align: center;">Add Role</td> </tr> </tbody> </table>	Role Browser		Role Name	Actions	accessmgr		admin		appdev	Delete	audit	Delete	cas	Delete	cli		datasec-exempt		dba	Delete	diag		infosec	Delete	inv		netadm	Delete	optim-audit	Delete	review-only		user		Add Role	
Role Browser																																						
Role Name	Actions																																					
accessmgr																																						
admin																																						
appdev	Delete																																					
audit	Delete																																					
cas	Delete																																					
cli																																						
datasec-exempt																																						
dba	Delete																																					
diag																																						
infosec	Delete																																					
inv																																						
netadm	Delete																																					
optim-audit	Delete																																					
review-only																																						
user																																						
Add Role																																						

Figure 6-6 Default and sample roles that use the User Role Browser

6.3.1 Custom role

Custom roles can be added to users (accounts) as with group users with the same operational or organizational role. This makes it easier to publish or share queries, reports, or workflow with users having the same role.

For example, there are four accounts, joe, bob, mary, and jane, who are all assigned the user role.

However, joe and mary are DBA managers; whereas bob and jane are DBAs.

Bob and jane can be added to the dba role so they can share reports with each other or anyone else with the dba role.

Joe and mary are added to the dba_mgmt role so they can receive reports that are distributed by the Guardium Audit process to the dba_mgmt role.

All four users have two roles: user and dba or dba_mgmt.

6.3.2 Adding a role

Complete the following steps to add a role:

1. Browse to the Access Management tab and click **User Role Browser**. Click the **Add Role** link at the bottom of the roles listing.
2. Enter the name of the new role and click **Add Role**, as shown in Figure 6-7.




Figure 6-7 Adding a custom role dba_mgmt'

To add a custom role (for example dba_mgmt) to an existing account, follow the steps that are described in 6.2.3, “Modifying a user’s role” on page 233, except that the user now has two roles that are selected: user and dba_mgmt.



Ongoing operations

In this chapter, we describe certain critical ongoing operations of Guardium system, outline optimization strategies, and provide maintenance suggestions, tools, and tips to keep your environment functioning smoothly and efficiently.

This chapter includes the following topics:

- ▶ Performance optimization and tuning
- ▶ Maintenance and updates
- ▶ Diagnostic tools
- ▶ Restoring audit data for forensic analysis

7.1 Performance optimization and tuning

In this section, we describe the built-in tools that are available to monitor and analyze the health of all InfoSphere Guardium components. The ability to use these tools is not only an essential part of ensuring that your Guardium environment is running smoothly, but can also be used when you are planning future expansion.

Configuring and interpreting these built-in tools means little if you do not know what corrective steps can be taken when issues are identified. Therefore, this section devotes considerable time describing the various strategies that are available to get things running smoothly.

7.1.1 S-TAP optimization and tuning

Beginning in Guardium V9, the primary tool for determining the overall health and performance of S-TAP is the S-TAP Statistics report. The S-TAP Statistics report is based on information that is collected from S-TAP and K-TAP on the host server and can be displayed as an interactive report or as part of an audit process. The information is collected at a configurable interval by the `guard_stap` process.

Note: As of v9, S-TAP Statistics functionality is only available on Linux and UNIX S-TAPs. Where indicated, however, some of the tuning parameters that are described in this section also apply to Windows S-TAPs.

Configuring S-TAP Statistics

The S-TAP Statistics functionality is not enabled by default. The initial configuration is done from the S-TAP configuration file (`guard_tap.ini`).

Complete the following steps to configure S-TAP Statistics reports:

1. Open the `guard_tap.ini` file on the server by using a text editor, such as Vi.
2. Browse the `stap_statistic` parameter.
3. Specify the polling interval.

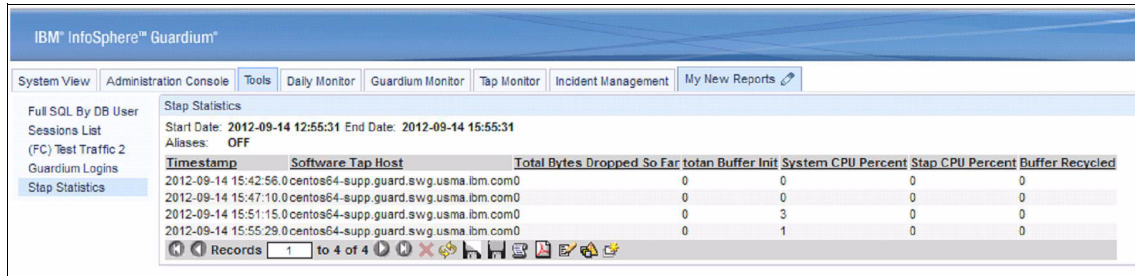
Values greater than 0 set the polling interval in hours. Values less than 0 set the polling interval in minutes. For example, setting `stap_statistic=-5` prompts `guard_stap` to collect performance information every 5 minutes.

For normal operations, set `stap_statistic=1`.

Use values less than 0 only when you are troubleshooting specific issues or when directed to do so by Guardium support.

Creating S-TAP Statistics report

There are no predefined reports for S-TAP statistics. However, there is a new S-TAP statistics domain for creating custom reports. (The process of creating custom reports is not described in this section.) Figure 7-1 shows an example report to get you started.



The screenshot shows the IBM InfoSphere Guardium interface. The main content area displays a 'Step Statistics' report for the period from 2012-09-14 12:55:31 to 2012-09-14 15:55:31. The report includes a table with the following data:

Timestamp	Software Tap Host	Total Bytes Dropped So Far	total Buffer Init	System CPU Percent	Stap CPU Percent	Buffer Recycled
2012-09-14 15:42:56.0	centos64-supp.guard.svg.usma.ibm.com0	0	0	0	0	0
2012-09-14 15:47:10.0	centos64-supp.guard.svg.usma.ibm.com0	0	0	0	0	0
2012-09-14 15:51:15.0	centos64-supp.guard.svg.usma.ibm.com0	0	3	0	0	0
2012-09-14 15:55:29.0	centos64-supp.guard.svg.usma.ibm.com0	0	1	0	0	0

Figure 7-1 Sample S-TAP Statistics report

Although there are more fields that can be added to an S-TAP Statistics report, the sample that is shown in Figure 7-1 includes the most commonly used parameters. The following fields are featured:

- ▶ **Timestamp:** Indicates when the record was created.
- ▶ **Software Tap Host:** The host system where data is collected.
- ▶ **Total Bytes Processed so Far:** The number of bytes that are logged by K-TAP.
- ▶ **Total Buffer Init:** The number of times the K-TAP buffer was reinitialized. Reinitializing the K-TAP buffer might be required if the contents become corrupted.
- ▶ **System CPU Percent:** Total CPU usage percentage on the host system for all processes.
- ▶ **S-TAP CPU Percent:** Total S-TAP CPU utilization on the host system. This value represents the overall S-TAP CPU utilization. For example, if the system has 10 cores, and S-TAP is using 30% of one, the overall S-TAP CPU usage is about 3%. The maximum CPU S-TAP can ever use on a server is 100% of one core because the guard_stap process is single threaded.
- ▶ **Buffer Recycled:** The number of times the S-TAP buffer overflowed.

In addition, there are other parameters that provide K-TAP information that also can be added to the S-TAP Statistics report. However, the data in these parameters is cumulative, so they might not accurately represent K-TAP performance for any specific point.

The following fields should be used when you are troubleshooting specific issues and only after the values are reset from the S-TAP side:

- ▶ **Total Bytes Dropped so Far:** Total number of bytes that are dropped by K-TAP. This value should be taken as a delta between two given points. If this number consistently grows, there might be insufficient resources on the host system for S-TAP to read data from the K-TAP buffer quickly enough.
- ▶ **Total Bytes Ignored:** Total number of bytes that are ignored by K-TAP as a result of any IGNORE STAP SESSION rules that might be implemented.
- ▶ **Total Response Bytes Ignored:** Total number of bytes that are ignored by K-TAP as a result of any IGNORE RESPONSES PER SESSION rules that might be implemented.

Interpreting the S-TAP Statistics report

To properly interpret the S-TAP Statistics report, it first helps to separate the fields in the report that are cumulative versus the fields that are real time. Cumulative means that the values in the fields are not reset until they are done manually by using a command on the host side, which is similar to the odometer in your car. Real-time fields are dynamic and are not required to be reset, which is similar to your vehicle's speedometer.

Real-time values

The following real-time value fields are available:

- ▶ **System CPU Percent**

The System CPU Percent field shows the CPU usage by all processes on the host database server. It is useful for showing how busy the host server is overall.
- ▶ **S-TAP CPU Percent**

The S-TAP CPU Percent field shows the overall CPU usage of S-TAP for the entire system. It is calculated by using the **pcpu** option from the **ps** command. S-TAP CPU usage might be indicating an issue in the following cases:

 - Usage is consistently at or near 100%. Such a condition might indicate that the **guard_stap** process is stuck in a loop and using all of the resources on one core. Run the **guard_diag** command when you encounter such cases.

- Overall usage is abnormally high. This number depends on the total number of cores that are running on the system. For example, consider a consistent S-TAP CPU usage of 5% on a system with 16 cores. In this case, 5% indicates that S-TAP is consistently using 80% of one core. If S-TAP is consistently running that high, it leaves little overhead to accommodate any other spikes in traffic. Worse, S-TAPs that are running close to 100% of one core might introduce performance degradation on the host server because it can make S-TAP unresponsive to K-TAP requests.

Cumulative values

The following cumulative value fields are available:

- ▶ **Total Bytes Processed so Far**

The Total Bytes Processed so Far value indicates the total number of bytes that were processed by K-TAP since the last reset of these values. This means that to come to any meaningful conclusions with this data, you must reset the values first. The values can be reset only directly from the database server by running the following command:

```
<S-TAP Shell Install
```

```
Directory>/guard_stap/ktap/current/guard_ktap_stat reset
```

or

```
<S-TAP GIM Install Directory>/modules/KTAP/current/guard_ktap_stat  
reset
```

Note: The reset command resets all of the K-TAP statistics.

Note: The value in Total Bytes Processes so Far rolls back over to 0 after reaching 4294967296 bytes (2^{32}). Therefore, if it was reset for some time, the value that is displayed might be a value that was rolled over several times.

On its own, there is little that can be learned from looking at only the total bytes processed value. Its delta over time can be used to estimate the volume of traffic that is processed by S-TAP if K-TAP is the only driver that is used to intercept traffic. For this purpose, it is not necessary to first reset the counter.

Total Bytes Processed is most helpful when it is used as baseline for some of the other statistics that are described next.

► Total Buffer Init

S-TAP can reinitialize the K-TAP buffer if any corruption of buffered data is detected. The number of times the K-TAP buffer is reinitialized is shown in the Total Buffer Init field.

► Buffer Recycled

Buffer Recycled is an important statistic. It indicates the number of times the S-TAP buffer overflowed, which is highly indicative of S-TAP performance issues on the host server. The S-TAP buffer can overflow for several reasons, including the following examples:

- Insufficient network bandwidth to accommodate the volume of data that is sent by the S-TAP to the Guardium appliance. This issue is most prevalent when the Guardium appliance and host database server are not in the same data center or LAN.
- Guardium collector is too busy to handle the volume that is sent from S-TAP (this is a rare case).

► Total Bytes Dropped so Far

The Total Bytes Dropped so Far value indicates the total number of bytes that were dropped by K-TAP. The value is cumulative, so any data drops that are shown in this field might not necessarily be recent. Before it is used in any analysis of S-TAP performance, the values should be reset by using the **guard_ktap_stat reset** command.

By default, K-TAP uses a 4 MB buffer file, which is configurable from the `guard_tap.ini`. If the `guard_stap` process cannot read data quickly enough from this buffer, K-TAP begins to drop data and the drops are reflected in this field. The significance of any drops that are shown here should be put in the context of Total Bytes Processed so Far. An excessive number of bytes dropped can be indicative of issues, including the following examples:

- Insufficient resources on host server for S-TAP (`guard_stap` process) to read data from the K-TAP buffer in a timely manner.
- Before Guardium V9, K-TAP does not support packet sizes larger than 64 KB. Data in packets that exceed this size (not common) is dropped and the value is reflected under Total Bytes Dropped so Far.

► Total Bytes Ignored

The Total Bytes Ignored value displays the amount of database traffic that is ignored at the K-TAP level when IGNORE STAP SESSION rules are implemented in the Guardium policy. It is useful for estimating how effectively the policy is ignoring traffic. As with other cumulative values, the value for Total Bytes Ignored should be considered only after a reset and in the context of Total Bytes Processed so Far.

- ▶ Total Response Bytes Ignored

The Total Response Bytes Ignored value displays the amount of database response traffic that is ignored at the K-TAP level as the result of any IGNORE RESPONSES PER SESSION rules that are implemented in the Guardium policy. This is a cumulative value and should be reset before it is used in any analysis.

Factors that affect S-TAP performance

There are many factors that directly and indirectly affect the performance of S-TAP on your database server. Many performance issues are caused by heavy traffic. In these cases, the best remedy is to ignore unimportant sessions through the Guardium Policy. In this section, the focus is specifically on S-TAP `guard_tap.ini` parameters that can have a negative effect on performance.

S-TAP options that have a negative effect on performance

The following S-TAP options might have a negative effect on performance:

- ▶ Transport Layer Security (TLS) encryption of S-TAP traffic by using the `use_tls` parameter Linux, UNIX, and Windows.

S-TAP traffic to the Guardium appliance can be encrypted by setting the `guard_tap.ini` parameter `use_tls=1`. The higher the volume of traffic that is monitored by S-TAP, the higher the processing power that is required to encrypt the data.

In general, it makes sense to encrypt only S-TAP traffic if the data that is sent to an appliance on a different network, or if the database traffic that is monitored is network encrypted.

- ▶ Compression of S-TAP traffic by using the `compression_level` parameter (Linux, UNIX, and Windows)

S-TAP traffic that is sent to the Guardium appliance can be compressed by setting the `compression_level` of the `guard_tap.ini` parameter 0 - 9. The default value is 0 and corresponds to no compression. A value of 1 corresponds to the lowest compression level and a value of 9 corresponds to the highest compression level. The higher the compression level that is requested, the more it affects S-TAP performance, which is observed as higher S-TAP CPU usage.

Compression is sometimes necessary to decrease the S-TAPs network bandwidth requirements, which is important on busy networks or when you are attempting to transmit S-TAP traffic to a Guardium appliance across a wide area network (WAN).

► Use of S-TAP firewall (S-GATE functionality) (Linux, UNIX, and Windows)

Enabling S-GATE firewall can have various effects on the performance of the S-TAP and the database server, depending on the exact configuration. The following different firewall options are available:

– Firewall in open mode

Firewall is enabled in open mode by setting the `guard_tap.ini` parameters as shown in the following examples:

- `firewall_installed = 1`
- `firewall_default_state=0`

The `firewall_installed` parameter enables firewall protection, which means that S-TAP requests a verdict from the Guardium appliance for each request packet that is monitored (a verdict is decision to allow or end the connection). Usually, this results in increased S-TAP CPU usage to a degree that is commensurate with the volume of traffic that is monitored. It can also result in slightly higher network bandwidth usage.

The `firewall_default_state=0` setting specifies that the firewall should operate in open mode, which means that while it is waiting for a verdict from the appliance, S-TAP does not hold up the database connections or traffic. Therefore, in open mode, users should not experience any latency when they are connecting to the database or running SQL statements.

The `firewall_default_state` parameter is the default state. Guardium users can still program “S-GATE Attach” rules in the policy to override this default and monitor specific sessions in closed mode.

– Firewall in closed mode

Firewall is enabled in closed mode by setting the `guard_tap.ini` parameters, as shown in the following examples:

- `firewall_installed=1`
- `firewall_default_state=1`

Setting `firewall_installed=1` enables S-TAP firewall protection and causes increased S-TAP CPU and network bandwidth usage.

In addition, setting `firewall_default_state=1` instructs the S-TAP firewall to operate in closed mode by default. This means that S-TAP holds database connections and traffic until it receives “allow” verdicts from the appliance.

While this configuration is the safest way to ensure that no unauthorized activities can take place, it also incurs the most significant performance penalty because of the following reasons:

- When you are connecting to the database, S-TAP holds sessions until it receives an allow verdict from the Guardium appliance. The speed at which this verdict is received depends on many factors, such as network speed, number of rules in the Guardium policy, and overall usage level of the collector.
- Individual SQL statements are also held by S-TAP until an allow verdict is received from the appliance. As such, they are subject to the same delays.

As with open mode, setting the `firewall_default_state` to closed mode defines how S-TAP should protect all new sessions. Individual sessions can still be monitored in open mode by using S-GATE Detach rules in the policy. Such sessions can still experience delays when they are connecting to the database while waiting to be detached, but performance is similar to open mode after the detach verdict is received.

New to Guardium V9 are the following `guard_tap.ini` firewall parameters that are used with open or closed mode:

- `firewall_force_watch`: This parameter is used with open mode (`firewall_default_state = 0`). It is a comma-separated list of IPs or networks (that use network masks) that specify which connections should be attached (watched) immediately. It is similar to attaching sessions from the policy without the inherent delay.
 - `firewall_force_unwatch`: This parameter is used with closed mode (`firewall_default_state = 1`). As with the force watch parameter, it is a comma-separated list of IPs or networks that are specifying which connections should be detached immediately. It is similar to detaching sessions from the policy, without the inherent delay. Users often employ this when they want to use closed mode for all connections except those that are originating from application servers, which have no tolerance for connection delays.
- Enabling UID Chain by using the `hunter_trace` parameter (Linux and UNIX only)

User ID (UID) chain is a mechanism that allows S-TAP to track the operating system user accounts in a switch user (su) chain before a database connection. It is enabled by configuring the `hunter_trace` parameter in the `guard_tap.ini` as shown in the following example:

```
hunter_trace=1
```

The other processing that is required to get UID chain information means users might observe increased S-TAP CPU usage when it is enabled.

- ▶ Running S-TAP in debug mode (Linux, UNIX, and Windows)

For certain S-TAP issues, you might be directed by technical support to run S-TAP diagnostics, which puts the S-TAP in debug mode. While it is running in debug mode, S-TAP uses more processor cycles and I/O resources on the database server. S-TAP debug can be run from S-TAP Control by clicking the send command icon (as shown in Figure 7-2), and then selecting **Run Diagnostics** from the drop-down menu.




Figure 7-2 Send command to S-TAP

S-TAP performance enhancing parameters

In the previous section, we focused on parameters that add functionality but can negatively affect S-TAP performance. In this section, we describe the following S-TAP parameters that help improve S-TAP performance:

- ▶ `ktap_fast_tcp_verdict` (Linux and UNIX only)

The `ktap_fast_tcp_verdict` parameter was first introduced in Guardium V8.01. The current default is `ktap_fast_tcp_verdict=0`.

To understand its role in performance, it is important to first understand that K-TAP does not have any prior knowledge of the Inspection Engines that are configured in the `guard_tap.ini`. Therefore, when a new TCP session is connected to the database, K-TAP must poll the `guard_stap` process to determine whether the session should be monitored, based on the TCP ports that are configured in the Inspection Engines.

In environments where there are hundreds or thousands of sessions per hour, this constant polling can translate into significant CPU time. In cases where the `guard_stap` process is busy (or becomes stuck), K-TAP can be kept waiting too long, and this can contribute database performance degradation and users might experience slow database connections.

For these reasons, we recommend setting `ktap_fast_tcp_verdict=1`. When set to 1, the TCP port information is loaded into K-TAP when S-TAP starts. The result is that K-TAP is no longer dependent on S-TAP to determine which TCP connections should be monitored, which reduces the likelihood of experiencing database performance degradation if S-TAP becomes slow or stuck.

When `ktap_fast_tcp_verdict=1` is set, the following limitation might result:

- `ktap_fast_tcp_verdict` disregards the IPs that were specified under “networks” or “Exclude IPs” in the Inspection Engines. All database traffic that matches the TCP ports that are configured are monitored and sent to the collector. Policy rules Ignoring Traffic still apply.
- As the name of the parameter implies, `ktap_fast_tcp_verdict` applies only to TCP connections. K-TAP still must poll S-TAP for other connection types, such as Oracle Bequeath, Shared Memory, and IPC.

▶ `ktap_request_timeout` (Linux and UNIX only)

The `ktap_request_timeout` parameter is used to configure how long KTAP waits for S-TAP to respond to a polling request. The current default is `ktap_request_timeout=5`, which should not be changed in most implementations.

The timeout does not apply to TCP connections when `ktap_fast_tcp_verdict=1`, but it applies to other connection types.

▶ `buffer_file_size` (Linux, UNIX, and Windows)

All S-TAP implementations use a buffer to accommodate temporary spikes in database traffic. Starting in V9, the default buffer size is decreased 100 - 50 MB, which is sufficient in most cases.

Users who are contemplating changing the default `buffer_file_size` should consider the following points:

- Environments with heavy database server traffic or many S-TAPs that are connecting to single collectors should use the default 50 MB or smaller buffer size. This might seem counter-intuitive because it decreases the amount of database traffic that S-TAP can continuously monitor if its connection with the collector is severed in any way. However, in busy database environments, there is little practical difference between a 50 MB and 100 MB buffer when the connection between S-TAP and collector is interrupted for a prolonged period.

In addition, increasing the buffer size to large values in busy environments can cause problems on the collector side. Upon reconnect with the collector, huge amounts of data can be flushed from the buffer, potentially overwhelming and creating drops on the collector side. The issue is exacerbated when there are multiple S-TAPs with large buffers reconnecting to the same appliance, as is the case during an unplanned inspection-core restart on the collector.

- There often is never a good reason for increasing the S-TAP buffer size more than 100 MB. Such increases often serve only as temporary fixes for bigger issues, such as insufficient network bandwidth or server resources for S-TAP to perform its work properly. It should be done only under the direction of Guardium support.
- ▶ `buffer_mmap_file` (Linux, UNIX, and Windows)

The `buffer_mmap_file` parameter governs how the S-TAP buffer is created. The following options are available:

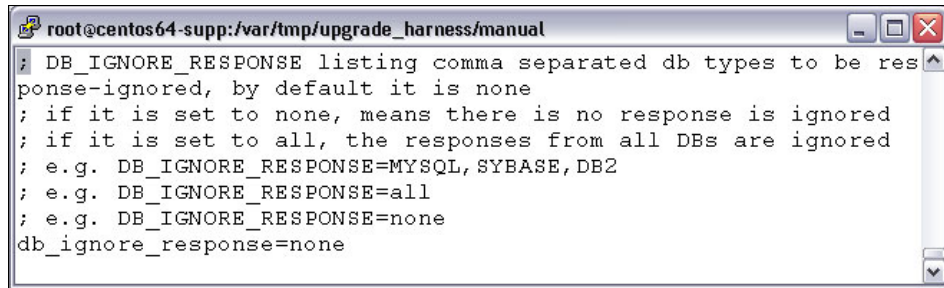
 - `buffer_mmap_file = 0` is the default configuration for new S-TAP installations starting from V8.2. This configuration offers the best performance because the S-TAP buffer exists only in memory. The drawback to this option is that the contents of the buffer are not retained if the S-TAP process is restarted.
 - `buffer_mmap_file = 1` was the default setting for V8.1 and prior versions. In this configuration, the buffer exists in memory and as a memory mapped file on disk. The operating system automatically handles the periodic synchronization of data between the buffer in memory and file on disk, which leads to an I/O performance penalty. The benefit to this configuration is that the buffer file is retained in the event of an S-TAP restart.
- ▶ `db_ignore_response` (Linux, UNIX, and Windows)

Ignoring responses is a good way to reduce the amount of traffic that S-TAP must intercept and can significantly improve its performance. Responses from the database include result sets, database exceptions (such as SQL errors), and failed login messages. If you do not need monitoring responses, you can increase performance by using one of the following methods:

 - The use of the `IGNORE RESPONSES PER SESSION` action from the policy is the recommended method because it allows the user to specify criteria, such as specific IPs, database users, and source applications for which to ignore responses. When ignoring responses by using policy rules, only failed logins are sent to the collector. Result sets from SQL statements are not sent, so be cautious if you plan to implement pattern matching by using extrusion rules.

Ignoring responses is effective for improving S-TAP performance in environments with long-running database sessions, but is less effective in environments with many short database sessions. This is because there is a small delay between the ignore rule firing and the S-TAP that is receiving the ignore instructions. If the database session ended before these operations occurred, the rule provides no benefit.

- The `db_ignore_responses` parameter was added to the S-TAP configuration in V9 and is the second option for ignoring responses. The default setting is `db_ignore_responses=none`. The configuration options are shown in Figure 7-3. Similar to IGNORE RESPONSES from the policy, this parameter instructs S-TAP to stop sending responses. Different from the policy action, the ignore takes effect immediately so it is much more effective for ignoring responses on short sessions. However, it does not provide the granular configuration options that are available in the policy rules and it ignores failed log ins. This parameter should be used only under the guidance of Guardium support.



```
root@centos64-supp:/var/tmp/upgrade_harness/manual
; DB_IGNORE_RESPONSE listing comma separated db types to be res
ponse-ignored, by default it is none
; if it is set to none, means there is no response is ignored
; if it is set to all, the responses from all DBs are ignored
; e.g. DB_IGNORE_RESPONSE=MYSQL, SYBASE, DB2
; e.g. DB_IGNORE_RESPONSE=all
; e.g. DB_IGNORE_RESPONSE=none
db_ignore_response=none
```

Figure 7-3 `db_ignore_responses` parameter

► `unix_domain_socket_marker`

The `unix_domain_socket_marker` parameter in the Inspection Engines configuration is used to configure the domain sockets (interprocess communication socket) for Oracle, MySQL, and PostgreSQL databases. In most case, the default setting of NULL properly collects all IPC traffic. However, in cases such as Oracle RAC environments, leaving the `unix_domain_socket_marker` setting to NULL forces S-TAP to monitor node-to-node traffic and impose an unnecessary performance penalty. In such environments, the `unix_domain_socket_marker` should be set to the KEY value of the IPC that is defined in the `tnsnames.ora` file.

► `intercept_types` parameter

The `intercept_types` parameter in the Inspection Engines is used to define the protocols that should be intercepted by S-TAP. The default setting is NULL, which captures all supported protocols. In some cases, this parameter is useful for determining whether performance issues are caused by the volume of specific traffic that uses a certain protocol. Figure 7-4 on page 250 shows the various configuration options for the `intercept_types` parameter.

Oracle	TCP, Named Pipe, PIPE
Informix	TCP, Shared Memory, TLI(SOLARIS only)
DB2	TCP, Shared Memory
MySQL	TCP, Named Pipe
Sybase	TCP, TLI(SOLARIS only)
PostgreSQL	TCP, Named Pipe

Figure 7-4 *intercept_types parameters*

► `participate_in_load_balancing`

The `participate_in_load_balancing` parameter allows you to balance the traffic that is intercepted by S-TAP across two or more appliances. Although this is less important for S-TAP performance, it can help resolve performance problems on the collectors by splitting the load.

Native S-TAP load balancing splits the traffic by database session, which sends each new session to a different appliance in the pool. Guardium also supports other load balancing methods that employ Cisco Global Site Selector (GSS) or an F5 Load Balancer. These methods are used to help distribute many S-TAPs across available collector resources, which is different from the session-based load balancing that is native to S-TAP.

7.1.2 Inspection Core performance and unit usage

The Inspection Core (sniffer) is the heart of the Guardium Collector. It receives all data that is sent from S-TAPs, Network TAPs, and SPAN ports. It is composed of the following components that perform various tasks of transforming network packets into data that can be stored in the internal MySQL database of the collector:

- **Sniffer Engine:** The function of this component is to reassemble packets that are coming from a SPAN port or Network TAP, or from S-TAPs that are using the packet capture (pcap) driver to capture data. It is not used to process data that is captured by the native S-TAP drivers, such as `lhmon` (in the case of Windows S-TAP) or `KTAP` (in the case of UNIX S-TAP).
- **Analyzer/Parser:** The analyzer determines the database type, protocol, and packet structure that is used for each monitored session. It then passes this information to the Parser, which, as the name implies, parses the SQL statements into their constituent parts (VERB, OBJECT, FIELD, and so on).

- **Logger:** The parsed data is then passed to the logger, which stores this data into the collector's database.

Each of these Inspection Core components feature dedicated buffers to cope with temporary spikes in traffic. When these buffers overflow, data loss occurs. When the appliance loses packets, you might notice data missing from Guardium reports or reports with missing fields (such as Database Username).

Therefore, managing the performance of the Inspection Core comes down to doing what is necessary to keep the various buffers from overflowing. The most efficient way to do this varies with the size of your Guardium environments.

In small environments of only a few collectors, you can monitor the Inspection Core performance by using the Buffer Usage Monitor report, as shown in Figure 7-5. In this chapter, we describe the most important parameters in this report.

Buff Usage Monitor																		
Start Date: 2013-08-17 15:33:03 End Date: 2013-08-18 15:33:03																		
Aliases: OFF																		
Timestamp	% CPU Sniffer	% Mem Sniffer	% CPU Mysql	% Mem Mysql	Mem Sniffer	Time Sniffer	Free Buffer Space	Analyzer Rate	Logger Rate	Analyzer Queue Length	Analyzer Total	Logger Queue Length	Logger Total	Session Queue Length	Session Total	Handler Data	Extra Info	ALP
2013-08-18 15:32:41.0	71	13	10	16	10133280818:153240100		825	690	0	308470	0	257125	3	78	52253	3	12	0
2013-08-18 15:31:40.0	53	13	8	16	10010360818:153138100		763	636	0	257311	0	214302	2	77	43541	2	10	0
2013-08-18 15:30:38.0	43	13	11	15	10010360818:153037100		691	576	0	210737	0	175494	2	77	35646	2	10	0
2013-08-18 15:29:36.0	59	12	6	14	10000120818:152935100		705	588	0	167891	1	139771	2	77	28390	2	7	0
2013-08-18 15:28:34.0	81	12	5	14	998988 0818:152833100		679	566	0	124141	1	103260	2	77	21081	2	3	0
2013-08-18 15:27:33.0	68	12	9	13	998988 0818:152731100		757	633	0	82031	0	68151	2	77	13797	2	6	0
2013-08-18 15:26:31.0	61	11	9	13	995916 0818:152630100		401	335	0	35845	1	29532	2	77	25952	2	5	0
2013-08-18 15:25:29.0	30	11	6	13	984648 0818:152528100		166	138	0	10938	0	8726	1	76	11171	1	2	0
2013-08-18 15:24:28.0	0	11	0	13	971332 0818:152427100		0	0	0	801	0	280	1	76	0	0	2	0

Figure 7-5 Buffer Usage Monitor Report

For medium to large Guardium environments, the best way to monitor Inspection Core performance across the environment is by using the Operational Dashboard functionality that was introduced in V9. The Operational Dashboard is accessible only from the Central Manager, as shown in Figure 7-6 on page 252.

Unit Utilization	
Start Date: 2012-09-01 15:44:35 End Date: 2012-09-25 15:44:35	
Using Merge Period Between 2012-09-11 and 2012-09-25.	
Aliases: ON	
Host Name	Max Overall Unit Utilization Level
blade01.guard.swg.usma.ibm.com	Low
blade02.guard.swg.usma.ibm.com	Low
blade03.guard.swg.usma.ibm.com	Low
blade04.guard.swg.usma.ibm.com	Low

Records 1 to 4 of 4

- Record Details
- Unit Utilization details
- Alias Definition

Figure 7-6 Operational Dashboard

The Operational Dashboard provides simple Low, Medium, and High utilization levels for each of the collectors in the environment. It calculates the utilization level that is based on several parameters in the Buffer Usage Monitor report data that is downloaded from the collectors. The Operational Dashboard is meant to provide a quick indication of potential performance issues across the estate, while the Buffer Usage Monitor report is still used to identify specific issues.

Small Guardium environments: Buffer Usage Monitor report

For environments with three or fewer collectors, the Buffer Usage Monitor report is your primary source of information regarding Inspection Core performance (although the report is not limited to this information only). This report is automatically updated by an internal script every minute, so the information that is contained is the most recent. By default, the appliance stores two weeks of data.

The Buffer Usage Monitor report consists of 47 or more columns, most of which you might never use in day-to-day monitoring of the appliance. One of the first things you should do is to create a simpler version of this report that contains the following columns:

- ▶ **Timestamp:** Shows when the data was collected.
- ▶ **% CPU Sniffer:** Shows a normalized representation of sniffer CPU usage. For example, 50% sniffer usage on an 8-core appliance means that the sniffer is using 400% CPU (4 cores).
- ▶ **% CPU Mysql:** Shows a normalized representation of MySQL CPU usage.
- ▶ **% Memory Mysql:** Shows the percentage of total system memory that is used by the MySQL database.
- ▶ **Free Buffer Space:** The percentage of free sniffer engine buffer space. The sniffer buffer engine is only used in implementations that use SPAN ports, Network TAPs, or S-TAP pcap. If the native S-TAP drivers are used, this value should always remain at 100%.

- ▶ **Mem Sniffer:** Shows sniffer memory usage in kilobytes (kB). This is an important parameter because on 32-bit appliances, the sniffer runs out of memory it can allocate at around 2.5 GB and restarts. Any data that is stored in any of the sniffer buffers during a restart is lost.
- ▶ **TID:** This is the sniffer process ID. Changes in the sniffer process ID indicate that the sniffer restarted, which results in the loss of any buffered data.
- ▶ **Analyzer Rate:** Provides a rough representation of the amount of data that is processed by the Analyzer/Parser per minute. The unit of data that is represented here is an internal structure that is closely analogous to a packet. The maximum analyzer rate a specific appliance can handle is a function of several variables, such as the appliance hardware, the type of data that is analyzed and parsed, and the type of rules that are used in the policy. Therefore, analyzer rate alone is not a good indicator of sniffer load, but it can be a good way to identify the busiest times of the day.
- ▶ **Analyzer Queue Length:** Indicates the amount of data that is in the Analyzer/Parser buffer. This value is one of the most direct indicators of sniffer performance. Ideally, the value here should remain at or close to zero. The analyzer queue might grow temporarily during temporary periods of high traffic, but should never remain elevated for more than five or six rows (5 - 6 six minutes) in the Buffer Usage Monitor report. The Analyzer/Parser Buffer is circular; therefore, after it fills, it begins to drop data as indicated in the Analyzer Lost Packets column.
- ▶ **Analyzer Lost Packets (ALP):** Data loss due to Analyzer Buffer overflow is shown here. The value in this field is cumulative, so it might continue to display lost data even after any issues are resolved. The value is reset after the sniffer is restarted.
- ▶ **Logger Rate:** Provides a rough representation of the amount of data that is processed by the logger per minute. The units here represent the parsed components of the SQL traffic that is inserted into the appliance's internal MySQL database. As with analyzer rate, the logger rate an appliance can handle depends on many factors, such as the appliance hardware, size of SQL statements that are logged, type of policy, and overall load on MySQL imposed by reports, and alerts.
- ▶ **Logger Queue Length:** Shows the amount of SQL data that is in the logger buffer and waiting to be inserted into the collector's database. Similar to the analyzer queue, consistently high amounts of data in the logger queue indicates that the appliance is unable to cope with the amount of traffic that is monitored. Temporarily spikes in buffered data are normal, as long as the buffer is flushed within several minutes.

Different from the Analyzer buffer, the Logger buffer is not circular. It grows until the sniffer cannot allocate more memory (2.5 GB maximum) and then the sniffer restarts. As data in the logger queue increases, so too does the overall sniffer memory usage. Any data in the logger buffer is lost during the sniffer restart.

- ▶ Sessions Queue Length: The total number of open sessions that are monitored by the sniffer. This information is important because the sniffer must allocate a certain amount of memory for each session that is monitored, and it cannot monitor more than 4000 simultaneous session at any time.
- ▶ Session Total: The overall number of sessions that are monitored since the last sniffer restart.

Other fields to include

Although the following fields do not relate to Inspection Core performance, they should be included in your simplified Buffer Usage Monitor report.

- ▶ Mysql Disk Usage: Displays the percentage MySQL disk space that is used.
- ▶ System CPU Load: Shows a normalized representation of total system CPU usage.

To create the simplified Buffer Usage Monitor report, use the Sniffer Buffer Usage Tracking domain. An example of the simplified report is shown in Figure 7-7.

Buff Usage Monitor																			
Start Date: 2013-08-17 15:33:03 End Date: 2013-08-18 15:33:03																			
Aliases: OFF																			
Timestamp	% CPU Sniffer	% Mem Sniffer	% CPU Mysql	% Mem Mysql	Mem Sniffer	Time Sniffer	Free Buffer Space	Analyzer Rate	Logger Rate	Analyzer Queue Length	Analyzer Total	Logger Queue Length	Logger Total	Session Queue Length	Session Total	Handler Data	Extra Info	ALP	
2013-08-18 15:32:41.0	71	13	10	16	10133280818:153240100			825	690	0	308470	0	257125	3	78	3 3 52253 12			0
2013-08-18 15:31:40.0	53	13	8	16	10010360818:153138100			763	636	0	257311	0	214302	2	77	2 2 43541 10			0
2013-08-18 15:30:38.0	43	13	11	15	10010360818:153037100			691	576	0	210737	0	175494	2	77	2 2 35646 10			0
2013-08-18 15:29:36.0	59	12	6	14	10000120818:152935100			705	588	0	167891	1	139771	2	77	2 2 28390 7			0
2013-08-18 15:28:34.0	81	12	5	14	998988 0818:152833100			679	566	0	124141	1	103260	2	77	2 2 21081 3			0
2013-08-18 15:27:33.0	68	12	9	13	998988 0818:152731100			757	633	0	82031	0	68151	2	77	2 2 13797 6			0
2013-08-18 15:26:31.0	61	11	9	13	995916 0818:152630100			401	335	0	35845	1	29532	2	77	2 2 5952 5			0
2013-08-18 15:25:29.0	30	11	6	13	984648 0818:152528100			166	138	0	10938	0	8726	1	76	1 1 1714 2			0
2013-08-18 15:24:28.0	0	11	0	13	971332 0818:152427100			0	0	0	801	0	280	1	76	0 0 2 0			0

Figure 7-7 Simplified Buffer Usage Monitor report

The simplified report that is shown in Figure 7-7 shows a lightly used collector. The Analyzer Queue values remain at or close to zero, which indicates that the Analyzer/Parser is keeping up with the traffic and there are no Analyzer Lost Packets. In addition, the Logger Queue is remaining at or close to zero, which indicates that it is keeping up as well. In the following section, we describe some common scenarios depicting various performance problems.

Performance Issue 1: Analyzer queue overflow

Figure 7-8 shows a scenario in which a sudden increase in the analyzer rate at 8:43 a.m. begins to overwhelm the sniffer. Even before this increase, the Analyzer/Parser appeared to be buffering some data, indicating that this appliance was already operating near its maximum performance.

During this spike in traffic, the Analyzer must start buffering large amounts of data, as shown by the increasing values in the Analyzer Queue Length. At approximately 8:48 a.m., the Analyzer/Parser buffers are full, and the sniffer begins to drop data, as shown in the Analyzer Lost Packets column. This kind of performance issue does not appear to be an isolated incident on this machine, as indicated by the large number of lost packets that existed before this event.

Timestamp	% CPU Sniff	% CPU Mysql	Mem Sniff	Sniffer Process ID	Analyzer Ra	Analyzer Queue Leng	ALP
8/12/13 8:59 AM	4	0	1292256	22868	13263	100851	22888011
8/12/13 8:58 AM	4	0	1292256	22868	13226	123175	22888011
8/12/13 8:57 AM	4	0	1292256	22868	13096	149336	22888011
8/12/13 8:56 AM	4	0	1292256	22868	12864	164330	22888011
8/12/13 8:55 AM	4	0	1292256	22868	13476	167818	22881334
8/12/13 8:54 AM	4	0	1292256	22868	13787	171253	22842433
8/12/13 8:53 AM	4	0	1292256	22868	13736	172766	22782747
8/12/13 8:52 AM	7	0	1292256	22868	14485	172526	22747439
8/12/13 8:51 AM	6	0	1292256	22868	14368	170622	22717531
8/12/13 8:50 AM	8	0	1292256	22868	14374	170330	22687476
8/12/13 8:49 AM	8	0	1292256	22868	13994	171600	22654768
8/12/13 8:48 AM	7	0	1292256	22868	13894	170095	22639334
8/12/13 8:47 AM	8	0	1292256	22868	13884	137133	22638400
8/12/13 8:46 AM	5	0	1292256	22868	13792	109393	22638400
8/12/13 8:45 AM	7	0	1292256	22868	13720	77099	22638400
8/12/13 8:44 AM	7	0	1292256	22868	13812	62526	22638400
8/12/13 8:43 AM	6	0	1292256	22868	13802	10417	22638400
8/12/13 8:42 AM	6	0	1292256	22868	4189	379	22638400
8/12/13 8:41 AM	5	0	1292256	22868	4094	93	22638400
8/12/13 8:40 AM	8	0	1292256	22868	4174	131	22638400
8/12/13 8:39 AM	7	0	1292256	22868	4173	89	22638400
8/12/13 8:38 AM	6	0	1292256	22868	4024	33	22638400
8/12/13 8:37 AM	6	0	1292256	22868	4213	48	22638400
8/12/13 8:36 AM	8	0	1292256	22868	3667	223	22638400
8/12/13 8:35 AM	6	0	1292256	22868	4204	1946	22638400
8/12/13 8:34 AM	1	0	1292256	22868	4864	0	22638400

Figure 7-8 Analyzer queue overflow

There are several reasons for issues with the Analyzer Queue overflowing, but the most common reason is that the sniffer cannot cope with the high rate of traffic that is monitored. In these cases, you must reduce the amount of traffic that is monitored by the appliance by using one of the following strategies:

- ▶ Moving some of the S-TAPs to less busy appliances.
- ▶ Introducing rules to filter more traffic. The most effective rule action to achieve filtering is the IGNORE STAP SESSION rule because the sessions are ignored by the S-TAP instead of being sent across the network to the appliance.
- ▶ S-TAP load balancing. Sometimes, a busy database server alone can overwhelm a collector. In these cases, it might help to load balance the traffic from this database to two or more collectors (for more information, see the `participate_in_load_balancing` parameter in “S-TAP options that have a negative effect on performance” on page 243).
- ▶ Consider using a Selective Audit policy. By default, the collector logs all data that is sent to it from S-TAPs or Hardware TAPs. A Selective Audit policy changes this behavior by monitoring only the database traffic that is specified in the policy rules.

Performance Issue 2: Logger queue overflow

Figure 7-9 shows a sudden increase in the logger queue. The analyzer queue is also high but recovers after two minutes. The logger queue is different from the analyzer queue in that it is not circular and continues to allocate memory until the sniffer reaches the 2.5 GB limit. In this report, observe that as the logger queue starts to grow at 8 p.m., the memory that is used by the sniffer (which is shown in the memory column) begins to increase by a proportional amount. When the sniffer reaches the 2.5 GB limit, it restarts.

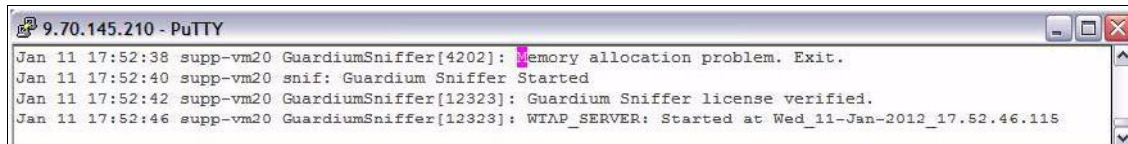
A	B	C	D	E	F	G	H	I
TIMESTAMP	% CPU Snif	% CPU MYSQL	PID	MEMORY	ANALYZE_RATE	LOG_RATE	ANALYZER_QUEUE	LOGGER_QUEUE
11/29/2012 20:03	78	0	11507	934176	2476	741	0	0
11/29/2012 20:02	78	8	9163	2573244	2476	741	649	106967
11/29/2012 20:01	60	2	9163	1881788	3852	508	59568	68615
11/29/2012 20:00	63	2	9163	1075132	239	54	17429	4999
11/29/2012 19:59	67	0	9163	934176	8406	399	0	0

Figure 7-9 *Logger queue overflow*

After the sniffer allocates memory, it does not release it even if the logger queue recovers. Therefore, it is possible to have a high sniffer memory usage even if the logger queues are not holding any data.

Sniffer restarts because of logger queue overflow is also shown in the collector’s syslog file (`/var/log/messages`). These messages come in two varieties. The first

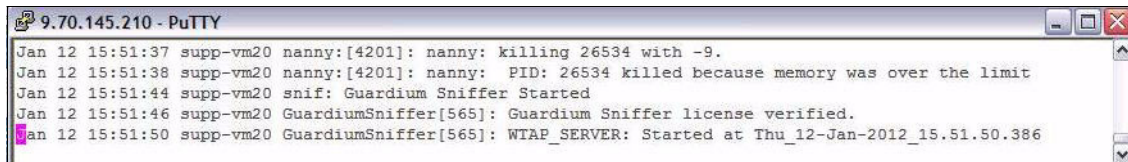
is a sniffer Memory Allocation Problem, which happens when the logger queues grow quickly, as shown in Figure 7-10.



```
9.70.145.210 - PuTTY
Jan 11 17:52:38 supp-vm20 GuardiumSniffer[4202]: Memory allocation problem. Exit.
Jan 11 17:52:40 supp-vm20 snif: Guardium Sniffer Started
Jan 11 17:52:42 supp-vm20 GuardiumSniffer[12323]: Guardium Sniffer license verified.
Jan 11 17:52:46 supp-vm20 GuardiumSniffer[12323]: WIAP_SERVER: Started at Wed_11-Jan-2012_17.52.46.115
```

Figure 7-10 Sniffer memory allocation problem

The second type of restart because of logger queue overflow happens when the Guardium “nanny” process, which monitors sniffer memory usage, detects that the sniffer is dangerously close to the 2.5 GB limit and restarts it, as shown in Figure 7-11.



```
9.70.145.210 - PuTTY
Jan 12 15:51:37 supp-vm20 nanny:[4201]: nanny: killing 26534 with -9.
Jan 12 15:51:38 supp-vm20 nanny:[4201]: nanny: PID: 26534 killed because memory was over the limit
Jan 12 15:51:44 supp-vm20 snif: Guardium Sniffer Started
Jan 12 15:51:46 supp-vm20 GuardiumSniffer[565]: Guardium Sniffer license verified.
Jan 12 15:51:50 supp-vm20 GuardiumSniffer[565]: WIAP_SERVER: Started at Thu_12-Jan-2012_15.51.50.386
```

Figure 7-11 Nanny process killing the sniffer

Usually, both types of restarts are caused by the same issues, the only difference being the speed at which the sniffer memory grows. Memory allocation problems happen when the sniffer memory grows quickly before the nanny process can react.

The logger queue can grow for the following reasons:

- ▶ Too much traffic or an overly aggressive policy with many heavy rules, such as Log Full Details. Though the solutions for Analyzer Queue issues can also apply here, most times it might be sufficient to reduce the number of Log Full Details or policy violation rules in the policy, or make such rules less inclusive.
- ▶ The logger might be competing for MySQL resources if there are an excessive number of reports, correlation alerts, or other internal processes that are running in the background. If your environment includes an Aggregator, consider running daily reports on that appliance instead.

The 2.5 GB maximum on sniffer memory allocation applies only to appliances running 32-bit Linux. This limitation no longer exists on 64-bit appliances, where the sniffer can allocate up to 33% of the physical memory that is installed.

Medium and large Guardium environments: Operational Dashboard

For environments larger than three collectors, it becomes impractical to monitor Inspection Core performance by individually analyzing the Buffer Usage Monitor report on each collector. Enter the Operational Dashboard.

The Operational Dashboard was introduced into the V9 Central Manager as a means of providing an enterprise-level view of collector usage. It employs a simple Low, Medium, or High indicator system to quickly show which collectors are over or under used. The analysis that is performed by the Operational Dashboard is based on Buffer Usage Monitor Data that is downloaded from the managed collectors. By using this data, it calculates the following statistics in 1-hour increments over a 24-hour period:

- ▶ Number of Sniffer restarts
- ▶ Sniffer Memory Utilization
- ▶ Percent MySQL Memory
- ▶ Free Buffer Space
- ▶ Analyzer Queue: Highest value of the Analyzer Queue for each hour in the test period
- ▶ Logger Queue: Highest value of the Logger Queue for each hour in the test period
- ▶ MySQL Disk Usage: Current MySQL disk usage
- ▶ System CPU Load: Highest System CPU load for each hour in the test period
- ▶ System Var Disk Usage: The utilization of the /var partition, which is where most of files that are generated by the appliance are stored
- ▶ Number of Requests: Number of SQL requests that are observed for each hour in the test period

The overall utilization level for any collector is equivalent to the highest level of the parameters measured. For example, if the number of Sniffer restarts tests as High while all other parameters test Low, the overall unit utilization is marked as High for the report.

Configuring the Operational Dashboard

Complete the following steps to configure the Operational Dashboard:

1. Upload the Buffer Usage Monitor Data from the managed collectors to the Central Manager.

This step is done from the Custom Table Builder under the Tools menu. Select **CM Buffer Usage Monitor** and then click **Upload Data**, as shown in Figure 7-12.




Figure 7-12 Custom table builder

By using the Upload Data menu, you can schedule the upload of Buffer Usage Monitor from the managed units and run the process immediately. When the operation completes, it displays the number of rows that are uploaded from each managed unit, as shown in Figure 7-13 on page 259.

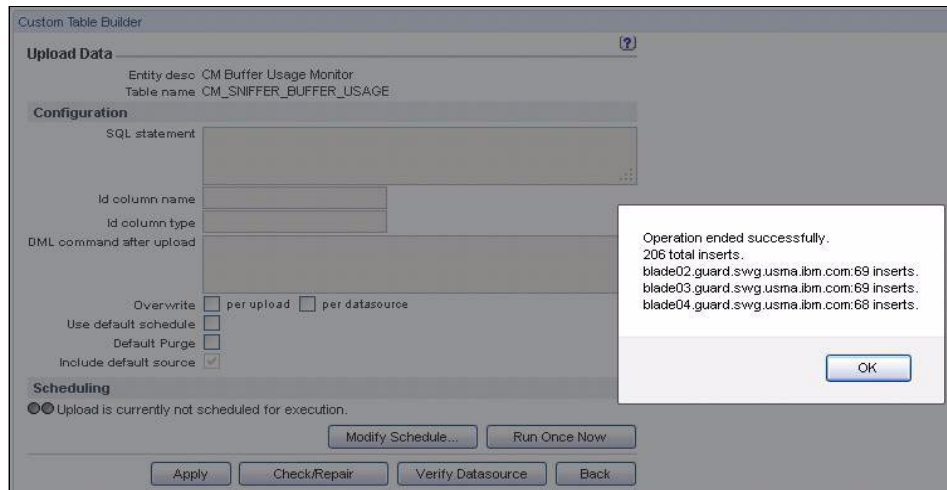


Figure 7-13 Data upload scheduler

- Click **Administration Console** → **Configuration** → **Unit Utilization Levels**. You might schedule this process or Run Once Now. The Unit Utilization Levels process takes the raw Buffer Usage Monitor data uploaded from the Managed Units and calculate hourly statistics for the parameters. For the first run of Unit Utilization Levels, the process calculates hourly statistics for the last 24 hours of data. This can be reconfigured later on, as shown in Figure 7-14.

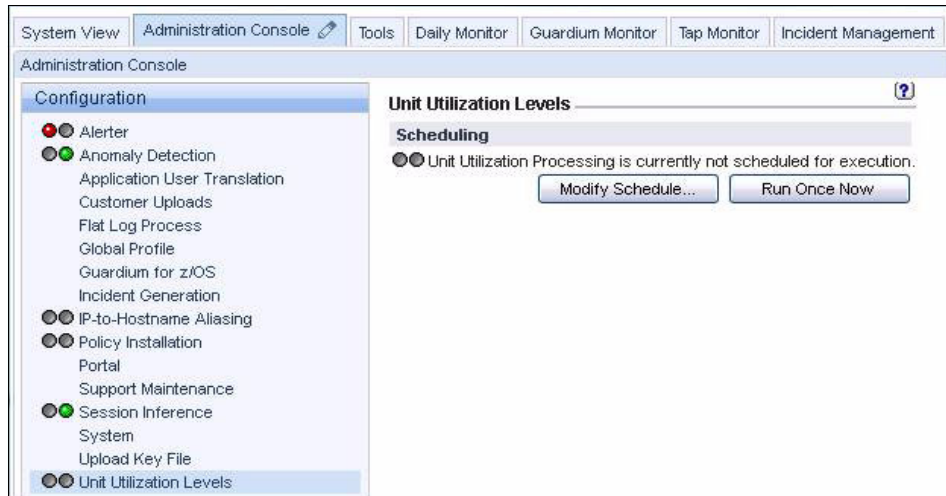


Figure 7-14 Unit Utilization Levels

Working with the Operational Dashboard

Complete the following steps to view the results of the Unit Utilization analysis:

- Click **Guardium Monitor** → **Units Utilization**. Overall collector utilization is displayed in the Unit Utilization report, as shown in Figure 7-15.



Figure 7-15 Unit Utilization report

Each managed collector is displayed with its corresponding utilization level. There are three levels: Low (green), Medium (yellow), and High (red). Double-click any managed unit and select the **Unit Utilization Details** report for an hourly breakdown of performance statistics for that unit.

- The Unit Utilization Details report provides a breakdown of the individual parameters that are used to calculate the overall collector utilization. The overall level is equal to the highest level of any of the individual parameters. For example, if the number of sniffer restarts reaches a level of Medium (but all other parameters are Low), the Overall Utilization level for that period is increased to Medium. Figure 7-16 on page 261 shows a collector with Low overall utilization.

IBM® InfoSphere™ Guardium®

Report **Unit Utilization details**
 hostName **blade02.guard.swg.usma.ibm.com**

Period Start	Overall Unit Utilization Level	Number Of restarts	Number Of restarts Level	Sniffer Memory	Sniffer Memory Level	Percent Mysql Memory	Percent Mysql Memory Level
2012-09-24 16:00:00.0	Low	0	Low	0	Low	7	Low
2012-09-24 17:00:00.0	Low	0	Low	0	Low	7	Low
2012-09-24 18:00:00.0	Low	0	Low	0	Low	7	Low
2012-09-24 19:00:00.0	Low	0	Low	0	Low	7	Low
2012-09-24 20:00:00.0	Low	0	Low	0	Low	7	Low
2012-09-24 21:00:00.0	Low	0	Low	0	Low	7	Low
2012-09-24 22:00:00.0	Low	0	Low	0	Low	7	Low
2012-09-24 23:00:00.0	Low	0	Low	0	Low	7	Low

Figure 7-16 Unit Utilization details

- The Low, Medium, and High utilization levels are established against a predefined but configurable set of thresholds, as shown in Figure 7-17 on page 262. The predefined thresholds should be adequate in most cases. Nevertheless, you might fine-tune the predefined parameters by double-clicking a particular threshold and starting the `update_utilization_thresholds` menu. Threshold 1 defines the level at which a particular parameter goes from Low to Medium. Threshold 2 defines the level at which a particular parameter is at High utilization.

Utilization Thresholds		
Aliases: OFF		
Parameter Name	Threshold 1	Threshold 2
Number Of Restarts	2	4
Sniffer Memory	2200000	2400000
Percent Mysql Memory	35	50
Free Buffer Space	75	90
Analyzer Queue	30000	50000
Logger Queue	30000	50000
Mysql Disk Usage	40	65
System CPU Load	45	85
System Var Disk Usage	40	70
Number of Requests	1000000	5000000
Number of Full SQL	500000	2500000
Number of Exceptions	1000	5000
Number of Policy Violations	1000	5000

Figure 7-17 Utilization Thresholds

- To clear the Unit Utilization data, open the Unit Utilization Distribution window as shown in Figure 7-18. Double-click the start icon in the report, and select **reset_unit_utilization_data** from the menu.

Unit Utilization Distribution			
Start Date: 2012-09-24 16:08:55 End Date: 2012-09-25 16:08:55			
Using Merge Period Between 2012-09-11 and 2012-09-25.			
Aliases: OFF			
Host Name	Low	Medium	High
blade01.guard.svg.usma.ibm.com	100	0	0
blade02.guard.svg.usma.ibm.com	100	0	0
blade03.guard.svg.usma.ibm.com	100	0	0
blade04.guard.svg.usma.ibm.com	100	0	0

Figure 7-18 Resetting Unit Utilization Analysis

5. You can reset the analysis data for a single, multiple, or for all managed units. Select the check box next to the units you want to reset, as shown in Figure 7-19. You can also specify how far back to analyze data. Recall the first run goes back 24 hours only.

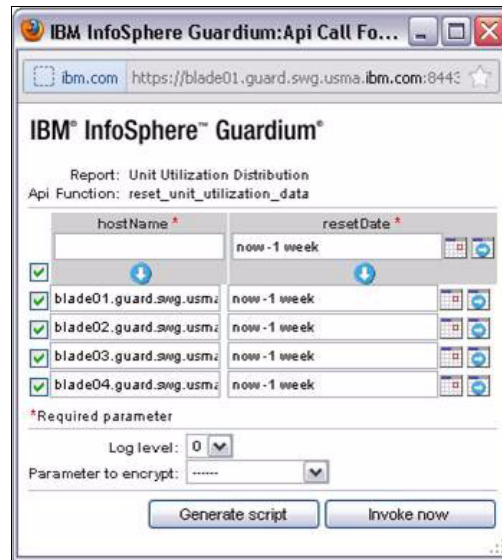


Figure 7-19 Specify which units to reset

Understanding Operational Dashboard output

Figure 7-20 shows an instance of a highly utilized collector. It is clear that the collector is high utilized because it is shown as red in the Unit Utilization window.

The screenshot displays the Operational Dashboard interface with the following sections:

- Unit Utilization:** Start Date: 2013-04-21 10:14:12 End Date: 2013-04-22 10:14:12. Aliaes: ON. Host Name: x12.guard.swg.usma.ibm.com. Max Overall Unit Utilization Level: High.
- Unit Utilization Distribution:** Start Date: 2013-04-21 10:14:12 End Date: 2013-04-22 10:14:12. Aliaes: OFF. Host Name: x12.guard.swg.usma.ibm.com. Low: 91, Medium: 5, High: 5.
- Utilization Thresholds:** Aliaes: OFF. Table with columns: Parameter Name, Threshold 1, Threshold 2.
- Unit Utilization Daily Summary:** Start Date: 2013-04-12 10:14:11 End Date: 2013-04-22 10:14:11. Aliaes: OFF. HostNameLike: LIKE %.

Host Name	Period Start	Max Number Of requests	Max Number Of requests Level	Number Of Requests % Increase	Max System Var Disk Usage	Max System Var Disk Usage Level	System Var Disk Usage % Increase	Max Mysql Disk Usage	Max Mysql Disk Usage Level	Mysql Disk Usage % Increase	Max Overall Utilization Level
x12.guard.swg.usma.ibm.com	2013-04-18 00:00:00	0	0	0	9	1	0	0	1	0	1
x12.guard.swg.usma.ibm.com	2013-04-19 00:00:00	7015317	3	0	9	1	0	0	1	0	3
x12.guard.swg.usma.ibm.com	2013-04-20 00:00:00	7137384	3	1	9	1	0	0	1	0	3
x12.guard.swg.usma.ibm.com	2013-04-21 00:00:00	6143719	3	-13	9	1	0	0	1	0	3
x12.guard.swg.usma.ibm.com	2013-04-22 00:00:00	0	1	0	9	1	0	0	1	0	1

Figure 7-20 Example of highly utilized collector

To identify the reason for high utilization on the collector, double-click the host and examine the Unit Utilization Report, as shown in Figure 7-21.

The screenshot shows the Unit Utilization report for the host x12.guard.swg.usma.ibm.com, which is highlighted in red. A context menu is open over the record, showing the following options:

- Record Details
 - Unit Utilization details
 - Alias Definition
 - Add API Mapping

Figure 7-21 Drill down to Unit Utilization details

The Unit Utilization report in Figure 7-22 shows a row during which the collector reached a Medium level of utilization because of many requests during that hour. Recall from Figure 7-17 on page 262 that Threshold 1 for number of requests is set to 1000000 and Threshold 2 is set to 5000000.

Period Start	Overall Utilization Level	Number Of requests	Number Of requests Level	Softer Memory	Softer Memory Level	Percent MySQL Memory	Percent MySQL Memory Level	Free Buffer Space	Free Buffer Space Level	Analyzer Status	Analyzer Status Level	Logger Status	Logger Status Level	MySQL Disk Usage	MySQL Disk Usage Level	System CPU Load	System CPU Load Level	System Var Disk Usage	System Var Disk Usage Level	Number Of requests	Number Of requests Level	Number Of Full SQLs	Number Of Full SQLs Level	Number Of Exceptions	Number Of Exceptions Level	Number Of Policy Violations	Number Of Policy Violations Level
2013-04-21 11:00:00.0	Low	0	Low	1089260	Low	21	Low	100	Low	0	Low	0	Low	0	Low	2	Low	9	Low	71	Low	0	Low	0	Low	0	Low
2013-04-21 12:00:00.0	Low	0	Low	1089260	Low	21	Low	100	Low	0	Low	0	Low	0	Low	2	Low	9	Low	71	Low	0	Low	0	Low	0	Low
2013-04-21 13:00:00.0	Low	0	Low	1089260	Low	21	Low	100	Low	0	Low	0	Low	0	Low	2	Low	9	Low	71	Low	0	Low	0	Low	0	Low
2013-04-21 14:00:00.0	Low	0	Low	1089260	Low	21	Low	100	Low	0	Low	0	Low	0	Low	2	Low	9	Low	71	Low	0	Low	0	Low	0	Low
2013-04-21 15:00:00.0	Low	0	Low	1089260	Low	21	Low	100	Low	0	Low	0	Low	0	Low	2	Low	9	Low	71	Low	0	Low	0	Low	0	Low
2013-04-21 16:00:00.0	Low	0	Low	1089260	Low	21	Low	100	Low	0	Low	0	Low	0	Low	2	Low	9	Low	71	Low	0	Low	0	Low	0	Low
2013-04-21 17:00:00.0	Low	0	Low	1089260	Low	21	Low	100	Low	0	Low	0	Low	0	Low	2	Low	9	Low	71	Low	0	Low	0	Low	0	Low
2013-04-21 18:00:00.0	Low	0	Low	1089260	Low	21	Low	100	Low	0	Low	0	Low	0	Low	2	Low	9	Low	71	Low	0	Low	0	Low	0	Low
2013-04-21 19:00:00.0	Low	0	Low	1089260	Low	21	Low	100	Low	0	Low	0	Low	0	Low	2	Low	9	Low	71	Low	0	Low	0	Low	0	Low
2013-04-21 20:00:00.0	Low	0	Low	1089260	Low	21	Low	100	Low	0	Low	0	Low	0	Low	2	Low	9	Low	71	Low	0	Low	0	Low	0	Low
2013-04-21 21:00:00.0	Medium	0	Low	1087875	Low	21	Low	100	Low	1	Low	35	Low	0	Low	57	Medium	9	Low	1996870	Medium	0	Low	440	Low	0	Low
2013-04-21 22:00:00.0	High	2	Low	1139131	Low	21	Low	100	Low	0	Low	0	Low	0	Low	70	Medium	9	Low	6143719	High	0	Low	911	Low	2	Low
2013-04-21 23:00:00.0	Low	0	Low	1148768	Low	21	Low	100	Low	0	Low	0	Low	0	Low	2	Low	9	Low	96	Low	0	Low	870	Low	1	Low
2013-04-21 00:00:00.0	Low	0	Low	1148768	Low	21	Low	100	Low	0	Low	0	Low	0	Low	2	Low	9	Low	96	Low	0	Low	0	Low	0	Low
2013-04-21 01:00:00.0	Low	0	Low	1148768	Low	21	Low	100	Low	0	Low	0	Low	0	Low	2	Low	9	Low	96	Low	0	Low	0	Low	0	Low
2013-04-21 02:00:00.0	Low	0	Low	1148768	Low	21	Low	100	Low	0	Low	0	Low	0	Low	2	Low	9	Low	96	Low	0	Low	0	Low	0	Low
2013-04-21 03:00:00.0	Low	0	Low	1148768	Low	21	Low	100	Low	0	Low	0	Low	0	Low	2	Low	9	Low	96	Low	0	Low	0	Low	0	Low
2013-04-21 04:00:00.0	Low	0	Low	1148768	Low	21	Low	100	Low	0	Low	0	Low	0	Low	2	Low	9	Low	96	Low	0	Low	0	Low	0	Low
2013-04-21 05:00:00.0	Low	0	Low	1195370	Low	21	Low	100	Low	0	Low	0	Low	0	Low	2	Low	9	Low	96	Low	0	Low	0	Low	0	Low
2013-04-21 06:00:00.0	Low	0	Low	1195204	Low	21	Low	100	Low	0	Low	0	Low	0	Low	2	Low	9	Low	96	Low	0	Low	0	Low	0	Low

Logger Queue	Logger Queue Level	MySQL Disk Usage	MySQL Disk Usage Level	System CPU Load	System CPU Load Level	System Var Disk Usage	System Var Disk Usage Level	Number Of requests	Number Of requests Level
0	Low	0	Low	3	Low	9	Low	71	Low
35	Low	0	Low	57	Medium	9	Low	1996870	Medium
3	Low	0	Low	70	Medium	9	Low	6143719	High
0	Low	0	Low	2	Low	9	Low	96	Low

Figure 7-22 Unit Utilization report

Also shown in Figure 7-22 is a second row in which the collector reached a High level of utilization. Once again, the reason is a high number of requests during that hour reaching over 6000000 and exceeding Threshold 2. Although the collector utilization is low during all other times, the Overall Utilization Level is always equal to the highest level in the Unit Utilization report.

Though this report results in overall High Utilization level, it might not necessarily be indicative of an issue. Because the machine reaches this level for only one hour in the test period, this is most likely an isolated event that can be ignored. If such events happen everyday and generate false alarms, it might make sense to modify the default thresholds from the Utilization Thresholds window that is shown in Figure 7-17 on page 262.

7.1.3 Interactive reports

To understand how to optimize reports, it is imperative to have an understanding of how reports work and how reported data is stored internally on the appliance. Make sure that you read Chapter 5, “Monitoring setup” on page 145 before you continue with the following performance process.

The term *report* is frequently used ambiguously to describe reports and queries. *Queries* are the definitions of what and how should be retrieved and *reports* are definitions of how the query results should be displayed. When we describe the performance of the reports, we are referring to the performance of the queries.

Each query that is associated with a particular predefined set of data is called a data domain. For example, Access domain for captured traffic, Exception domain for captured errors from the database server, or Guardium activity domain to monitor activities that are performed by Guardium users. There are approximately 40 different domains on the Guardium appliance today.

Figure 7-23 shows report configuration windows of Guardium.

The screenshot displays the Guardium report configuration window. The interface is divided into several sections:

- System View:** Administration Console, Tools, Daily Monitor, Guardium Monitor, Tap Monitor, Incident Management.
- Config & Control:** Report Building.
- Entity List:** A list of entities including Client/Server, Session, Server IP/Server Port, App User Name, Application Events, FULL SQL Value, FULL SQL, SQL, Access Period, Command, Object, Object/Command, Join, Field SQL Value, Qualified Object, Field, and Object/Field.
- Activity By Client IP:** Main Entity: Session. A table with columns Seq, Entity, and Attribute. The table contains three rows:

Seq	Entity	Attribute
<input type="checkbox"/>	1 Client/Server	Client IP
<input type="checkbox"/>	2 Command	SQL Verb
<input type="checkbox"/>	3 Object	Object Name
- Additional Options:** Addition mode: AND (selected), OR, HAVING.
- Table Header:** Entity, Agg., Attribute.

Figure 7-23 Report configuration

Most of the domains deal with small amount of data and are irrelevant for this performance description. The one domain that deals with large amount of data and is used often is the Access domain. This domain contains database activity data that is captured by appliance. This domain is composed of many tables with many millions of rows of data in each table. You should carefully consider your decisions when you are designing reports in this domain. In Figure 7-23 on page 266, the second column shows the entity list that is available in the Access domain. Each entity corresponds to an underlying table and has a list of attributes. Attributes correspond to fields in the table. When you select fields (attributes) from the Session entity and some other fields from the SQL entity, these two entities are joined to retrieve data that you request. The more entities that are joined in your query, the more complex the final query is and takes longer to produce the results.

To better understand how query builder works, you should have a more detailed understanding about how captured data is stored in the underlying tables (entities).

Guardium monitors and captures numerous details about database user activity. All of this information can be put into the following major categories:

- ▶ **Who:** Describes a connection to a database, who made a connection, and when the connection was made.
- ▶ **What:** Contains the SQL statements that were run on the database.

Connection details include attributes (fields), such as server IP address, database type, database user name. This information is recorded in the Client/Server and Session entities as shown in Figure 7-23 on page 266. Two entities are used because although the login information (IP, user name, and so on) of a user stays the same for every connection, there is unique information for every connection; for example, login time or client port. The relatively static, repeatable login information is stored in Client/Server entity and the unique, connection-specific information is stored in the Session entity. Splitting information between two entities helps to reduce data redundancy and saves disk space.

In addition to login information, Guardium captures the SQL statements that are issued by the user or an application. The SQL statements are recorded in SQL entity. To create queries with conditions on specific groups of tables or sets of commands, Guardium parses captured SQLs to commands, objects, and fields and places this information in three other entities: Commands, Objects, and Fields.

Therefore, if you want to create a report that shows only the activity on a particular table, you can create a query with a condition, such as where `OBJECT.OBJECT_NAME = 'myTable'`, or, if you want to create a report that shows only DML activity, you can create query condition, such as `WHERE COMMAND.VERB in group 'DML commands'` The redundancy here helps to create more efficient queries.

When you create a query, you enter a query name first then you select a main entity. It is important to select a Main Entity that tells query builder the focal point for the new report and how to construct a query. Ultimately, it might also affect query performance.

Consider the following examples.

Figure 7-24 and Figure 7-25 on page 269 show two queries that have identical fields that are selected. The only difference between these two queries is the main entity that is selected, one is Session (the better choice) and the other is Command.

SessionList MainEntitySession

Main Entity: Session Add Count Add

Seq.	Entity	Attribute	Field Mode	Order-by	Sort Ran
<input type="checkbox"/>	1	Session	Session Start	Value	<input type="checkbox"/>
<input type="checkbox"/>	2	Session	Session End	Value	<input type="checkbox"/>
<input type="checkbox"/>	3	Client/Server	Client IP	Value	<input type="checkbox"/>
<input type="checkbox"/>	4	Client/Server	Server IP	Value	<input type="checkbox"/>
<input type="checkbox"/>	5	Client/Server	DB User Name	Value	<input type="checkbox"/>
<input type="checkbox"/>	6	Client/Server	Source Program	Value	<input type="checkbox"/>

AND OR HAVING

Entity	Agg.	Attribute	Operator	Runtime Param.
--------	------	-----------	----------	----------------

Figure 7-24 Report with Session as the main entity




Figure 7-25 Report with Command as the main entity

The following query is generated by the query builder with the Session main entity:

```
select ... from GDM_ACCESS, GDM_SESSION where....
```

The following query is generated by the query builder with the Command main entity:

```
select ... from GDM_ACCESS, GDM_SESSION, GDM_CONSTRUCT_INSTANCE, GDM_SENTENCE where...
```

Both queries have the same columns. However, the first query joins two tables to produce the results and the second query has four tables that participated. The second query takes longer to complete. Even more important, most likely there are more records in a second report and some of the rows appear multiple times.

When Command is selected as the main entity, the report generator defines the report with the focus on “command”. There are most likely many commands in a session and each command appears on the report in a separate row, even if you do not have a command that is displayed on the report.

Main entities are organized hierarchically from high-level details to more granular. Thus, the main entity defines the level of details in the report. Selecting a main entity on too high a level in the list might limit your ability to select fields to report. An example is a single SQL statement with multiple fields. If you select SQL as a main entity, your level of detail is an SQL statement and each line in the report is dedicated to one SQL statement. This means that you cannot display fields in the same line because there is no space for multiple fields.

However, you can use the count function to display total count of the fields in a SQL statement or the max function to display the highest field value. Figure 7-26 shows a report definition with SQL as the main entity for database activity. You cannot add the SQL Verb field from the Command entity because the Command entity is positioned lower than SQL entity in the entity list. Therefore, you cannot use the field value directly but you can apply the value to one of the math functions, such as count, min, or max.




Figure 7-26 Report with SQL as the main entity

When you are designing a new query, consider the relationships between entities to avoid data redundancy in reports. Figure 7-27 shows another example with Field as the main entity and a few columns selected.




Figure 7-27 Report with Field as the main entity

Figure 7-28 shows a snippet of the report that is generated by this query.

Field level activity		
Start Date: 2012-09-08 11:48:31 End Date: 2013-07-08 11:48:31		
Aliases: OFF		
Server IP	DB User Name	Sql
155.180.58.42N087764		UPDATE msdb.dbo.sysjobactivity SET start_execution_date = @startExecutionDate WHERE job_id = ? and session_id = ?
155.180.58.42N087764		UPDATE msdb.dbo.sysjobactivity SET start_execution_date = @startExecutionDate WHERE job_id = ? and session_id = ?
155.180.58.42N087764		UPDATE msdb.dbo.sysjobactivity SET start_execution_date = @startExecutionDate WHERE job_id = ? and session_id = ?
155.180.58.42N087764		UPDATE msdb.dbo.sysjobactivity SET start_execution_date = @startExecutionDate WHERE job_id = ? and session_id = ?
155.180.58.42N087764		UPDATE sysjobactivity SET last_executed_step_date = @lastExecuteStepDate, last_executed_step_id = ? WHERE job_id = ? AND session_id = ?
155.180.58.42N087764		UPDATE sysjobactivity SET last_executed_step_date = @lastExecuteStepDate, last_executed_step_id = ? WHERE job_id = ? AND session_id = ?
155.180.58.42N087764		UPDATE sysjobactivity SET last_executed_step_date = @lastExecuteStepDate, last_executed_step_id = ? WHERE job_id = ? AND session_id = ?
155.180.58.42N087764		UPDATE sysjobactivity SET last_executed_step_date = @lastExecuteStepDate, last_executed_step_id = ? WHERE job_id = ? AND session_id = ?
155.180.58.42N087764		UPDATE sysjobactivity SET last_executed_step_date = @lastExecuteStepDate, last_executed_step_id = ? WHERE job_id = ? AND session_id = ?

Figure 7-28 Report snippet

The report shows the same line repeatedly because selecting Field as the main entity instructs the report generator to dedicate one line in the report to a field. If the Field attribute is added to report, we can see that after we select the Field domain as a main entity, each line of the report is associated with one Field and the rest of information is repeated as needed, as shown in Figure 7-29.

Field level activity			
Start Date: 2012-09-08 11:48:31 End Date: 2013-07-08 11:48:31			
Aliases: OFF			
Server IP	DB User Name	Sql	Field Name
155.180.58.42N087764		UPDATE msdb.dbo.sysjobactivity SET start_execution_date = @startExecutionDate WHERE job_id = ? and session_id = ?	start_execution_date
155.180.58.42N087764		UPDATE msdb.dbo.sysjobactivity SET start_execution_date = @startExecutionDate WHERE job_id = ? and session_id = ?	@startExecutionDate
155.180.58.42N087764		UPDATE msdb.dbo.sysjobactivity SET start_execution_date = @startExecutionDate WHERE job_id = ? and session_id = ?	job_id
155.180.58.42N087764		UPDATE msdb.dbo.sysjobactivity SET start_execution_date = @startExecutionDate WHERE job_id = ? and session_id = ?	session_id
155.180.58.42N087764		UPDATE sysjobactivity SET last_executed_step_date = @lastExecuteStepDate, last_executed_step_id = ? WHERE job_id = ? AND session_id = ?	last_executed_step_date
155.180.58.42N087764		UPDATE sysjobactivity SET last_executed_step_date = @lastExecuteStepDate, last_executed_step_id = ? WHERE job_id = ? AND session_id = ?	@lastExecuteStepDate
155.180.58.42N087764		UPDATE sysjobactivity SET last_executed_step_date = @lastExecuteStepDate, last_executed_step_id = ? WHERE job_id = ? AND session_id = ?	last_executed_step_id
155.180.58.42N087764		UPDATE sysjobactivity SET last_executed_step_date = @lastExecuteStepDate, last_executed_step_id = ? WHERE job_id = ? AND session_id = ?	job_id
155.180.58.42N087764		UPDATE sysjobactivity SET last_executed_step_date = @lastExecuteStepDate, last_executed_step_id = ? WHERE job_id = ? AND session_id = ?	session_id

Figure 7-29 Report with Field attribute added

Certain database operations (such as GROUP BY, DISTINCT, ORDER BY, or HAVING clauses) provide flexibility to the report builder. However, these operations might take more processor time. If you have report performance issues, consider revising your report to limit the usage of these database operations.

In general, the data volume that is stored on the appliance is the major factor that can affect the report performance. When you tune the report performance, consider the following points:

- ▶ Define the purge process to run nightly.
- ▶ Configure the data retention period to the minimum that is allowed by your business requirements.
- ▶ Record Full SQL only when it is necessary (for example, when monitoring sensitive objects or when monitoring privileged users). Full SQL tables can add data volume quickly.
- ▶ Reduce the period of the report to have a positive effect on the report run time.
- ▶ Analyze MySQL database performance once a month to update the index cardinality by clicking **Perform Maintenance Actions** → **TURBINE analyze**.

In the federated environments, reports frequently run on the aggregators instead of the collectors. Data that is exported from the collectors nightly in one-day chunks is transferred to the aggregator. On an aggregator, data from multiple collectors and multiple days is merged for reports. For efficiency reasons, not all of the data that is presented on an aggregator is merged and made available for reports. The default merge period is 14 days. You can change the merge period by using the `store aggregator merge_period X` CLI command, where X is the number of merging days.

By reducing the merging days, you reduce the data volume that query works with and improve query performance. (This is true only for interactive reports.) If you run reports in the background by using the ad hoc option, the background report process creates its own set of merged data to match the report period. Background reports run the same way as Audit Process reports.

7.1.4 Audit process reports

As the name indicates, the interactive reports are run interactively with the users waiting for the results in front of the monitor. Therefore, the run time of the interactive reports is limited to three minutes after which the report is ended automatically. Many reports might require much longer time to run. These reports should run as ad hoc background reports or as part of the audit process. Both options use the same internal mechanism. The audit process that is described in this section also is applicable to the ad hoc background reports.

Audit process is a mechanism that allows you to submit a group of tasks to run asynchronously, on a predefined schedule, or on demand and forward the run results to a group of predefined receivers automatically.

There are a few factors that can influence the performance of the audit process. First, all of the considerations that are described in 7.1.3, “Interactive reports” on page 266 about how to build efficient queries are also applicable to the audit report process. These considerations become even more important because the reports that are running as tasks in the audit process typically deal with large data volume. However, there are a few significant differences in the way the audit process reports are run.

The results of the interactive reports are kept in the memory and last only while they are displayed on the monitor. When you move to another page, the results are gone. If you want to return to the report results, you must rerun the report.

The results of the audit process reports are not deleted at the end of the run and are stored in the internal database for later view, as needed. Each line of the audit process report is stored as a row in the table in the internal database. This table can grow large quickly and might eventually fill up all available space in the database and affect system performance. An internal process is built in to keep the table size under control. This process is based on some simple rules that users must follow to make sure that the old data is deleted in time and the report result table does not grow too large.

First, configure how long you want to keep your results under the Audit Process Definition, as shown in Figure 7-30 on page 274.

Audit Process Builder

Audit Process Definition ?

Description: test

Active: There is no schedule associated with this process

Archive Results:

Keep for a minimum of: 0 days or 5 runs

CSV/CEF File Label: test Zip CSV for mail

Email Subject: _____

View Run Once Now Modify Schedule...

Receiver Table

Figure 7-30 Defining the report result retaining period

The default setting of five runs means that only the last five runs of the report results are kept and when you run the report again, the new results are recorded and the oldest run results are deleted. You can use number of days (for example, 30 days) instead of number of runs. Configure this parameter to a value that meets your business needs but not too big to help keeping the report results table size down.

The second setting is related to the designated result receivers and their required action. The report results must be reviewed (and optionally signed) by receivers before they can be purged. Figure 7-31 shows that the results are sent to admin user for review.

Receiver	Action Req.	To-Do List	Email Notif.	Cont.	Appv. if Empty
<input checked="" type="checkbox"/> admin	<input checked="" type="radio"/> Review <input type="radio"/> Sign	<input checked="" type="checkbox"/>	<input checked="" type="radio"/> No <input type="radio"/> Link <input type="radio"/> Full Results	<input checked="" type="checkbox"/>	<input type="checkbox"/>
<input checked="" type="checkbox"/> (admin admin)					

Add Receiver

Receiver name: _____ Search users

Action Required: Review Sign

To-Do List: Add

Email Notification: None Link Only Full Results

Continuous:

Approve if Empty: Yes

Figure 7-31 Setting report receivers

Note: Unsigned or not reviewed report results can fill up the database and affect report performance.

Specifying task result receiver was mandatory on the Guardium version before V9.0. This caused some issues with report results not being purged in time. In Version 9.0 and later, specifying a task receiver is optional. If you upgrade your appliances from the version before 9.0, review your audit process definitions and verify that task results receivers are defined only when review of the results is required.

Some customers use the audit process reports to transfer data to other systems in their environment for integration and reporting purposes. These types of reports tend to be large with hundreds of thousands of rows in a single report and can take some time to complete. To improve the performance of these reports, avoid the use of ORDER BY, GROUP BY, or HAVING statements wherever possible. In particular, the ORDER BY statement might affect performance dramatically on large reports (over 100000 rows).

Audit process can run on the collector or aggregator on any appliance. However, audit processes are run on aggregators typically.

On aggregators, the audit process does not run on the default database but on the ad hoc database subset that is created simultaneously. This database subset includes only data from days that are required by the report. This significantly reduces the data volume that the audit process reports work with and reduces report run time.

Data that is used by reports is transferred from collectors by the daily import and merge process. Reports should be run after the daily import and merge process so the report can use the most current data.

Typically, the import and merge process runs shortly after midnight. Allow enough time for the import and merge process to complete and schedule the audit processes to start at early morning hours to avoid process congestion.

7.1.5 Back up and purge

Backups are performed to save vital data and configuration information for retrieval when a Guardium Server becomes unusable, in need of repair, or is being replaced.

You can have a full backup or a “snapshot” of the Guardium server. In virtualized environments, a backup also can be done by making an actual snapshot of the Guardium machine. Restoring data from the Guardium system backup replaces all existing data that is stored on the appliance with the data from the backup file. As a result, all activity that is collected after the last backup is lost. For more information about the granular (day-by-day) restoration of the data, see Chapter 8, “Disaster recovery” on page 313.

It is best to run full system backup at least once a month or more that is based on system activity, restoration frequency, and business requirements.

Backup files are generated in the /var partition before it is transferred to an outside storage facility. To avoid an out-of-space failure, you must ensure that the /var partition has enough disk space to complete the backup. It is best to have threshold alerts set to monitor system and MySQL disk usage.

In general, the largest component of the backup is data that directly correlates to the MySQL database usage. The higher the MySQL database usage (% used), the larger a system backup file size is.

One method to check the MySQL database and /var partition usage is through the Buffer Usage report, as shown in Figure 7-32.

MySQL on Disk Usage	MySQL Is Up	Promiscuous Received	Sniffer Connections Ended	Sniffer Connections Used	Sniffer Packets Dropped	Sniffer Packets Ignored	Sniffer Packets Throttled	System Cpu Load	System Memory Usage	System Root Disk Usage	System Uptime	System Var Disk Usage
0	0	106	0	0	0	0	0	0	90	18	48415512	
0	0	126	0	0	0	0	0	1	90	18	48414902	
0	0	84	0	0	0	0	0	0	90	18	48414282	

Figure 7-32 Buffer usage of MySQL and /var partition usage

Alternatively, you can check the space utilization through System View, as shown in Figure 7-33 on page 277.

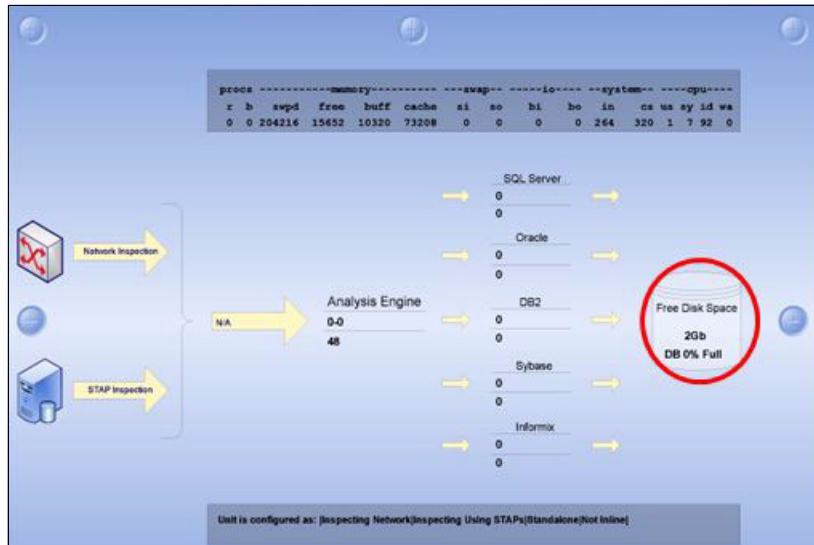


Figure 7-33 System view

To verify that the backup completed successfully, complete the following steps to check the Aggregation archive log:

1. Log in to the Guardium graphic user interface (GUI) as Admin.
2. Select **Guardium Monitor** → **Aggregation/archive Log**.
3. Sort by “backup” type.

The log file is also accessible through the `fileserver` utility. The corresponding log file is `turbine_backup.log`.

The purge operation should be done on regular basis to free up space and speedup access operations on the internal database. The default purge value is 60 days. By default, purge is scheduled at 5:00 a.m. daily.

The amount of data that can be stored on the appliance depends on many criteria, including appliance type, disk space, and policy. The purge period must be adjusted to reflect the optimal balance between data accessibility and quick response time of the system process. In a typical implementation, best practices are to keep the data for 30 days or less.

You can reset the purge period through Data Archive and Data Export, as shown in Figure 7-34 on page 278.

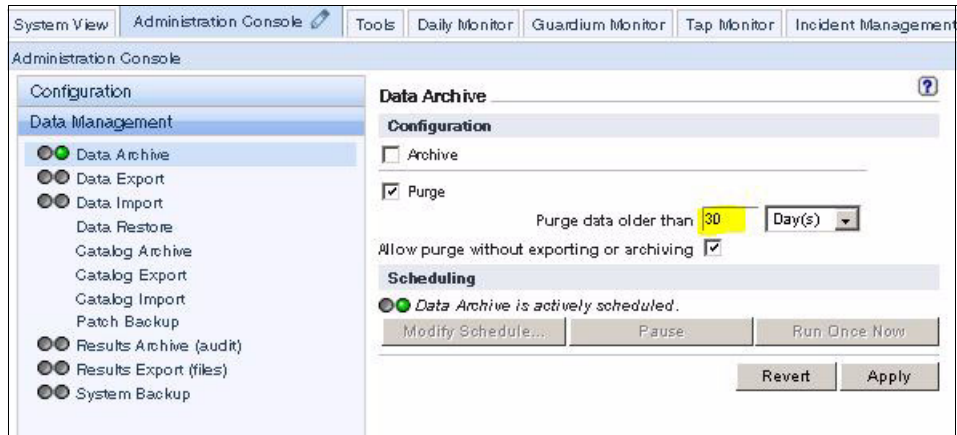


Figure 7-34 Setting purge period

There is no warning when you purge data that was not archived or exported. The purge operation does not purge restored data whose age is within the “do not purge restored data” time frame that is specified on a restore operation window.

To verify that the purge is running on your system and to check its completion status, click **Aggregation/Archive Log** under the Guardium Monitor tab in Guardium GUI, as shown in Figure 7-35.

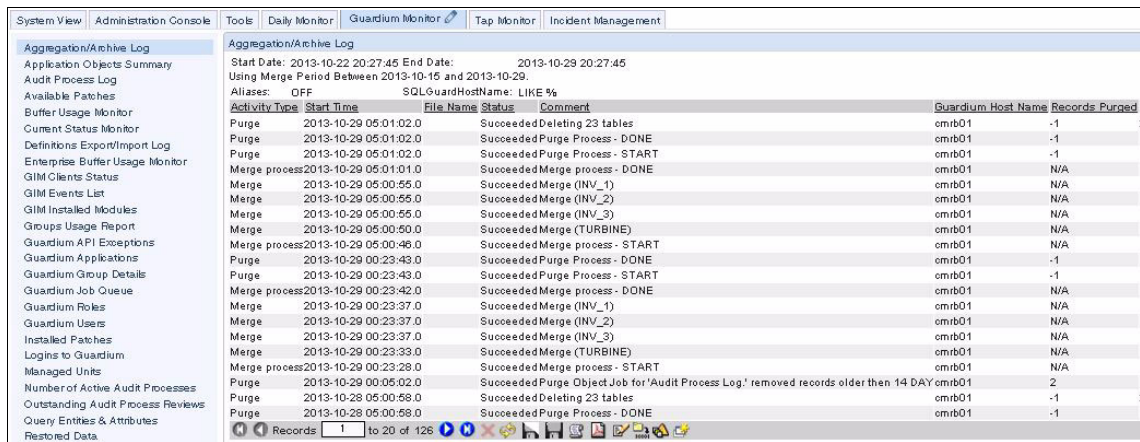


Figure 7-35 Aggregation/Archive Log

Note: A successful data purging operation is critical to maintain a healthy and efficient system.

7.1.6 Central management

Central Manager is the Guardium appliance that monitors and controls other Guardium units in a federated environment. Unmanaged Guardium appliances are referred to as stand-alone units.

The smooth and efficient operation of a Central Manager unit is critical to the overall Guardium system performance. In this section, we describe some considerations about Central Manager efficiency and maintenance.

Guardium definitions

Central Manager houses most of the definitions of all of the units that report to it. When users submit any report, query, or audit process on any managed unit in a federated (centrally managed) environment, definitions of this activity are retrieved directly from Central Manager. Therefore, latency between Central Manager and its managed units can be a contributing factor for potential user interface slowness on the corresponding managed units.

Users can use Central Manager or any of its managed units to modify those definitions. Regardless of the appliance where the definition changes were made, updated content (with an exception of Policies and Groups) is immediately available on all the appliances across the federated environment.

Note: In distributed federated environments with high latency between Central Manager and its managed units, it is best to apply all definition changes directly on the Central Manager appliance.

Portal synchronization and remote policy installation

Policy rules and Groups are considered a special type of definitions. Similar to the other Guardium definitions, the master copy of Policy and Groups are kept on the Central Manager unit. However, to provide a high level of operation efficiency, managed units maintain a local copy of those definitions that are synchronized with the master copy on Central Manager regularly.

The Portal user synchronization process controls the Group synchronization operation, as shown in Figure 7-36 on page 280. The default (which also is the best) synchronization frequency is 30 minutes.



Figure 7-36 Managed Unit Portal User Synchronization

You can deploy a policy by using the installation policy option by clicking **Data Management** → **Central Manager** → **Central Manager**, as shown in Figure 7-37.

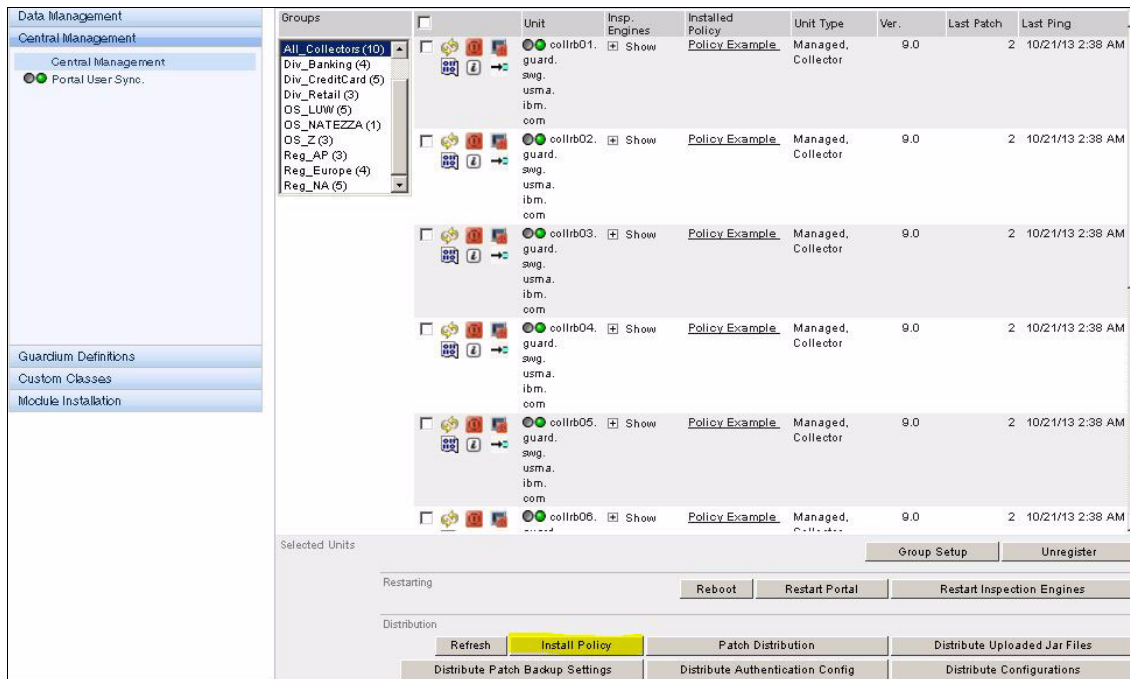


Figure 7-37 Deploying a policy

Schedule the Remote policy installation options (from Central Manager) on a permanent basis. The frequency and specific time of the push depends on the “quiet time” window and the rate of policy changes in a particular organization.

For enterprise customers with many managed units, it is best to split the policy installation into smaller subsets of units and schedule the installation at different times to avoid connection bottlenecks on Central Manager.

Remote source invocation

The Central Manager appliance also can schedule audit tasks to run on aggregators and collectors through the remote source function. Though the tasks run on the managed unit, this function requires Central Manager to dedicate significant memory resources to handle the result output data of the audit tasks. In certain extreme cases, this can cause “out of memory” exceptions and report failure. To prevent this issue, it is best to spread the task invocation schedule throughout the day to minimize processing multiple task results simultaneously, especially those with larger result sets.

Guardium Version 9.0 patch 50 and later increase the memory threshold for Tomcat server that helps increasing virtual memory capacity of the Central Manager appliance to handle larger result sets. This is especially important for the 64-bit version.

7.2 Maintenance and updates

In this section, we provides insights about Guardium product updates and patches. We describe methods and considerations for upgrading Guardium appliances and agents and where to download the Guardium updates and patches.

7.2.1 Appliance updates

Guardium software updates are issued regularly and typically include new enhancements, bug fixes, and security and vulnerability updates. Appliance updates are managed by and available only through Guardium internal patch installation mechanisms.

Bug fixes, which are committed into the current version, are always merged into all future versions. Depending on the critical nature of the issues, some fixes also are back-ported into supported downward versions. Back porting decisions are typically made on case-by-case basis. Because not all the changes are automatically back ported, it is important to keep up with latest versions and Guardium Patch Update (GPU) patching.

Combined fixes applies to fixes that require software changes on agent and appliance sides. In those cases, implementation must allow the server and client to have different version. Release notes of the appropriate agent provides specific reference to the patch level of the corresponding appliance that are required to address a particular issue.

Guardium software changes within the same major version that are delivered by one of two packaging methods: Ad hoc patches or GPU fix packs. Though the content and purpose of those two packages are different, they are produced by using the same Guardium patching mechanism.

All Guardium patches are encrypted and can be applied for corresponding version or Guardium software only.

Ad hoc patches

Ad hoc patches provide temporary relief to the customers with urgent, critical, or prevalent Guardium software issues. In addition, ad hoc patches are used to address configuration and data manipulation requests for specific customer needs.

Ad hoc patches contain only the modules that fix a specific customer issue. Ad hoc patches often depend on the latest GPU version of the major code version. Modules, which are included in ad hoc packages, contain all of the modifications that were previously made to the same module within the same version.

The distribution vehicle that is most often used for the ad hoc patches is Support PMR systems. Ad hoc patches that are uploaded to the corresponding PMR automatically send email notification to the customers.

GPU patches

GPU is a cumulative fix pack that contains all the software modifications, database changes, and security updates that are committed into the corresponding version since the GA release version. GPU must be installed on the corresponding main version of Guardium. For example, GPU V9.0p50 can be installed only on V9.0 appliance. However, within the version, any latest GPU patch can be installed on the top of any Guardium patch level.

In GPU patches, database changes and operating system upgrades are kept to a minimum to keep GPU upgrades relatively quick and an easy way to update Guardium software.

Guardium GPU fix packs are released to IBM Fix Central quarterly. To benefit from latest version of Guardium software, install GPU patches in timely manner.

IBM software downloading sites

Guardium releases and fix packs are provided through IBM Passport Advantage or IBM Fix Central. All new GA releases of Guardium software are deployed directly to Passport Advantage. All of the post GA release software modifications, documents, and fix packs are released through Fix Central.

Fix Central contains the latest versions of Guardium Agents (such as S-TAP, Configuration Audit systems, Guardium Installation Manager), GPU fix packs, Database Protection Knowledgebase Subscription (DPS) updates, and Appliance upgrade bundles. Those changes often become available a few months after the GA release. When new versions of a software component are released to Fix Central, the version in Passport Advantage of the same component becomes obsolete. Therefore, it is important to always download the latest versions of Guardium software directly from Fix Central, which is available at this website:

<http://www-933.ibm.com/support/fixcentral/>

Note: Fixes for certain Guardium components are not released to Fix Central. These fixes include customer software licenses, appliance ISO images, upgrade patches, and release documentation. These components are available from the Passport Advantage at IBM Passport Advantage site, which is available at this website:

<http://www-01.ibm.com/software/lotus/passportadvantage/>

Patch deployment with Central Manager

The Central Manager appliance provides a quick and easy way to distribute Guardium patches to all managed units that are connected to this environment. The managed units must be registered to this Central Manager appliance first.

Before you install patches or distribute patches to the managed units through Central Manager, patches must be uploaded to the Central Manager appliance. You can upload patches to the Guardium system by using one of the following methods:

- ▶ The **fileserv** utility
- ▶ Central Manager distribution
- ▶ CLI patch installation command

The **fileserv** utility is a Guardium CLI command that provides access to internal logs and patch uploading services through a stand-alone web server instance.

To upload patches using the **fileserv** utility, complete the following steps:

1. Open an SSH client and log in to the Central Manager appliance as a CLI user.
2. Run the **fileserv** command.
3. Use the displayed link to open the file server page in a web browser.
4. Click **Upload a patch**.
5. Browse to the location of the patch and click **Upload**, as shown in Figure 7-38.




Figure 7-38 Uploading patches

6. Press Enter to stop file server, as shown in Figure 7-39.

```
Welcome cli - your last login was Thu Apr 25 02:50:01 2013
cmrb01.guard.swg.usma.ibm.com> fileserv
Creating the index file.

Starting the file server. You can find it at http://cmrb01.guard.swg.usma.ibm.com
The timeout has been set to 600 seconds and it may timeout during the uploading.

Press ENTER to stop the file server.

Stopping process
ok
```

Figure 7-39 Stopping file server

You also can distribute the patches through the Guardium Patch Distribution function in the Administration Console by clicking **Administration Console** → **Central Management** → **Central Management**.

Figure 7-40 shows list of the managed units that are available through the Administration Console.

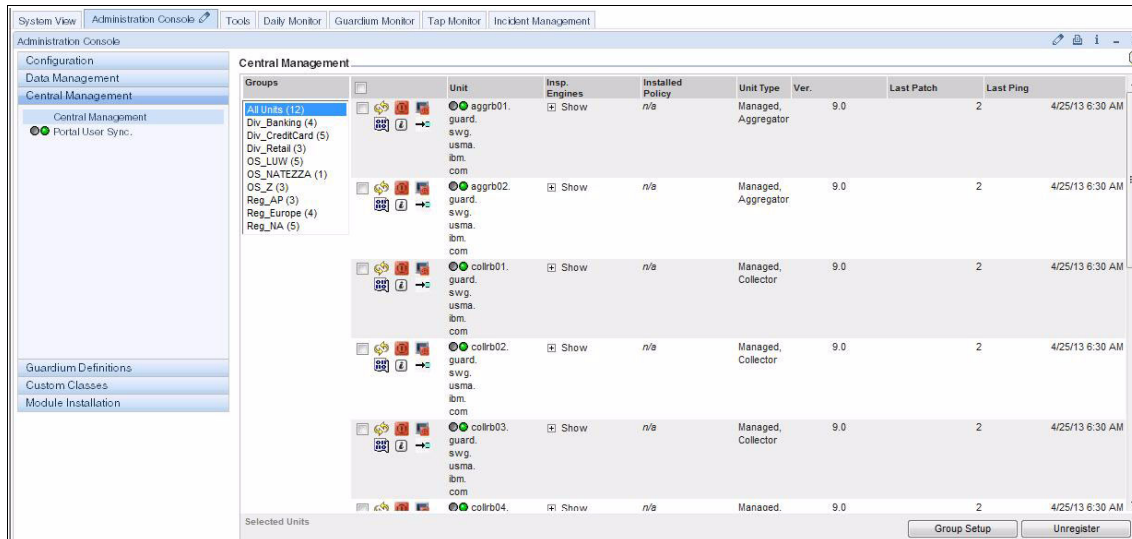


Figure 7-40 Managed unites in Administration console

Guardium can distribute patches to the target individual appliance and all managed units simultaneously. However, for a Guardium environment that has large number of appliances, distribution management can be a challenging task and requires a more systematic approach. In a large enterprise environment, appliances and managed units can be grouped by geographical or business-related areas by using the Guardium Group Setup in Central Manager.

A managed unit can belong to multiple groups. The default ALL Units group on the Central Manager contains all of the managed units.

An example is a customer with three major divisions. Each division has separate operational schedules and maintenance time frames. Guardium agents are on Linux and z/OS systems that are in two different geographic regions. To accommodate all aspects into the enterprise solution, appliance grouping can include considerations for associated division, region, and agent operational system.

To create a group, click **Group Setup** at the bottom of the Central Management window, as shown in Figure 7-41 on page 286. Enter group name and click **Add**.



Figure 7-41 Group Setup

To assign the new unit to a group, on Central Management window, select the managed units that you want to assign and click **Group Setup** to open the Group Setup window. A list of the selected units and groups is displayed. Select the group the units that you want to add and click **Update groups**, as shown in Figure 7-42.



Figure 7-42 Update group

By using the patch distribution function, you can apply remote installation instantaneously by using the Run Once Now option or scheduling it for a later time.

Patch deployment with CLI

Guardium also provides an option to upload and install patches directly through a CLI. This option is mostly usable for stand-alone appliances and is an alternative to the Central Manager distribution method.

You can locate the patches to be uploaded by running the **show system patch available** command, as shown in Figure 7-43. This command allows users to see and validate available patches, including md5sum of the patch to make sure that the patch was uploaded successfully.

```
collrb01.guard.swg.usma.ibm.com> show system patch available
Attempting to retrieve the patch information. It may take time. Please wait.
P#      Description                               Version Md5sum
      Dependencies
2       Guardium Patch Update (GPU) for Version 9.0  9.0
ok
```

Figure 7-43 Showing available patches

Note: The **show system patch available** command displays patches that are uploaded by using the **filesaver** utility or through Central Manager distribution only. If the patch is uploaded by using a different method (directly from CLI command), run the **store system patch install** command to see and install the patch.

Note: Make sure to close the **filesaver** command after you upload the patch, as shown in Figure 7-39 on page 284. If file server is not closed properly or timed out during the uploading process, the patch might not be visible by running the **show system patch available** command.

To install the uploaded patches, run the **store system patch install sys now** command, as shown in Figure 7-44.

```
collrb01.guard.swg.usma.ibm.com> store system patch install sys now
The backup profile is not set for saving the backup file if patch installation fails.
If you want to save the backup file, please answer "NO" to the question
and run CLI command "store backup profile" to set up the parameters.
yesyou want to continue (yes or no)?
y
List the files in the patches directory:
1. SqlGuard-9.0p9997.tgz.enc
Please choose patches to install (1-1, or multiple numbers separated by ",", or q to quit):
```

Figure 7-44 Installing patches

Patches that are uploaded to Central Manager can be distributed to the managed units without installing them locally on Central Manager. This condition works only for patches that are uploaded through fileserver. All patches that are installed locally on the Central Manager (regardless of the upload method) are also available for distribution.

To upload and install patch directly from CLI without pre-uploading, run the **store system patch install** command with the file transfer flags. This command includes SCP and FTP patch transfer methods and provides interfaces to upload patches directly from CD and DVD.

Figure 7-45 shows an example of uploading the Pre-GPU Health Check patch 9.0p 9997 to our lab collector by using SCP transfer method. The CLI allows users to abbreviate words if it identifies a unique value. In this example, “install” was abbreviated as “in”.

```
collrb01.guard.swg.usma.ibm.com> store system patch in scp

Please enter the following information for file transfer:
Host to import patch from: myhost
User on myhost: myaccount
Full path to the patch, including name (file name may use wildcard *): SqlGuard-9.0p9997.tgz.enc
Password:

Enter the SCP port if you need to use a special port.
Enter "0" or press "Enter key" to use the default port: 0

The file transfer process can take a while to complete.
Leave the terminal open and do not answer any questions until the transfer is complete.

Starting transfer, please wait.
spawn /usr/bin/scp myaccount@myhost:SqlGuard-9.0p9997.tgz.enc /var/log/guard/patches/
```

Figure 7-45 Beginning of the **store system patch installed scp** command

Upon successful file transfer, the user is prompted to install the uploaded patch to the current system.

7.3 Diagnostic tools

Supportability and diagnostic testing of distributed systems always pose significant challenges. Adding to it the complexity and sophistication of such a versatile product as Guardium can make support and problem diagnosis a demanding task. In the past few years, IBM made great strides in the development of supportability and diagnostic tools to ease your effort in supporting a Guardium environment.

In this section, we describe enhancements, recommendations, and best practices concerning troubleshooting Guardium.

7.3.1 Appliance diagnostic tools

MustGather is the diagnostic tool that Guardium provides for gathering the appliance information for troubleshooting. MustGather was first introduced in Version 8.2. In the subsequent version, new functions were added and existing functions were enhanced.

MustGather procedures

MustGather includes procedures for gathering specific areas of diagnostic information. The MustGather of Version v9.0p50 contains the following diagnostic functions:

- ▶ **agg_issues**: For issues that are related to aggregation processes
- ▶ **alert_issues**: For issues that are related to alerts
- ▶ **app_issues**: For issues that are related to GUI
- ▶ **audit_issues**: For issues that are related to audit processes
- ▶ **backup_issues**: For issues that are related to the backup process
- ▶ **cm_issues**: For issues that are related to central management functionality
- ▶ **miss_dbuser_prog_issues**: For issues that are related to missing database user and source programs
- ▶ **purge_issues**: For issues that are related to purge process
- ▶ **scheduler_issues**: For issues that are related to scheduler functionality
- ▶ **sniffer_issues**: For issues that are related to sniffer functionality
- ▶ **system_db_info**: For issues that are related to appliance space and the performance of databases and operating systems

Guardium patch v9.0p50 introduces the following MustGather commands for network and patch installation issues:

- ▶ **patch_install_issues**: For issues that are related to patch installation and upgrades
- ▶ **network_issues**: For issues that are related to network architecture

To see a list of available MustGather commands, run the **com must** command in CLI, as shown in Figure 7-46.

```
collrb01.guard.swg.usma.ibm.com> com must
support must_gather ?
support must_gather agg_issues
support must_gather alert_issues
support must_gather app_issues
support must_gather audit_issues
support must_gather backup_issues
support must_gather cm_issues
support must_gather miss_dbuser_prog_issues
support must_gather purge_issues
support must_gather scheduler_issues
support must_gather sniffer_issues
support must_gather system_db_info
```

Figure 7-46 List available MustGather commands

You can run MustGather commands from the CLI. Figure 7-47 shows running the **support must_gather sniffer_issues** command.

```
collrb01.guard.swg.usma.ibm.com> support must_gather sniffer_issues

This operation may take several minutes to complete.

9.0.0_r45260_v90_1-e158-20121102_1546
Created file /var/log/guard/must_gather/sniffer_logs/sniffer.20131023.tgz.
ok
collrb01.guard.swg.usma.ibm.com> █
```

Figure 7-47 Running MustGather command

When the MustGather command is run through CLI, the MustGather command displays an output files location. The default location of the output files is the `./must_gather/<issue>_logs` directory.

Part of the MustGather output is an archived file (with a `.tgz` extension), which includes all of the collected diagnostic results. This file can be downloaded easily and analyzed by experienced Guardium users and administrators through the **fileserver** utility. If the issue requires further review by the Guardium support team, upload this `.tgz` file to the RETAIN® (PMR) system.

In Figure 7-47, you can see that the **sniffer_issues** MustGather produces the `sniffer_20130531.tgz` archive file.

Figure 7-48 on page 291 shows the **fileserver** view of the **sniffer_issues** MustGather output, including the `.tgz` file.

ID	Date	Permissions	File Name
4096	Fri May 31 11:03:23 2013	rw-xr-xr-x	must_gather/sniffer_logs
10362	Fri May 31 11:03:04 2013	rw-r--r--	must_gather/sniffer_logs/snif_buf_usage.txt
4080	Fri May 31 11:03:14 2013	rw-r--r--	must_gather/sniffer_logs/messages
5564	Fri May 31 11:02:54 2013	rw-r--r--	must_gather/sniffer_logs/sniffer_info_output.txt
2739	Fri May 31 11:03:14 2013	rw-r--r--	must_gather/sniffer_logs/snif_status.txt
1090	Fri May 31 11:03:14 2013	rw-r--r--	must_gather/sniffer_logs/snif_nanny_messages.txt
4510638	Fri May 31 11:03:14 2013	rw-r--r--	must_gather/sniffer_logs/snif_stderr.txt
264	Fri May 31 11:03:14 2013	rw-r--r--	must_gather/sniffer_logs/sniff_buffer_usage.log
2416	Fri May 31 11:03:14 2013	rw-r--r--	must_gather/sniffer_logs/db_output.txt
67429	Fri May 31 11:03:23 2013	rw-r--r--	must_gather/sniffer_logs/sniffer.20130531.tgz
25443	Fri May 31 11:03:16 2013	rw-r--r--	must_gather/sniffer_logs/SOFTWARE_TAPS.txt
258	Fri May 31 11:03:14 2013	rw-r--r--	must_gather/sniffer_logs/iptraf_output.txt
33629	Fri May 31 11:03:14 2013	rw-r--r--	must_gather/sniffer_logs/SNIFFER_BUFFER_USAGE.csv
75802	Fri May 31 11:03:22 2013	rw-r--r--	must_gather/sniffer_logs/SOFTWARE_TAP_DBS.txt

Figure 7-48 Viewing the .tgz file

The .tgz file includes diagnostic files, logs, and reports that are relevant to the sniffer operation and the generic information about the health of the appliance, operating system (system_output.txt) and onboard MySQL database (db_output.txt).

Figure 7-49 shows the content of the **sniffer_issues** MustGather output .tgz file.

Name	Size	Type	Date Modified
autopdzip		File Folder	5/31/2013 11:11 AM
db_output.txt	3 KB	Text Document	5/31/2013 11:03 AM
iptraf_output.txt	1 KB	Text Document	5/31/2013 11:03 AM
messages	4 KB	File	5/31/2013 10:25 AM
snif_buf_usage.txt	11 KB	Text Document	5/31/2013 11:03 AM
snif_nanny_messages.txt	2 KB	Text Document	5/31/2013 11:03 AM
snif_output.txt	1 KB	Text Document	5/31/2013 8:15 AM
snif_status.txt	3 KB	Text Document	5/31/2013 11:03 AM
snif_stderr.txt	4,405 KB	Text Document	5/31/2013 11:03 AM
sniff_buffer_usage.log	1 KB	Text Document	5/31/2013 8:15 AM
SNIFFER_BUFFER_USAGE.csv	33 KB	Microsoft Office Exc...	5/31/2013 11:03 AM
sniffer_info_output.txt	6 KB	Text Document	5/31/2013 11:02 AM
SOFTWARE_TAP_DBS.txt	75 KB	Text Document	5/31/2013 11:03 AM
SOFTWARE_TAPS.txt	25 KB	Text Document	5/31/2013 11:03 AM
system_output.txt	11 KB	Text Document	5/31/2013 11:02 AM

Figure 7-49 Content of .tgz file

The generic information about the appliance is collected in the `system_output.txt` and `db_output` files and is included in every MustGather command. The `system_output.txt` file contains unit type, a list of the installed patches, the results of the `top` command, system interfaces, memory, and I/O operation statistics.

Figure 7-50 shows a snippet of the content of `system_output.txt`.

```
#####
=====Output of the Unit Type:=====
#####
Managed Netinsp stap
#####
=====Output of the Show build:=====
#####
Build: 9.0
Release: 9.0.0_r45260_v90_1-e158-20121102_1546
Snif version: manual-gmachine-v90-r45260-20121102_1546
#####
=====Output of the Installed patches:=====
#####
P#      Who      Description      Request Time      Status
2      CLI      Guardium Patch Update (GPU) for 2013-03-14 20:10:13  DONE: Patch installation Succeeded.
#####
=====Output of the Net macs:=====
#####

eth0:   00:50:56:8A:3C:D4
eth1:   00:50:56:8A:3C:D5
eth2:   00:50:56:8A:3C:D6
eth3:   00:50:56:8A:3C:D7

#####
=====Output of the top command:=====
#####
ESC[ESC[2Jtop - 05:27:29 up 83 days, 13:29, 1 user, load average: 0.36, 0.13, 0.04
Tasks: 75 total, 1 running, 73 sleeping, 0 stopped, 1 zombie
Cpu(s): 0.5%us, 0.5%sy, 0.0%ni, 99.0%id, 0.0%wa, 0.0%hi, 0.0%si, 0.0%st
Mem: 4147600k total, 3772080k used, 375520k free, 290040k buffers
Swap: 1437776k total, 0k used, 1437776k free, 1098792k cached

  PID   USER   PR    NI  VIRT   RES    SHR  S    %CPU  %MEM     TIME  COMMAND
    1    root   15     0  2160   652    560  S    0.0   0.0   0:03.16  init
    2    root   RT    -5     0     0     0  S    0.0   0.0   0:00.00  migration/0
    3    root   34    19     0     0     0  S    0.0   0.0   0:00.44  ksoftirqd/0
    4    root   10    -5     0     0     0  S    0.0   0.0   120:19.48  events/0
    5    root   10    -5     0     0     0  S    0.0   0.0   0:00.00  khelper
    6    root   10    -5     0     0     0  S    0.0   0.0   0:00.00  kthread
```

Figure 7-50 `system_output.txt`

The database diagnostic file `db_output.txt` contains the size and total usage of the database, currently running queries, and a list of the most often used tables.

Figure 7-51 shows content of db_output.txt.

```

#####
=====Output of the Top 20 tables:=====
#####

```

Table Size (M)	I/D %	Est. Rows	Name
278822	0	N/A	GDM_CONSTRUCT_TEXT
37056	0	N/A	MONITOR_VALUES
3159	7	2438244	REPORT_RESULT_DATA_ROW
2883	0	N/A	GDM_EXCEPTION
2819	0	N/A	GDM_SESSION
2328	0	N/A	GDM_CONSTRUCT_INSTANCE
1363	0	7241730	GDM_FIELD
394	0	N/A	MONITOR_VALUES_CHANGED_COLUM
326	0	1550725	GDM_OBJECT
98	0	984512	GDM_SENTENCE
53	0	76052	GDM_CONSTRUCT
31	0	102814	GDM_ACCESS
21	0	N/A	GDM_CONSTRUCT_VALUES
7	18	63517	DB_ERROR_TEXT
6	0	N/A	DATASOURCE_VERSION_HISTORY
6	0	N/A	GDM_FLAT_LOG
4	0	N/A	MESSAGE
3	4	17632	CUSTOM_TABLE_RUN
3	0	N/A	GUARD_USER_ACTIVITY_AUDIT
2	0	1723	SNIFFER_BUFFER_USAGE

Figure 7-51 Content of db_output.txt

This diagnostic information is helpful in preventing bottlenecks of the system or database operation. As an example, in Figure 7-51, the large size of the GDM_CONSTRUCT_TEXT table (Full SQL) can be an alert to the Guardium administrator as a potential need to further tighten the filtering of the Log Full Details policy rules.

In addition to the generic information, MustGather collects issue-specific diagnostic information. For example, in the case of sniffer issues, MustGather collects buffer usage report output, syslog output, nanny process messages, and basic S-TAP configuration.

Note: MustGather collects log data for the last two days only. Therefore, it is important to run the diagnostic gathering procedure during or right after the issue occurrence.

Starting MustGather through Guardium GUI

Guardium v9.0p50 introduces an option to run MustGather procedure directly from Guardium GUI. You can access this option by clicking **Administration Console** → **Support Information Gathering**, as shown in Figure 7-52 on page 294.




Figure 7-52 Running MustGather from GUI

By using the Support Information Gathering window, users can collect the diagnostic data and send an email with the output directly to a specified recipient. To use this option, the user must choose an email address in the “email” field.

GUI users can also schedule diagnostic gathering to run at a specific preset date and time. This option should simplify collecting diagnostic information during off hours and weekends.

By using the GUI MustGather window, users also can submit several diagnostic commands together. (The diagnostics procedures are run sequentially.) All output files are sent to the email recipient upon completion of entire process.

The email includes the output packages of all submitted diagnostic procedures and optional other parameters, such as PMR number and Issue Description.

Apart from email, all diagnostic result outputs are available for download by clicking **Administration control** → **Support Information Results**, as shown in Figure 7-53.




Figure 7-53 Support Information Results window

To download any of the MustGather outputs, click the corresponding diskette line and then click **Save** in the pop-up window.

You also can start MustGather remotely from Central Manager appliance. This option allows you to submit or schedule diagnostic testing remotely on single or multiple managed units simultaneously, as shown in Figure 7-54 on page 296.

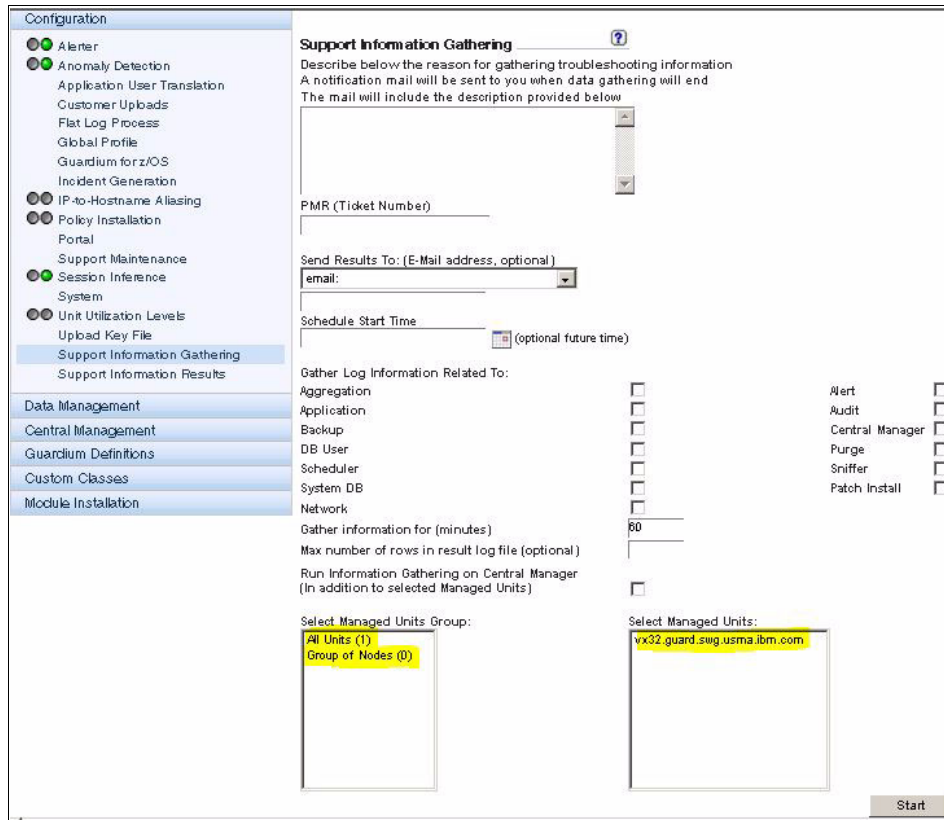


Figure 7-54 Starting MustGather remotely

Starting MustGather by using the `grdapi` command

Guardium provides a standard `grdapi` command interface for scripting diagnostic collection, as shown in the following command syntax:

```
grdapi <must_gather_api> [parameter1=value1] [parameter2=value2] ...
```

The following parameters are available:

```
commandsList (String) - MustGather type, for example, "sniffer_issues".
description (String)
duration (Integer)
emailDestination (String)
invokingUser (String)
maxLength (Integer)
pmrNumber (String)
start (Date)
timestamp (Date)
api_target_host (String)
```

For example:

```
grdapi must_gather_api commandsList=sniffer_issues invokingUser=JohnD
description=SnifferRestarts emailDestination=john@mydomain.com
duration=10 maxLength=10 timestamp="2013-09-10" start="2013-09-10"
pmrNumber=1231231234
```

Support analyze commands

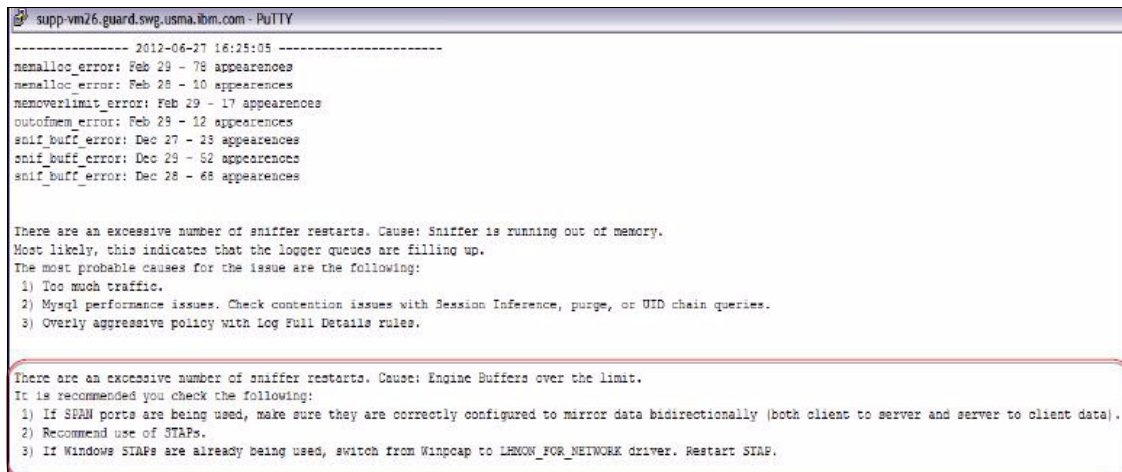
Guardium introduced the following class of the commands recently to automatically check potential issues and generate warnings to the user diagnostic testing to proactively prevent potential issues:

- ▶ **support analyze tap_property**
- ▶ **support analyze sniffer**

support analyze tap_property

This command analyzes the value of fields and specific field combinations for SOFTWARE_TAP_PROPERTY and SOFTWARE_TAP_DB_SERVER tables to identify potential issues with S-TAP configuration.

Figure 7-55 shows an example of the output of the **support analyze tap_property** command.



```
supp-vm26.guard.svg.usma.ibm.com - PuTTY
----- 2012-06-27 16:25:05 -----
memalloc_error: Feb 29 - 78 appearances
memalloc_error: Feb 28 - 10 appearances
memoverlimit_error: Feb 29 - 17 appearances
outofmem_error: Feb 29 - 12 appearances
snif_buff_error: Dec 27 - 28 appearances
snif_buff_error: Dec 29 - 52 appearances
snif_buff_error: Dec 28 - 68 appearances

There are an excessive number of sniffer restarts. Cause: Sniffer is running out of memory.
Most likely, this indicates that the logger queues are filling up.
The most probable causes for the issue are the following:
1) Too much traffic.
2) Mysql performance issues. Check contention issues with Session Inference, purge, or UID chain queries.
3) Overly aggressive policy with Log Full Details rules.

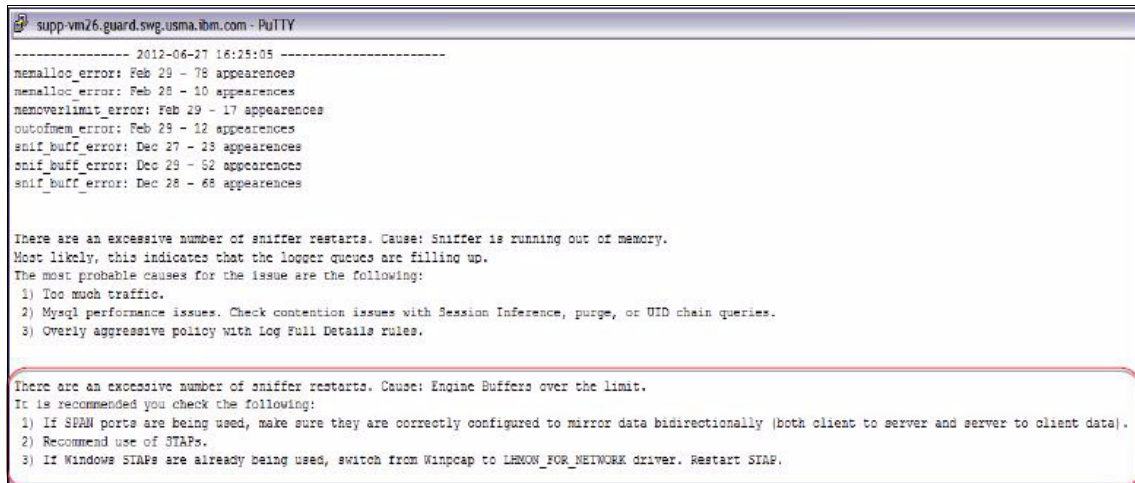
There are an excessive number of sniffer restarts. Cause: Engine Buffers over the limit.
It is recommended you check the following:
1) If SIEM ports are being used, make sure they are correctly configured to mirror data bidirectionally (both client to server and server to client data).
2) Recommend use of STAPs.
3) If Windows STAPs are already being used, switch from Winpcap to LINUX_FOR_NETWORK driver. Restart STAP.
```

Figure 7-55 The support analyze tap_property command output

support analyze sniffer

This command scans recent sniffer output messages for potential issues that are related to the sniffer operation and stability.

Figure 7-56 shows an example of the output of the **support analyze sniffer** command.



```
supp-vm26.guard.svg.usma.ibm.com - PuTTY
----- 2012-06-27 16:25:05 -----
memalloc_error: Feb 29 - 78 appearances
memalloc_error: Feb 28 - 10 appearances
memoverlimit_error: Feb 29 - 17 appearances
outofmem_error: Feb 29 - 12 appearances
sniff_buff_error: Dec 27 - 28 appearances
sniff_buff_error: Dec 29 - 52 appearances
sniff_buff_error: Dec 28 - 68 appearances

There are an excessive number of sniffer restarts. Cause: Sniffer is running out of memory.
Most likely, this indicates that the logger queues are filling up.
The most probable causes for the issue are the following:
1) Too much traffic.
2) Mysql performance issues. Check contention issues with Session Inference, purge, or UID chain queries.
3) Overly aggressive policy with Log Full Details rules.

There are an excessive number of sniffer restarts. Cause: Engine Buffers over the limit.
It is recommended you check the following:
1) If SPAN ports are being used, make sure they are correctly configured to mirror data bidirectionally (both client to server and server to client data).
2) Recommend use of STAPs.
3) If Windows STAPs are already being used, switch from Winpcap to LEMON_FOR_NETWORK driver. Restart STAP.
```

Figure 7-56 The support analyze sniffer command output

7.3.2 Agents diagnostic tests

The **guard_diag** (for UNIX system) and **diag.bat** (for Windows system) are utilities that facilitate the collection of diagnostic information for S-TAP issues. You can run **guard_diag** and **diag.bat** directly from the database server as script or from Guardium appliance GUI.

UNIX and Linux S-TAP diagnostic tests

You can run **guard_diag** on all Guardium supported, UNIX based operating system platforms. Introduced in v8.2, **guard_diag** is included in S-TAP installation packages for all consecutive versions.

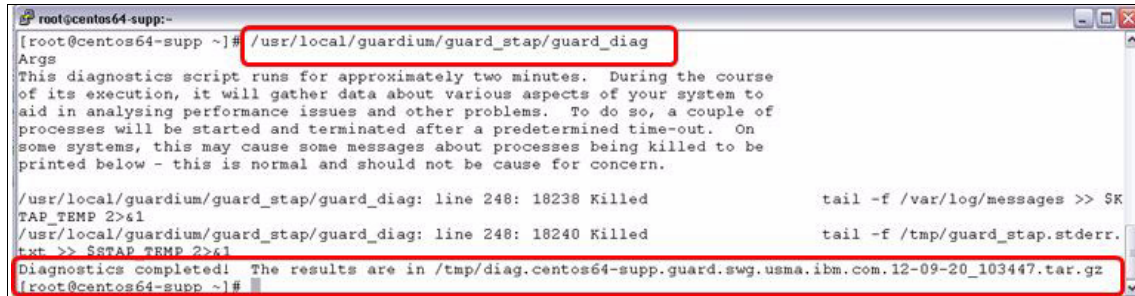
The **guard_diag** command is SQL Guard version independent. It does not require S-TAP to run or even to be installed on the system. However, the diagnostic information that is collected without active S-TAP is limited to corresponding system information only.

Note: In most cases, Guardium support requires **guard_diag** output that is collected while S-TAP is up and running. It is best to collect this diagnostic data during the time frame when the specific issue was reproduced.

By default, the **guard_diag** script is in the S-TAP installed directory. From the database server, you can run **guard_diag** as shown in the following example:

```
./guard_diag [output_dir] [-ktap_verbose_logging] [-l level] [-s duration]
```

Figure 7-57 shows running **guard_diag** from a database server directly.



```
root@centos64-supp:~# /usr/local/guardium/guard_stap/guard_diag
Args
This diagnostics script runs for approximately two minutes. During the course
of its execution, it will gather data about various aspects of your system to
aid in analysing performance issues and other problems. To do so, a couple of
processes will be started and terminated after a predetermined time-out. On
some systems, this may cause some messages about processes being killed to be
printed below - this is normal and should not be cause for concern.

/usr/local/guardium/guard_stap/guard_diag: line 248: 18238 Killed          tail -f /var/log/messages >> $K
TAP_TEMP 2>41
/usr/local/guardium/guard_stap/guard_diag: line 248: 18240 Killed          tail -f /tmp/guard_stap.stderr.
.txt >> $STAP_TEMP 2>41
Diagnostics completed! The results are in /tmp/diag.centos64-supp.guard.swg.usma.ibm.com.12-09-20_103447.tar.gz
root@centos64-supp ~]#
```

Figure 7-57 Running **guard_diag**

By default, the script output is placed in the `/tmp` directory. You can specify an output directory as an argument. All collected data is combined into a single `.tar` file, as shown in the following example:

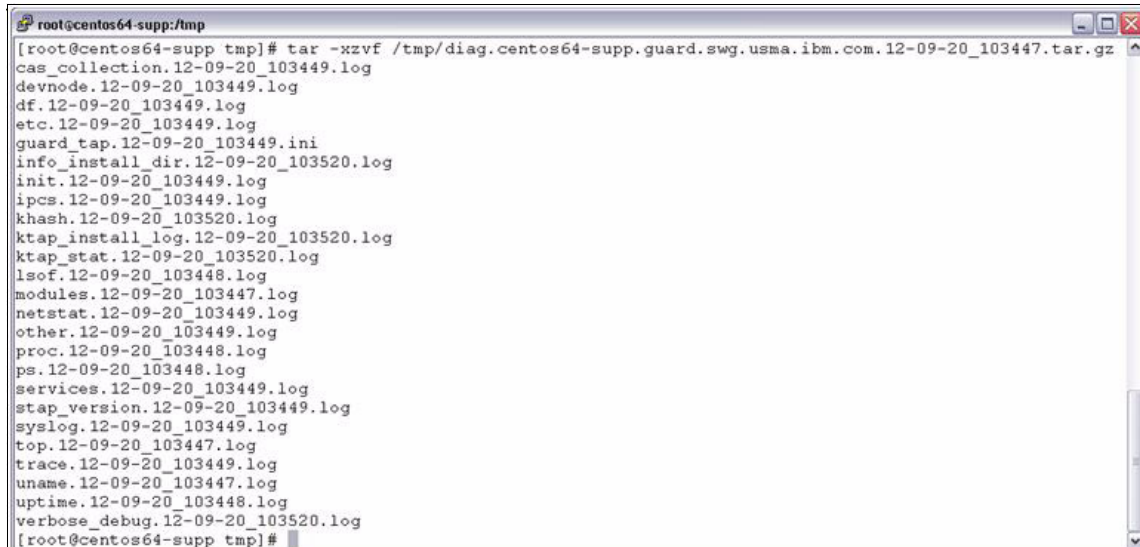
```
/var/tmp/diag.<hostname>.<date>.tar.gz
```

Typical output of the **guard_diag** utility contains the following Guardium and system operational logs, configuration details, and some platform-specific information:

- ▶ The **uname -a**
- ▶ List of installed kernel modules
- ▶ The **top** output (or its equivalent)
- ▶ Processor number and type
- ▶ The **Lsof** output
- ▶ The **netstat** output
- ▶ Disk free statistics
- ▶ The **uptime** output
- ▶ The **ps -ef** output
- ▶ Copy of `/etc/services`
- ▶ Contents of `/etc/inittab`
- ▶ Platform-specific information (release, memory, boot information, and so on)

- ▶ S-TAP specific data (version, guard_tap.ini, KTAP statistics, KTAP hash contents, and so on)
- ▶ GIM log and configuration files (if GIM is installed)
- ▶ CAS-specific data (if CAS installed)

Figure 7-58 shows the content of a `guard_diag` output file.



```
root@centos64-supp:/tmp
[root@centos64-supp tmp]# tar -xzf /tmp/diag.centos64-supp.guard.swg.usma.ibm.com.12-09-20_103447.tar.gz
cas_collection.12-09-20_103449.log
devnode.12-09-20_103449.log
df.12-09-20_103449.log
etc.12-09-20_103449.log
guard_tap.12-09-20_103449.ini
info_install_dir.12-09-20_103520.log
init.12-09-20_103449.log
ipcs.12-09-20_103449.log
khash.12-09-20_103520.log
ktap_install_log.12-09-20_103520.log
ktap_stat.12-09-20_103520.log
lsdf.12-09-20_103448.log
modules.12-09-20_103447.log
netstat.12-09-20_103449.log
other.12-09-20_103449.log
proc.12-09-20_103448.log
ps.12-09-20_103448.log
services.12-09-20_103449.log
stap_version.12-09-20_103449.log
syslog.12-09-20_103449.log
top.12-09-20_103447.log
trace.12-09-20_103449.log
uname.12-09-20_103447.log
uptime.12-09-20_103448.log
verbose_debug.12-09-20_103520.log
[root@centos64-supp tmp]#
```

Figure 7-58 `guard_diag` output

You also can start diagnostic testing directly from Guardium GUI on the appliance. When you are starting from the GUI, the output is placed in the `/var/tmp` directory.

Figure 7-59 shows the `guard_diag` starting from GUI.




Figure 7-59 Starting `guard_diag` from GUI

Windows S-TAP diagnostics

The S-TAP diagnostic script, `diag.bat`, for Windows is packaged and distributed with S-TAP installers, starting with Guardium V9. You can run this diagnostic routine on any supported Guardium platform. This routine is Guardium version independent.

To run the Windows diagnostic routine on a database server, double-click `diag.bat` in the installation directory, as shown in Figure 7-60 on page 302.




Figure 7-60 Starting Windows diagnostic routine

The output of the **diag.bat** is placed in the `diag` subdirectory. The output is compressed into a single `.zip` file. Diagnostic data is logically grouped among following files:

- ▶ `stap.txt`
- ▶ `tasks.txt`
- ▶ `system.txt`
- ▶ `evtlog.txt` or `evtlog2008.txt`
- ▶ `reg.txt`

Similar to `guard_diag`, the output of the `diag.bat` utility contains Guardium and system operational logs, configuration details, and platform-specific information. The follow diagnostic information is collected by `diag.bat`:

- ▶ Content of `guard_tap.ini`
- ▶ The Guardium S-TAP installation log
- ▶ All running tasks
- ▶ List of all installed kernel drivers

- ▶ Operating system information that is collected from the system information utility
- ▶ The **ipconfig /all**
- ▶ The **netstat -nao**
- ▶ The **ping** and **tracert** results from the database server to the Guardium appliance
- ▶ Processor usage for `guardium_stapr`, overall system processor usage
- ▶ The `guardium_stapr` process handle count and memory usage
- ▶ Event log messages that are generated by S-TAP
- ▶ System event log messages of error and warning type

The diagnostic information that is collected by the **diag.bat** utility is most efficient when it is collected during problem occurrence or reproduction. When an issue cannot be reproduced, it is best to run the **diag.bat** utility immediately after the issue occurs.

Note: The **diag.bat** output is required by support for all S-TAP related issues. Providing this output ahead of time also helps to speed up support's investigation.

7.4 Restoring audit data for forensic analysis

The audit data is recorded in the appliance internal database constantly. Customers are advised to keep only a short history of the audit data on the appliance. The archive and purge mechanism is configured and scheduled during implementation the stage to run nightly to archive the last day's data and purge data that is older than certain number of days. This practice helps to keep only certain number of days' data on appliance to save disk space and boost appliance performance.

Note: The Guardium archive function creates signed, encrypted files that cannot be altered. Do not change the names of the generated archive files. The archive and restore operations depend on the file names that are created during the archiving process.

There are times when there is a need to review historical data. When such a need arises, archive files can be restored on the appliance to restore data.

The audit data is stored in a normalized way in an internal database of an appliance. The audit data that changes constantly is referred as *dynamic* data and the audit data that stays relatively constant is referred as *static* data. The user login time is a good example of dynamic data. It is unique for every user for each login. User name is an example of static data. It stays the same for every login of this user to the database.

To save space on storage servers, Guardium uses an incremental archive strategy. The dynamic audit data is archived when it is observed. Static audit data is archived only when the data is observed in the first time. For example, the login time is archived every time that each user logs in to the database server, but not the user name. This incremental approach to move data to data storage reduces the size of archive files dramatically. The trade-off is that a single archive file might not contain all of the audit data that is needed to be restored back to the appliance. To compensate for this trade-off, the archive process generates a full (not incremental) archive file the first time the archive process runs and then the first day of every month.

Note: On Guardium version 8.2 and later, the incremental archive strategy is used only for the archive that is taken from a collector. Full archive is always taken from an aggregator for static data to simplify the archive restore process.

Restored audit data can be viewed as the regular audit data by using interactive or audit process reports.

7.4.1 Restore strategies

Note: The Guardium archive function creates signed, encrypted files that cannot be tampered with. Do not change the names of the generated archive files. The archive and restore operations depend on the file names that are created during the archiving process.

The archive and export activities use an operating system encryption algorithm to create encrypted data files. The restore system must have the encryption algorithm that the archive system used to restore the archived data.

The archived files can be restored by retrieving them through the archive catalog. The Guardium catalog tracks where every archive file is sent so that the archive files can be retrieved and restored with minimal effort at any point in the future. A separate catalog is maintained on each appliance. A new record is added to the catalog whenever the appliance archives data or results.

To restore data to a new appliance, you must first add entries of the files to be restored to the catalog manually or by using catalog export and import function.

Scenario 1: Restoring a few days of recent data

Assume that you have two weeks of data on an appliance and you want to review a couple of days' worth of data from a prior month. Table 7-1 lists the date and status of the archived data.

Table 7-1 Archived data status

Date	Status
May 17	Currently on the appliance (today)
May 16	Currently on the appliance
.....
May 4	Currently on the appliance
May 3	Oldest day on the appliance
.....
April 27	Need to restore
April 28	Need to restore

To restore the data of April 27 and April 28, click **Administration Console** → **Data Management** → **Data Restore**, as shown in Figure 7-61. Specify the date range that you want to restore and optionally the host name of appliance where archive files were taken, and click **Search**.




Figure 7-61 Enter dates range to search catalog

The Data Restore function displays a list of available archive files in the searching range, as shown in Figure 7-62.

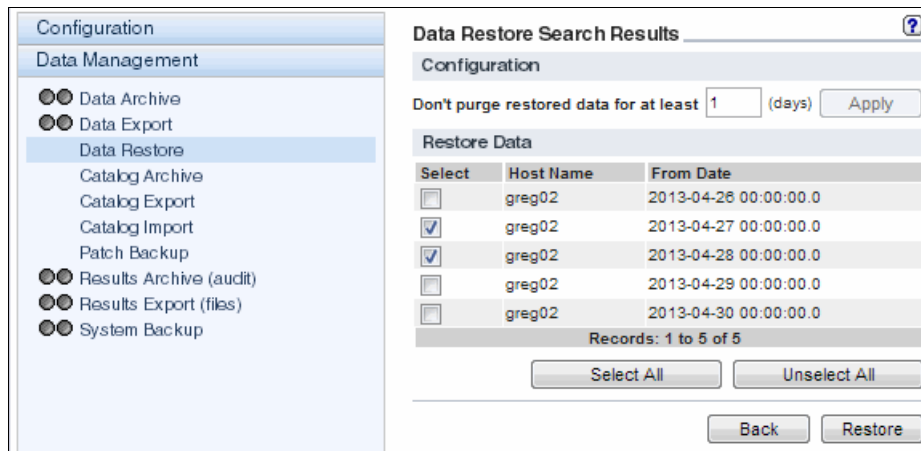


Figure 7-62 Archive file list

Select the files that you want and then click **Restore**.

Scenario 2: Restoring archived data on an empty appliance

To avoid interference with current data that is operating, you can restore archived data to a stand-alone appliance that has no other audit data on it and is designated for restoring and reviewing historical data.

Because the target appliance is not where the archive files are taken, the archive catalog does not have the archive file entries. To add records manually to the catalog, click **Administration Console** → **Data Management** → **Catalog Archive**, specify date range, and then click **Search**.

Review the returned results, as shown in Figure 7-63 on page 307. If the archive files that you want are not in the list, click **Add** to add entries manually.

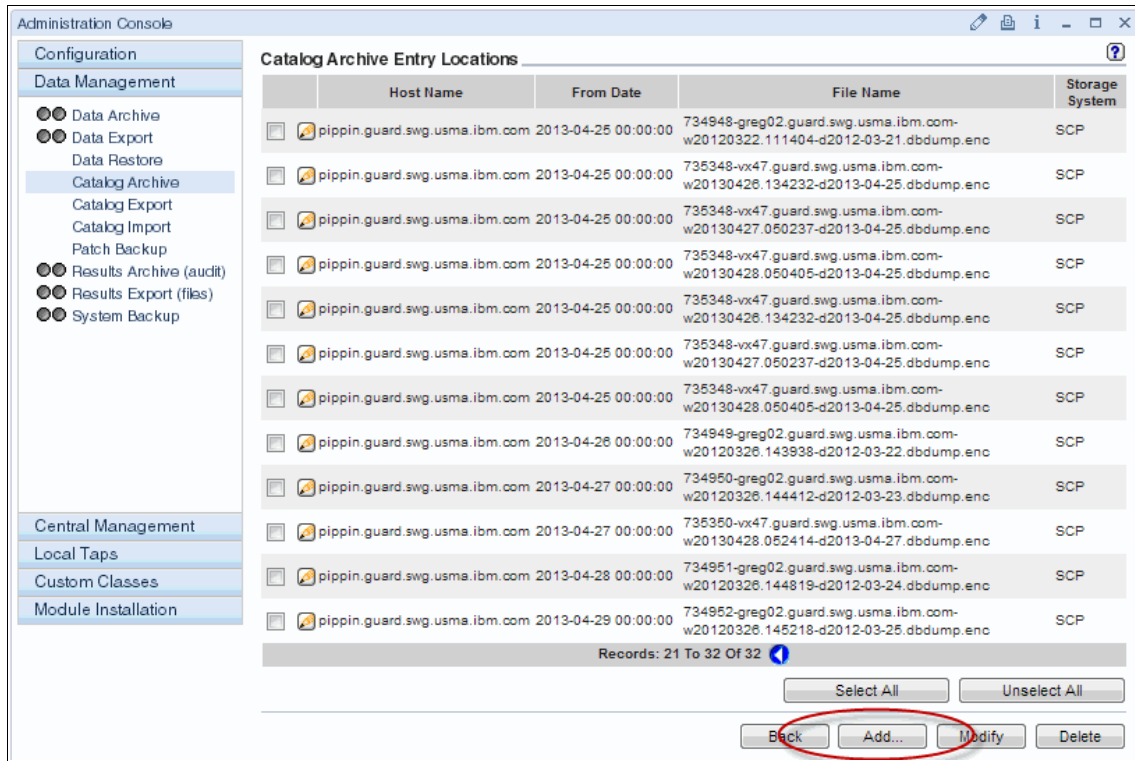


Figure 7-63 Existing catalog entries

In the Add location panel (see Figure 7-64), enter the required information and click **Save** to add the new entry to catalog.

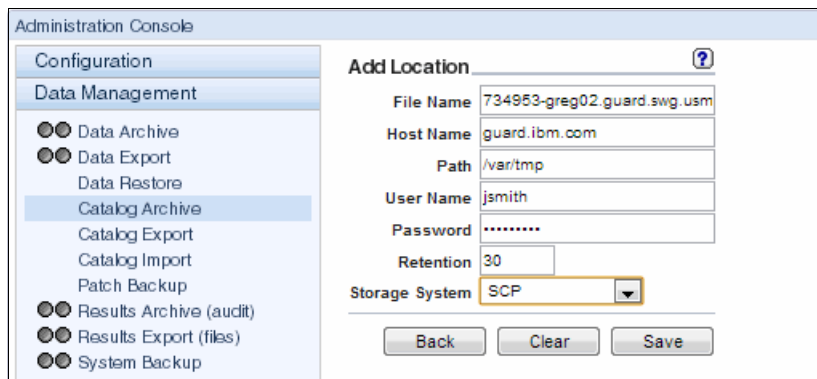


Figure 7-64 Adding new entry to catalog

Another method to add entries to the archive catalog is by using the catalog export and import function.

To export catalog entries from the appliance where the archive files were taken, click **Administration Console** → **Data Management** → **Catalog Export**, select file entries that you want to import, and then click **Export**, as shown in Figure 7-65.

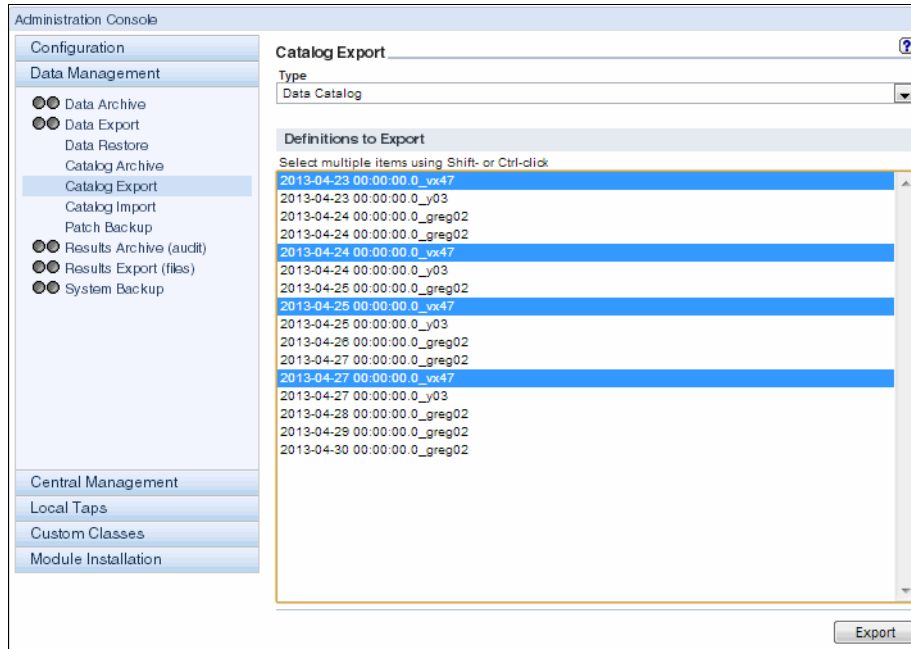


Figure 7-65 Exporting the archived file entries

The tool generates a file that can be imported to the appliance where data is restored. To import the file, click **Catalog Import**, upload the catalog entry file, and select the green check box to import the previously exported entries, as shown in Figure 7-66 on page 309.

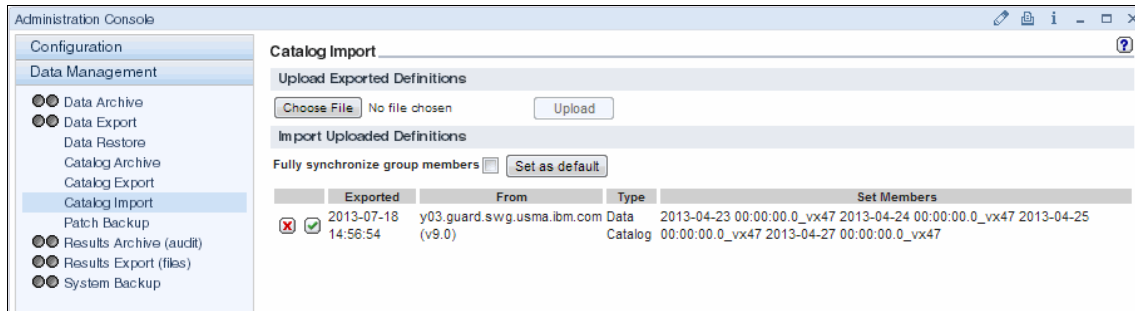


Figure 7-66 Importing entries

After the entries are added to catalog, use the data restore procedure that is described in “Scenario 1: Restoring a few days of recent data” on page 305 to restore the archive files.

Note: The archive files from collectors or from Guardium system before Version 9.0 are archived by using incremental archive strategy. When you are restoring such archive files, you must always start from the first of the month and work your way up to the days you need. For example, if you want to restore data from 13, 14, and 15 of May from a collector, start with May 1, 2, 3, and so on until May 15.

If you restore archived files from aggregator from V8 or later, you can restore only the days that you want; that is, 13, 14 and 15 of May.

Scenario 3: Restoring a few months worth of data on an empty appliance

When you want to restore data for a longer period (for example, three months), restoring file by file as described in “Scenario 1: Restoring a few days of recent data” on page 305 can be time-consuming. An alternative method is the use of the backup files if they are available.

Restoring data from a backup file gives you data only on the period that the data was backed up. For example, if the data retention period of the appliance is two months, each backup file contains only two months worth of data. If you want three months worth of data, you must restore the extra month data that you need from the archived files after you restore the backup data. If the data retention period is one month, you must restore the other two months of data that you need from the archived files after the backup data is restored.

Note: We can use only one backup file in such a restore scenario because restoring from backup file overrides all the existing data.

7.4.2 Restore audit process result sets

The following archive operations are available on the Administration Console under Data Management:

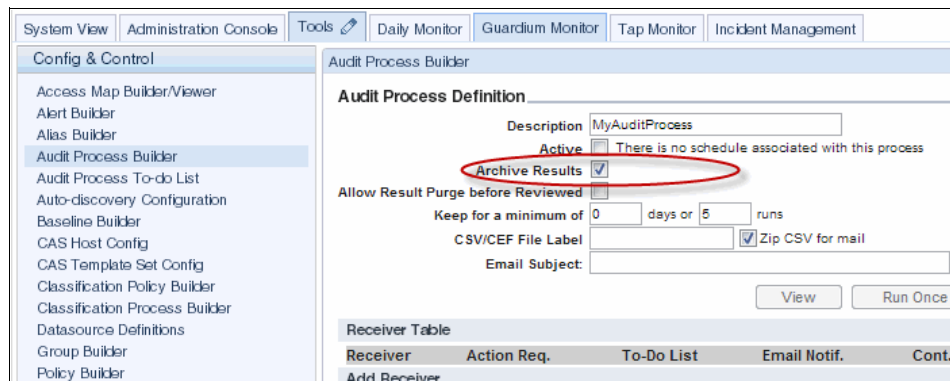
- ▶ Data archive
- ▶ Results archive

The results archive that is taken through the Audit Process Build includes the following information:

- ▶ Audit tasks results from the following sources:
 - Reports
 - Assessment tests
 - Entity audit trail
 - Privacy sets
 - Classification processes
- ▶ View
- ▶ Sign-off trails
- ▶ Accommodated comments from workflow processes

The results sets are purged from the system according to the workflow process definition.

To archive the results set for a particular audit process, select **Archive Results** in Audit Process Definition of the Audit Process Builder, as shown in Figure 7-67.



The screenshot shows the 'Audit Process Builder' interface. The left sidebar contains a list of tools, with 'Audit Process Builder' selected. The main area is titled 'Audit Process Definition' and contains the following fields and controls:

- Description: MyAuditProcess
- Active: There is no schedule associated with this process
- Archive Results: (circled in red)
- Allow Result Purge before Reviewed:
- Keep for a minimum of: 0 days or 5 runs
- CSV/CEF File Label: Zip CSV for mail
- Email Subject:
- Buttons: View, Run Once

Below the configuration is a 'Receiver Table' with the following columns: Receiver, Action Req., To-Do List, Email Notif., and Cont. The first row is partially visible with the text 'Act Receiver'.

Figure 7-67 Setting result archiving

To have all selected results archived regularly, click **Administration Console** → **Data Management** → **Results Archive**, configure your archive process, and set up a schedule.

Note: The audit process results that must be signed are not archived until they are reviewed and signed.

The results sets can be restored only into the investigation center. You can set up an investigation center by creating a special investigation user account on a Guardium appliance.

Remember the following rules to correctly configure an investigation center for a user account:

- ▶ The user's last name must be set to the name of one of the three available investigation centers: INV_1, INV_2, or INV_3 (case-sensitive).
- ▶ The user account must have the INV role assigned to it.

Figure 7-68 shows setting an investigation account.

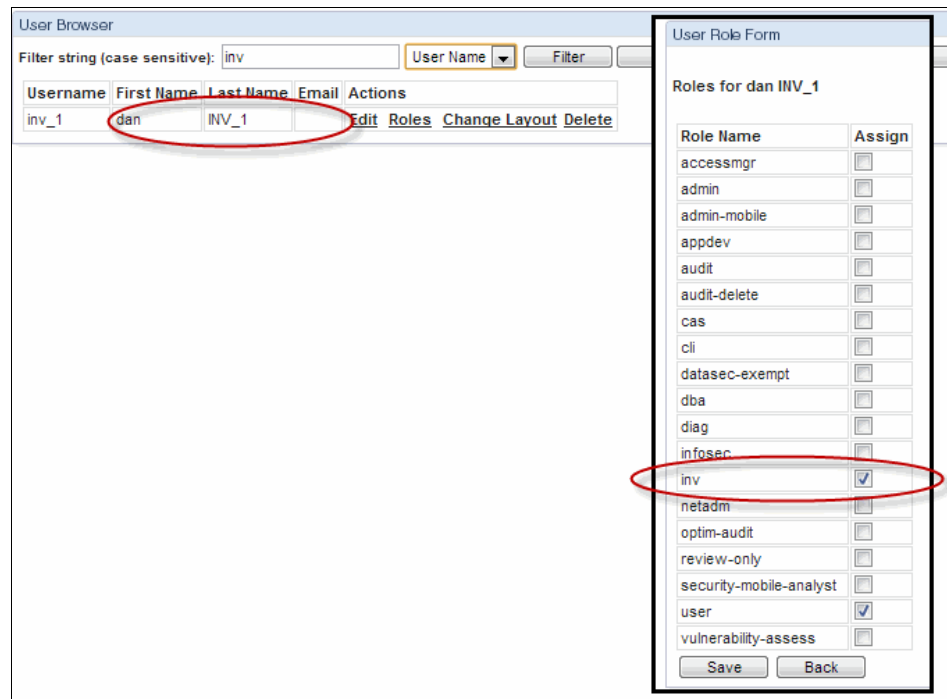


Figure 7-68 Setting an investigation account

Up to three investigation centers can be defined simultaneously on an appliance. Each center has its own internal database and does not interfere with each other. Also, data in the investigation center is not mixed up with the regular audit data. The investigation center also can be used to restore data archives.

For more information about setting up and using an investigation center, see the *Investigation Center* article on the online help at Guardium GUI, which is available at this website:

http://pic.dhe.ibm.com/infocenter/igsec/v1/index.jsp?topic=%2Fcom.ibm.guardium.software.app.install.doc%2FtopicsV90%2Fsoftware_appliance_installation_guide.html



Disaster recovery

In this chapter, we summarize the steps for disaster recovery in different scenarios.

This chapter includes the following topics:

- ▶ Overview
- ▶ Appliance recovery
- ▶ Appliance recovery steps

8.1 Overview

Guardium provides the following methods that can be used to back up data from appliance to a designated location:

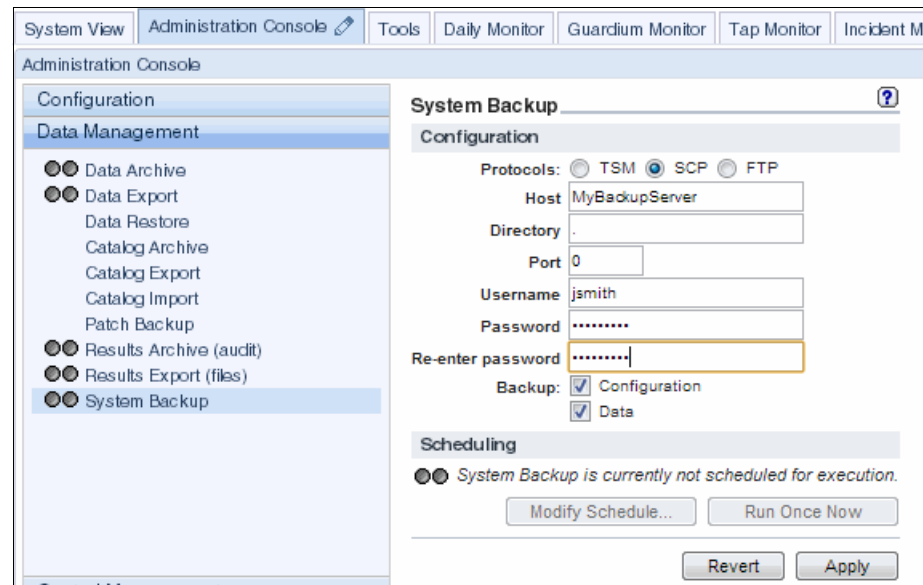
- ▶ Backup
- ▶ Archive

8.1.1 System backup and restore

Guardium internal database contains audit data and all the appliance system definitions, such as groups, queries, reports, audit processes, alerts, and policies. System backup is a snapshot of an existing internal database. When restored, it overrides the existing data.

Backup can be started from the user interface (UI) or from a command-line interface (CLI).

Figure 8-1 shows the System Backup user interface.



The screenshot displays the Administration Console interface for System Backup. The top navigation bar includes tabs for System View, Administration Console (active), Tools, Daily Monitor, Guardium Monitor, Tap Monitor, and Incident M. The left sidebar shows a tree view under Data Management, with System Backup selected. The main content area is titled 'System Backup' and contains the following sections:

- Configuration:**
 - Protocols: TSM, SCP, FTP
 - Host: MyBackupServer
 - Directory: .
 - Port: 0
 - Username: jsmith
 - Password:
 - Re-enter password:
 - Backup: Configuration, Data
- Scheduling:**
 - System Backup is currently not scheduled for execution.
 - Buttons: Modify Schedule..., Run Once Now

At the bottom right, there are buttons for Revert and Apply.

Figure 8-1 System Backup user interface

Figure 8-2 shows the CLI of the system backup utility.

```
collrb01.guard.swg.usma.ibm.com> backup system
  1. DATA
  2. CONFIGURATION

Please enter the number of your choice: (q to quit) 1

  1. SCP
  2. CONFIGURED DESTINATION
  3. Snapshot
  4. Existing Backup Copy to Snapshot location

Please enter the number of your choice: (q to quit) 1
Enter the destination backup host: MyBackupServer
Enter the destination username: jsmith
Enter the destination directory: /bck
Enter the password for jsmith@MyBackupServer ? *****

Enter the SCP port if you need to use a special port.
Enter "0" or press "Enter key" to use the default port:
Gathered all the information needed to perform the backup.
We are now going into auto pilot mode, please do not enter any
passwords or commands until this operation finishes.

Please wait, this may take some time.
gpg: checking the trustdb
gpg: 3 marginal(s) needed, 1 complete(s) needed, classic trust model
gpg: depth: 0 valid: 1 signed: 0 trust: 0-, 0q, 0n, 0m, 0f, 1u
spawn /usr/bin/scp /var/dump/turbine_backup.tgz jsmith@MyBackupServer:/bck/2013-07-23-1837-collrb
01.guard.swg.usma.ibm.com-SQLGUARD_DATA.tgz
```

Figure 8-2 System backup CLI

System backup includes the following variations:

- ▶ Data backup: Contains a snapshot of the entire database, audit data, and all system definitions.
- ▶ Configuration backup: Contains only system definitions and no audit data; therefore, it is smaller in size.

Backup should be configured and scheduled during the implementation stage to run on regular bases.

To restore backup file, run the **system restore** CLI command, as shown in Figure 8-3.

```
collrb01.guard.swg.usma.ibm.com> restore system

1. SCP
2. FTP
3. TSM
4. CENTERA
5. Snapshot

Please enter the number of your choice: (q to quit) 1
Enter the backup host: MyBackupServer
Enter the backup host username: jsmith
Enter the remote directory: /bck
Enter the remote file name (file name may use wildcard *): /bck/2013-07-23-1837-collrb01.guard.sw
g.usma.ibm.com-SQLGUARD_DATA.tgz
Enter the password for jsmith@MyBackupServer ? *****

Enter the SCP port if you need to use a special port.
Enter "0" or press "Enter key" to use the default port:

Attempting to retrieve file. It may take time. Please wait.
During the transfer, please do not enter the password or answer any questions.

Starting transfer, please wait.
spawn /usr/bin/scp jsmith@MyBackupServer:/bck//bck/2013-07-23-1837-collrb01.guard.swg.usma.ibm.co
m-SQLGUARD_DATA.tgz /var/dump/restore/tmp/
```

Figure 8-3 Restore command-line interface

Restoring system backup overrides existing data on an appliance. Backup is important and often is used when a lost appliance must be recovered.

8.1.2 Data archive and restore

Customers usually keep only a small amounts of historical data on the appliance to save disk space and boost performance. Typically, customers keep no more than one or two months worth of audit data on an appliance (in some cases, as little as a few days). At the same time, auditors might require to keep audit data available for a few years. Guardium provides the mechanism to back up audit data in chunks of one day and store them on a remote location. This is known as *archiving*.

The archive process should be configured during implementation to run on daily. The archive files can be used for data restoration for forensic purposes when a limited number of days data needs to be restored. Restoring the archive file does not override the existing data and it can be used to restore data on an appliance where some data exists.

For more information about data archive and restore processes, see 7.4, “Restoring audit data for forensic analysis” on page 303.

8.1.3 Storage location

Guardium supports EMC Centera storage system and IBM Tivoli Storage Manager. Centera or Tivoli Storage Manager as a storage destination can be enabled by running the following CLI command:

```
>store storage-system Centera <backup|archive> on
```

or

```
>store storage-system TSM <backup|archive> on
```

For more information about Tivoli Storage Manager or Centera configuration instructions, see the appliance help (select ? from upper right side when you are logged in to the GUI) Archive, Purge, and Restore article.

Before you restore from Tivoli Storage Manager, the `dsm.sys` configuration file must be uploaded to the Guardium appliance by using the CLI command.

Before you restore from EMC Centera, a PEA configuration file must be uploaded the Guardium appliance through the Data Archive panel.

8.2 Appliance recovery

To complete the appliance recovery process, the following information is required:

- ▶ New physical or virtual appliance.
- ▶ ISO image of Guardium software.
- ▶ The same patches that were installed on the appliance when the last system backup was taken.
- ▶ Latest system backup files.
- ▶ Daily archive files.
- ▶ License, SSL certificates, and any settings that must be set manually.

The following items are not backed up and must be installed manually to complete disaster recovery process:

- ▶ License: License is not installed by backup restore; therefore, it must be installed manually.
- ▶ SSL Certificate (optional): SSL Certificate is not backed-up; therefore, it must be installed manually.

- ▶ Language (optional): Use the CLI command **store language** to change from English (default).
- ▶ Network Time Protocol (NTP) settings (recommended): Use the CLI commands **store system ntp server** and **store system ntp state** to complete NTP server configuration.
- ▶ Time zone (recommended): Use the CLI command **store system clock timezone** to configure system time zone.
- ▶ If you are using Tivoli Storage Manager or Centera as a storage location, configure the appliance to support it before you attempt to restore the backup files.

8.3 Appliance recovery steps

Appliances of different types (collectors, aggregators, and central manager) have different purposes and store data in different ways. Use a different recovery strategies for different appliances. The recovery strategy that is used also depends on whether the appliance that must be recovered is a stand-alone appliance or is used with other appliances.

In this section, we describe the detail disaster recovery steps for each appliance type and usage scenario.

8.3.1 Stand-alone collector

Complete the following steps for a single, stand-alone collector:

1. Use the appropriate ISO image to build an appliance.
2. Apply all of the patches to bring the appliance to the same patch level as it was when the last backup was taken.
3. Restore the data backup.
4. Restore the configuration backup.
5. Restore the archive files for missing days, as needed.
6. Apply the license.

8.3.2 Collector with Aggregator no CM

When you must recover a collector that aggregates data daily to the aggregator, complete the following steps:

1. Use the appropriate ISO image to build an appliance.
2. Apply all of the patches to bring the appliance to the same patch level as it was when the last backup was taken.

Data recovery is optional because all audit data exists on an aggregator.

3. Restore the configuration backup.

Restore archive files is optional because data exists on aggregator.

4. Apply the license.

8.3.3 Aggregator: Not centrally managed

Use the following recover steps when recovering an aggregator that is not centrally managed:

1. Use the appropriate ISO image to build an appliance.
2. Apply all of the patches to bring the appliance to the same patch level as it was when the last backup was taken.
3. Restore the data backup.
4. Restore the configuration backup.
5. Restore the archive files for missing days as needed.
6. Apply the license.

Alternatively, to restore data backup, you might consider aggregating data again from collectors, which depends on the number of collectors and the retention period requirements on the collectors and aggregator.

8.3.4 Centrally managed collector

For the collector that is centrally managed, the configuration data is stored on central manager and restore data on the collector is not required. Complete the following steps to complete disaster recovery for a managed collector:

1. Use the appropriate ISO image to build an appliance.
2. Apply all of the patches to bring the appliance to the same patch level as it was when the last backup was taken.

3. Restore the data backup. This step might be optional if there is an aggregator in this environment and data from the collector exists on the aggregator.
4. It is not required to restore the configuration backup because the definitions are pooled from central manager.
5. Shared secret must be set to enable communication with central manager. Shared secret is used to encrypt communication of the appliance with central manager. Use CLI command **store system shared secret** to complete this step.
6. (Optional) Restore archive files for missing days, as needed. This step is not required if you are not restoring data backup.
7. Register the newly built collector with central manager (licenses are pooled from central manager).

8.3.5 Centrally managed aggregator

For the centrally managed aggregator, the configuration data is stored on central manager and restoring the data on the aggregator is not required during disaster recovery. Complete the following steps to complete the disaster recovery for managed aggregator:

1. Use the appropriate ISO image to build an appliance.
1. Apply all of the patches to bring the appliance to the same patch level as it was when the last backup was taken.
2. Restore the data backup.
It is not required to restore the configuration backup because the definitions are pooled from central manager.
3. Shared secret must be set to enable communication with central manager. Shared secret is used to encrypt communication of the appliance with central manager. Use CLI command **store system shared secret** to complete this step.
4. Restore archive files for missing days, as needed.
5. Register with central manager (licenses are pooled from central manager).

8.3.6 Dedicated central manager (no data aggregation)

For recovery of dedicated central manager, complete the following steps:

1. Use the appropriate ISO image to build an appliance.
2. Apply all of the patches to bring the appliance to the same patch level as it was when the last backup was taken.
3. Restore the configuration backup.
4. Apply the license.



Upgrade best practices

This chapter provides methods and strategies of Guardium software upgrades for appliances and agents.

We describe the best practices for Guardium Version 8.2 to Version 9.1p100 upgrade process for 32-bit and 64-bit versions. We also describe the differences in the approaches for each environment.

This chapter includes the following topics:

- ▶ Content delivery for 32-bit and 64-bit systems
- ▶ Compatibility considerations for 32-bit versus 64-bit systems
- ▶ RAM memory considerations
- ▶ Enterprise upgrade strategy
- ▶ Appliance upgrade methods
- ▶ Agents upgrade methods
- ▶ Upgrading bundle monitoring and status validation

9.1 Content delivery for 32-bit and 64-bit systems

The Guardium Version V9.0p50 and V9.1p100 releases contain the following infrastructure changes for underline Red Hat operational system and MySQL database:

- ▶ Red Hat operational system changes
Guardium V9.1p100 provides customer with option to install Guardium appliance software on 32-bit and 64-bit Red Hat platforms.
- ▶ MySQL database changes
On 64-bit platforms, the internal MySQL database version is upgraded from V4 to V5 (V5.6.11).

Red Hat nor MySQL provide an upgrade path to move to their newer software versions. This limitation dictates the differences in upgrade approaches that are used for the following 32-bit versus 64-bit platforms:

- ▶ On 32-bit platform:
 - For V9.0 and V9.0p02 appliances, Guardium content is delivered through Guardium Patch Update (GPU) patch to provide an upgrade path to the new V9.1p100 (32-bit) version.
 - For V8.2 appliances, Guardium provides V8.2 to V9.1p100 (32-bit) upgrade bundle patch, which is applicable on V8.2 version appliance regardless of the patching level. The bundle path allows direct upgrade to V9.1p100 (32-bit) version level.
- ▶ On 64-bit platform:
 - For V9.0p50 64-bit appliance, Guardium content is delivered through V9.1p100 (64-bit) GPU patch to provide an upgrade path.
 - For new appliances or appliances before V9.0p50, content is delivered through V9.1p100 (64-bit) product ISO, which requires a rebuild of the appliances. No upgrade path is available for new 64-bit system installation.
 - Any system backup that is generated on V8.2 or later version can be restored on V9.1p100 32-bit and 64-bit system.

9.2 Compatibility considerations for 32-bit versus 64-bit systems

In addition to the infrastructure changes, to meet new stricter MySQL field naming conventions, Guardium must make several changes in the existing internal database table fields.

These changes introduce more Central Management and Aggregation limitations during upgrade transitional periods when managed or aggregated appliances have different Guardium versions or patch levels (mixed environment).

In this section, we describe the mixed environment compatibility limitations.

9.2.1 Central management

A V9.1p100 (64-bit) Guardium Central Manager cannot manage appliances that are built with a lower version than V9.0p50.

Figure 9-1 shows the compatibility between different versions of V8.2 and V9.0.

Appliance	Central Manager			
	64-bit 9.0p50 can manage		32-bit 9.0p50 can manage	
64-bit 9.0p50	yes		yes	
32-bit 9.0p50	yes		yes	
9.0	no		yes	
8.2	no		yes	

Figure 9-1 Central management compatibility

9.2.2 Aggregation

A 64-bit Guardium appliances cannot aggregate to any 32-bit Guardium aggregators. It can aggregate only to 64-bit level aggregator, as shown in Figure 9-2.

Appliance	Aggregator			
	64-bit 9.0p50 can aggregate		32-bit 9.0p50 can aggregate	
64-bit 9.0p50	yes		no	
32-bit 9.0p50	yes		yes	
9.0	yes		yes	
8.2	yes		yes	

Figure 9-2 Aggregation compatibility

9.2.3 Central management and aggregation combo appliances

In cases where the Central Manager also serves as an aggregator appliance, both limitations of Central Manager and Aggregator functions apply. Figure 9-3 shows a summary of the limitations of a Central Manager and Aggregation combo environment.









Appliance	Central Manager / Aggregator	
	64-bit 9.0p50 can manage and aggregate	32-bit 9.0p50 can manage and aggregate
64-bit 9.0p50	yes 	no 
32-bit 9.0p50	yes 	yes 
9.0	no 	yes 
8.2	no 	yes 

Figure 9-3 Combo environment compatibility

For customers that are upgrading to 32-bit Guardium systems, the limitations that are shown in Figure 9-3 are not relevant. This limitation is also not relevant to customers who do not use Central Manager and Aggregation combo appliance. However, customers who use Central Manager and Aggregation combination appliance and are planning to upgrade to 64-bit Guardium system, this limitation dictates the upgrade strategy for the entire corresponding enterprise environment. To maintain uninterrupted communication between appliances on different versions, customers are required to proceed with a two-step upgrade approach in which all of the appliances must be upgraded to V9.1p100 32-bit level before you proceed with the 64-bit implementation.

Figure 9-4 on page 327 shows another view of mixed environment compatibility limitations for Central Manager and Aggregator functionality.




Figure 9-4 Compatibility chart

The blue lines in Figure 9-4 represent communication compatibility between Central Manager and its managed units for different versions of Guardium solution.

9.3 RAM memory considerations

The following virtual memory requirements for post-V9.0p50 systems are different between 32-bit and 64-bit architectures:

- ▶ 32-bit architecture: Memory must be in the range of 8 GB - 16 GB.
- ▶ 64-bit architecture: Minimum memory requirement is 16 GB.

Note: Some new features are memory-intensive. To use these features, we strongly encourage customers to upgrade memory capacity within the recommended range. In a 64-bit architecture, consider at least 24 GB of RAM.

9.4 Enterprise upgrade strategy

Upgrading large enterprise environments requires thorough planning and preparation. During the initial stages of preparation, the Guardium administrator often must decide the following issues:

- ▶ Upgrade strategy and logistics
- ▶ Change control management
- ▶ Required personal availability
- ▶ Contingency planning

Choosing the right upgrade strategy is one of the most important decisions to ensure a smooth and successful upgrade. Strategy often must be adjusted to fit an operational flow of your organization and work around multiple known and unpredictable constraints. We also strongly suggest choosing the strategy that allows the quickest way to complete the upgrade of all of the appliances across the entire estate.

9.4.1 Change control management

More organizations today adopt extremely strict control policies over their test and production environments. These policies must be taken into the account when you are preparing for upgrade.

Implementation and upgrade of the Guardium solution usually requires involvement of Guardium administrators, DBAs, IT systems, and security personnel. Their presence and full cooperation are important during the implementation.

The upgrade process usually cannot be done simultaneously on all appliances and all S-TAPs. Therefore, it requires a multi-staged upgrade approach. During the transition period, the Guardium environment operates in hybrid mode with Version 8.2 and Version 9.0 Guardium software (mixed version mode).

Note: Hybrid stage refers to the transition period during the upgrade where customers have a hybrid Version 8.2 and Version 9.x Guardium solution.

9.4.2 Top-to-bottom rollout order

Upgrade the IBM InfoSphere Guardium environment in a top-to-bottom order. Top-to-bottom order is an important requirement to ensure continuous connectivity between different components of the Guardium solution. This means that the upgrade starts with Central Manager, then Aggregators, Collectors, and then all agents, as shown in Figure 9-5. This order is enforced to always keep corresponding higher-level appliances on the same or higher level of Guardium software.




Figure 9-5 To-to-bottom rollout order

9.4.3 Strategies

The general strategies of the Guardium enterprise solution can be represented by two models, Horizontal and Vertical, or some variations of mixture of these two.

Horizontal model

Through horizontal model, users often upgrade all of the appliances of the same type before moving forward with other appliances. Upgrade all of the Central Managers first, then upgrade all the aggregators, then all the collectors, and finish with agents upgrade, as shown in Figure 9-6 on page 330.




Figure 9-6 Horizontal approach

Vertical model

Figure 9-7 on page 331 shows the vertical upgrade approach. This method often requires the upgrade of single Central Manager (1 in the figure), followed by an aggregator that is managed by this Central Manager, then followed by all the collectors exporting to this aggregator. You then move on to the next aggregator and its collectors until all of the managed units are upgraded.

When all of the managed units of the Central Manager (1 in the figure) are upgraded, the Guardium administrator should proceed with upgrade of Central Manager (2 in the figure) and its managed units following the same upgrade order.




Figure 9-7 Vertical approach

The vertical model is a preferable approach because it allows users to minimize the time that is required to have the central manager and its managed units staying in the hybrid mode, which might associate with several policy limitations.

This approach also minimizes the time that aggregators and collectors are required to be the hybrid mode. Staying in hybrid mode might introduce aggregation performance affects that are related to the continuous dynamic conversion of the data into the newer version of database structures.

64-bit transition strategy

Guardium V9.1p100 (64-bit) has no upgrade path. It requires rebuilding the appliances by using new V9.1p100 64-bit ISO image and restoring the data by using the **restore db-from-prev-version** CLI command.

To maneuver between various hybrid mode compatibility limitations, you should decide on one of the following approaches:

- ▶ Two-step approach
- ▶ Direct approach

Two-Step approach

The Two-Step procedure approach includes the following steps:

1. Upgrade all of the appliances to V9.1p100 32-bit patch level by using the suggested vertical model.
2. Rebuild the Central Manager. Following the vertical model to rebuild all of the appliances to V9.1p100 64-bit level by using v9.1p100 64-bit ISO image.

- Restore the backed up data by using the **restore db-from-prev-version** CLI command.

Note: For more post-9.1p100 patches that might be required before the **restore db-from-prev-version** command is run, see the Flashes and Alerts section of IBM Guardium Customer Support website at:

http://www-947.ibm.com/support/entry/portal/alerts/software/information_management/infosphere_guardium?productContext=-168397159

Figure 9-8 shows the two-step upgrade approach.

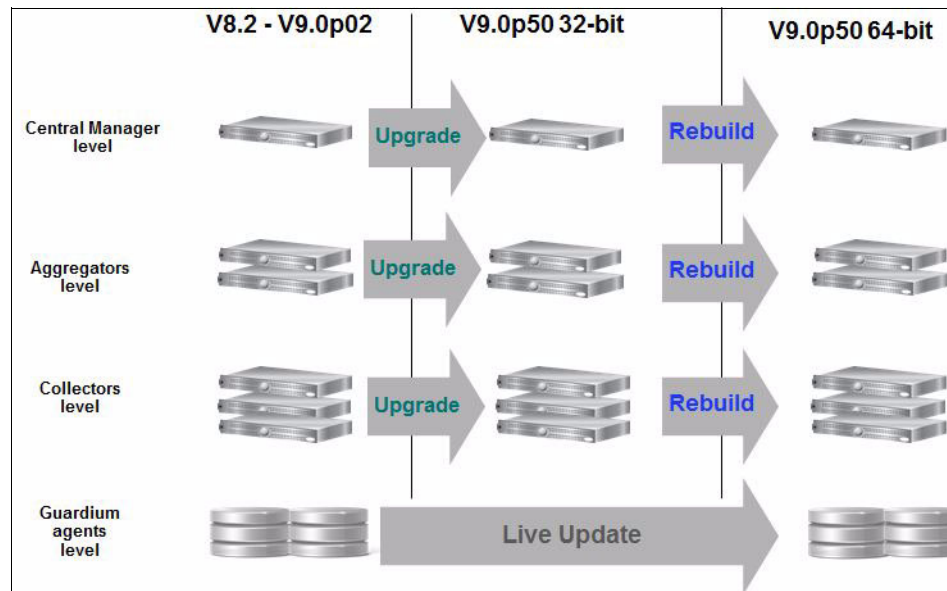


Figure 9-8 Two-step upgrade approach

The two-step approach includes the following advantages and disadvantages:

- ▶ Advantages:
 - Provides a straightforward transition path and is applicable for any type of Guardium enterprise architecture.
 - After completion of the first step, the entire environment is brought to the stable homogeneous level with most of the new features of version V9.1p100 already available for the user.

► Disadvantage:

The two-step approach is generally much longer upgrade route that requires users to complete upgrade and then to rebuild all of the appliances across the estate. This duplication can be avoided with the direct approach.

Direct approach

The direct approach (as shown in Figure 9-9) provides a logistical way to work around V9.1p100 compatibility with an earlier version constrains. In fact, the first action of the direct approach is the same as the two-step approach. It requires upgrading the central manager to V9.1p100 32-bit level. However, at this point, the rest of the managed unit appliances can be rebuilt directly to the V9.1p100 64-bit level.




Figure 9-9 Direct approach

The direct approach includes the following advantages and disadvantages:

- Advantage: This approach does not require transitional 32-bit managed unit upgrade; therefore, it is expected to have a shorter upgrade time frames.
- Disadvantage: This approach is not applicable for environments with combined central manager and aggregator appliances.

9.4.4 Enterprise upgrade summary

The Guardium upgrade can be summarized by the following points:

- ▶ **Limitations:** Although the hybrid mode is supported by Guardium, many functions are limited until all components are at the same version.
- ▶ **RAM memory upgrade:** Some new features are memory-intensive. To use these features, consider upgrading RAM memory capacity to 24 GB.
- ▶ **Appliances upgrade:** Choose the upgrade strategy that minimizes time duration when appliances operate in mixed V8.2 and V9.0 or in mixed 32-bit and 64-bit architectures.
- ▶ **Purging:** Purging unnecessary data from the appliance significantly decreases the duration of the upgrade process. The pre-upgrade health check enforces the database size to be below 50% capacity. We suggest having the databases disk space usage under 20% before the upgrade, if possible. For V9.1p100, GPU requirement is to have the database size below 80% usage.
- ▶ **Duration:** The following factors contribute to the upgrade process:
 - **Size:** Usage and data distribution of internal MySQL database tables
 - **Capacity** of the appliance (virtual memory, processor, and so on)
- ▶ **Downtime:** A typical upgrade of a Guardium appliance is expected to last for several hours. During this time, the appliance might not be accessible and does not perform any data collection activity.

Note: To avoid loss of data collection, you can redirect S-TAPs to different collectors. However, it is required that both of these collectors are exporting to the same aggregator.

9.5 Appliance upgrade methods

In this section, we describe the methods to upgrade the Guardium appliance to version V9.1p100. We also highlight the more complicated path that starts from V8.2 appliance.

In this section, we also describe the following most common upgrade methods:

- ▶ Upgrade appliances by using Guardium V8.2 to V9.0p100 (32-bit) upgrade bundle.
- ▶ Rebuild appliances with V9.1p100(64-bit) ISO image and restore the V8.2 or V9.0 backup.

- ▶ Build new separate V9.1p100(64-bit) appliances and gradually retire the old ones.

9.5.1 Resource download locator

Guardium upgrade packages are available in two sites: IBM Fix Central and IBM Passport Advantage. Figure 9-10 shows the upgrade bundle of V8.2 to V9.1p100.

Target Source	32-bit → 32-bit	32-bit → 64-bit	64-bit → 64-bit
V8.2	V8.2 to V9.1p100 bundle patch (IBM Fix Central)	Rebuild with V9.1p100 64-bit ISO (IBM PA)	N/A
V9.0 (GA)	V9.1p100 32-bit GPU patch (IBM Fix Central)	Rebuild with V9.1p100 64-bit ISO (IBM PA)	N/A
V9.0p02	V9.1p100 32-bit GPU patch (IBM Fix Central)	Rebuild with V9.1p100 64-bit ISO (IBM PA)	N/A
V9.0p50	V9.1p100 32-bit GPU patch (IBM Fix Central)	Rebuild with V9.1p100 64-bit ISO (IBM PA)	V9.1p100 64-bit GPU patch (IBM Fix Central)
V9.1p100	N/A	Rebuild with V9.1p100 64-bit ISO (IBM PA)	N/A

Figure 9-10 Resource download locator for transition to V9.1p100

9.5.2 Upgrading appliances by using V8.2 to V9.0p50 upgrade bundle

The suggested approach to upgrade appliance from V8.2 to the latest V9.0 (32-bit) is through the V9.1p100 upgrade bundle patch.

The Guardium upgrade bundle patch serves as a special wrapper that implements the sequential installation of several upgrade patches that are required to bring the system to the wanted Guardium version with the most recent GPU level.

The V8.2 to V9.1p100 upgrade bundle contains the following patches:

- ▶ V8.2 to V9.0 upgrade patch
- ▶ V9.0 Health Check for GPU and Upgrade
- ▶ V9.1p100 GPU patch (32-bit)

Therefore, the bundle size is large and expected to take longer than regular upgrade patches.

Note: During upgrade bundle installation, the appliance is expected to restart few times automatically. This is normal behavior of the upgrade patch installation process.

Note: Do not restart the appliance manually until the installation of all patches is complete.

The Pre-upgrade Health Check patch is a mandatory prerequisite of the upgrade bundle. This is the same health check patch that is used for regular V9.0 upgrade. The installation sequence includes the following steps:

1. Run Health Check patch for V8.2 to V9.0 upgrade, as shown in Figure 9-11.




Figure 9-11 Health Check patch

2. Install the upgrade bundle for V8.2 to V9.1p100 (32-bit) upgrade. The fix pack uses the following name:

InfoSphere_Guardium_v8.2_to_9.1p100_Upgrade_Bundle_<Timestamp>

For the appliances that are already on V9.0, upgrading to V9.1p100 requires the installation of the following patches:

- ▶ V9.0 Health Check for GPU and Upgrade
- ▶ V9.0p50 GPU patch (32-bit)

The V9.0p50 GPU patch is applicable on the appliances with any V9.0 patch level.

Both patches are available for download through the following IBM Fix Central website:

<http://www-933.ibm.com/support/fixcentral/>

Step-by-step upgrade procedure

When you are upgrading the appliances, make sure that you cover all of the required steps. It is also imperative to prepare the fallbacks and recovery plan. The following procedure provides you a quick reference of the tasks that are required to have a smooth upgrade:

1. Check system requirements.

For the system requirements, see InfoSphere Guardium system requirements, which are available at this website:

<http://www.ibm.com/support/docview.wss?&uid=swg27039049>

2. Purge the systems.

The upgrade requires the appliance internal MySQL database to be below 50% usage. To speed up the upgrade process and minimize the risks that are related to the upgrade of large amounts of data, lower the database usage to 20% or less.

For GPU upgrades, Guardium enforces a maximum of 80% internal database usage.

3. Upload and install pre-upgrade Health Check patch.

Pre-upgrade Health Check is mandatory prerequisite for upgrade and GPU patches.

4. Run the pre-upgrade system backup for current version.

System backup is an important step for disaster recovery and the alternative upgrade route.

5. If enabled, turn off high-availability.

If you have a high availability configured unit, turn off this functionality through the CLI and reboot the appliance before the upgrade is done.

6. Decide on the GUI layout.

By default, the upgrade process applies the V9.0 appearance. For customers with large amounts of panel customizations in the previous versions, Guardium includes an option to maintain the customer's current GUI layout by using the `keep_psm1 s` CLI command. Run the command before the upgrade.

7. Stop the export process.

Temporarily pause any export activity for the period of appliance upgrade.

8. Upload and install upgrade bundle patch.

9. Post-upgrade activity.

Turn on high availability (if required), catch up on all missed and current exports, and install the new V9.1p100 or later Accelerators patches.

9.5.3 Rebuilding appliances with V9.1p100 (64-bit) ISO image and restoring the V8.2 or V9.0 backup

As an alternative upgrade method, you can rebuild Guardium appliances by using Guardium V9.1p100 product ISO image and then restoring data by using the **restore db_from_previous_version** command.

For 32-bit architecture, this method implies rebuilding by using the original V9.0 ISO image, followed by the installation of V9.1p100 GPU upgrade patch.

This is the only method to upgrade to Guardium version with 64-bit architecture. Guardium provides direct V9.1p100 64-bit product ISO image. Authorized customers can find 9.1p100 64-bit ISO image on Passport Advantage.

Note: For data restore and recovery purposes, make sure to generate a system backup before you begin rebuilding the appliance.

Complete the following steps to rebuild and restore the appliance:

1. Run the pre-upgrade Health Check patch.
2. Create a data and configuration backup of the Guardium V8.2 appliance.
3. Install the Guardium V9.1p100 64-bit ISO image.
4. Copy the data and configuration backup file to the appliance by using the CLI **import file** command.
5. Restore and upgrade data from V8.2 or V9.0 backup by using the CLI **restore db-from-prev-version** command.
6. Generate a new V9.1p100 system backup.

Note: Another post-9.1p100 patch might be required before you run the restore command. For more information, see the IBM Guardium Customer Support page (Flashes and Alerts section), which is available at this website:

http://www-947.ibm.com/support/entry/portal/product/information_management/infosphere_guardium?productContext=-168397159

9.5.4 Building new separate V9.1p100(64-bit) appliances and gradually retiring the old ones

The following procedure is used to perform the upgrade by using a parallel environment:

1. Install Guardium V9.1p100 64-bit with an ISO image.
2. Restore the latest pre-V9.1p100 system backup or import definitions.
3. Configure the S-TAPs to report to the new appliance.

Note: All of the reports and queries that require pre-upgrade data must be maintained on old appliances for the time that is equivalent to the purge period.

This method is attractive especially for virtualized environments in which the introduction of more virtual machine appliances is easy and cheap. This method also requires a limited restoration procedure (usually definitions only), there is low risk of data corruption, and no monitoring downtime is required.

However, this method requires at least the temporary expansion of the Guardium environment with newly configured physical or virtual appliances, which requires more maintenance efforts through the transition period.

9.6 Agents upgrade methods

In this section, we describe the upgrade methods that can be used for Guardium agents.

9.6.1 GIM upgrades

Guardium Installation Manager (GIM) agent installation or upgrade must be completed before the installation of other Guardium agents.

The following GIM installation options are available:

- ▶ GIM shell installer
This option is required on systems without GIM installed.
- ▶ GIM bundle upgrade (used on the system with GIM)
This option is applied through Guardium appliance GUI. The GIM bundle upgrade is applicable only on systems with previously installed GIM.

Note: GIM software installation does not require reboots because GIM has no kernel modules that are associated with it.

9.6.2 UNIX S-TAP upgrade

UNIX S-TAP can be installed through the following two major installation options:

- ▶ S-TAP shell installer (from the command line on the database server)
 - This upgrade option can be used in the following modes:
 - Live update by using the **guard-stap-update** utility (preferable)
 - Uninstall, reboot, and install the new version (required during the major upgrade of host operating system)
- ▶ GIM STAP bundle (from through Guardium appliance UI)
- ▶ GIM S-TAP bundle is often used with the live update option (KTAP_LIVE_UPDATE=Y) that allows S-TAP to upgrade without rebooting the host database server.

Note: Starting with version V8.x, Guardium supports live (boot-less) KTAP upgrade, which does not require a reboot of the database server after installation of new version of S-TAP is complete. The live update mechanism is controlled through the GUI with the KTAP_LIVE_UPDATE parameter or BUNDLE-STAP/KTAP installers by using the **guard-stap-update** utility.

9.6.3 Windows S-TAP upgrade

The following installation options are available for Windows S-TAP:

- ▶ STAP silent installer (from the command prompt on the database server):
 - Live update by using **setup.exe /s /z "UPGRADE"** (preferable)
 - Uninstall, reboot, and install the new version (required during the major upgrade of host operating system)
- ▶ GIM S-TAP bundle (from Guardium Appliance UI)
 - GIM S-TAP bundle often is used with the live update option, which allows S-TAP to be upgraded without rebooting the system.

9.7 Upgrading bundle monitoring and status validation

In this section, we describe how to monitor the upgrade process and validate the status.

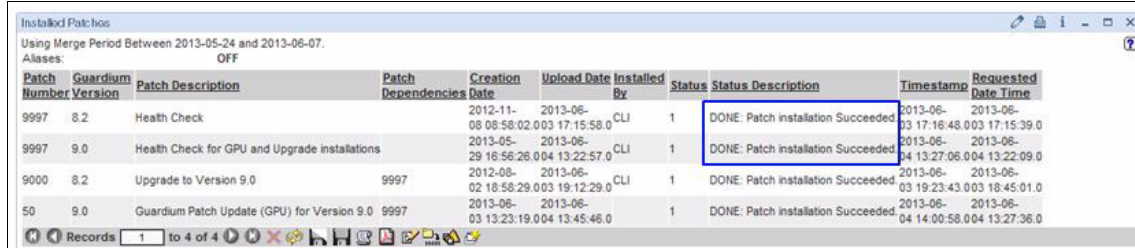
9.7.1 Health check monitoring

For 32-bit architecture, pre-upgrade Health Check is a mandatory requirement that is enforced during the upgrade patch installation.

Starting with V9.0p50, pre-upgrade Health Check is also a prerequisite for any GPU patch installation.

For 64-bit architecture, pre-upgrade Health Check cannot be enforced because this version requires a system rebuild. However, in the post-rebuild stage, data restore is required. We advise that you run the Health Check patch before you generate a pre-upgrade system backup.

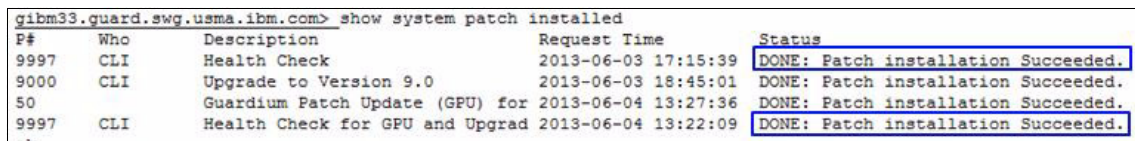
To validate the successful installation of pre-upgrade Health Check patch, click **Guardium appliance GUI** → **Guardium Monitor** → **Installed patches**, as shown in Figure 9-12.



Patch Number	Guardium Version	Patch Description	Patch Dependencies	Creation Date	Upload Date	Installed By	Status	Status Description	Timestamp	Requested Date Time
9997	8.2	Health Check		2012-11-08 08:58:02	2013-06-03 17:15:58	CLI	1	DONE: Patch installation Succeeded.	2013-06-03 17:16:48	2013-06-03 17:15:39
9997	9.0	Health Check for GPU and Upgrade installations		2013-05-29 16:56:26	2013-06-04 13:22:57	CLI	1	DONE: Patch installation Succeeded.	2013-06-04 13:27:06	2013-06-04 13:22:09
9000	8.2	Upgrade to Version 9.0	9997	2012-08-02 18:58:29	2013-06-03 19:12:29	CLI	1	DONE: Patch installation Succeeded.	2013-06-03 19:23:43	2013-06-03 18:45:01
50	9.0	Guardium Patch Update (GPU) for Version 9.0	9997	2013-06-03 13:23:19	2013-06-04 13:45:46	CLI	1	DONE: Patch installation Succeeded.	2013-06-04 14:00:58	2013-06-04 13:27:36

Figure 9-12 Health Check GUI status

Alternatively, you can also access Guardium CLI and run the `show system patch installed` command to check status, as shown in Figure 9-13.



```
gibm33.guard.swg.usma.ibm.com> show system patch installed
```

P#	Who	Description	Request Time	Status
9997	CLI	Health Check	2013-06-03 17:15:39	DONE: Patch installation Succeeded.
9000	CLI	Upgrade to Version 9.0	2013-06-03 18:45:01	DONE: Patch installation Succeeded.
50		Guardium Patch Update (GPU) for	2013-06-04 13:27:36	DONE: Patch installation Succeeded.
9997	CLI	Health Check for GPU and Upgrad	2013-06-04 13:22:09	DONE: Patch installation Succeeded.

Figure 9-13 Health Check CLI status

In both cases, the successful completion status is indicated by the “DONE: Patch installation succeeded” message.

You can access the Health Check log files through the CLI file server facility, as shown in Figure 9-14.

```
gibm33.guard.swg.usma.ibm.com> fileserver
Creating the index file.

Starting the file server. You can find it at http://gibm33.guard.swg.usma.ibm.com
The timeout has been set to 600 seconds and it may timeout during the uploading.

The upload will only be accessible from the IP you are logged in from: 9.50.19.227

Press ENTER to stop the file server.

Stopping process
ok
```

Figure 9-14 File server facility

Health check logs are in the `diag/current` folder and with a file naming convention, such as `health_check_<timestamp>`. Figure 9-15 shows sample log list.

```
Fri Jun 7 00:27:23 2013 rwxr-xr-x  diag/current
Fri Jun 7 00:27:23 2013 rw-r--r--  diag/current/health_check.20130604132320.log
Fri Jun 7 00:27:23 2013 rw-r--r--  diag/current/health_check.20130603171623.log
```

Figure 9-15 Health Check logs

Figure 9-16 shows the content of the Health Check log for a successful upgrade. This indicates that all checks passed the tests and the appliance is ready for upgrade.

```
/var/log/guard/diag/current/health_check.20130603171623.log

There is NO issue with DB size.
There is NO issue with disk space.
There is NO issue with hostname.
There is NO issue with /var/guardium/jakarta-tomcat-4.1.30/webapps/ROOT/WEB-INF/conf/local-portlets.xreg.
Appliance is ready for upgrade.
```

Figure 9-16 Health Check log content

Starting from V9.1p100, in addition to upgrade, Guardium introduces the mandatory pre-GPU Health Check patch, as shown in Figure 9-17 on page 343. This patch is embedded in the upgrade bundle to run automatically during V8.2 to V9.0 and V9.1p100 GPU patch installation.

```

/var/log/guard/diag/current/health_check.20130604132320.log

There is NO issue with DB size.
There is NO issue with disk space.
There is NO issue with hostname.
Appliance is ready for GPU installation/upgrade.

```

Figure 9-17 Pre-GPU Health Check patch

9.7.2 Upgrade process monitoring

You can check the upgrade process completion from GUIs and CLIs. Figure 9-18 shows a sample patch installation status from the GUI.

Patch Number	Guardium Version	Patch Description	Patch Dependencies	Creation Date	Upload Date	Installed By	Status	Status Description	Timestamp	Requested Date Time
9997	8.2	Health Check		2012-11-08 08:58:02	2013-06-03 17:15:58	CLI	1	DONE: Patch installation Succeeded.	2013-06-03 17:16:48	2013-06-03 17:15:39
9997	9.0	Health Check for GPU and Upgrade installations		2013-05-29 16:56:26	2013-06-04 13:22:57	CLI	1	DONE: Patch installation Succeeded.	2013-06-04 13:27:06	2013-06-04 13:22:09
9000	8.2	Upgrade to Version 9.0	9997	2012-08-02 18:58:29	2013-06-03 19:12:29	CLI	1	DONE: Patch installation Succeeded.	2013-06-03 19:23:43	2013-06-03 18:45:01
50	9.0	Guardium Patch Update (GPU) for Version 9.0	9997	2013-06-03 13:23:19	2013-06-04 13:45:46	CLI	1	DONE: Patch installation Succeeded.	2013-06-04 14:00:58	2013-06-04 13:27:36

Figure 9-18 Installing patches completed: GUI

Figure 9-19 shows a sample patch installation status from the CLI.

```

gibm33.guard.swg.usma.ibm.com> show system patch installed
P#      Who      Description      Request Time      Status
9997    CLI      Health Check      2013-06-03 17:15:39  DONE: Patch installation Succeeded.
9000    CLI      Upgrade to Version 9.0  2013-06-03 18:45:01  DONE: Patch installation Succeeded.
50      CLI      Guardium Patch Update (GPU) for Version 9.0  2013-06-04 13:27:36  DONE: Patch installation Succeeded.
9997    CLI      Health Check for GPU and Upgrad  2013-06-04 13:22:09  DONE: Patch installation Succeeded.
ok

```

Figure 9-19 Installing patches completed: CLI

You also can access the upgrade installation log by using the CLI file server command, as shown in Figure 9-20.

```

Fri Jun 7 00:27:23 2013 rw-rw-r--  diag/depot/patch-9.0p50 GPU May 2013 20130604134546.log
Fri Jun 7 00:27:23 2013 rw-rw-r--  diag/depot/check-8.2p9000 Upgrade to Version 9.0 20130603191229.log

```

Figure 9-20 Accessing installation log by using CLI command

Figure 9-21 on page 344 shows a successful installation message in the installation log. The “DONE: Installation Successful” message at the end of the installation log indicates a successful completion of the V9 upgrade process.

```
run_install.sh, warning, STEP: Execute Post Install Actions. Version: '8.2' Patch: '9000'
run_install.sh, info, * starting tomcat...
run_install.sh, warning, STEP: Making copy of inittab. Version: '8.2' Patch: '9000'
run_install.sh, info, Turning recover the following inittab service(s):
run_install.sh, warning, STEP: Recovering inittab. Version: '8.2' Patch: '9000'
run_install.sh, warning, STEP: Update the Activity Audit for: 'Upgrade to Version 9.0'. Version: '8.2'
run_install.sh, warning, DONE: Installation Successful. Version: '8.2' Patch: '9000'
```

Figure 9-21 Upgrade installation log

9.7.3 Troubleshooting upgrade issues

In this section, we describe some common upgrade issues and how to fix them.

Upgrade error handling

Figure 9-22 and Figure 9-23 show the error message of a failure during the upgrade process from the GUI and CLI.

Patch Number	Guardium Version	Patch Description	Patch Dependencies	Creation Date	Upload Date	Installed By	Status	Status Description	Timestamp
9997	8.2	Health Check		2012-05-22 10:10:16.001	2012-10-01 14:14:00.0	CM	1	DONE: Patch installation Succeeded.	2012-10-01 14:14:31.001
9000	8.2	Upgrade to Version 9.0		2012-08-02 18:58:29.001	2012-10-01 14:23:58.0	CM	0	ERROR: Patch Installation Failed - Recovery from pre patch backup Succeeded	2012-10-01 14:24:59.001

Figure 9-22 Patch installation failure: GUI

```
qa-vm18.guard.swg.usma.ibm.com> show system patch insta
P# Who Description Request Time Status
9997 CM Health Check 2012-10-01 14:12:02 DONE: Patch installation Succeeded.
9000 CM Upgrade to Version 9.0 2012-10-01 14:17:02 ERROR: Patch Installation Failed - Recovery from pre patch backup Succeeded
ok
```

Figure 9-23 Patch installation failure: CLI

Note: The “Recovery from pre-patch backup succeeded” message refers to the auto recovery of several critical internal tables and files.

Error status often requires a detailed review of the log’s content to determine the root cause of failure. You can access the patch upgrade log file through the file server, as shown in Figure 9-24.

```
-rw-rw-r-- 1 root tomcat 10463 Oct 1 14:03 patch-8.2p9997_20121001140331.log
-rw-rw-r-- 1 root tomcat 10463 Oct 1 14:14 patch-8.2p9997_20121001141400.log
-rw-rw-r-- 1 root root 476686 Oct 1 14:24 check-8.2p9000 Upgrade to Version 9.0 20121001142358.log
-rw-rw-r-- 1 root root 141368 Oct 1 14:24 patch-8.2p9000 Upgrade to Version 9.0 20121001142358.log
-rw-rw-r-- 1 root tomcat 29918 Oct 1 14:25 patch_installer.log
```

Figure 9-24 Upgrade log file

Search for the “ERROR” or “fatal” phrases to identify installation error messages. Figure 9-25 shows that the particular error indicates an issue with one of the internal MySQL tables.

```
012-10-01 14:24:42, .run_install.sh, info, sql file name /var/tmp/tmp.AYLLj20949/SqlGuard-8.2p9000_Upgrade_to_Version_9.0/update_files/
fixed_files.dir/076_AuditTask_2.sql
012-10-01 14:24:42, .run_install.sh, info, sql revert file name /var/tmp/tmp.AYLLj20949/SqlGuard-8.2p9000_Upgrade_to_Version_9.0/update
files/fixed_files.dir/076_AuditTask_2.sql.revert
012-10-01 14:24:42, .run_install.sh, warning, STEP: SQL EXE - DML update using '076_AuditTask_2.sql'. Version: '8.2' Patch: '9000'
RROR 1060 (42S21) at line 1: Duplicate column name 'WRITE AXIS'
012-10-01 14:24:42, .run_install.sh, fatal, ERROR: DML File '076_AuditTask_2.sql' executing SQL. Version: '8.2' Patch: '9000'
-----
012-10-01 14:24:42, .run_install.sh, info, DML File /var/tmp/tmp.AYLLj20949/SqlGuard-8.2p9000_Upgrade_to_Version_9.0/update_files/fixed
files.dir/076_AuditTask_2.sql return status 1
```

Figure 9-25 Searching for error

If an upgrade fails, the end of the log file indicates the failure, as shown in Figure 9-26.

```
2012-10-01 14:24:50, .run_install.sh, warning, STEP: Recovering inittab. Version: '8.2' Patch: '9000'
2012-10-01 14:24:56, .run_install.sh, fatal, ERROR: Patch Installation Failed. Version: '8.2' Patch: '9000'
-----
2012-10-01 14:24:56, .run_install.sh, info, rm_patch_file - not manager unit
2012-10-01 14:24:56, .run_install.sh, info, rm_patch_file - /var/log/guard/patches/SqlGuard-8.2p9000_Upgrade_to_Version_9.0.tgz.enc
2012-10-01 14:24:58, .run_install.sh, info, Patch file remove - Could not find /var/log/guard/patches/SqlGuard-8.2p9000_Upgrade_to_Vers
on_9.0
2012-10-01 14:24:58, .run_install.sh, info, Ended .run_install.sh
```

Figure 9-26 Failed upgrade

Upgrading MustGather

In many cases, upgrade issues require the attention of support personnel. As is the standard procedure, support engineers require that the basic diagnostic information that is related to the issue be attached to the PMR. The information that is gathered includes basic version, patch level, system information, and installation log files.

Starting with V9.0p50, Guardium introduces a new diagnostic gathering procedure for upgrade and patch installation issues. For more information, see Chapter 7, “Ongoing operations” on page 237.)

This procedure can be started through the CLI or GUI. Figure 9-27 shows the GUI option.




Figure 9-27 MustGather patch installation from the appliance GUI

Starting GUI MustGather is also available through Central Manager.

The output of MustGather is archived in to a single .tar file for the convenience of downloading the file from the appliance and uploading the file to the corresponding PMR.



Use cases

In this chapter, we describe the following common use cases:

- ▶ How to automate the tracking process for database changes
This tracking process integrates external ticket information into the reporting process to streamline the database administrator's (DBA) activity.
- ▶ How to configure the connection profile to help implement security best practices for database security
Connection profiling is the process by which you identify connections to the database and identify if they should be authorized, investigated, or prevented.

This chapter includes the following topics:

- ▶ Creating an integrated change management report
- ▶ Connection profiling and security best practices

10.1 Creating an integrated change management report

Change management is part of typical operational procedures. It is important to track database changes to the change management process. The ability of Guardium to import and correlate information dramatically reduce the manual time in reconciling change requests within the database environment.

In this section, we describe how to integrate and reduce the time that is required to reconcile database changes with the overall change management process.

Because each change management system is different, we present this information based on a generic system; however, this process can be applied to all types of systems. Figure 10-1 shows the overall process to obtain an integrated change management report.




Figure 10-1 Change management integration process

10.1.1 Importing change tickets with Enterprise Integrator

Importing change tickets with Enterprise Integrator consists of the following steps:

1. Define or import the table definition structure.
2. Upload the data from the external table.
3. Define the custom domain.
4. Define the custom query with the new external domain entities.
5. Place the query on the portal.

In many customers' environments, a table or view is defined within a database for the Guardium system to upload change tickets. For this example, we assume that a database table that contains the change ticket number and description is in place. We run an Oracle script (as shown in Example 10-1) to simulate the external change management table.

Example 10-1 Sample script to add a table to upload into the Guardium appliance

```
-- Create ChangeRequest table to simulate troubleticketing system
-- The specific table you will need depends on your ticketing system.
-- Please consult the Change Ticketing system vendor for details

Create table ChangeRequest (
  ChangeID varchar2(30),
  NAME varchar2(30),
  REQDATE varchar2(8),
  EXPECTED varchar2(8),
  DESCRIPTION varchar2(100),
  AFFECTED varchar2(10),
  APPROVED varchar2(1),
  COMPLETED varchar2(8) );

-- Insert some data into the newly created table
-- this data will be imported into the Guardium system

Insert into
ChangeRequest(ChangeID,NAME,REQDATE,EXPECTED,DESCRIPTION,AFFECTED,APPROVED,COMPLETED)
VALUES('1279','BILL SMITH','05-21-10','05-23-10','Modify Schema to
include new product sales',
'REVENUES','Y','05-23-10');

Insert into
ChangeRequest(ChangeID,NAME,REQDATE,EXPECTED,DESCRIPTION,AFFECTED,APPROVED,COMPLETED)
```

```
VALUES('1280','BILL SMITH','05-22-10','05-23-10','Modify Schema to
include net sales',
'REVENUES', 'Y','05-23-10');
```

```
Insert into
ChangeRequest(ChangeID,NAME,REQDATE,EXPECTED,DESCRIPTION,AFFECTED,APPRO
VED,COMPLETED)
```

```
VALUES('1281','BILL SMITH','05-23-10','06-03-10','Rollup
calculations',
'REVENUES', 'Y','06-03-10');
```

Defining or importing the table definition structure

Complete the following steps to define the table definition structure within the Guardium appliance so that the Guardium appliance knows what format to use to upload the database table information:

1. Select **Monitor/Audit** → **Build Reports** → **Custom Reporting**, as shown in Figure 10-3 on page 352.

This menu provides access to the tools that are required to import custom data from external sources. Select the custom table builder to define the structure of the external information and the ability to upload this information into Guardium. Custom domain builder is used to define how Guardium users access this information when the external information is uploaded. By using the custom query builder, you can create reports on the domain from the information that was imported.

We use the following items at the bottom of the window for the complete process:

- Custom table builder
- Custom domain builder
- Custom query builder

Select **Custom Table Builder**.

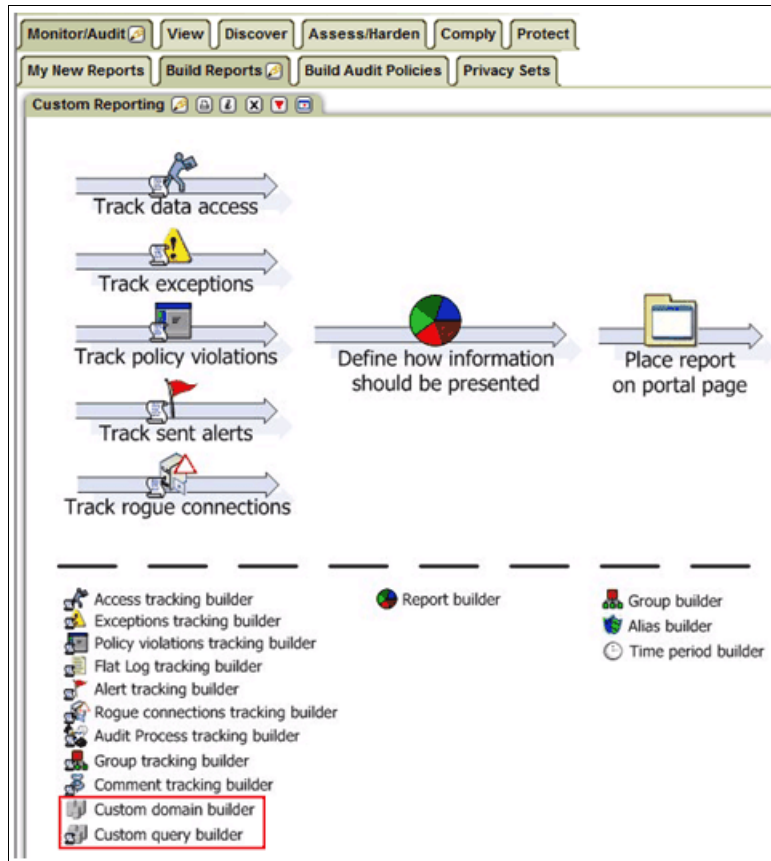


Figure 10-3 Custom Reporting

2. In the Custom Tables tab, select **Upload Definition**, as shown in Figure 10-4 on page 353).

The goal is to retrieve the table definition by connecting to the external database table and retrieving the table definition. The table definition is similar to the SQL “DESCRIBE” statement for Oracle.

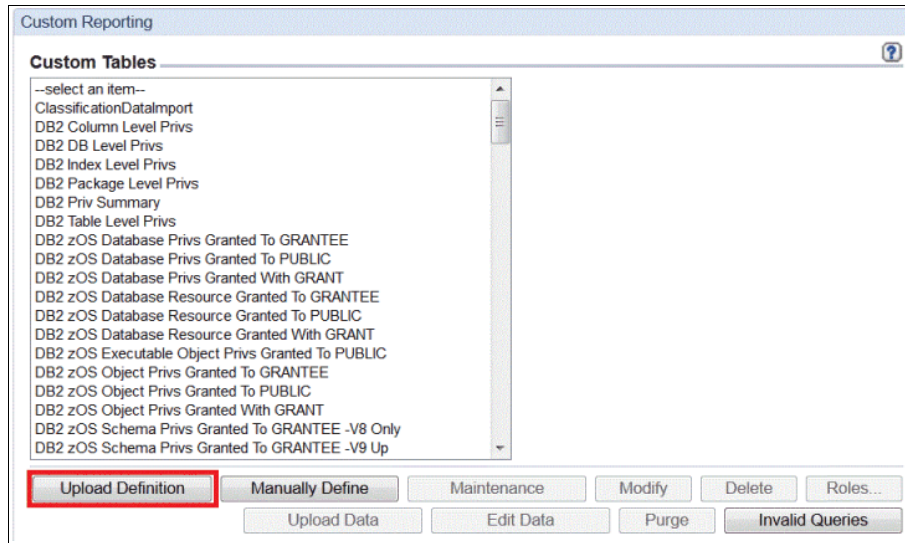


Figure 10-4 Uploading definition

3. In the Import Table Structure tab (see Figure 10-5 on page 354), enter the following information:
 - Entity description

This is the entity name in the query that is used later in the process. This entity allows you to refer to the external ticket table information when you join this with the DBA's SQL activity.
 - Table Name

This is a name for the internal Guardium table to be used in the report.
 - SQL Statement

This is the SQL statement that is used to retrieve the table structure. When the table structure is retrieved, you can report any element within the table.

This SQL statement references the database table that was created in the prerequisite script. In production implementations, this can be a table within a staging database or one that was taken directly from the ticketing system.

Select **Add Datasource**.




Figure 10-5 Import Table Structure tab

4. In the Datasource Finder tab, select **New**, as shown in Figure 10-6.
We define where the external database is and the credentials that are needed to retrieve the table definition and content later in the process.

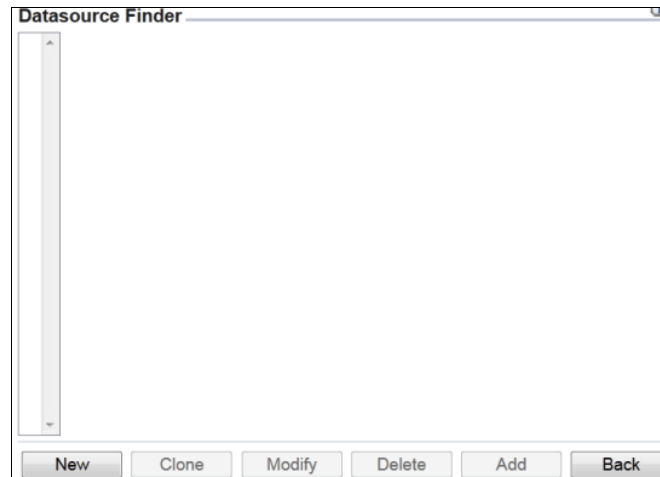


Figure 10-6 Datasource Finder

Note: To retrieve information from the ChangeRequest, you need appropriate permissions and access to read the table in the database.

5. In the Datasource Definition tab (see Figure 10-7), define the following specific information for the external data source:
 - Datasource Definition:
 - Name
 - Database Type
 - Authentication Credentials: Save Password
This means that we save the user name and password within Guardium.
 - Location Information:
 - Host Name/IP
 - Port
 - Service Name

The screenshot shows the 'Datasource Definition' window. The 'Name' field is 'oracle56-joe'. The 'Database Type' is 'Oracle (DataDirect)'. The 'Severity classification' is 'NONE'. The 'Share Datasource' checkbox is checked. In the 'Authentication' section, 'Save Password' is checked, 'Login Name' is 'joe', and the 'Password' is masked with dots. In the 'Location' section, 'Host Name/IP' is '10.10.9.56', 'Port' is '1521', and 'Service Name' is 'XE'. There are also empty fields for 'Informix Server', 'Schema', 'Connection Property', and 'Custom Url'. The 'CAS' section has empty fields for 'Database Instance Account' and 'Database Instance Directory'. The 'Roles' section shows 'No roles have been assigned to this datasource' and a 'Roles...' button. At the bottom are 'Apply' and 'Back' buttons.

Figure 10-7 Datasource Definition tab

It is always a good idea to test the connection to make sure that the credentials are correct, and to make sure that you have IP connectivity to the database server.

Select **Apply**, then select **Test Connection** to verify database connection.

If you are successful, you see the “Datasource can be successfully connected” message.

Select **Apply** and then **Back** after all of the information is completed.

6. Highlight the **Oracle56_joe_ORACLE(Custom Domain)** data source and select **Add** in the Datasource Finder tab, as shown in Figure 10-8. The Datasource Finder tab is where you can define external connections to the Guardium system.




Figure 10-8 Adding external connections to the Guardium system

- Retrieve the table definition structure. In the Import Table Structure tab (as shown in Figure 10-9), select **Retrieve** to run the `Select * from ChangeRequest` statement and retrieve the table structure.

The screenshot shows the 'Import Table Structure' tab in a software interface. The interface includes a navigation menu at the top with options like 'View', 'Quick Start', 'Monitor/Audit', 'Discover', 'Assess/Harden', 'Comply', and 'Privacy Sets'. Below the menu, there are sections for 'Custom Reporting' and 'Import Table Structure'. The 'Import Table Structure' section contains three input fields: 'Entity desc' with the value 'ExternalTickets', 'Table Name' with the value 'ExternalTickets', and 'SQL statement' with the value 'Select * from ChangeRequest'. Below these fields is a 'Datasources' section with a table listing a data source named 'oracle56_joe_ORACLE(Custom Domain)' with type 'ORACLE', host '10.10.9.56', and user 'joe'. At the bottom of the form are 'Retrieve' and 'Back' buttons.

Figure 10-9 Import Table Structure tab

Note: You can use Guardium database activity monitoring to monitor the activity when the connection is made. The SQL request originates from the Appliance (10.10.9.248) to the external database (10.10.9.56) in this example.

If the SQL statement runs successfully, you see the table structure information (ExternalTickets) from the database that created an internal table on the Guardium appliance, as shown in Figure 10-10 on page 358.




Figure 10-10 Newly created custom table ExternalTickets

8. Highlight **ExternalTickets** and select **Modify** to validate the table structures that were imported, as shown in Figure 10-11.




Figure 10-11 Modify the ExternalTickets structure

Figure 10-12 shows the table definition that was imported into the system. The table structure definition must link to the internal Guardium table structure. By linking the Application Event Value String with the ID column, you can join the ticket ID with the SQL Session from the DBA.

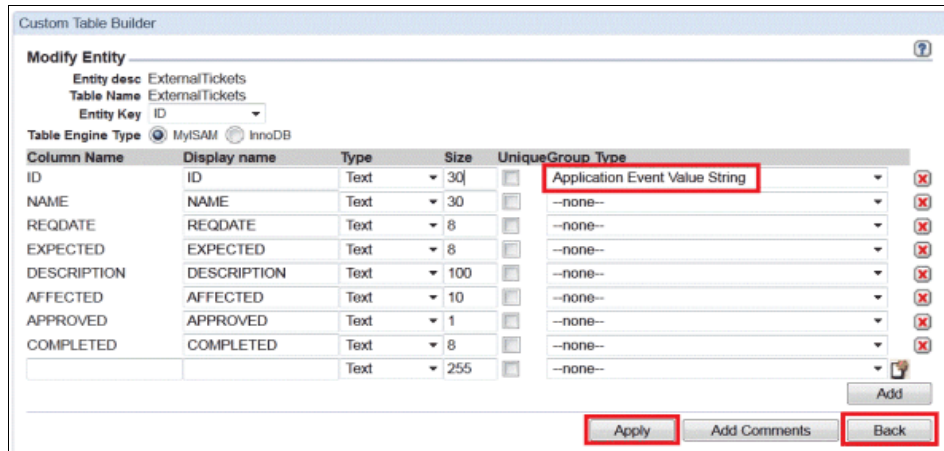


Figure 10-12 Table structure definition

We use the ID column to link with the Guardium Application Event Value String for the unified report. Change the Group Type in the upper right column so that it matches the Application Event Value String. You can now modify any of the column display names if you want to change any of this information. In our example, we leave these alone to keep them as they were imported.

Select **Apply**, and then **Back**.

This table structure is good for our purposes. We completed the process to define the table structure. The next step is to upload the data within the external table.

The other selection in the Maintain Custom Table menu is Manage Table Index. Click **Insert** to open Table Index Definition. The pop-up window suggests columns in the table to add to indexes that are based on columns that are used on custom domains as Join conditions. Select the columns and click **Save**. Indexes are created (or re-created).

The table engine types for custom tables and entitlements (InnoDB and MyISAM) appear for all predefined custom databases because the data that is stored on the Guardium internal database is MYSQL-based at the time of this writing.

The two major types of table storage engines for MySQL databases are InnoDB and MyISAM. These two MySQL table engine types feature the following major differences:

- ▶ InnoDB is more complex; MyISAM is simpler.
- ▶ InnoDB is more strict in data integrity; MyISAM is looser.
- ▶ InnoDB implements row-level lock for inserting and updating; MyISAM implements table-level lock.
- ▶ InnoDB has transactions; MyISAM does not.
- ▶ InnoDB has foreign keys and relationship constraints; MyISAM does not.
- ▶ Changing the engine type is disallowed (and the selection grayed out) if the row number in the table is greater than 1 MB.

10.1.2 Uploading the data from the external table

You must upload data into the newly created custom table after the table definition is defined or imported. You also can use this method to integrate heterogeneous database entitlement information for z/OS and Informix databases.

Complete the following steps to upload the data from the external table:

1. Select **Upload Data** from the Custom Tables tab, as shown in Figure 10-13 on page 361.

By selecting Upload Data, the Guardium appliance connects to the external database table and retrieves the content within this table after all of the configuration information is provided. The information that is retrieved from the database is then stored within Guardium. This information can be provided in a separate report or combined with other information to join more meaningful data in the reporting process.

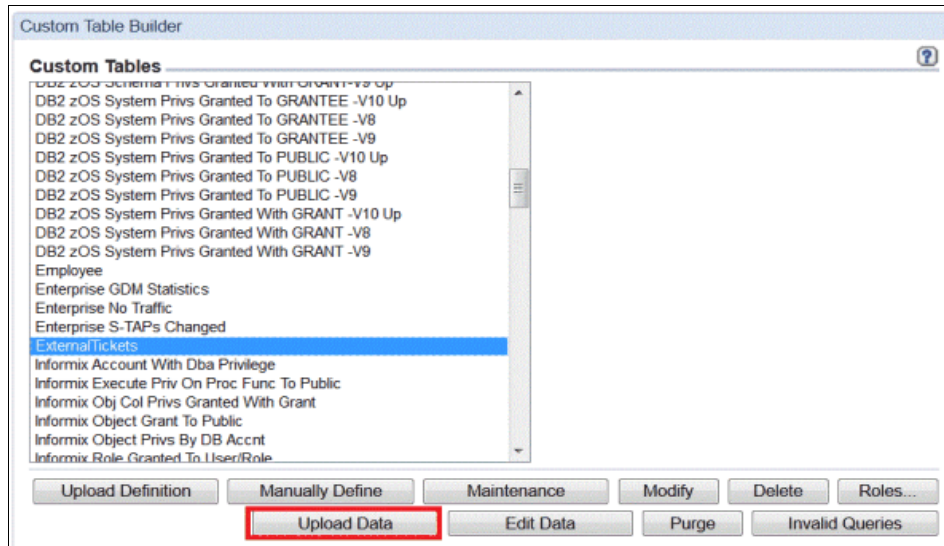


Figure 10-13 Uploading data

To connect to the database, you must provide a data source, which has the appropriate credentials to log in to the database server.

Select **Add Datasource**, and the Datasource Finder window opens, as shown in Figure 10-14. We want to add the data source `oracle56-joe_oracle(Custom Domain)`, which was defined as shown in Figure 10-7 on page 355. Select this data source and then click **Add**.




Figure 10-14 Add the data source `10.10.9.56-joe_ORACLE(Custom Domain)`

2. Select **Check/Repair** (as shown in Figure 10-15 on page 362) to see whether the SQL statement is valid when it is run on the external Oracle database.

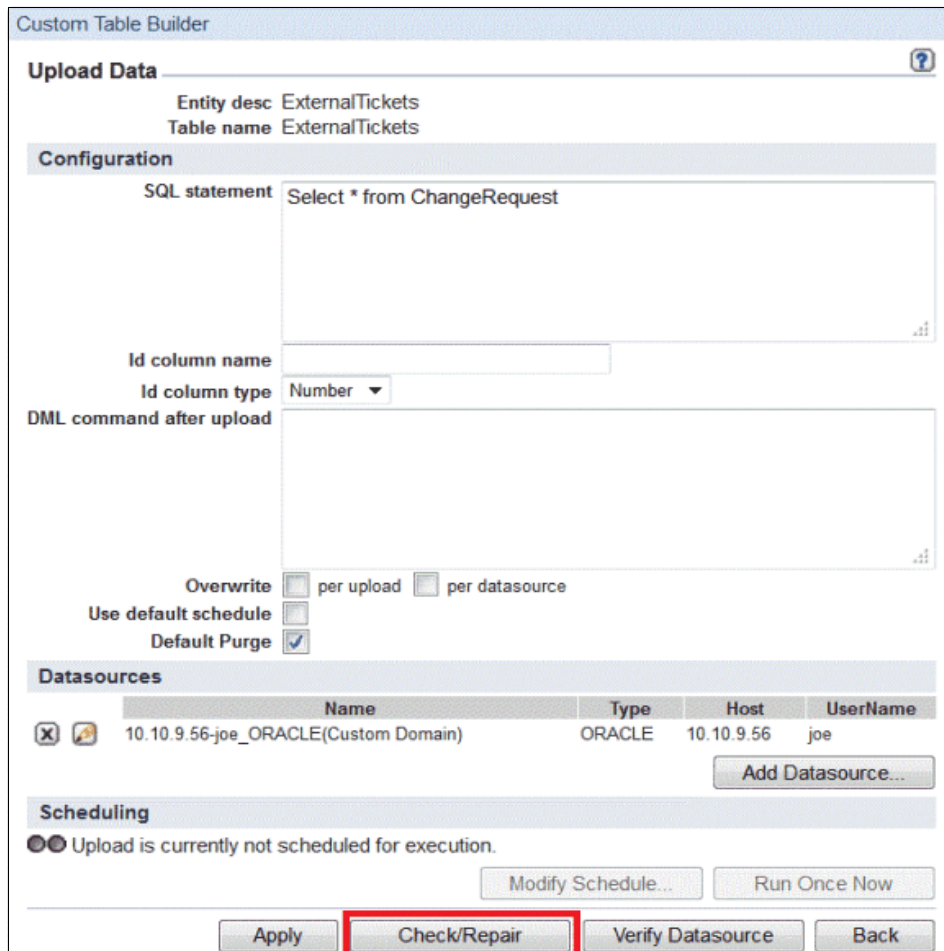


Figure 10-15 Validate the table structure and definitions

You should see the “Operation ended successfully” message if the new table structure is valid.

Select **Apply**.

Note: You have many options for how you want to clean up the external database table. One option is to use a SQL statement to delete the contents of the external table by selecting **DML command after upload**. This option gives you tremendous flexibility. You can also overwrite all the previously imported Guardium data when you upload by selecting **Overwrite**. Finally, you can schedule this process to happen automatically by selecting **Modify Schedule**.

3. Select **Run Once Now** to populate the information within the table structure that is defined, as shown in Figure 10-16.

The screenshot shows the 'Custom Table Builder' window. The 'Upload Data' section is active, showing 'Entity desc: ExternalTickets' and 'Table name: ExternalTickets'. The 'Configuration' section includes a text area for the 'SQL statement' containing 'Select * from ChangeRequest', an empty 'Id column name' field, and a 'Number' dropdown for 'Id column type'. Below this is a 'DML command after upload' text area. There are checkboxes for 'Overwrite' (with 'per upload' and 'per datasource' options), 'Use default schedule', and 'Default Purge' (checked). The 'Datasources' section contains a table with one entry: '10.10.9.56-joe_ORACLE(Custom Domain)' of type 'ORACLE' on host '10.10.9.56' with user 'joe'. A 'Scheduling' section shows 'Upload is currently not scheduled for execution.' and a 'Run Once Now' button highlighted with a red box. At the bottom are 'Apply', 'Check/Repair', 'Verify Datasource', and 'Back' buttons.

Name	Type	Host	UserName
10.10.9.56-joe_ORACLE(Custom Domain)	ORACLE	10.10.9.56	joe

Figure 10-16 Run to populating information

You should see the “Operating ended successful” message when the table information is populated. This action also validates how many entries were in the table. In our example, there were 12 entries in the ChangeRequest table.

Select **Back**.

Data is now uploaded to the Guardium Appliance.

10.1.3 Defining the custom domain

Now we need to define the custom domain for unifying the external data and the internal data that is captured from the DBA activity. Complete the following steps to define the custom domain:

1. Select **Monitor/Audit** → **Build Reports** → **Custom Reporting**.

Select **Custom Domain Builder** to define how you want to get access to the new information, which was uploaded in the previous step.

In Domain Finder, select **New**, as shown in Figure 10-17 on page 364.

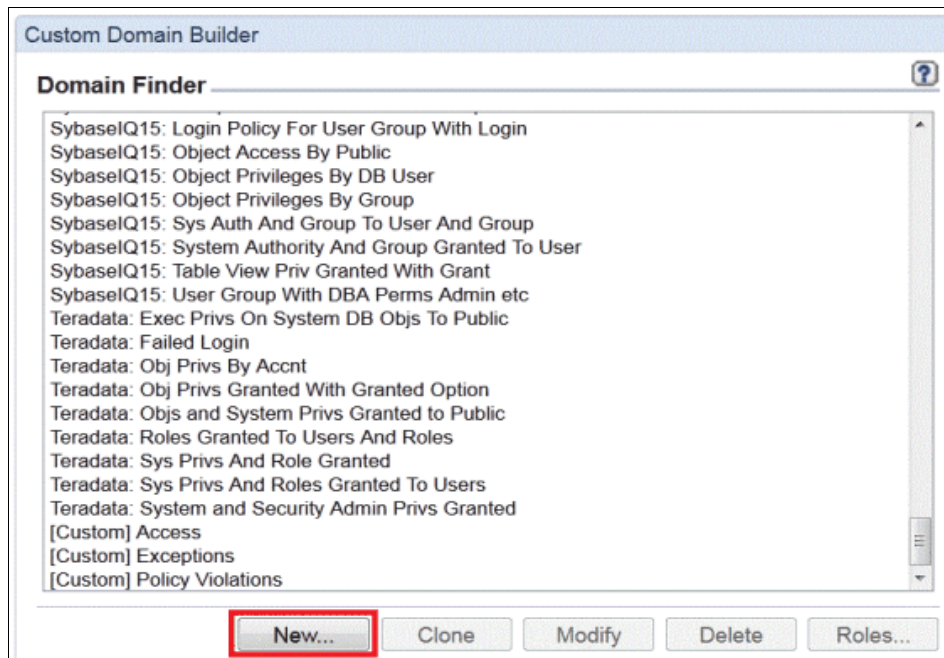


Figure 10-17 Creating a domain

The following special domains are found at the bottom of the window:

- [Custom] Access: This domain captures database traffic.
- [Custom] Exceptions: This domain captures SQL Errors, Failed Logins, and so on.
- [Custom] Policy Violations: This domain contains security violations that were triggered by the Policy Rules.

By linking the uploaded information with these custom domains, you can unify external reporting information to make the reports more meaningful to your stakeholders (auditors, information security, DBAs, and so on).

We link information with the [Custom] Access domain.

The goal is to create a domain of information for a report with the newly uploaded data within Guardium.

We create two domains. The first domain consists of only the imported ticket information. The second domain consists of the imported ticket information and the monitoring structure for the DBA activity when they place their marker within the SQL transactions.

2. Complete the following steps in the Custom Tables Domain (as shown in Figure 10-18):
 - a. Enter ExternalTickets in the Domain name section.
 - b. Highlight **ExternalTickets** from the Available Entities.
 - c. Select >>.

There are many Guardium domains in which information can be found. In this example, we are creating a domain, but there are other valuable domains, such as Enterprise No Traffic Alert, which can be used to alert if there is no audit data coming from a database. Enterprise S-TAP Changed reports about the changes on the configuration of your S-TAPs and more.

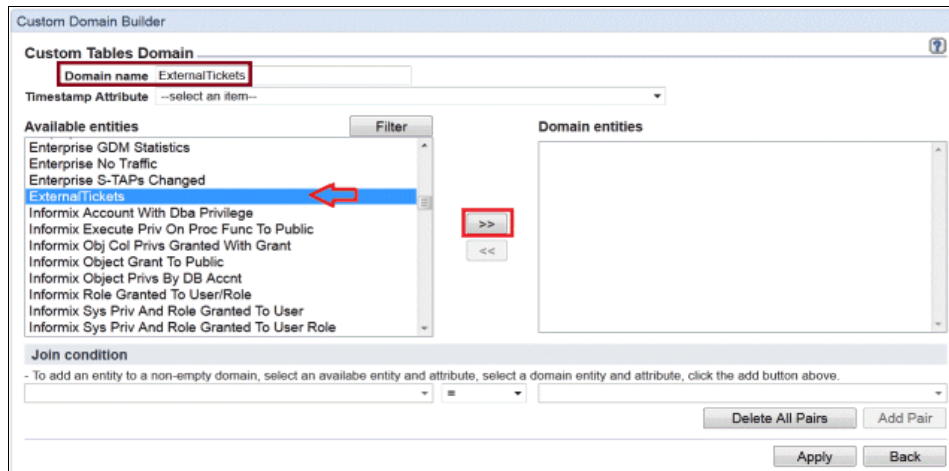


Figure 10-18 Custom Tables Domain

The ExternalTickets was moved to the right under Domain entities, as shown in Figure 10-19 on page 366.




Figure 10-19 Custom Domain with Timestamp added

3. Select **External_Tickets.SqlGuard_timestamp** from the drop-down list for the Timestamp Attribute. This is the time that the data is uploaded into the SQL Guard Appliance.

Select **Apply** and the custom Domain definition is completed. Select **Back**.

You successfully created the custom domain that includes the external ticket information only. Now you can create the second domain of information to report on. This consists of the imported ticket information and the monitoring structure for the DBA activity when they place their marker within the SQL transactions.

Complete the following steps to create the second domain:

1. Clone the [Custom] Access domain so that we can link the ExternalTickets to create a unified reporting information domain.

Select a domain to be cloned. Scroll to the bottom and select **[Custom] Access**, as shown in Figure 10-20 on page 367.

In this process, we link the native Guardium Access domain with the ExternalTickets domain.




Figure 10-20 Clone the [Custom] Access domain

2. Change the Domain Name from Copy of [Custom] Access to External Tickets and DBA Activity, as shown in Figure 10-21.

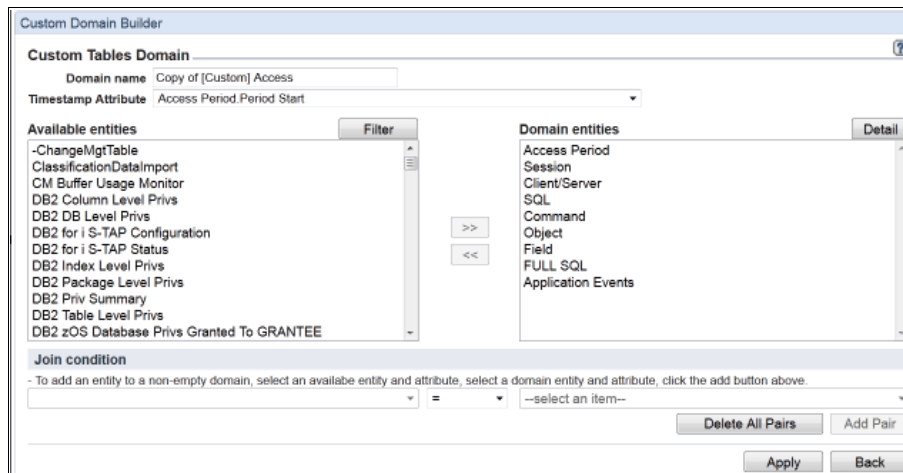


Figure 10-21 Clone of the [Custom] Access domain: change domain name

If you are familiar with Guardium, the Domain entities on the right side are entities that are shown in the Query builder where you create your reporting.

3. Click **Apply**. You see the successful message when the Access Domain is cloned successfully.

Notice the Domain Name changed at the top of the window, as shown in Figure 10-22. We are going to perform an outer join from the ExternalTickets table and the Guardium Application Events table. By using this join, we can see information whether a DBA entered a change request ticket by including information without the Application Events.

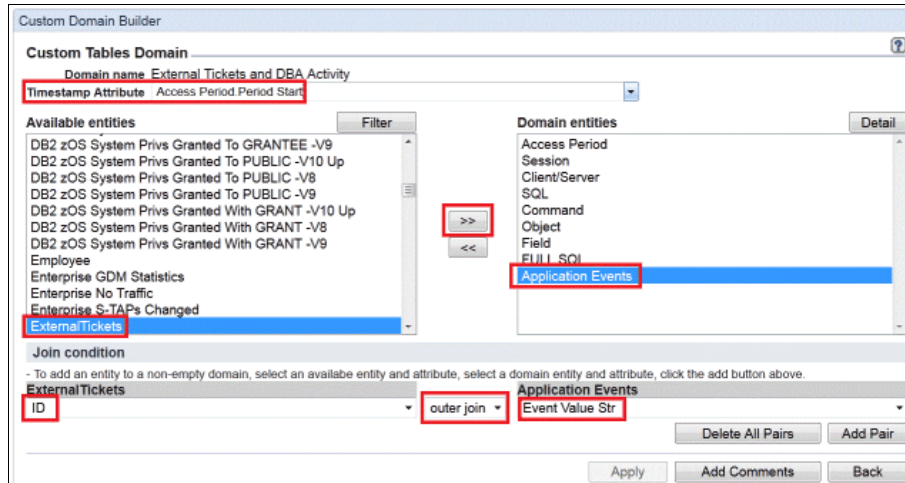


Figure 10-22 Joining to Domains of Information

Complete the following steps:

- a. Timestamp Attribute: Select **Access Period.Period Start** as the time stamp.
- b. Available entities: Select the **ExternalTickets**.
- c. Select **ID** under ExternalTickets at the lower left.
- d. Select **Application Events Entity** for the Domain Entities on the right side.
- e. Select **Event Value Str** as the field to join.
- f. Under Join Condition, select **outer join** in the middle drop-down list.
- g. Select **>>**, which moves the ExternalTickets entity into the Domain Entities on the right side. This move links ExternalTickets with Application Events Entity within the Custom Domain of External Tickets and DBA Activity.

Note: It is critical to have the outer join in this linkage so that if the DBA does not enter a change ticket or enters an incorrect ticket, these items show up in this report.

In the Join Condition fields, if you selected **ExternalTickets** for table to add a “=” instead of **outer join** in domain Application Events in the linkage, only items that matched show up in this report. This might be a wanted goal; however, there are other ways to accomplish this in the query condition portion of the report.

Click **Apply**.

The ExternalTickets moved to the Domain Entities, as shown in Figure 10-23.

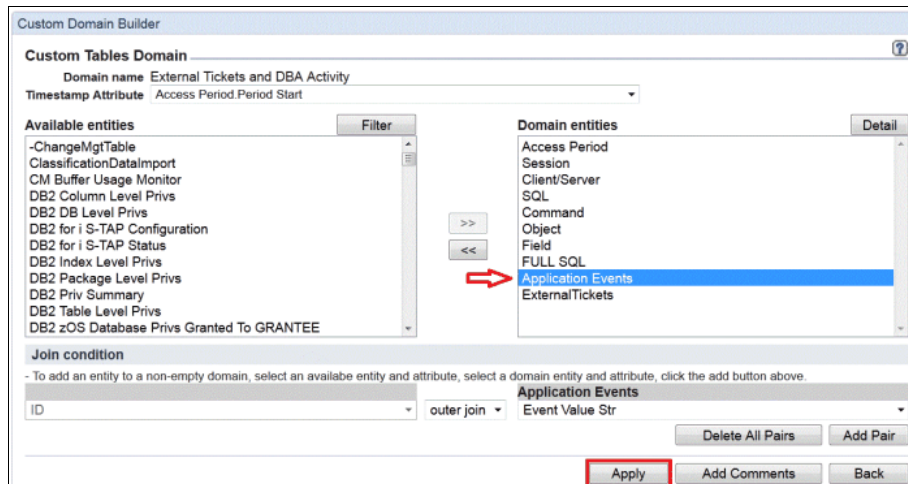


Figure 10-23 ExternalTickets linked with Application Events Entity

4. Select **Back**. Figure 10-24 shows that two New Domains, “External Tickets” and “External Tickets and DBA Activity”, are successfully created.

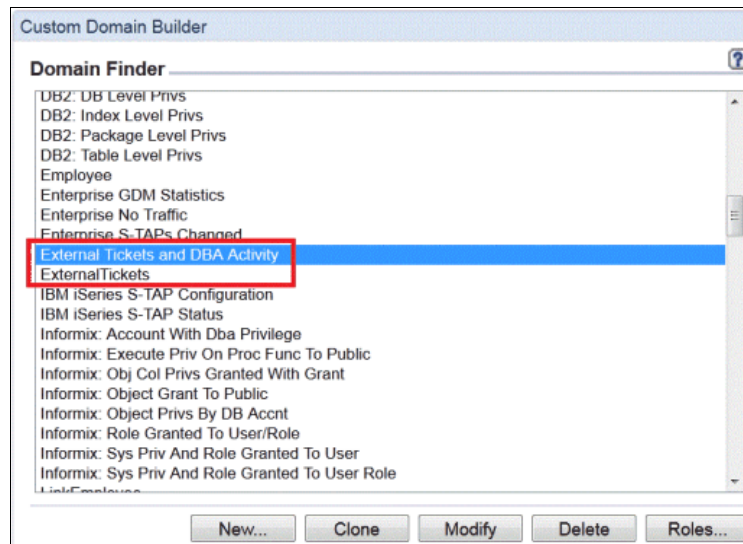


Figure 10-24 Two domains successfully created

You successfully created two domains of information. Now you must create a query for each domain. Then, you can put these reports on the portal.

10.1.4 Defining custom query with new external domain entities

Complete the follow steps to create the custom query for each report:

1. Select **Monitor/Audit** → **Build Reports** → **Custom Reporting**. Select **Custom query builder** (as shown in Figure 10-25 on page 371) to create reports and put these reports on the GUI (Portal).

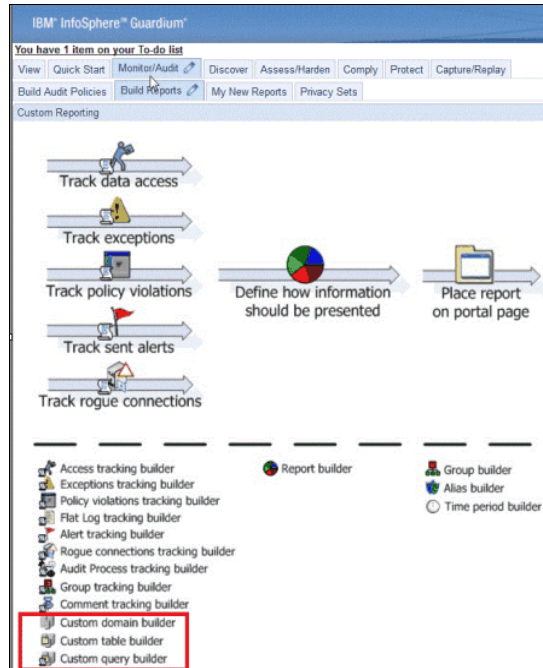


Figure 10-25 Selecting Custom query builder

2. Search the ExternalTickets domain that we created in the previous step, as shown in Figure 10-26.

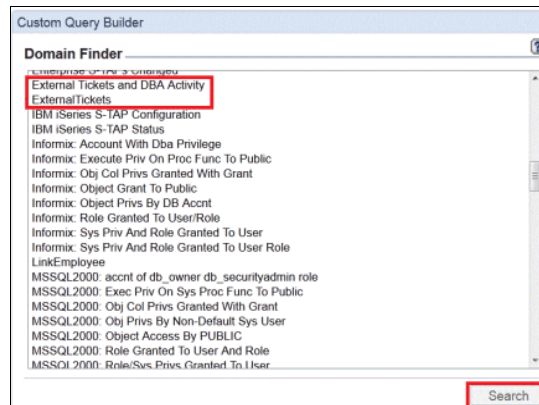


Figure 10-26 Searching ExternalTickets domain

3. Select **New** to create a query, as shown in Figure 10-27.




Figure 10-27 New Custom Query from Custom Query Builder

4. Enter the following information in the New Query - Overall Details tab, as shown in Figure 10-28:
 - Query Name: **ExternalTickets**
 - Main Entity: **ExternalTickets**




Figure 10-28 Defining Query Name and Main Entity

- Click **ExternalTickets** main entity so it expands the list of attributes. Add the attributes to the query as shown in Figure 10-29. Upon completion, select **Save**.

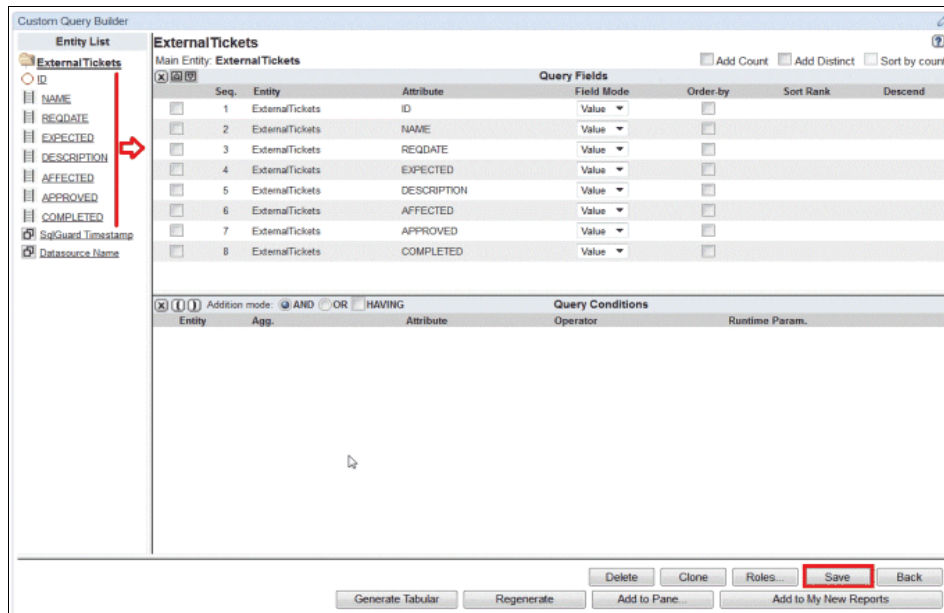


Figure 10-29 Adding attributes to query

- Select **Add to My New Reports**. When this is completed, you receive the “Report added to My New Reports pane” message that confirms that the new report ExternalTickets is added to My New Reports pane in the GUI. Select **Back**.

By adding the report to the Portlet, you can now view the uploaded information.

Repeat the following steps to define the query for the second domain, External Tickets and DBA Activity:

- Select **Custom query builder** from the Custom Report tab.
- Search the **External Tickets and DBA Activity** domain that we created, as shown in Figure 10-30 on page 374.

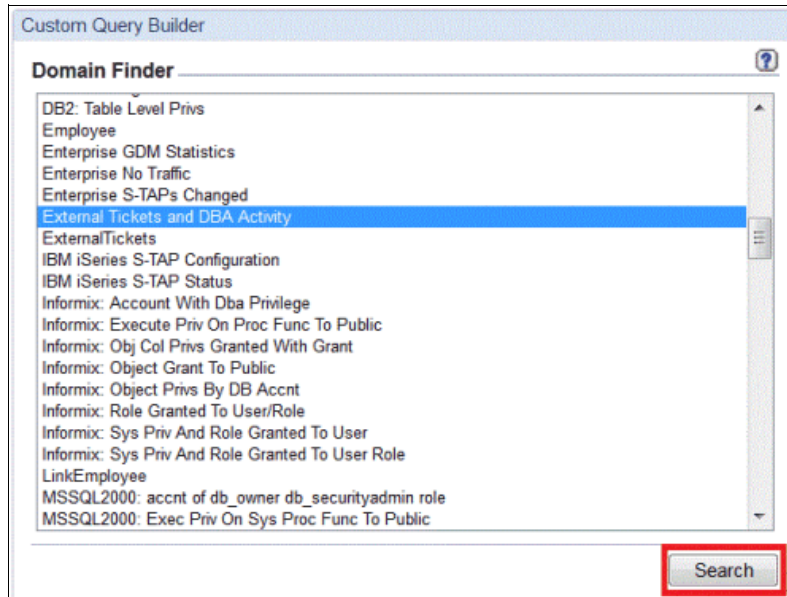


Figure 10-30 Custom Query Builder for External Tickets and DBA Activity

3. Select **Search**, then select **New** from the Query Finder tab, as shown in Figure 10-31.


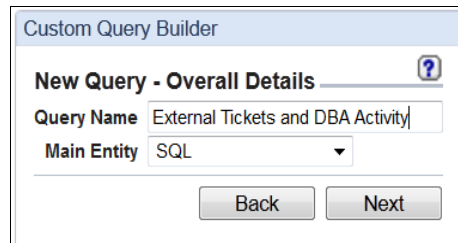


Figure 10-31 Query Finder

4. In the New Query - Overall Details tab (see Figure 10-32), enter the following information:
 - Query Name: External Tickets and DBA Activity
 - Main Entity: SQL



The screenshot shows a dialog box titled "Custom Query Builder" with a sub-tab "New Query - Overall Details". Inside the dialog, there is a text input field for "Query Name" containing the text "External Tickets and DBA Activity" and a dropdown menu for "Main Entity" currently set to "SQL". At the bottom of the dialog, there are two buttons: "Back" and "Next".

Figure 10-32 External Ticket and DBA Activity definition

5. Select the appropriate entities and add the attributes to the query, as shown in Figure 10-33 on page 376. This report filters for only DDL changes (Create, Alter, and Drop). We use the following entities:
 - Access Period: Timestamp
 - ExternalTickets: Description and ID
 - Application Events: Event Value Str and Event User Name
 - Client/Server: Client IP, Server IP, and DB User Name
 - SQL: SQL

Click **Save** → **Generate Tabular** → **Add to My New Reports** → **Done**.

This report shows the external ticket information and the DBA SQL activity.




Figure 10-33 Query definition for External Tickets and DBA Activity

You successfully added two new reports to the “My New Reports” pane.

10.1.5 Customizing the reports

Complete the following steps to customize the reports in the My New Reports pane:

1. Click **Monitor/Audit** → **My New Reports**, as shown in Figure 10-34. Select **ExternalTickets**.



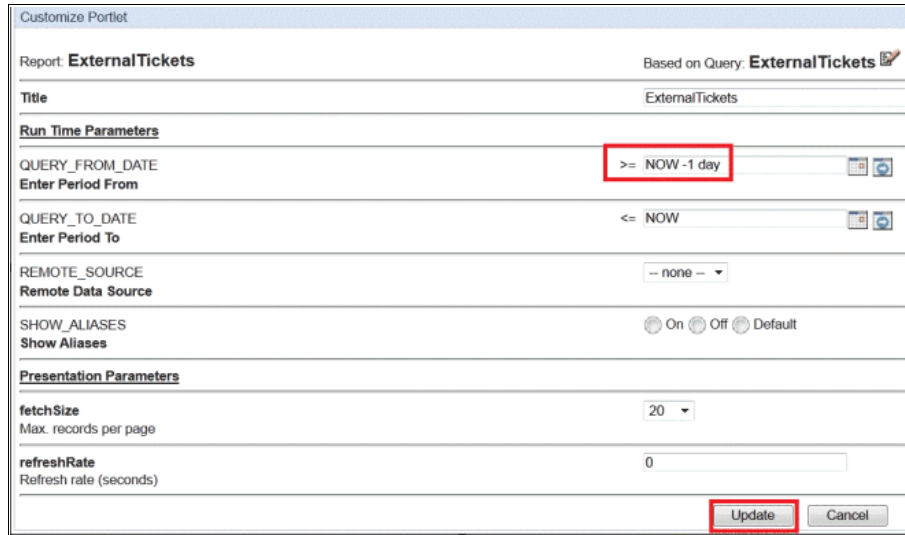


Figure 10-34 Customize the date range to show information in the report

- Click the pencil icon  at the upper right of the External Tickets report to define the date range of the report. When we imported this report, we selected **Access Period.Period Start timestamp**, so this is when the data range should be selected to see the appropriate information, as shown in Figure 10-35. Most customers synchronize their ticketing systems daily for these types of reports. In this example, we use Now -1 day to obtain the information in our report.





Customize Portlet



Report: **ExternalTickets** Based on Query: **ExternalTickets**

Title: ExternalTickets


Run Time Parameters

QUERY_FROM_DATE:  

Enter Period From

QUERY_TO_DATE:  

Enter Period To


REMOTE_SOURCE: 

Remote Data Source

SHOW_ALIASES: On Off Default

Show Aliases

Presentation Parameters

fetchSize: 

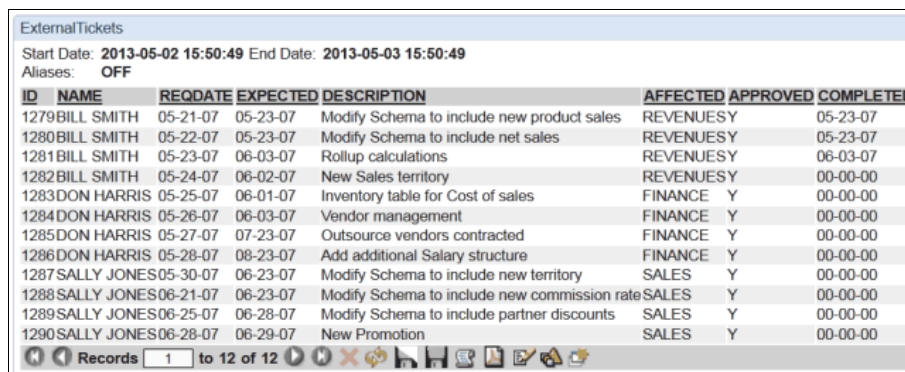
Max. records per page

refreshRate:

Refresh rate (seconds)

Figure 10-35 Updated the time frame to show data in the report

Figure 10-36 shows that the external tickets were imported correctly.



ExternalTickets

Start Date: 2013-05-02 15:50:49 End Date: 2013-05-03 15:50:49

Aliases: OFF

ID	NAME	REQDATE	EXPECTED	DESCRIPTION	AFFECTED	APPROVED	COMPLETED
1279	BILL SMITH	05-21-07	05-23-07	Modify Schema to include new product sales	REVENUES	Y	05-23-07
1280	BILL SMITH	05-22-07	05-23-07	Modify Schema to include net sales	REVENUES	Y	05-23-07
1281	BILL SMITH	05-23-07	06-03-07	Rollup calculations	REVENUES	Y	06-03-07
1282	BILL SMITH	05-24-07	06-02-07	New Sales territory	REVENUES	Y	00-00-00
1283	DON HARRIS	05-25-07	06-01-07	Inventory table for Cost of sales	FINANCE	Y	00-00-00
1284	DON HARRIS	05-26-07	06-03-07	Vendor management	FINANCE	Y	00-00-00
1285	DON HARRIS	05-27-07	07-23-07	Outsource vendors contracted	FINANCE	Y	00-00-00
1286	DON HARRIS	05-28-07	08-23-07	Add additional Salary structure	FINANCE	Y	00-00-00
1287	SALLY JONES	05-30-07	06-23-07	Modify Schema to include new territory	SALES	Y	00-00-00
1288	SALLY JONES	06-21-07	06-23-07	Modify Schema to include new commission rate	SALES	Y	00-00-00
1289	SALLY JONES	06-25-07	06-28-07	Modify Schema to include partner discounts	SALES	Y	00-00-00
1290	SALLY JONES	06-28-07	06-29-07	New Promotion	SALES	Y	00-00-00

Records 1 to 12 of 12

Figure 10-36 External Ticket information was successfully displayed in Guardium

3. Generate some SQL traffic (as shown in Example 10-2) to display the next report.

Example 10-2 Integrating ticket information with SQL activity

```
[root@osprey ~]# sqlplus joe/guard
```

```
SQL*Plus: Release 11.2.0.2.0 Production on Tue May 7 16:46:35 2013
```

```
Copyright (c) 1982, 2011, Oracle. All rights reserved.
```

```
Connected to:
```

```
Oracle Database 11g Express Edition Release 11.2.0.2.0 - 64bit Production
```

```
SQL> select 'GuardAppEvent:Start', 'GuardAppEventType:RECONCILE',  
2 'GuardAppEventUserName:joed', 'GuardAppEventStrValue:1279' from dual;
```

```
'GUARDAPPEVENT:STAR 'GUARDAPPEVENTTYPE:RECONCIL 'GUARDAPPEVENTUSERNAME:JOE  
-----  
'GUARDAPPEVENTSTRVALUE:127  
-----
```

```
GuardAppEvent:Start GuardAppEventType:RECONCILE GuardAppEventUserName:joed  
GuardAppEventStrValue:1279
```

```
SQL> create table NewProductSales(  
2 ProductID integer,  
3 Name varchar2(20),  
4 cost integer,  
5 revenue integer );
```

```
Table created.
```

```
SQL> SELECT 'GuardAppEvent:Released' from dual;
```

```
'GUARDAPPEVENT:RELEASE  
-----  
GuardAppEvent:Released
```

```
SQL>
```

Figure 10-37 on page 379 shows the captured audit data that is unified with the ticket information in a single report. Notice that the ID and Description fields are from the ticketing system, so the auditor has an easy reference to the authorized ticket that is associated with this SQL activity.

In our example, the DBA “joe” was tasked to modify the schema to include the new product sales table. He created the table NewProductSales with the appropriate fields, according to ticket number 1279.

External Tickets and DBA Activity							
Start Date:		2013-05-07 13:58:12		End Date:		2013-05-07 16:58:12	
Aliases:		OFF		DBUserLike:		LIKE joe	
FULLSQLLike:		LIKE %					
Timestamp	ID	DESCRIPTION	Event Value Str	Client IP	Server IP	DB User Name	Full Sql
2013-05-07 16:51:07.0				10.10.9.56	10.10.9.56	JOE	SELECT 'GuardAppEvent:Released' from dual
2013-05-07 16:50:56.0	1279	Modify Schema to include new product sales	1279	10.10.9.56	10.10.9.56	JOE	create table NewProductSales(ProductID integer, Name varchar2(20), cost integer, price integer) select 'GuardAppEvent:Start', 'GuardAppEventType:RECONCILE', 'GuardAppEventUserName:joed', 'GuardAppEventStrValue:1279' from dual
2013-05-07 16:49:57.0	1279	Modify Schema to include new product sales	1279	10.10.9.56	10.10.9.56	JOE	

Figure 10-37 The DBA activity that is unified with the Ticket ID information

The following fields can help in the reporting information:

- GuardAppEventType: String value
- GuardAppEventUserName: String value
- GuardAppEventNumValue: Numeric value

The key here is to have the DBA run the Guardium API so that their SQL activity tags the session with the appropriate change ticket number. In our example, we used 1279, as shown in the following select statement:

```
select 'GuardAppEvent:Start', 'GuardAppEventType:RECONCILE',
      2 'GuardAppEventUserName:joed', 'GuardAppEventStrValue:1279' from
dual;
```

4. Now that you have the general report working, you can enhance the report by using different colors to identify if the DBA does not enter any ticket information or if they enter an incorrect ticket number.

Select **Monitor/Audit** → **Build Reports** → **Custom Reporting** → **Define how information should be presented**, as shown in Figure 10-38 on page 380.

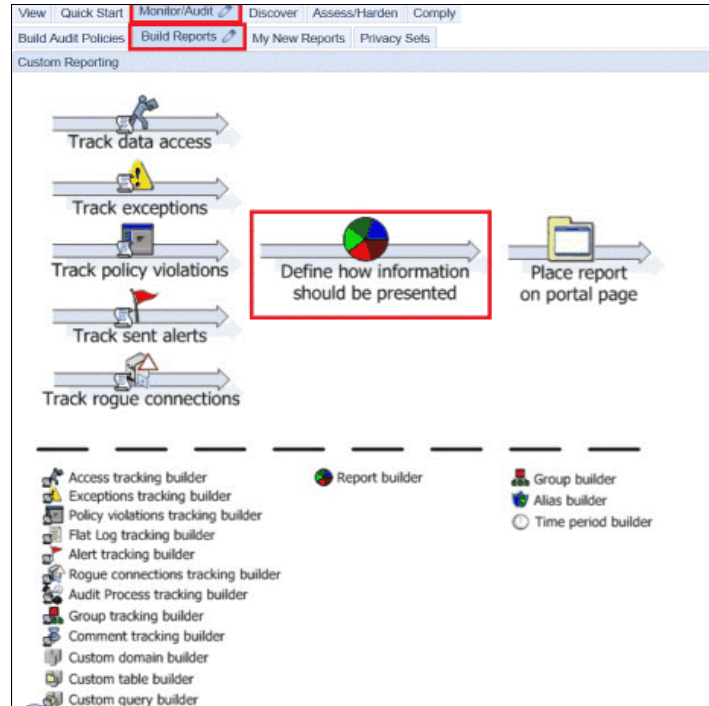


Figure 10-38 Customize the report information

5. Select the **External Tickets and DBA Activity** report from the drop-down list of the Report Title field, as shown in Figure 10-39. This report should be at the top of this window because we started the name of this report with a “-” so that it appears at the top.

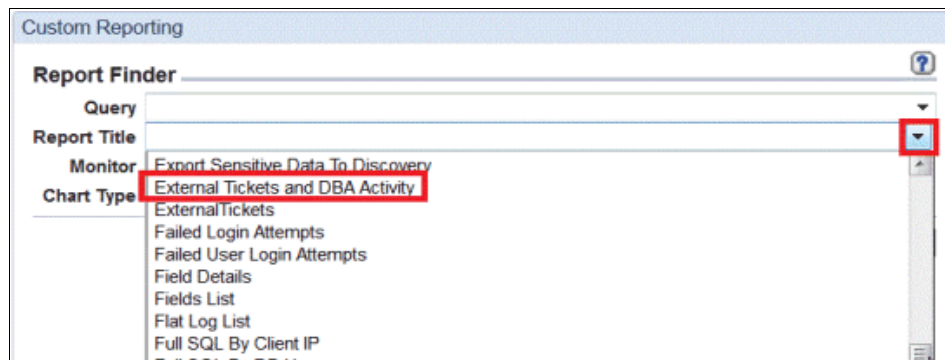


Figure 10-39 Find and modify the External tickets and DBA Activity report

6. Select **Modify** to modify certain header information and add colors, depending on the wanted conditions, as shown in Figure 10-40.




Figure 10-40 Modify information

7. Modify the column headings in the report from Event Value Str to Ticket Entered (this is what the DBA entered), as shown in Figure 10-41.




Figure 10-41 Modifying the column headings in the report

8. In the Report Parameter Description tab, select **Next**, as shown in Figure 10-42.

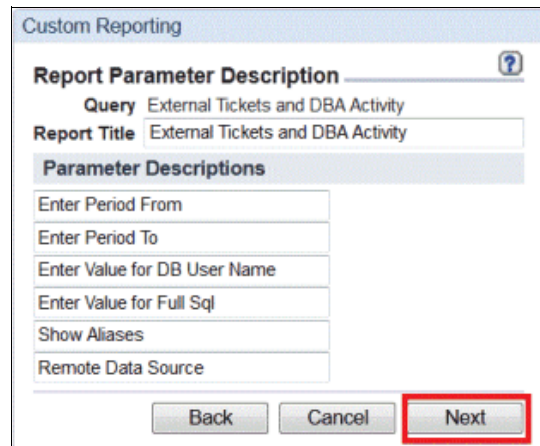


Figure 10-42 Custom Reporting Parameter Description

9. In the Report Attributes tab, select **Next**, as shown in Figure 10-43.

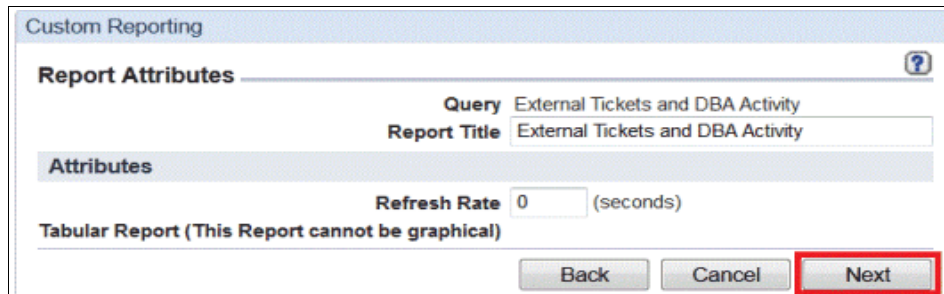


Figure 10-43 Reporting Refresh Rate

10. In the Report Color Mapping window (see Figure 10-44 on page 383), enter Ticket Entered in the Column field and select the color red. Click **Add** to add this entry into the background color section.

Select **CHANGEID** column and select the color yellow. Click **Add** to add this into the background color section.

If someone does not enter a ticket number in their SQL session, the entry is highlighted in red.

If someone enters a ticket number that is not within the Change Ticket system that was imported to Guardium, the entry is highlighted in yellow.

Custom Reporting

Report Color Mapping ?

Query External Tickets and DBA Activity

Report Title External Tickets and DBA Activity

Background Color

Column	Operator	Value	Color
Ticket Entered	=		X
ID	=		X

Add

Back Cancel **Next**

Figure 10-44 Custom color mapping for Ticket Entered and ID from ticketing system

11. Select **Save** to submit the report, as shown in Figure 10-45.

Custom Reporting

Submit Report ?

Query External Tickets and DBA Activity

Report Title External Tickets and DBA Activity

Back Cancel Comments... Roles... **Save**

Figure 10-45 Complete the customization of the report by saving the parameters

Complete the following steps to test these scenarios:

1. Generate traffic, as shown in Figure 10-46.

External Tickets and DBA Activity						
Start Date:	2013-05-07 16:51:55	End Date:	2013-05-07 22:51:55			
Aliases:	OFF	DBUserLike:	LIKE joe			
FULLSQLLike:	LIKE create%					
Timestamp	ID	DESCRIPTION	Ticket Entered	Client IP	Server IP	DB User Name Full Sql
2013-05-07 22:50:53.0	1279	Modify Schema to include new product sales	1279	10.10.9.56	10.10.9.56	JOE create table joe1(i int)

Figure 10-46 Generating traffic to test the scenarios

Because the test ticket number 7777 is not in the ticketing system, this ticket is highlighted in yellow in the report, as shown in Figure 10-47.

External Tickets and DBA Activity							
Start Date:		2013-05-07 14:56:11		End Date:		2013-05-07 20:56:11	
Aliases:		OFF		DBUserLike:		LIKE joe	
FULLSQLLike: LIKE create%							
Timestamp	ID	DESCRIPTION	Ticket Entered	Client IP	Server IP	DB User Name	Full Sql
2013-05-07 20:54:27.0			7777	10.10.9.56	10.10.9.56	JOE	create table bogus (i int) create table NewProductSales(ProductID integer, Name varchar2(20), cost integer, revenue integer)
2013-05-07 16:50:56.0	1279	Modify Schema to include new product sales	1279	10.10.9.56	10.10.9.56	JOE	

Figure 10-47 Invalid ticket that is entered by DBA is highlighted as yellow

2. Try the same procedure, but without any Guardium API. The entry without any ticket is highlighted in red, as shown in Figure 10-48.

External Tickets and DBA Activity							
Start Date:		2013-05-08 04:04:23		End Date:		2013-05-08 10:04:23	
Aliases:		OFF		DBUserLike:		LIKE joe%	
FULLSQLLike: LIKE %							
Timestamp	ID	DESCRIPTION	Ticket Entered	Client IP	Server IP	DB User Name	Full Sql
2013-05-08 10:04:13.0				10.10.9.56	10.10.9.56	JOE	create table no_ticket(i int)
2013-05-08 10:03:58.0			7777	10.10.9.56	10.10.9.56	JOE	create table bogus (i int)

Figure 10-48 DBA activity that is highlighted in red without any approved ticket

There are other ways to automate this process to make it easier for the DBA to enter in ticket information. The `login.sql` script (as shown in Example 10-3) is one way to help enter information that can be used for this purpose.

Example 10-3 The `login.sql` file automates sending `GuardAppEvents`

```
[joe@osprey ~]$ more login.sql

select to_char(sysdate,'YYYY-MM-DD HH24:MI:SS') dcol from dual;

accept EventUserName char prompt "Enter Business Owner:  "
accept TicketNumber char prompt "Enter Change Request (ticket) Number:
"

set feedback on
set termout off
select 'GuardAppEvent:Start', 'GuardAppEventType:RECONCILE',
```

```
'GuardAppEventUserName:&EventUserName',
'GuardAppEventNumValue:&TicketNumber',
'GuardAppEventStrValue:&TicketNumber' from dual;
undefine EventUserName
undefine TicketNumber
[joe@osprey ~]$
[joe@osprey ~]$ sqlplus joe/guard
```

SQL*Plus: Release 11.2.0.2.0 Production on Tue May 7 22:50:18 2013

Copyright (c) 1982, 2011, Oracle. All rights reserved.

```
Connected to:
Oracle Database 11g Express Edition Release 11.2.0.2.0 - 64bit
Production
```

```
DCOL
-----
2013-05-07 22:50:18
```

```
Enter Business Owner: RetailBanking
Enter Change Request (ticket) Number: 1279
SQL> create table joe1(i int);
```

Table created.

```
SQL>
```

The script automatically prompts for the ticket information and another field GuardAppEventType, which can be used in the reporting to reconcile these database changes. This is shown next.

The result is that you are prompted automatically to enter a ticket number for your session. This greatly automates the ticket integration process for DBAs. Figure 10-49 shows the result from an audit report.

External Tickets and DBA Activity							
Start Date:	2013-05-07 16:51:55	End Date:	2013-05-07 22:51:55				
Aliases:	OFF	DBUserLike:	LIKE joe				
FULLSQLLike:	LIKE create%						
Timestamp	ID	DESCRIPTION	Ticket Entered	Client IP	Server IP	DB User Name	Full Sql
2013-05-07 22:50:53.0	1279	Modify Schema to include new product sales	1279	10.10.9.56	10.10.9.56	JOE	create table joe1(i int)

Figure 10-49 Resulting audit report for the login.sql ticket automation

10.2 Connection profiling and security best practices

Access to the database server should follow security best practices and must be treated with great care because most of your sensitive information is within this data store. There should be appropriate controls to identify who should have access and where and when this access can occur.

Connection profiling is the process by which you identify connections to the database and identify whether they should be authorized, investigated, or prevented.

For more information about different components of this process, see the following YouTube videos:

- ▶ *Connection Profiling Part 1 of 3*: A demonstration of how connection profiling works in Guardium V9 GPU 50:
<http://youtu.be/yRoRkAExVz0>
- ▶ *Connection Profiling Part 2 of 3*: A how to guide about configuring connection profiling:
<http://youtu.be/bm6nnATDzeU>
- ▶ *Connection Profiling Part 3 of 3*: This shows the audit process approval on how to authorize new connections to the database, which includes the application owner's involvement in the process:
<http://youtu.be/NwndWdCmAic>

From a security best practices perspective, it is a good idea to block anything that is not authorized until someone can review the details for the new connection. For example, in Figure 10-50 on page 387, the application server IP address of 10.10.9.240 that uses the database user “apps” is an authorized connection to the database server. Any other connections (that is, not 10.10.9.240), should be blocked. This is a simple but effective approach to securing the database.

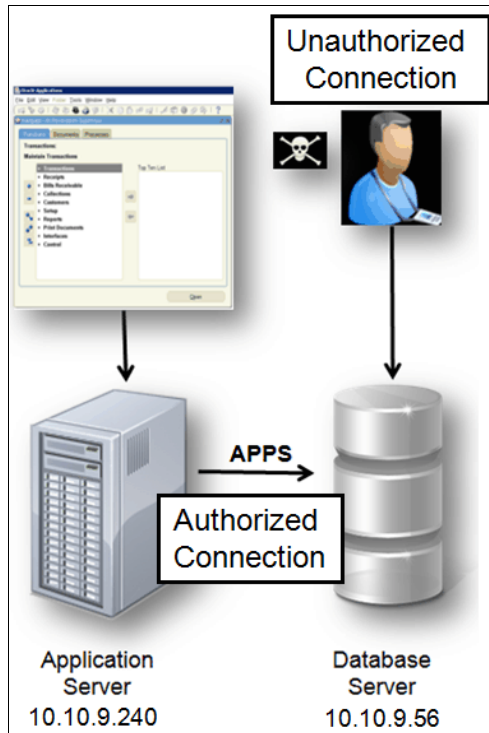



Figure 10-50 Database security

The Guardium security policy that is required for the configuration that is shown in Figure 10-50 is shown in Figure 10-51 on page 388. You can create a group “connection profiling list”, for example, that has all of the details of the authorized connection.

Simple Policy To Eliminate Unnecessary Risk

- Drop connections NOT in the Authorized Group “Connection Profiling List”



1 Access Rule: Block Unauthorized Connections Not in Group AllowedConnectionsToDatabase (Installed)

Cat.	Classif.	Sev.	Client IP	Server IP	Src App.	DB Name	DB User	App. User	Client IP/Src App./DB User/Server IP/Svc. Name					
Data Security	Priv Users	1	ANY	ANY	ANY	ANY	ANY	ANY	Not Connection Profiling List					
Svc. Name	OS User	Net Protocol	Field	Pattern	XML Pattern	Client MAC	DB Type							
ANY	ANY	ANY	ANY	ANY	ANY	ANY	ANY							
Object	Command	Object/Command Group	Object/Field Group	Records Affected Threshold	Trigger Once Per Session	Masking Pattern	Replacement Character	Min. Ct.	Reset Int.	Quarantine Min.	Rec. Vals.	Cont.	Period	Action
ANY	ANY	ANY	ANY	0	<input type="checkbox"/>	ANY	*	0	0	0	<input checked="" type="checkbox"/>	<input type="checkbox"/>	ANY	S-GATE TERMINATE ALERT PER MATCH
App Event Exists	Event Type	App Event Num. Val.	App Event Date	Event User Name	App Event Text Val.									
<input type="checkbox"/>	ANY	ANY	ANY	ANY	ANY									

Figure 10-51 Security Policy to prevent unauthorized connections not in the connection profiling list

If you want to add a new connection to the database server, you add the new connection details into the Connection Profiling List group. Similarly, if there is a connection that is not authorized, you remove this connection from the group. For more information, see the YouTube video *Connection Profiling Part 2 of 3*, which is available at this website:

<http://youtu.be/bm6nnATDzeU>

How do you determine whether a connection should be authorized? In the ideal world, the information security and DBA group work with the application owners to help them identify the risk that is associated with a new connection to the database server. For example, they can help them identify the following information:

- ▶ Database user: Is the new connection from an authorized database user who can directly access information inside the database tables, which can bypass all of the application controls?
- ▶ IP address: Is this connection started from within a secure zone and to be trusted, or is this new connection coming from an unknown IP address?

- ▶ Source program: Is the source program of the database connection an authorized program, such as WebSphere®, or is it a developer application, such as SQL/Plus or Toad?

In many cases, too much information that is provided to the application owners overwhelms them and they do not have a productive conversation with the information security group. By providing a coarse level of information, such as a connections report, it allows the information security team to have a productive conversation with the application owners and bridge the IT-to-Business gap by showing them the risk that is associated with some of the new connections.

In many cases, predefined service accounts have access to the database server so that they can fix issues after going through a change control process. In other cases, a new connection can be seen, which is someone probing to see their boundaries of their privileges and access control. These are the types of connections that should be blocked if they have no business need to see the information. In other cases, this type of implementation can prevent zero day attacks because you proactively block unauthorized connections that might come from a worm or virus that is trying to spread within your environment.

For more information about this conversation between information security, DBA, and the business owners, see the YouTube video, *Connection Profiling Part 3 of 3*, which is available at this website:

<http://youtu.be/NwndWdCmAic>

This video shows the audit process approval cycle on how to authorize a new connection.



Integration with other IBM products

In this chapter, we describe some of the integration that the InfoSphere Guardium has with the other IBM products.

The worlds of security, auditing, and data touch many aspects of information technology. Customers are storing and analyzing a tremendous amount of data. Some studies suggest that we are generating as much data in two days as we did from the dawn of man through 2003. It is hard for us to comprehend the effect of that statement, as about everything is generating data today. It is no wonder that identifying sensitive information in the vast amount of data and how people access the data becomes a huge challenge.

This is why it is important to integrate with other products to help give visibility and synergy from a security and operational perspective. Figure 11-1 on page 392 shows a sampling of Guardium integration points with other IBM products, from the Information Management to Security portfolio, to help customers secure and audit their environment.

11.1 Security and audit

In this section, we describe some of the product integration of Guardium within the broader IBM portfolio regarding security and audit integration. The first set of IBM products (BigInsights™, DB2, Netezza®, Informix, PureData™, and so on) allow for S-TAP to be loaded to intercept the transactions to provide an audit trail of the activity. Depending on the platform, some advanced features, such as blocking and masking, can also be accomplished with Guardium integration without changes to the application or database.

11.1.1 BigInsights integration

Big Data is all around us and is gaining a tremendous amount of momentum. Customers are implementing Hadoop clusters and storing a vast array of information. It can be log files, tweets, social media sentiment, and a host of other information that help the business gain valuable insights into improving their products, services, and offerings. One of the challenges with gathering all of this data is to understand what is stored in these servers and what controls are in place to secure this data.

The following questions can be asked to gauge your level of comfort and maturity of understanding within this environment:

- ▶ Is there any sensitive information that is stored in these nodes?
- ▶ Who is accessing this information?
- ▶ How many MapReduce jobs were run against this data?
- ▶ How are privilege users are defined, and what controls are in place to assure the corporation they are not abusing their privileges?

In the world of Big Data, one of the most critical items to be concerned about is securing the information that you store in your Hadoop clusters. Figure 11-2 on page 394 shows a simplistic view of the Big Data comparison.

Database	Big Data
<ul style="list-style-type: none"> • Database commands <ul style="list-style-type: none"> - SQL - DB Utilities - Scripts - Batch jobs • Database Users <ul style="list-style-type: none"> - Privilege Users - Applications - Test & QA - Developers • Database Clusters <ul style="list-style-type: none"> - a collection of servers to provide scalability and availability at low cost • Data that is stored <ul style="list-style-type: none"> - High valued data 	<ul style="list-style-type: none"> • Big Data commands <ul style="list-style-type: none"> - HBASE - Alter, count, create, drop, get, put, list, etc - HDFS – rmr,ls, tail, chmod, chown, expunge, & files, permissions - MapReduce operations - Job, operation, target, permissions • Big Data Users <ul style="list-style-type: none"> - Privilege Users - Applications - Test & QA - Developers • Big Data Architecture <ul style="list-style-type: none"> - a collection of commodity components to provide scalability and availability at low cost • Data that is stored <ul style="list-style-type: none"> - “unknown” value of data - Looking for valuable insights with VAST amount of data

Figure 11-2 Simplistic view of database components and Big Data components

In database security and auditing, you can monitor the SQL access. For example, you can monitor select statements to understand who is accessing sensitive data. In Big Data, you monitor HBASE GET commands or HDFS cat commands. Grant statements in SQL are roughly similar to **chmod** and **chown** commands within HDFS.

Figure 11-3 on page 395 shows that Guardium can audit a BigInsights environment. A MapReduce job was submitted to the name node in the Hadoop cluster. S-TAP is loaded on the name node to copy this information and send it to the Guardium appliance in real time. Base on your security policy on the Guardium appliance, you can proactively alert your Security Information and Event Manager (SIEM) that an unauthorized user submitted a MapReduce job.




Figure 11-3 Alerting on unauthorized activity is critical in Big Data environments

With Version 9 of Guardium, there are predefined reports to help monitor and audit this environment, as shown in Figure 11-4. The Unauthorized MapReduce jobs report within InfoSphere Guardium helps monitor who is using the Cluster.

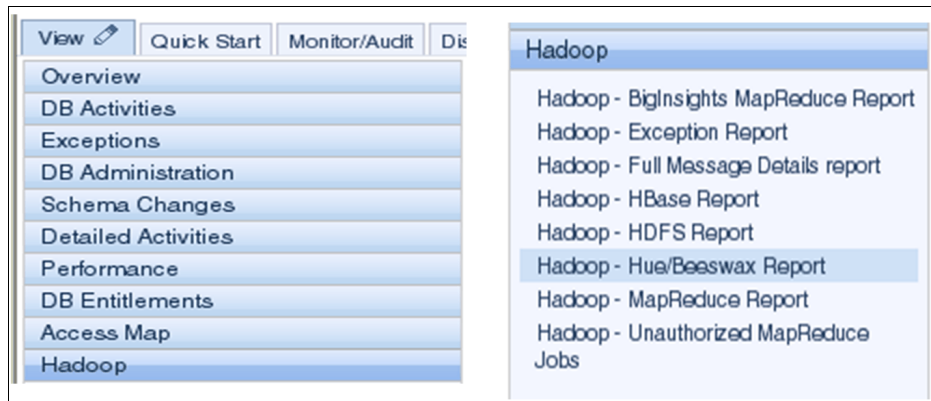


Figure 11-4 MapReduce jobs report

When a Hadoop environment is monitored, the following items should be considered:

- ▶ Separation of duties

In the architecture that is shown in Figure 11-3 on page 395, there are proper separation of duties because a non-Hadoop administrator is responsible for the audit policy creation and reporting level.

Privilege users cannot tamper with the audit logs.

- ▶ HDFS and HBASE Auditing

You must be able to capture relevant HDFS and HBASE information to understand who is accessing or potentially modifying your information.

- ▶ Store audit logs outside of Hadoop system

Guardium streams the transaction information off the servers to reduce overhead and protect the audit data from tampering. If the audit files are on the server, the administrator or someone that hacks into the system potentially can tamper with the audit logs.

This tampering is not possible if the audit logs are not inside the Hadoop cluster.

- ▶ Reporting

Guardium reporting can display and easily filter log information and messages between Hadoop infrastructures to provide meaningful reports to information security or auditing personnel.

Guardium also reports MapReduce jobs for easy identification for both security and operational benefits.

- ▶ Integration with SIEM and real time alerting

Guardium fully integrates with the leading SIEM systems (QRadar, ArcSight, Envision, and so on) by using syslog and real time alerting. This is critical to provide security of your Hadoop infrastructure with your other standard processes.

- ▶ Heterogeneous support and centralized policy administration

Because customers choose products for many reasons, Guardium includes heterogeneous support for multiple distributions of Hadoop.

Policies within the Guardium solution do not require expertise of Hadoop. This allows for easy definition in a heterogeneous Hadoop infrastructure.

The audit collection policy can be defined through a single window for these different Hadoop distributions. All of reports are normalized, which provides the same information from different Hadoop systems in a common format.

- ▶ NoSQL support
Guardium supports the collection of audit data and reporting from the NoSQL databases, which are often deployed within the Hadoop infrastructure.
- ▶ Integration into business process
Guardium can create a process to review collected reports and distribute these reports for compliance sign-off. In addition, the system tracks the sign off from these reports.

11.2 Databases and data warehouses

Guardium is well-known for heterogeneous database activity monitoring. In this section, we describe the integration and functionality of the different database platforms, such as DB2, Informix, Netezza, and PureData. We also describe the following functionality categories:

- ▶ Basic database activity monitoring (DAM)
- ▶ Advanced database activity monitoring
- ▶ Basic vulnerability assessment
- ▶ Advanced vulnerability assessment

11.2.1 Basic database activity monitoring

At the time of this writing, Guardium defines basic DAM as listed in Table 11-1.

Table 11-1 Basic DAM functionality

Automation for system administration	Compliance and security	Automation and integration
User and roles management	Query builder	User and Role integration (LDAP, Radius, and so on)
Data level security	Report builder	Archiving integrations (Tivoli Storage Manager, Centera)
Database discovery through network scanning	Compliance workflow (audit process)	SIEM integration (QRadar, ArcSight, Envision, and so on)
Database instance discovery using agent on database server	Privacy sets	Administrative integration (GuardAPI)

Automation for system administration	Compliance and security	Automation and integration
Group management (white/black list)	Incident manager	Reports export (CEF, CSV, AXIS, SCAP, and so on)
Portal management	Predefined reports	Real-time export of security events
Self monitoring	Security Policies & Predefined Policies	Automation-Click Integration
Internal audit trail	Forensics and drill-down	Aliases

At a high level, the basic DAM functionality includes the following abilities:

- ▶ Discover new databases in the network, as shown in Figure 11-5. Guardium can scan the network to discover database servers in the network. This function helps you identify what resources can have sensitive data that should be protected.

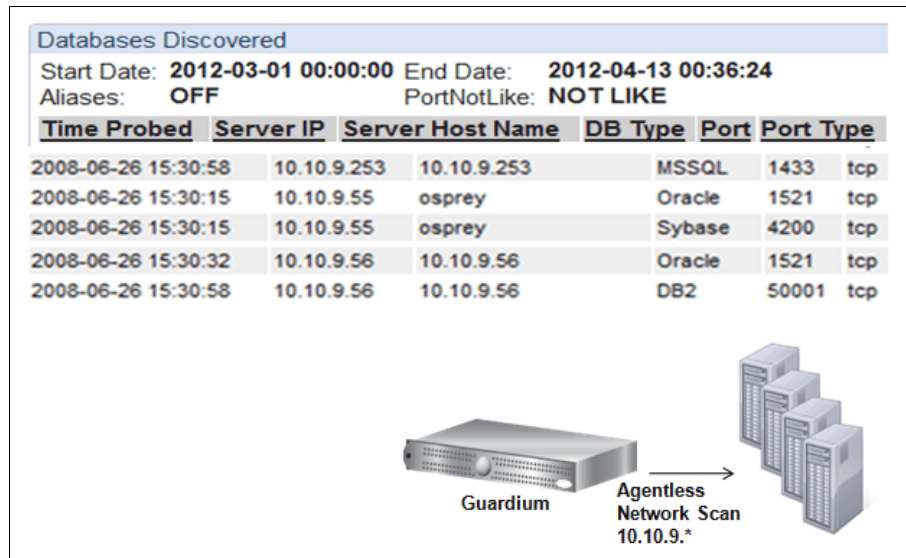


Figure 11-5 Database discovery

- ▶ Classify what sensitive data is within these databases, as shown Figure 11-6 on page 399.
- ▶ Add this sensitive information database table into the security policy to audit and monitor who has access to this information
- ▶ Send real-time alerts that are based on the policies that are defined.

- ▶ Create reports and audit policies on the information that is collected.
- ▶ Create an audit process by using these reports to validate compliance procedures are properly enforced. These audit processes include comments and escalations so that your documentation trail can be used for internal and external audits.
- ▶ Create an incident to be resolved if something violates a security policy.
- ▶ Provide forensics and drill-down reports if you must perform a postmortem of a security breach.
- ▶ Provide reports and alerts for self-monitoring the Guardium system to ensure its health and operational status.
- ▶ Provide for automation of tasks through guardAPI and click integration. This feature can link reports to automatically populating guardAPI commands and audit processes for totally automating tasks.
- ▶ Archiving audit data.

After you discover unknown databases in the network, you can review these databases to identify the location of your sensitive information. Figure 11-6 shows several ways to locate your sensitive information, from a catalog search, by permissions, or searching the actual data in the database.

- **Discover database instances on network**
- **Catalog Search: Search the database catalog for table or column name**
 - Example: Search for tables where column name is like "%card%"
- **Search by Permission: Search for the types of access that have been granted to users or roles**
- **Search for Data: Match specific values or patterns in the data**
 - Example: Search for objects matching `guardium://CREDIT_CARD` (a built-in pattern defining various credit card patterns)
- **Search for Unstructured Data: Match specific values or patterns in an unstructured data file (CSV, Text, HTTP, HTTPS, Samba)**

The figure consists of two screenshots of the 'Classification Policy Builder' interface. The top screenshot shows the 'Rule Type' dropdown menu with 'Search For Data' selected. The bottom screenshot shows the 'Search For Data' configuration page with various search criteria like 'Table Name Like', 'Data Type', and 'Search Expression'.

Figure 11-6 Searching for sensitive data

All of this information is considered basic database activity monitoring. In many organizations, there is a separate group that is responsible for the data and might feature one of the following names:

- ▶ Data architecture
- ▶ Data steward
- ▶ Application developer
- ▶ Application owner

These types of roles identify where the sensitive information is within the application and the tables that are inside the database. It is a difficult task to constantly keep this information updated and have the synergies to share this with the database administrators or security officers that must put the appropriate controls around this data.

The ability to search for data in various ways to validate that these applications are monitored correctly is a key element to reduce risk for the organization. The best security model includes with checks and balances and this should not be any different when it comes to the most value data asset in the organization, such as intellectual property, customer, credit card, and personally identifiable information (PII). This is where you might want to put more security controls around the sensitive information.

11.2.2 Advanced database activity monitoring

DAM includes the ability to have a proactive security policy to prevent unauthorized access to sensitive information. A good security policy can proactively block unwanted access and quarantine the users until the information security group can validate their intentions. For example, some organizations might outsource DBA activity. In this case, you can proactively block access to credit card information, as shown in Figure 11-7 on page 401.




Figure 11-7 Proactive blocking of unauthorized access to sensitive information

In Figure 11-7, an outsource DBA that is named Joe uses `sqlplus` to gain access to credit card information. There is a proactive security policy that prevents this access because it is beyond his job responsibilities to view information that is inside the database. In this example, you can also add another element to quarantine this individual for a predefined period. This method can prevent the user from accessing anything within the database until security can validate that their intention is not malicious. The ability to block access, even though you are the system account in Oracle, `db2inst1` in DB2, `SA` in SQL Server, and so on, is a powerful tool to help ensure that the proper security policies are in place to protect your data. This can be done with a single security policy across your heterogeneous environment, without changing your application or database configuration.

Some organizations are comfortable blocking access, while other organizations might decide to mask or redact sensitive data. If you are an organization that must comply with Payment Card Industry (PCI) regulations, you might use Dynamic Data Masking (DDM) to mask personal account numbers (PAN), as shown in Figure 11-8 on page 402. DDM masks sensitive data as needed when an unauthorized individual tries to access sensitive information. The information can be stored in plain text in the database. But when it is displayed to the unauthorized user, the sensitive information is redacted.




Figure 11-8 Dynamic Data Masking for credit card information

In Figure 11-8, the unauthorized user tries to access credit card information, but the policy that is enforced with S-TAP masks the result set so that only a partial PAN is displayed. This masking feature is an important part of the overall security policy that can be implemented to protect your sensitive data.

Blocking and masking are part of an advanced database activity monitoring strategy. After you have these components built into your security policy, you want to assess the other risks that might be available to your database servers because you now have the visibility to understand what servers have sensitive information. Guardium's Security Assessment is the next phase to this lifecycle.

11.2.3 Basic vulnerability assessment

Vulnerability assessment is the ability to identify what risks are associated with the configuration and usage of your database servers, and to provide recommendations on how to close the gaps on these identified risks.

Figure 11-9 on page 403 shows three areas of the assessment process. Vulnerability assessment architecture contains three components: database layer assessment tests, operating system assessment tests, and behavior activity assessment tests.

- **Based on industry standards: DISA STIG, CIS Benchmark and CVE**
- **Extensive Library of pre-built tests for all supported platforms**
- **Customizable tests to address your specific corporate security policies**
 - Via custom scripts, SQL queries, environment variables, etc.
- **Combination of tests ensures comprehensive coverage:**
 1. Database settings
 2. Operating system
 3. Observed behavior

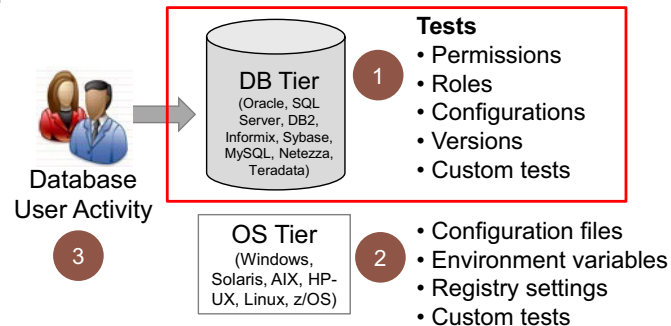


Figure 11-9 Vulnerability assessment architecture

The basic assessment of the database is done within the first tier of the assessment at the database tier. This is where database permissions, roles, configuration parameters, and the database version are defined. An example of permission might be to grant the DBA role to someone. This is a powerful permission and should not be given lightly because this person has all access to the information that is inside the database. Another common configuration is the database version. In many large organizations, it is important to understand what version and configuration of the database is running in production. The vulnerability assessment module allows you to verify this information.

11.2.4 Advanced vulnerability assessment

The more advanced vulnerability assessment includes extra elements within the assessment process.

The second level of vulnerability assessment occurs at the operating system tier. The database is like any application that runs on the operating system. There are parameters at the operating system that control the security of the database application.

These configurations must be monitored at the operating system level to ensure proper security. A good example of this is the `listener.ora` file in Oracle. If there are any changes to this file, such as `chmod 666 listener.ora` that allows anyone that can log in to the system to read and write to this file, the integrity of database might be compromised because this file identifies connection information to the Oracle database.

The third area of vulnerability assessment is concerned with the behavior or user activity of the database server. In this area, it is important to understand usage patterns to identify any potential area of compromise or misuse of the database system, as shown in the following examples:

- ▶ Some customers configure their systems so that only a single IP address can log in to the database server with DBA privileges. If a database administrator logs in to the system from many different IP addresses, this issue is a concern that highlights this account might be shared by many individuals or their security policy is compromised.
- ▶ Database activity is using privileged accounts for most of the application work. This situation violates the concept of least privileges, which states that you need only enough privilege to perform your task without having too many privileges to exceed your job responsibility.
- ▶ Excessive number of SQL Errors is another key behavioral test. This can indicate whether the application is running poorly or might be compromised by a SQL injection type of attack.

Some other advanced functionality includes entitlement reports that identify who has what permissions and roles within the database. Many auditors look to identify who is a privilege user in the database. The entitlement report allows you to easily determine this information.

Figure 11-10 on page 405 shows the heterogeneous support (DB2, Informix, MS SQL Server, MySQL, Netezza, Oracle, PostgreSQL, Sybase, and Teradata) for entitlement reports within Guardium. These reports can be automatically distributed for audit review by using the Audit Process facility to validate appropriate permissions within the database.

IBM InfoSphere® Guardium 02.24 | Edit Account

My New Reports | Standard Reports | Discover | Assess/Harden | Comply | Protect | Quick Start | Sarbanes-Oxley Accelerator | PCI Accelerator | Data Privacy Accelerator

Overview

DB Activities

Exceptions

DB Administration

Schema Changes

Detailed Activities

Performance

DB Entitlements

DB2

Informix

MS-SQL

MySQL

Netezza

Oracle

PostgreSQL

Sybase

Teradata

Access Map

ORA Obj And Columns Priv

Start Date: 2010-08-25 01:35:38 End Date: 2010-08-30 01:35:38
Aliases: ON Grantable: LIKE %

Grantee	Privilege	Table Name	Owner	Grantor	Grantable	Datasource Name	SqlGuard Timestamp	Count of ORA Obj And Columns Priv
AQ_ADMINISTRATOR_ROLE	EXECUTE	DBMS_AQ	SYS	SYS	NO	10.10.9.59-system	2010-08-27 15:02:06.0	1
AQ_ADMINISTRATOR_ROLE	EXECUTE	DBMS_AQADM	SYS	SYS	NO	10.10.9.59-system	2010-08-27 15:02:06.0	1
AQ_ADMINISTRATOR_ROLE	EXECUTE	DBMS_AQELM	SYS	SYS	NO	10.10.9.59-system	2010-08-27 15:02:06.0	1
AQ_ADMINISTRATOR_ROLE	EXECUTE	DBMS_AQIN	SYS	SYS	NO	10.10.9.59-system	2010-08-27 15:02:06.0	1
AQ_ADMINISTRATOR_ROLE	EXECUTE	DBMS_AQJMS_INTERNAL	SYS	SYS	NO	10.10.9.59-system	2010-08-27 15:02:06.0	1

Records 1 to 5 of 23997

ORA Accnts of ALTER SYSTEM

Start Date: 2010-08-25 01:35:38 End Date: 2010-08-30 01:35:38
Aliases: ON

Grantee	Privilege	Admin Option	Datasource Name	SqlGuard Timestamp	Count of ORA Accnts of ALTER SYSTEM
XDB	ALTER SESSION	NO	10.10.9.59-system	2010-08-27 15:02:05.0	1
BI	ALTER SESSION	NO	10.10.9.59-system	2010-08-27 15:02:05.0	1
SYS	ALTER SYSTEM	NO	10.10.9.59-system	2010-08-27 15:02:05.0	1
SYS	ALTER SESSION	NO	10.10.9.59-system	2010-08-27 15:02:05.0	1
SH	ALTER SESSION	NO	10.10.9.59-system	2010-08-27 15:02:05.0	1

Records 1 to 5 of 14

ORA Accnts with BECOME USER

Start Date: 2010-08-25 01:35:38 End Date: 2010-08-30 01:35:38
Aliases: ON

Grantee	Privilege	Admin Option	Datasource Name	SqlGuard Timestamp	Count of ORA Accnts with BECOME USER
DBA	BECOME USER	YES	10.10.9.59-system	2010-08-27 15:02:05.0	1
SYS	BECOME USER	NO	10.10.9.59-system	2010-08-27 15:02:05.0	1
IMP_FULL_DATABASE	BECOME USER	NO	10.10.9.59-system	2010-08-27 15:02:05.0	1

Records 1 to 3 of 3

ORA Object privileges

Start Date: 2010-08-25 01:35:38 End Date: 2010-08-30 01:35:38
Aliases: ON

Grantee	Table Name	Owner	Privilege	Datasource Name	SqlGuard Timestamp	Count of ORA Object privileges
IX	DBMS_CAPTURE_ADM	SYS	EXECUTE	10.10.9.59-system	2010-08-27 14:58:28.0	1
BI	CUSTOMERS	OE	SELECT	10.10.9.59-system	2010-08-27 14:58:28.0	1
ORDSYS	EXP_PKG_OBJ	SYS	INSERT	10.10.9.59-system	2010-08-27 14:58:28.0	1
BI	BOMBAY_INVENTORY	OE	SELECT	10.10.9.59-system	2010-08-27 14:58:28.0	1
ORDSYS	EXP_DEP_OBJ	SYS	DELETE	10.10.9.59-system	2010-08-27 14:58:28.0	1

Records 1 to 5 of 93

ORA SYSDBA and SYSOPER Accnts

Start Date: 2010-08-25 01:35:38 End Date: 2010-08-30 01:35:38
Aliases: ON

Username	Is Sysdba	Is Sysoper	Is External Password	Datasource Name	SqlGuard Timestamp	Count of ORA SYSDBA and SYSOPER Accnts
SYS	TRUE	TRUE	FALSE	10.10.9.59-system	2010-08-27 15:02:04.0	1

Records 1 to 1 of 1

Figure 11-10 Heterogeneous database entitlement report

11.3 Data lifecycle management integration

Data lifecycle management is the process of moving inactive data that still has value to long-term storage so that your production database servers can run more efficiently, as shown in Figure 11-11.




Figure 11-11 Data lifecycle management

For more information about data lifecycle management, see *Implementing an InfoSphere Optim™ data Growth Solution*, SG24-7936.

11.3.1 Identifying archive candidates

DAM can help the overall process of data lifecycle management by helping identify which tables were not accessed or used within a certain period. These tables might be good archive candidates where you can move them off your production server so that it can run faster and be more efficient. Guardium features predefined reports in which you can populate the data warehouse groups to help you use the standard reports for this purpose.

Figure 11-12 on page 407 shows an example of the predefined groups for identifying archive candidates.

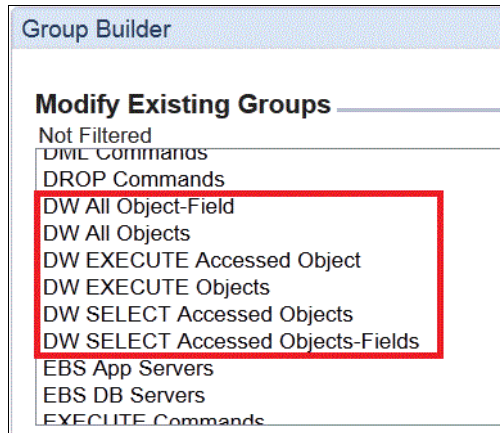


Figure 11-12 Predefined groups for identifying archive candidates

To access the predefined reports, click **View** → **Performance**, as shown in Figure 11-13. This report shows all of the objects that might be a candidate for archiving. The How-To Guide Overview of the Guardium Information Center has information about how to configure the data warehouse reports and the underlying groups, which this report is dependent upon. The guide is available at this website:

http://pic.dhe.ibm.com/infocenter/igsec/v1/topic/com.ibm.guardium91.doc/how_to/topics/how-to_guide_overview.html




Figure 11-13 Predefined data warehouse reports to help identify archive candidates

By selecting DW Select Object Access report, you can identify what objects were accessed by using a SQL Select command within the time frame on the report.

11.3.2 Auditing archive access

After the archive files are created (as shown in Figure 11-14), the InfoSphere Optim Archive application facilitates application access to the archive files. From a security perspective, it is important to audit who accesses these archive files because the same sensitive information is in the archive files and the production databases. When someone accesses these archive files, Optim Archive can send this information to Guardium, as shown in Figure 11-14.




Figure 11-14 Integrated with IBM InfoSphere Optim Data Growth Solution

There is a predefined role within the Guardium system that is called *optim-audit*. The Accessmgr can assign the optim-audit role to people who are interested in viewing the archive access reports, as shown in Figure 11-15.

The screenshot shows the 'Access Management' interface with the 'Data Security' tab selected. The 'User Role Form' is displayed for 'Jack Auditor'. A table lists various roles with checkboxes in the 'Assign' column. The 'optim-audit' role is highlighted with a red box and has its checkbox checked. Other roles with checked checkboxes include 'pci', 'sox', and 'user'. Roles with unchecked checkboxes include 'accessmgr', 'admin', 'appdev', 'audit', 'Baselll', 'cas', 'cli', 'DataPrivacy', 'datasec-exempt', 'dba', 'diag', 'infosec', 'inv', 'netadm', 'review-only', and 'SimpleView'. The interface also includes a 'Save' button and a 'Back' button at the bottom.

Role Name	Assign
accessmgr	<input type="checkbox"/>
admin	<input type="checkbox"/>
appdev	<input type="checkbox"/>
audit	<input type="checkbox"/>
Baselll	<input type="checkbox"/>
cas	<input type="checkbox"/>
cli	<input type="checkbox"/>
DataPrivacy	<input type="checkbox"/>
datasec-exempt	<input type="checkbox"/>
dba	<input type="checkbox"/>
diag	<input type="checkbox"/>
infosec	<input type="checkbox"/>
inv	<input type="checkbox"/>
netadm	<input type="checkbox"/>
optim-audit	<input checked="" type="checkbox"/>
pci	<input checked="" type="checkbox"/>
review-only	<input type="checkbox"/>
SimpleView	<input type="checkbox"/>
sox	<input checked="" type="checkbox"/>
user	<input checked="" type="checkbox"/>

Figure 11-15 Optim-audit role within Accessmgr account

The optim-audit role has a predefined report about who accessed the archive logs. If the user is granted access to this role, these predefined reports are available in the user's portal.

Figure 11-16 shows that the predefined Optim Archive reports are available when the optim-audit role is available to the user.

Details	Request Type	Request Name	Optim Server	User	Start Time	Errors	Elapsed Time
12	ODM	OPTIM_AUDIT	(LOCAL)	SYSTEM	2013-03-28 12:58:37.00	0	
10	ODM	OPTIM_AUDIT	(LOCAL)	NCOSENZ	2013-03-28 12:54:21.00	0	
8	ODM	OPTIM_AUDIT	(LOCAL)	SYSTEM	2013-03-28 12:51:16.00	0	
6	ODM	OPTIM_AUDIT	(LOCAL)	SYSTEM	2013-03-28 12:44:06.00	0	
4	ODM	OPTIM_AUDIT	(LOCAL)	NCOSENZ	2013-03-27 12:45:32.00	0	
2	ODM	OPTIM_AUDIT	(LOCAL)	NCOSENZ	2013-03-27 12:24:01.00	0	

Figure 11-16 Predefined Optim Archive access reports

For Optim Archive to send this information to Guardium, you must enable the auditing within the Optim Archive application by completing the following steps:

1. Select **Audit Selection...** from General tab of product options.
2. In Audit Facility window, select **Hosted by Guardium** or **Hosted by Optim/Guardium** from the Audit Status drop-down menu.
3. Click **Guardium Settings...**
4. In Advanced Setting (Guardium setting), provide the IP address and DNS name of the Guardium appliance to send the audit information to be stored.

Figure 11-17 shows the Optim Archive configuration to send audit information to Guardium.

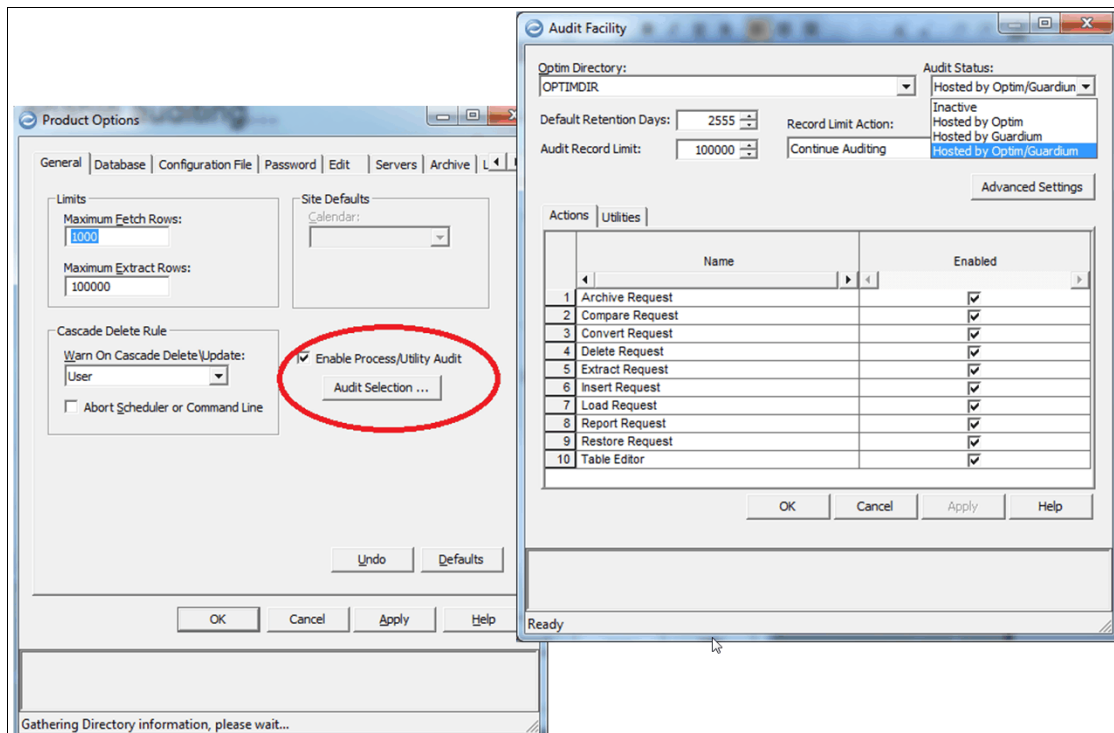


Figure 11-17 Optim Archive configuration to send audit information to Guardium

11.3.3 Test data management integration

Another aspect of data lifecycle management is test data management (TDM). This is the process to take a subset of information from your production server so that you can provide testing and application development data. Figure 11-18 on page 412 shows the test data management process.

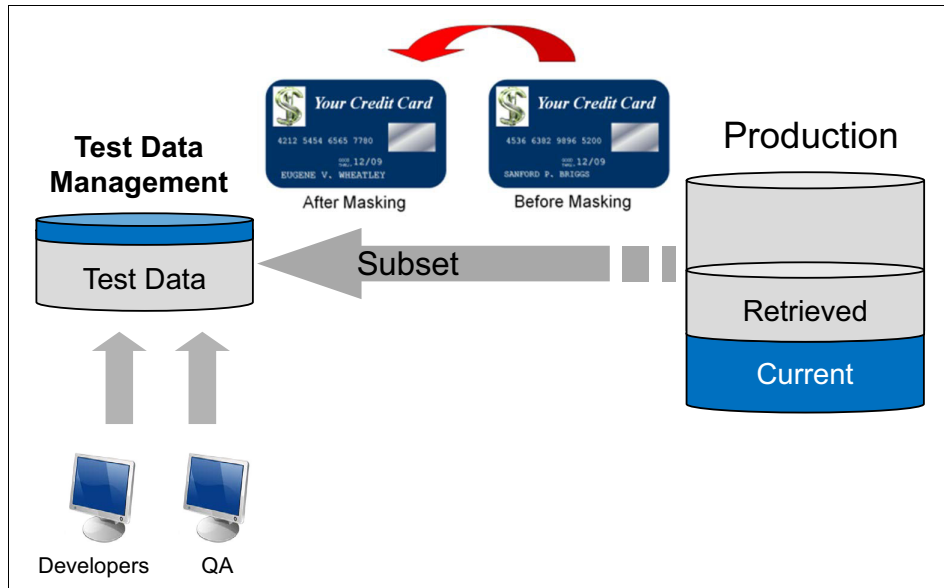


Figure 11-18 The test data management process

During this process of taking a subset of your production data, you can apply masking policies so that you can desensitize data. This process is sometimes referred to as static data masking (SDM) because you statically mask the data when you extract it from your production database server. This is important so that you do not have actual sensitive data in your test and application development environment where there might not be the strict security policies like your production environment.

IBM Optim Test Data Management Solution and Guardium can share these masking policies. If an unauthorized user accesses sensitive data on your production server, you can perform dynamic data masking (DDM) to redact the result set to protect this information. Guardium can be used for dynamic data masking so that this unauthorized access can be dynamically masked similar to your test and development environment. You can configure this feature by exporting and importing the policy from Guardium and IBM Optim with eXtensible Access Control Markup Language (XACML), which is an industry-standard access control policy language.

Figure 11-19 shows an example of importing and exporting masking policies with Guardium with XACML.




Figure 11-19 Import and export masking policies with Guardium

Masking policies assume that you understand where your sensitive data is. It can be difficult to track this information. This is why Guardium and InfoSphere Discovery can help automate the process of finding your sensitive information.

11.3.4 Exchanging the sensitive information location with InfoSphere Discovery

There is a section in the online Help for Guardium that describes the process to import and export the Guardium classification results, which identifies where your sensitive data is, as shown in Figure 11-20 on page 414. The sensitive data can be imported or exported from InfoSphere Discovery.

How to transfer sensitive data

Take sensitive data information, identified and classified in one InfoSphere® product and transfer that information to another InfoSphere product, such as InfoSphere Discovery.

Both InfoSphere Guardium® and InfoSphere Discovery have the capability to identify and classify sensitive data, such as Social Security Numbers or credit card numbers.

A customer of the InfoSphere Guardium product can use a bidirectional interface to transfer identified sensitive data information from one product to another. Those customers who have already invested the time in one InfoSphere product can transfer the information to the other InfoSphere product.

Note: In InfoSphere Guardium, the Classification process is an ongoing process that runs periodically. In InfoSphere Discovery, Classification is part of the Discovery process that usually runs once.

Note: The data will be transferred via CSV files.

The summary of Export/Import procedures is as follows:

- Export from Guardium - Run the predefined report (Export Sensitive Data to Discovery) and export as CSV file.
- Import to Guardium - Load to a custom table against CSV datasource; define default report against this datasource.

Follow these steps:

1. Export from Guardium - Export Classification Data from InfoSphere Guardium to InfoSphere Discovery
2. As an admin user in the Guardium application, go to Tools > Report Building > Classifier Results Tracking > Select a Report > Export Sensitive Data to Discovery (See screenshot).

Note: Add this report to the UI pane (it is not by default).

The screenshot shows the 'Report Building' interface in InfoSphere Guardium. The main window is titled 'Export Sensitive Data To Discovery'. It features a 'Query Fields' table and a 'Query Conditions' section.

Seq	Entity	Attribute	Field Mode	Order-by	Sort Rank	Descend
<input type="checkbox"/>	1	Classification Process Results	Data Source Type	Value		<input type="checkbox"/>
<input type="checkbox"/>	2	Classification Process Results	Host	Value		<input type="checkbox"/>
<input type="checkbox"/>	3	Classification Process Results	Port	Value		<input type="checkbox"/>
<input type="checkbox"/>	4	Classification Process Results	DB Name	Value		<input type="checkbox"/>
<input type="checkbox"/>	5	Classification Process Results	Schema	Value		<input type="checkbox"/>
<input type="checkbox"/>	6	Classification Process Results	Service Name	Value		<input type="checkbox"/>
<input type="checkbox"/>	7	Classification Process Results	Table Name	Value		<input type="checkbox"/>
<input type="checkbox"/>	8	Classification Process Results	Column Name	Value		<input type="checkbox"/>

Below the table, there are options for 'Addition mode' (AND, OR, HAVING) and a 'Query Conditions' section with a table:

	Entity	Attr	Attribute	Operator	Runtime Param.
<input type="checkbox"/>	WHERE	Classification Process Results	Table Name	LIKE	Parameter tableNameLike
<input type="checkbox"/>	AND	Classification Process Results	Schema	LIKE	Parameter schemaLike

Figure 11-20 How to transfer sensitive data

InfoSphere Discovery can then use this information to build the business object, which might be required to understand the data model for archiving database objects.

Figure 11-21 on page 415 shows that InfoSphere Discovery can exchange the sensitive data location with Guardium so that appropriate security and audit policies can be applied according to corporate security standards.




Figure 11-21 Exchange sensitive data with InfoSphere Discovery

Now that we understand where our sensitive data is, it is important to identify who is using it. In many cases, this can be a challenge, especially in a three-tier environment.

11.4 Identifying the user activity

After you review Figure 11-22, how can you tell who performed the transactions to the database, Joe or Bob?




Figure 11-22 Identifying the user in a three-tier environment

To answer the question, you need more information because the only piece of information you have is that Apps (the database user) performed some transactions to the database. So, you do not know whether it was Joe or Bob. This is a common scenario in which a single database user shares this connection to the database. Guardium addresses this problem by using the following methods:

- ▶ Custom identification procedures
- ▶ GuardAppEvents and GuardAppUser
- ▶ Set client user
- ▶ WebSphere application user information
- ▶ CICS® application user

11.4.1 Custom identification procedures

The use of custom identification procedures is an option if the application was written so that it calls a stored procedure when it opens a new connection. This stored procedure might be called to set the user name, time zone, and language, for example. If so, we can scrape the user information by identifying one of the parameters of the stored procedure when it is run. Figure 11-23 shows how to configure custom identification procedures for the stored procedure AppEndUser.

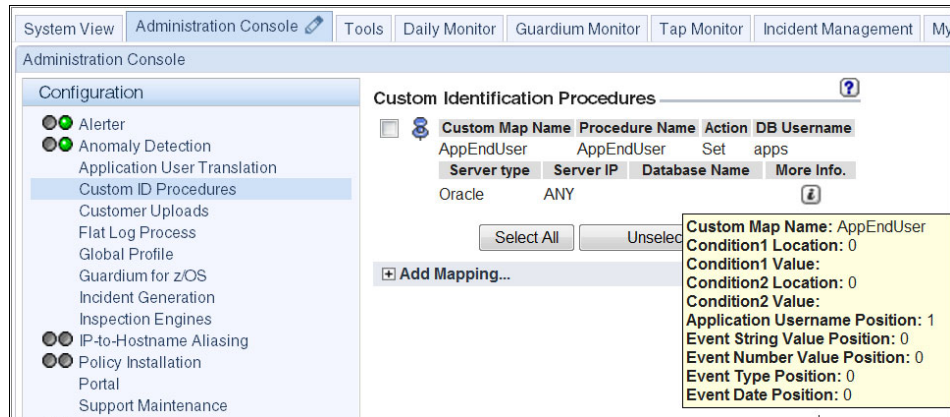


Figure 11-23 Custom Identification procedures

In this example, the stored procedure is called so that we extract the actual user in position 1 of the stored procedure execution. This is identified as Application Username Position:1 in Figure 11-23. The result of configuring custom identification procedures is that you can now uniquely and deterministically identify who performed the transactions.

In Figure 11-24 on page 418, under the Application User field in the audit report, “joe” performed some transactions and “bob” performed other transactions to the database based on the pooled application user “apps”. This is determined when the stored procedure, AppEndUser, is run and extracting the first parameter ‘joe’ or ‘bob’.

```

[root@osprey ~]# sqlplus apps/guardium
SQL*Plus: Release 11.2.0.2.0 Production on Fri Jul 5 17:41:01 2013
Copyright (c) 1982, 2011, Oracle. All rights reserved.

Connected to:
Oracle Database 11g Express Edition Release
11.2.0.2.0 Production on Fri Jul 5 17:41:01 2013

SQL> set serveroutput on
SQL> exec AppEndUser ('joe');
joe
PL/SQL procedure successfully completed.

SQL> create table joe (i int);
Table created.

SQL> insert into joe values(10);
1 row created.

SQL>
SQL> exec AppEndUser ('bob');
bob
PL/SQL procedure successfully completed.

SQL> create table bob (i int);
Table created.

SQL> insert into bob values(20);
1 row created.

SQL>

```

Client IP	Server IP	DB User Name	Application User	Full Sql
10.10.9.56	10.10.9.56	APPS	joe	BEGIN AppEndUser ('joe'); END;
10.10.9.56	10.10.9.56	APPS	joe	create table joe (i int)
10.10.9.56	10.10.9.56	APPS	joe	insert into joe values(10)
10.10.9.56	10.10.9.56	APPS	bob	BEGIN AppEndUser ('bob'); END;
10.10.9.56	10.10.9.56	APPS	bob	create table bob (i int)
10.10.9.56	10.10.9.56	APPS	bob	insert into bob values(20)

Figure 11-24 Extracting extra information to identify the unique application user

11.4.2 GuardAppEvents and GuardAppUser

If the application does not have a stored procedure, you can insert dummy SQL statements into your application to provide the extra context of who is performing the transaction. This SQL statement is placed where you open a new connection to the database. The following example is Oracle that needs the dual dummy database table name in the SQL statement:

```

SELECT 'GuardAppEvent:START', 'GuardAppEventUsername:Joe',
'GuardAppEventType:yourStringHere1',
'GuardAppEventStrValue:yourStringHere2', 'GuardAppEventNumValue:4321'
FROM Dual;

```

Figure 11-25 shows that GuardAppEvents can be used to provide user identity in a pooled database user environment.

```
SELECT 'GuardAppEvent:START', 'GuardAppEventUsername:Joe', 'GuardAppEventType:yourStringHere1',
'GuardAppEventStrValue:yourStringHere2', 'GuardAppEventNumValue:4321' FROM Dual;

drop table joe;

SELECT 'GuardAppEvent:START', 'GuardAppEventUsername:Bob', 'GuardAppEventType:yourStringHere1',
'GuardAppEventStrValue:yourStringHere2', 'GuardAppEventNumValue:1234' FROM Dual;

drop table bob;
```


-AppEventUser							
Start Date:	2013-07-05 15:21:42	End Date:	2013-07-05 18:21:42				
Aliases:	OFF	ClientPLike:	LIKE %				
DBUserLike:	LIKE apps	FullSQLLike:	LIKE %				
ServerPLike:	LIKE %	ServerTypeLike:	LIKE %				
SourcePrgLike:	LIKE %	netProtoLike:	LIKE %				
Client IP	Server IP	DB User Name	Event User Name	Event Type	Event Value Str	Event Value Num	Full Sql
10.10.9.240	10.10.9.56	APPS	Joe	yourStringHere1	yourStringHere2	4321	SELECT 'GuardAppEvent:START', 'GuardAppEventUsername:Joe', 'GuardAppEventType:yourStringHere1', 'GuardAppEventStrValue:yourStringHere2', 'GuardAppEventNumValue:4321' FROM Dual
10.10.9.240	10.10.9.56	APPS	Joe	yourStringHere1	yourStringHere2	4321	drop table joe
10.10.9.240	10.10.9.56	APPS	Bob	yourStringHere1	yourStringHere2	1234	SELECT 'GuardAppEvent:START', 'GuardAppEventUsername:Bob', 'GuardAppEventType:yourStringHere1', 'GuardAppEventStrValue:yourStringHere2', 'GuardAppEventNumValue:1234' FROM Dual
10.10.9.240	10.10.9.56	APPS	Bob	yourStringHere1	yourStringHere2	1234	drop table bob

Figure 11-25 GuardAppEvents can be used to provide extra context in a pooled database user connection

In the report that is shown in Figure 11-25, the following points are important:

- ▶ These transactions belong to the same database session ID, 1059 in this case. This is typical of a pooled database user connection.
- ▶ The DB User Name is “APPS”, which is the only database user that is defined for this transaction.
- ▶ Each of the database transactions can now be identified uniquely in the Event User Name field for “Joe” and “Bob”.
- ▶ You can add context through the GuardAppEventType, GuardAppEventStrValue, and GuardAppEventNumValue. Some customers add client IP, client host name, and other relevant information that is specific to their application.
- ▶ The following syntax is for IBM DB2:

```
SELECT 'GuardAppEvent:START', 'GuardAppEventUsername:Joe',
'GuardAppEventType:yourStringHere1',
'GuardAppEventStrValue:yourStringHere2',
'GuardAppEventNumValue:4321' FROM SYSIBM.SYSDUMMY1
```

- ▶ The following syntax is for Microsoft SQL Server and Sybase:

```
SELECT 'GuardAppEvent:START', 'GuardAppEventUsername:Joe',
'GuardAppEventType:yourStringHere1',
'GuardAppEventStringValue:yourStringHere2',
'GuardAppEventNumValue:4321'
```

GuardAppUser is similar, but you cannot add contextual information as you can with GuardAppEvents. The following syntax is for GuardAppUser DB2:

```
Select 'GuardAppUser:Joe' From Sysibm.Sysdummy1
```

Figure 11-26 shows a GuardAppUser DB2 example to help identify unique users in a pooled connection environment.

The terminal window shows the following commands and output:

```
[db2inst1@osprey ~]$ db2 connect to sample2 user apps using guardium
Database Connection Information
Database server = DB2/LINUX8664 10.1.0
SQL authorization ID = APPS
Local database alias = SAMPLE2

[db2inst1@osprey ~]$ db2 "Select 'GuardAppUser:Joe' From Sysibm.Sysdummy1"
1
-----
GuardAppUser:Joe
1 record(s) selected.

[db2inst1@osprey ~]$ db2 "drop table joe"
DB20000I The SQL command completed successfully.

[db2inst1@osprey ~]$ db2 "Select 'GuardAppUser:Bob' From Sysibm.Sysdummy1"
1
-----
GuardAppUser:Bob
1 record(s) selected.

[db2inst1@osprey ~]$ db2 "create table bob (i int)"
DB20000I The SQL command completed successfully.

[db2inst1@osprey ~]$ db2 terminate
DB20000I The TERMINATE command completed successfully.

[db2inst1@osprey ~]$
```

The GuardAppUser application log window shows the following session details:

Session ID	Client IP	Server IP	DB User Name	Application User	Full Sql	Timestamp
1152	10.10.9.56	10.10.9.56	APPS	Joe	Select 'GuardAppUser:Joe' From Sysibm.Sysdummy1	2013-07-09 10:17:18.0
1152	10.10.9.56	10.10.9.56	APPS	Joe	drop table joe	2013-07-09 10:17:31.0
1152	10.10.9.56	10.10.9.56	APPS	Bob	Select 'GuardAppUser:Bob' From Sysibm.Sysdummy1	2013-07-09 10:17:41.0
1152	10.10.9.56	10.10.9.56	APPS	Bob	create table bob (i int)	2013-07-09 10:18:06.0

Figure 11-26 GuardAppUser with DB2 example

GuardAppUser is an effective way to communicate who is the application user if you do not want to provide any other information. In the report, the application user is shown in the **Access Period** → **Application User** field.

11.4.3 Setting client user ID

Some applications are written so that they provide the extra information within the pooled connection; for example, PeopleSoft, SAP, or Siebel. In these cases, the extra contextual information for the application user is within the **Access Period** → **Application User** field, as shown in Figure 11-27.

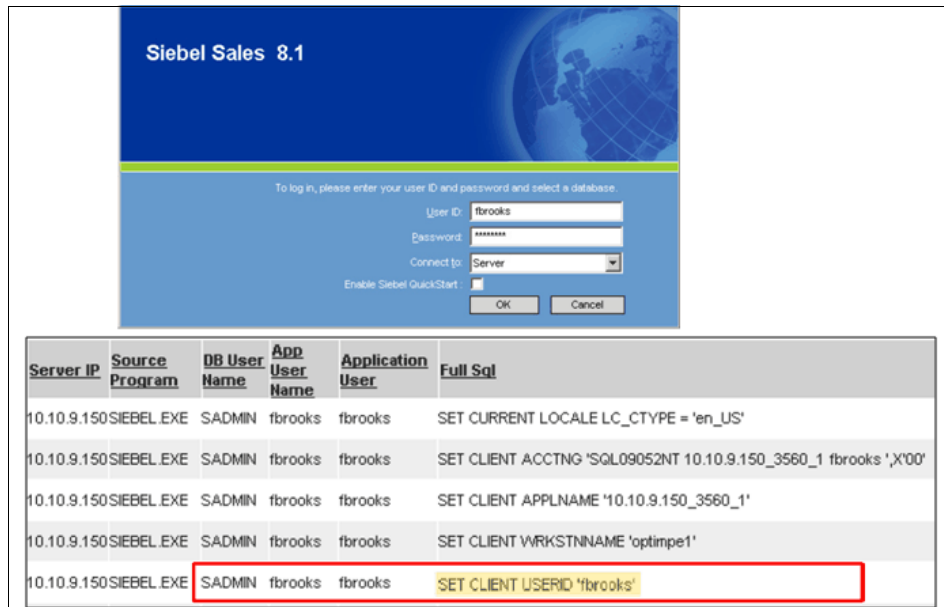


Figure 11-27 Application User information that is sent by the application

Figure 11-27 also shows that the Application User information is sent by the application during the new connection to the database. This depends on how the application is written to determine whether this other context is provided.

Depending on the SAP version, it also sends more information to help identify the SAP application user that is shown in Figure 11-28 on page 422. SAP application user DDIC performs an SU01 T-Code to add a user. This information can be audited by Guardium without any application changes.

For more information, see *Monitor database activity for application users with Guardium and WebSphere Application Server*, which is available at this website:

<http://www.ibm.com/developerworks/data/library/techarticle/dm-1208monitordbactivity/index.html>

If you are connecting to a DB2 database with WebSphere, you also might want to provide DB2 Client Info context. In this example, you can add code after you open the connection to the database server, as shown in Figure 11-29.

Timestamp	Server Type	Server IP	Client IP	Network Protocol	DB User Name	Full Sql	DB2 Client Info
2012-01-18 14:16:39.0	DB2	10.10.9.5	10.10.9.240	DRDA	SERVERIBMUSER	SELECT cardid,name from creditcard	WSUSER=Joe DiPietro,WRKSTN=IP=1.1.1.1,APPL=JoeApplication

2 – Execute the SQL Statement

```
// Execute SQL to send additional client information to database
System.out.println("Executing SQL Statement...");
try {
    stmt = conn.createStatement();
} catch (SQLException e1) {
    // TODO Auto-generated catch block
    e1.printStackTrace();
}
try {
    rs = stmt.executeQuery("SELECT cardid,name from creditcard");
} catch (SQLException e) {
    // TODO Auto-generated catch block
    e.printStackTrace();
}

// Loop through the result set printing each line
System.out.println("Before result sets\n");
System.out.println("Cardid NAME");
System.out.println("-----");

try {
    while (rs.next()) {
        Integer i = rs.getInt("CARDID");
        String name = rs.getString("NAME");
        System.out.println(i + " " + name);
    }
} catch (SQLException e) {
    // TODO Auto-generated catch block
    e.printStackTrace();
}
```

1 – Set the client info information

```
// Set the additional client information to be passed to DB2
System.out.println("Setting DB2 Client Info...");
try {
    conn.setClientInfo("ClientUser", "Joe DiPietro");
    conn.setClientInfo("ApplicationName", "JoeApplication");
    conn.setClientInfo("ClientHostName", "IP=1.1.1.1");
} catch (SQLClientInfoException e1) {
    // TODO Auto-generated catch block
    e1.printStackTrace();
}
System.out.println("Set DB2 Client Info Successful...");
```

Figure 11-29 DB2 Client Info parameters to provide extra application context

11.4.5 CICS application user

As you can see, every application and environment can be slightly unique when it comes to adding information to identify the true application user. This is no exception when it comes to the mainframe environment and CICS. There is a configuration for the DB2conn within the CEDDA utility where you can tell CICS to send more information to DB2 so that Guardium can extract the CICS application user. As shown in Figure 11-30 on page 424, the key to this configuration is to ensure that the AuthType is set for UserID or Group.

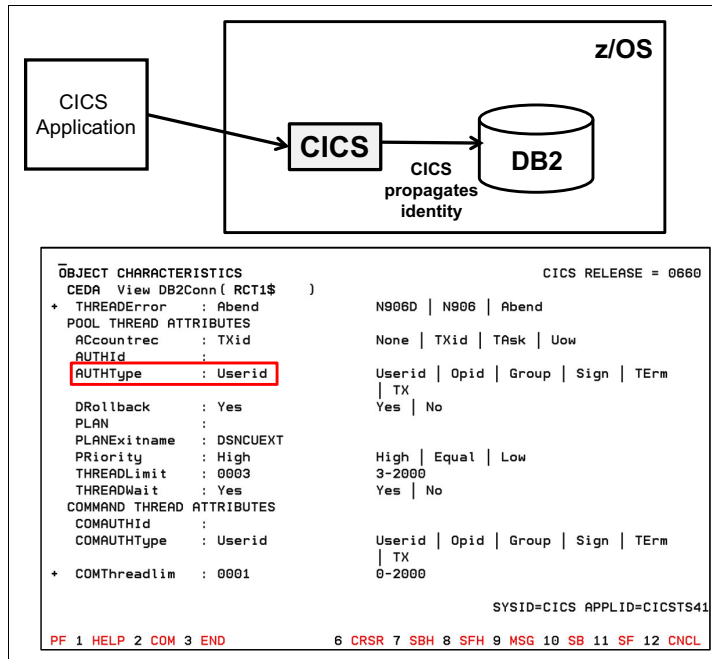


Figure 11-30 CICS configuration to send more information to identify the application user

In some cases, custom applications in the CICS environment might use Cobol to interact with DB2. These programs use more static SQL versus dynamic SQL. For these types of applications, you can also use the GuardAppEvent with Static SQL and bind variables. The key is that this information is sent to the database.

The following static SQL statement is used:

```
SELECT 'GuardAppEvent:Start','GuardAppEventBindvalue', cast(? as
char(100)), cast(? as char(100)), cast(? as char(100)) cast(? as
char(100)), cast(? as char(100)) FROM SYSIBM.SYSDUMMY1
```

Where the bind variables are mapped as follows (in this sequence), as shown in Figure 11-31 on page 425:

- ▶ Event User Name: JoeD
- ▶ Event Value Str: EventStrValue:RECONCILE
- ▶ Event Type: EventType:ChangeRequest
- ▶ Event Value Num:1281

GuardAppEvents with Bind Variables (Static SQL)

Full_Sql						
SELECT 'GuardAppEvent.Start', 'GuardAppEvent.Bindvalue', cast(? as char(100)), cast(? as char(100)), cast(? as char(100)), cast(? as char(100)) FROM SYSIBM.SYSDUMMY1						
Bind Variables Values	Event User Name	Event Value Str	Event Type	Event Value Num	Session Id	Source Program
'JoeD', 'EventStrValue.RECONCILE', 'EventType.ChangeRequest', '1281', JoeD	JoeD	EventStrValue.RECONCILE	EventType.ChangeRequest	1281	8	IBM-6CD8U7TIXYD.DB2JCC_APPLI

Figure 11-31 GuardAppEvents with static SQL and bind variables

In Figure 11-31, the extra application information is sent to the database, and Guardium extracts these bind variables into the Application Events fields (Event Username, Event Value Str, Event Type, Event Value Num). This information is carried through the session when the next SQL statement is run, (Select * from CreditCard where CardID=? and Name like?), as shown in Figure 11-32. The Application Events are carried throughout the same session (session ID 8).

Full_Sql						
SELECT 'GuardAppEvent.Start', 'GuardAppEvent.Bindvalue', cast(? as char(100)), cast(? as char(100)), cast(? as char(100)), cast(? as char(100)) FROM SYSIBM.SYSDUMMY1						
SELECT * FROM creditcard where CardID=? and Name like ?						
Bind Variables Values	Event User Name	Event Value Str	Event Type	Event Value Num	Session Id	Source Program
'JoeD', 'EventStrValue.RECONCILE', 'EventType.ChangeRequest', '1281', JoeD	JoeD	EventStrValue.RECONCILE	EventType.ChangeRequest	1281	8	IBM-6CD8U7TIXYD.DB2JCC_APPLI
0, 'User%'	JoeD	EventStrValue.RECONCILE	EventType.ChangeRequest	1281	8	IBM-6CD8U7TIXYD.DB2JCC_APPLI

Figure 11-32 GuardAppEvents with static SQL and bind variables Session ID 8

Providing more application information can help identify the unique users that provided the transaction. You also can have WebSphere propagate its user identity to CICS and then propagate this information to DB2. For more information about the WebSphere to CICS identity propagation, see this website:

http://pic.dhe.ibm.com/infocenter/cicsts/v4r1/index.jsp?topic=%2Fcom.ibm.cics.ts.doc%2Fdfht5%2Ftopics%2Fidprop_intro.html

Now that we described some of the background information, we describe the following other IBM products that feature built-in application user information.

11.4.6 InfoSphere MDM and Data Stage application user

As sensitive data is being used for the “single version of the truth”, it is important to understand who touches this information. Figure 11-33 shows the application user “barneyrubble” reviewed some sensitive data.

Full SQL	Timestamp	Application User	Client IP	DB User Name	Session Start	Source Program	Full Sql	
	1586	2010-04-12 15:53:46	barneyrubble	169.254.198.66	DB2ADMIN	2010-04-12 14:50:28	DB2JCC_APPLICATION	SELECT 'GuardAppUser.barneyrubble' FROM SYSIB...
	1587	2010-04-12 15:53:46	barneyrubble	169.254.198.66	DB2ADMIN	2010-04-12 14:50:28	DB2JCC_APPLICATION	SELECT 'GuardAppUser.barneyrubble' FROM SYSIB...
	1588	2010-04-12 15:53:46	barneyrubble	169.254.198.66	DB2ADMIN	2010-04-12 14:50:28	DB2JCC_APPLICATION	SELECT CONTACT.CONT_ID , CONTACT.ACCE_C...
	1589	2010-04-12 15:53:46	barneyrubble	169.254.198.66	DB2ADMIN	2010-04-12 14:50:28	DB2JCC_APPLICATION	SELECT CONTACT.CONT_ID , CONTACT.ACCE_C...
	1590	2010-04-12 15:53:46	barneyrubble	169.254.198.66	DB2ADMIN	2010-04-12 14:50:28	DB2JCC_APPLICATION	SELECT PERSONNAME.PERSON_NAME_ID , PE...
	1591	2010-04-12 15:53:46	barneyrubble	169.254.198.66	DB2ADMIN	2010-04-12 14:50:28	DB2JCC_APPLICATION	SELECT PERSONNAME.PERSON_NAME_ID , PE...
	1592	2010-04-12 15:53:46	barneyrubble	169.254.198.66	DB2ADMIN	2010-04-12 14:50:28	DB2JCC_APPLICATION	SELECT IDENTIFIER.IDENTIFIER_ID , IDENTIFIE...
	1593	2010-04-12 15:53:46	barneyrubble	169.254.198.66	DB2ADMIN	2010-04-12 14:50:28	DB2JCC_APPLICATION	SELECT IDENTIFIER.IDENTIFIER_ID , IDENTIFIE...
	1594	2010-04-12 15:53:46	barneyrubble	169.254.198.66	DB2ADMIN	2010-04-12 14:50:28	DB2JCC_APPLICATION	SELECT LOBREL.LOB_REL_ID , LOBREL.ENTIT...
	1595	2010-04-12 15:53:46	barneyrubble	169.254.198.66	DB2ADMIN	2010-04-12 14:50:28	DB2JCC_APPLICATION	SELECT LOBREL.LOB_REL_ID , LOBREL.ENTIT...
	1596	2010-04-12 15:53:46	barneyrubble	169.254.198.66	DB2ADMIN	2010-04-12 14:50:28	DB2JCC_APPLICATION	SELECT PRIVPREF.PPREF_ID , PRIVPREF.VALU...
	1597	2010-04-12 15:53:46	barneyrubble	169.254.198.66	DB2ADMIN	2010-04-12 14:50:28	DB2JCC_APPLICATION	Select default_src_val_id, entity_name, instance_pk...
	1598	2010-04-12 15:53:46	barneyrubble	169.254.198.66	DB2ADMIN	2010-04-12 14:50:28	DB2JCC_APPLICATION	SELECT 'GuardAppUserReleased' FROM SYSIBM...
	1599	2010-04-12 15:53:46	barneyrubble	169.254.198.66	DB2ADMIN	2010-04-12 14:50:28	DB2JCC_APPLICATION	SELECT 'GuardAppUserReleased' FROM SYSIBM...
	1600	2010-04-12 15:53:46		169.254.198.66	DB2ADMIN	2010-04-12 14:50:28	DB2JCC_APPLICATION	SELECT CONTACT.CONT_ID , CONTACT.ACCE_C...
	1601	2010-04-12 15:53:46		169.254.198.66	DB2ADMIN	2010-04-12 14:50:28	DB2JCC_APPLICATION	SELECT CONTACT.CONT_ID , CONTACT.ACCE_C...
	1602	2010-04-12 15:53:46		169.254.198.66	DB2ADMIN	2010-04-12 14:50:28	DB2JCC_APPLICATION	SELECT CONTACT.CONT_ID , CONTACT.ACCE_C...

Figure 11-33 MDM integration with application user

Within an IBM InfoSphere Master Data Management (InfoSphere MDM) environment, Guardium can also be used to identify the following unauthorized activities in InfoSphere MDM environments:

- ▶ Direct access to InfoSphere MDM database by privileged users, such as SysAdmins, DBAs, developers, and outsourced personnel
- ▶ Sharing of service account credentials
- ▶ Unusual number of SQL errors and failed logins to InfoSphere MDM database
- ▶ Configuration changes to InfoSphere MDM database (ports, privileges, accounts, and so on)

Figure 11-34 shows a high-level summary of the goal of integrated solution.

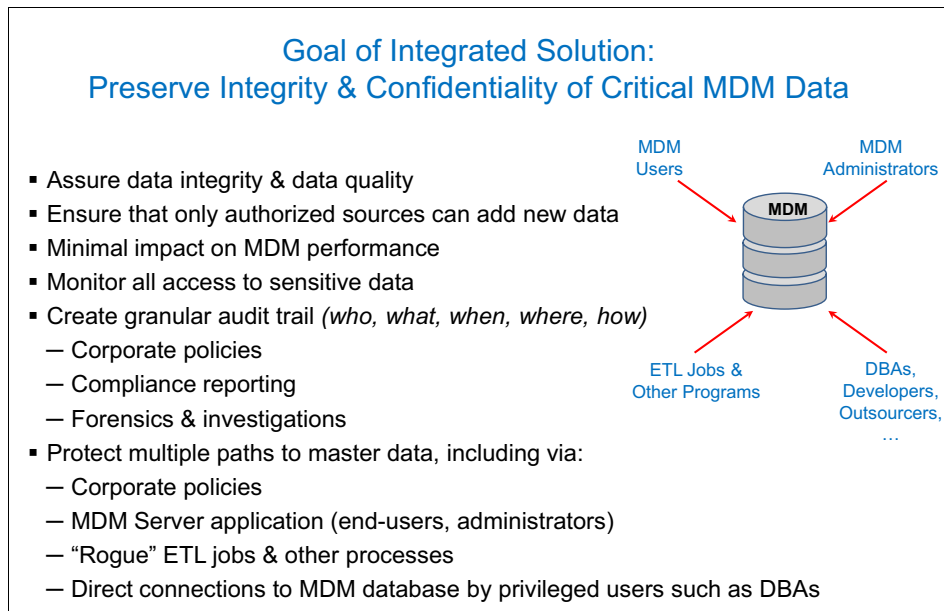


Figure 11-34 Preserve the integrity and confidentiality of InfoSphere MDM data

11.4.7 Cognos application user

Another sensitive data store that is used for powerful analytics is Cognos®. There is much information that must be protected in this environment, and you can review the Cognos application as shown in Figure 11-35 on page 428.

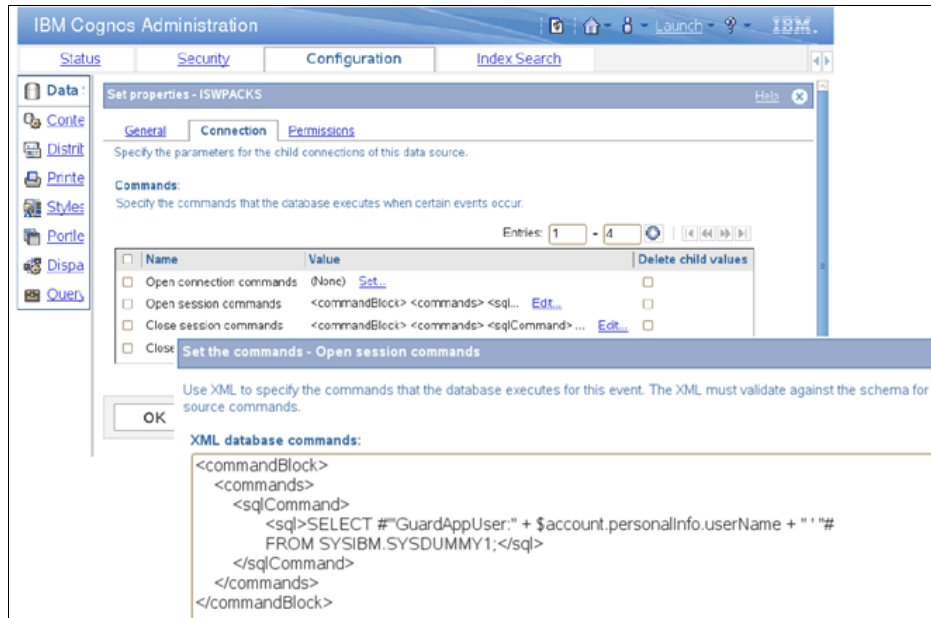


Figure 11-35 Guardium and Cognos application user configuration

We use the GuardAppUser SQL Statement with the Cognos `$account.personalInfo.userName` variable to send the extra application user information to the database so that Guardium can deterministically identify the user of the Cognos transactions. In this example, we have the luxury of a well-written application that allows you to send SQL statements after the pooled connection is opened. You can see this in the Open Session Commands configuration within the Connection parameters of the Cognos Administration window that is shown in Figure 11-35.

11.5 Security integration

There are many elements within the security intelligence framework that integrate with Guardium. In this section, we describe the QRadar Security Information Event Management (SIEM) integration.

11.5.1 Guardium and QRadar integration

For more information, see the following YouTube video:

<http://youtu.be/MOP12R2Kkjc>

QRadar features the following integration areas:

- ▶ Real-time alerts
- ▶ Asset Export Information Source (AXIS)
- ▶ Security Compliance Automation Protocol (SCAP)

Real-time alerts is the process of identifying a security policy violation within the database or Big Data environment that needs immediate security attention. These violations are defined within the policy on the Guardium system. After the policies are defined, the Guardium collector can send a real-time alert in the Log Event Extended Format (LEEF) through a syslog connection to the QRadar system, as shown in Figure 11-36.

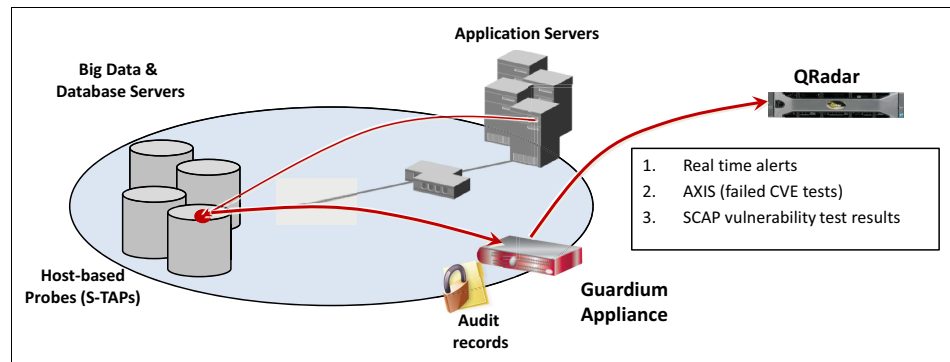


Figure 11-36 QRadar and Guardium integration

A SIEM system provides real-time analysis of security events that are generated by various devices on your network. Figure 11-37 on page 430 shows some of these types of devices and the information that is generated through a Guardium solution.




Figure 11-37 SIEM Systems

The challenge with a SIEM is to take millions of security events and reduce them to actionable incidents. In general, you do not want to forward all of the database and Big Data activities to the SIEM because doing so generates too much noise. However, you want to send a subset of information that violates a security policy, as shown in the following examples:

- ▶ SQL errors on a production database server
- ▶ Failed login to the database after five attempts within 5 minutes
- ▶ Unauthorized users who are accessing credit card information that is stored in the Big Data or database server

For these types of security events, the Guardium system sends a real-time alert to QRadar for investigation, as shown in Figure 11-38.

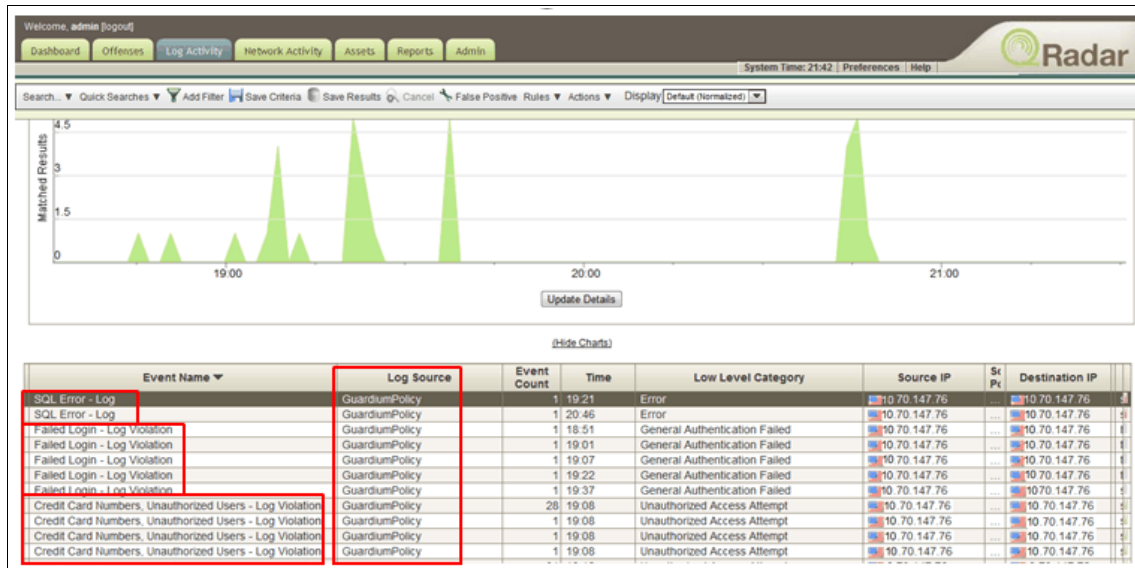


Figure 11-38 Real-time security policy violations that are sent from Guardium to QRadar

These events are sent to the QRadar system in the LEEF format, as shown in Figure 11-39. This allows QRadar to incorporate the Big Data and database security violations with the rest of the infrastructure that is monitored.

Guardium sends information in the LEEF Format so QRadar can parse correctly

"Log Source Type" is IBM Guardium

Figure 11-39 Guardium sends security events in the LEEF format for QRadar to parse correctly

Another aspect of security is the risk that is associated with devices that have known vulnerabilities. If a device has a known vulnerability, it might be used to obtain access to that device. After access to that device is obtained, the hacker can steal the data or hop to another device in the network to get access to their goal. There are industry best practices concerning managing vulnerabilities. For more information, see the following Common Vulnerabilities and Exposures (CVE) website:

<http://cve.mitre.org/>

Guardium can help close the gap to understand which database servers have vulnerabilities. For example, in addition to the proprietary vulnerability assessment tests that were developed within IBM, we use the Security Assessment application within Guardium that identifies the following industry-standard tests:

- ▶ Center for Internet Security (CIS). For more information, see this website:
<http://www.cisecurity.org/>
- ▶ Security Technical Implementation Guides (STIGS). For more information, see this website:
<http://iase.disa.mil/stigs/>

These results are important to understand which security tests failed and those that succeeded. For example, the Security Assessment test that is shown in Figure 11-40 on page 434 has a risk score of 35%. In the Result Summary section, there is a high-level summary of which tests passed or failed, depending on the category of the assessment test.

Also included is a section for recommendations that you can use to prioritize the remediation of failed tests. After you have this data, you have the option of exporting this to the QRadar system. In this particular example, we are only concerned about the failed CVE tests (the tests that include known vulnerabilities according to industry security best practices) that are sent to QRadar. We export this in the AXIS format, and QRadar associates these failed tests with the database asset.

In this example, the database asset has 128 failed CVE tests. This allows us to identify which CVE tests we want to remediate and to track the asset over time. Ideally, on the next security assessment run, we have less than 128 failed CVE tests because they are remediated.

The screenshot shows the QRadar Asset interface with a table listing assets. One asset is highlighted with a red box, showing its IP address (10.10.9.58), MAC, Name, User, Group, Network (Net_10_0_0_0), Weight (0), Risk Level (128), Vulnerabilities (128), and Last Seen (2012-09-12 02:43). Below this, a window titled 'Guardium Vulnerability Assessment' displays the results of an Oracle Assessment. The assessment shows 35% tests passing and 128 vulnerabilities. A bar chart shows the test passing percentage over time. A red box highlights the 'Vulnerabilities' column in the QRadar table and the '128' count in the assessment window. A red arrow points from the '128' count to the 'Export as SCAP xml' button. A text box on the right says 'Failed CVE's uploaded to QRadar via AXIS'.

Figure 11-40 Guardium Security Assessment and QRadar AXIS and SCAP integration

If you want to send the entire security assessment results to QRadar, you use the SCAP format. This format includes passed and failed tests where the CVE information is a subset of the overall Security Assessment tests. All of these tests can be automated by distributing these tests with the workflow system of Guardium called the Audit Process.

11.5.2 Tivoli Netcool

Another aspect to security is to receive alerts when something is wrong or a security violation occurred. The process of monitoring the network infrastructure matured during the last 20 years, and one aspect of this monitoring is with Simple Network Monitoring Protocol (SNMP). SNMP is an industry standard that allows a central manager to monitor many devices in the network.

The issue becomes how to monitor the unique events of each device. The standard allows for flexibility through the Management Information Base (MIB). The MIB allows each device to describe itself with the appropriate attributes that are unique to that device. There also is a set of common attributes with SNMP, such as processor and memory.

Guardium has an SNMP MIB that can be incorporated with Netcool®. You configure the SNMP parameters by clicking **Administration console** → **Alerter**, as shown in Figure 11-41.

The screenshot shows the Administration Console interface. The top navigation bar includes tabs for Product Demo, System View, Administration Console (active), Tools, Daily Monitor, Guardium Monitor, and Tap Monitor. The main content area is titled "Administration Console" and features a left-hand navigation tree under "Configuration". The tree includes items like Alerter, Anomaly Detection, Application User Translation, Custom ID Procedures, Customer Uploads, Flat Log Process, Global Profile, Guardium for z/OS, Incident Generation, Inspection Engines, IP-to-Hostname Aliasing, Policy Installation, Portal, Support Maintenance, Session Inference, System, Upload Key File, and Unit Utilization Levels. The "Alerter" item is selected. The main panel displays the configuration for the "Alerter". It includes a "Polling Interval (seconds)" field set to 60 and an "Active on startup" checkbox checked. Below this is the "SMTP" section with fields for IP Address/Host Name (10.10.9.56), Port (25), User Name (joe), Password, Re-enter password, Return Email Address (guardium@us.ibm.com), and Authentication method (None). A "Test Connection" button is next to the IP field. The "SNMP" section follows, with fields for IP Address/Host Name (10.10.9.240), "Trap" Community (masked with dots), and Retype Community. Another "Test Connection" button is next to the IP field. At the bottom of the configuration area are "Stop", "Restart", and "Apply" buttons.

Figure 11-41 SNMP parameters that are configured in the administration console

To poll the Guardium device or for Guardium to send an SNMP trap, you must configure the passwords. The passwords for SNMP are defined within the SNMP community strings, as shown in Figure 11-41.

There are two complementary methods for a user to receive SNMP information from a Guardium appliance: Traps and Polling.

Traps

Traps are unsolicited alerts that are generated by an appliance and sent to an SNMP manager, such as Netcool. In the following subsections, we describe how to use SNMP to poll the Guardium Appliance.

Definitions

The Guardium implementation of SNMP uses the following definitions:

- ▶ **Guardiumsnmp:** The SNMP community for Guardium appliances.
- ▶ **Port 161:** The open port to query Guardium appliances that uses SNMP.
- ▶ **UCD-SNMP-MIB:** A widely used MIB, which can be used to query for many SQL Guard metrics.
- ▶ **HOST-RESOURCES-MIB:** Another MIB that can be used to query Guardium appliance.

Polling

In a polling scenario, an SNMP management system or a user queries the Guardium appliance by using standard SNMP commands. The SNMP management system can then send alerts that are based on user-defined thresholds.

The Guardium appliance provides standard metrics (by using the MIBs, UCD-SNMP-MIB, and HOST-RESOURCES-MIB) to monitor the health of the machine and a set of custom metrics (by using extensions in the MIB and UCD-SNMP-MIB), which provides information that is specific to the Guardium appliance.

Standard metrics

Displaying data that is relevant to any server, these metrics measure key performance statistics, such as memory usage, disk usage, and CPU usage.

A full list of Guardium SNMP OIDs is available in the Monitoring via SNMP section of the Guardium Administration Help Book Guide, which is available at this website:

http://pic.dhe.ibm.com/infocenter/igsec/v1/topic/com.ibm.guardium91.doc/administer/topics/monitoring_via_snmp.html

The following examples provide other methods to query an appliance by using Net-SNMP.

Use the following command to retrieve information about one metric by using the numeric object identifier (OID):

```
#snmpget -v 1 -c guardiumsnmp supp8.guardium.com
.1.3.6.1.4.1.2021.9.1.7.1
UCD-SNMP-MIB::dskAvail.1 = INTEGER: 472296
```


The following command provides the same result by using a “human readable” version of the OID:

```
#snmpget -v 1 -c guardiumsnmp supp8.guardium.com dskAvail.1
UCD-SNMP-MIB::dskAvail.1 = INTEGER: 472296
```

Removing the 1 at the end of the line gives you the status on all of the available disks (use `snmpwalk` instead of `snmpget` to retrieve multiple metrics), as shown in the following example:

```
# snmpwalk -v 1 -c guardiumsnmp supp8.guardium.com dskAvail
UCD-SNMP-MIB::dskAvail.1 = INTEGER: 472296
UCD-SNMP-MIB::dskAvail.2 = INTEGER: 60494636
```

Finally, querying on disk provides all disk information in this subsection of the UCD-SNMP-MIB, as shown in the following example:

```
# snmpwalk -v 1 -c guardiumsnmp supp8.guardium.com dsk
UCD-SNMP-MIB::dskIndex.1 = INTEGER: 1
UCD-SNMP-MIB::dskIndex.2 = INTEGER: 2
UCD-SNMP-MIB::dskPath.1 = STRING: /
UCD-SNMP-MIB::dskPath.2 = STRING: /var
UCD-SNMP-MIB::dskDevice.1 = STRING: /dev/sda5...
```

The following `snmpwalk` commands also are useful:

```
snmpwalk -v 1 -c guardiumsnmp supp9.guardium.com memory
snmpwalk -v 1 -c guardiumsnmp supp9.guardium.com system
```

As you can see, there is much information that is available to monitor the health of the Guardium appliance.

11.5.3 IBM Security Access Manager for enterprise single sign-on

Another aspect to security is single sign-on. In this example, Guardium can be integrated with Tivoli Enterprise Single Sign-On, as shown in Figure 11-42 on page 438. The Guardium user installs the Tivoli Access Manager Enterprise Single Sign-On agent on their desktop. When the user logs in to the Guardium GUI from the web browser, the Tivoli Access Manager Enterprise Single Sign-On agent captures the login information and stores this in the Tivoli Access Manager Enterprise Single Sign-On wallet.




Figure 11-42 IBM Security Access Manager for Enterprise Single Sign-On

Figure 11-43 on page 439 shows the Tivoli Access Manager Enterprise Single Sign-On wallet with the new information for 10.10.9.248, which is the Guardium appliance.




Figure 11-43 IBM Security Access Manager for Enterprise Single Sign-On wallet

The next time the user accesses the Guardium appliance through the web browser, their credentials are automatically presented in the login form, as shown in Figure 11-44.

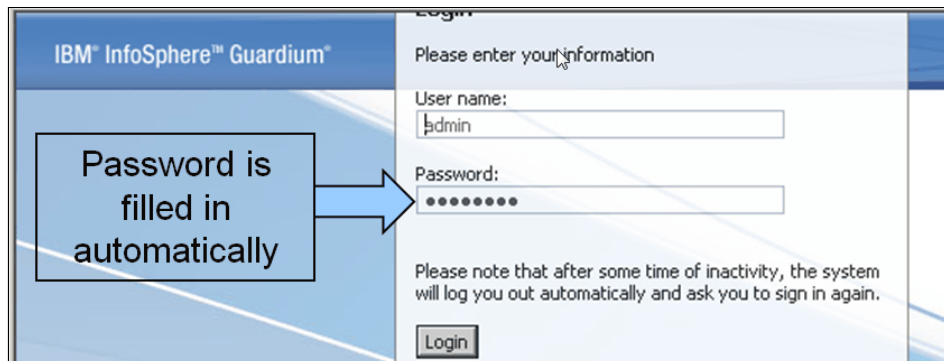


Figure 11-44 Credentials are automatically entered

The Single Sign On method is an excellent way to store all of the user name and password information that is required for applications in an enterprise environment.

11.5.4 Tivoli Directory: Lightweight Directory Access Protocol

The next area of integration for Guardium is with Lightweight Directory Access Protocol (LDAP). Many customers use LDAP to consolidate various information. One popular aspect of the use of LDAP with Guardium is to use this information to identify who can access the database server. Figure 11-45 shows how Guardium is integrated with Tivoli Directory.




Figure 11-45 Tivoli Directory integration with Guardium to help provide database access control

In this process, new users can be added or deleted from Tivoli Directory. These users are imported into an authorized group on the Guardium system regularly, as shown in step 2 of Figure 11-45.

Some customers might use Active Directory (AD) for this definition. In this case, SamAccountName is a common LDAP attribute that is used to import the Active Directory users from the Authorized Database Users group, as shown in Figure 11-46 on page 441. Users Joe and Joed were imported because they were members of the Authorized Database Users group within AD.

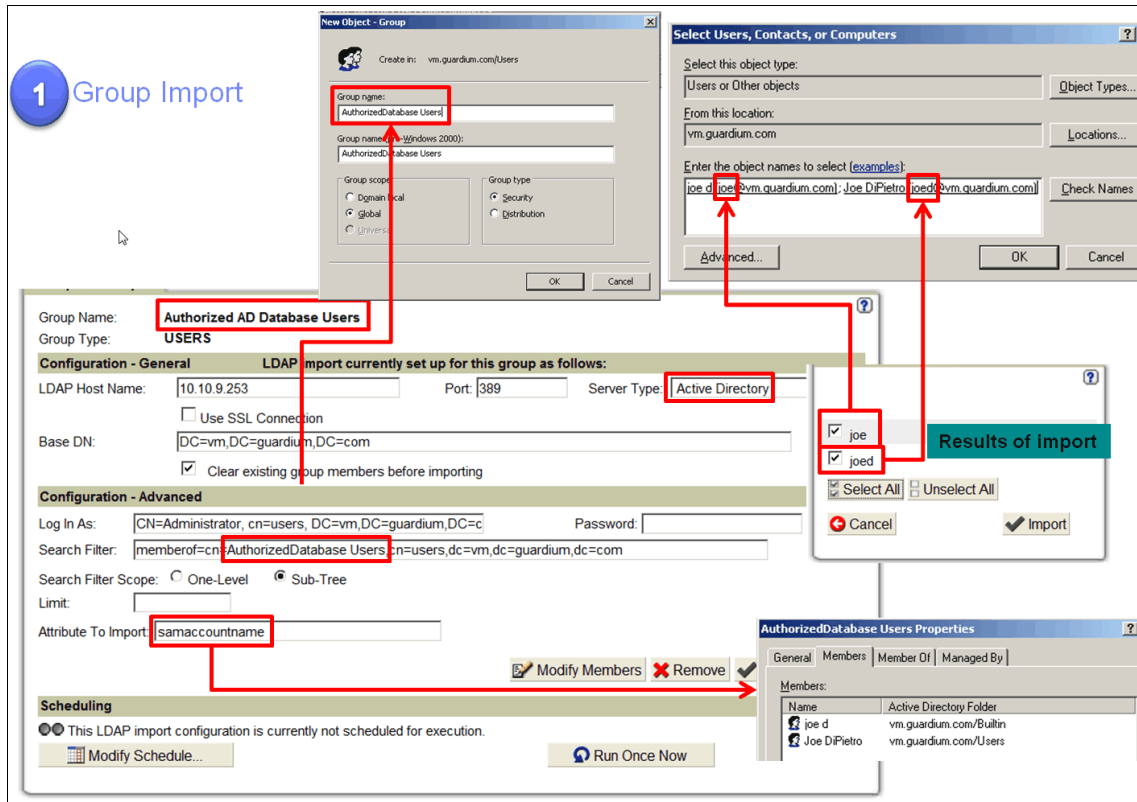


Figure 11-46 Active Directory Import into Guardium with SamAccountName attribute

The configuration of what to import from the LDAP server is flexible and can manage different LDAP attributes to satisfy a customer's unique requirements.

If the users are defined within these groups, they can access the database because of the policy definition. If the users are not defined within this group, they cannot access the database, as shown in Figure 11-47.

Block Unauthorized Access

Access Rule Definition

Rule #6 Description: Block Non-Active Directory Authorized Users

Category: Security Classification: Unauthorized Access Severity: HIGH

Not Server IP _____ and/or Group _____

Not Client IP _____ and/or Group _____

Not Client MAC _____ Net. Protocol _____ and/or Group _____

DB Type _____ Not Service Name _____ and/or Group _____

Not DB Name _____ and/or Group _____

Not DB User _____ and/or Group (Public) Authorized AD Database Users

Not App. User _____ and/or Group _____

Not OS User _____ and/or Group _____

Not Src App. _____ and/or Group _____

Not Field Name _____ and/or Group _____

Not Object _____ and/or Group _____

Not Command _____ and/or Group _____

Object/Command Group _____

Object/Field Group _____

Pattern _____ XML Pattern _____

Period _____

App Event Exists Event Type _____ Event User Name _____

App Event Values

Text _____ Numeric _____ Date _____

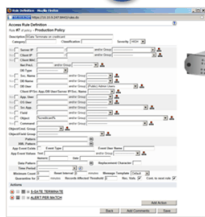
Min. Ct. [0] Reset Interval (minutes) [0]

Continue to next Rule Rec. Vals.

Action [S-GATE TERMINATE]

3

Group information updated in policy



Guardium heterogeneous DB access control policy

- Updated group content will be installed into the policy
- If not in the group, they will be prevented from executing any transactions in the database

Figure 11-47 Policy definition to allow authorized LDAP Users access to the database

This example shows how to use predefined information that is in the LDAP server to control access to the database. You can also use information that is defined in the LDAP server to enhance your audit reports. Some common examples are to add employee department, manager name, location, and other attributes into these audit reports.

11.5.5 IBM Endpoint Manager and Guardium Integration

In trying to assess your infrastructure and the overall risk to your environment, it is important to understand how end points (desktops, notebooks, mobile devices, and so on) affect this risk. End points access all kinds of sensitive data from applications and database servers. Guardium can help integrate the vulnerability assessment information of database servers into IBM Endpoint Manager so that you have a complete picture of this risk.

Figure 11-48 shows how Guardium uses the Security Compliance Automation Protocol (SCAP) to provide the vulnerability assessment information of the database servers, and consolidate that information within Tivoli Endpoint Manager. This is a powerful mechanism to consolidate and understand the risk of your database servers within the scope of your endpoint devices.




Figure 11-48 IBM Endpoint Manager integration with Guardium through SCAP

11.5.6 Tivoli Storage Manager and Guardium Integration

It is important to maintain archive logs of the audit trail. If you have a security breach, you want to review the archive logs to understand the details of when this potential breach occurred and how long the perpetrators were in your environment.

Tivoli Storage Manager helps centralize and automate data protection to help reduce the risks that are associated with data loss. This highly scalable software helps you manage more data with less infrastructure and simplified administration. You can save money, improve service levels, and comply with data retention regulations. Guardium can use Tivoli Storage Manager to archive audit data within the same framework as the rest of your backups by using Tivoli Storage Manager, as shown in Figure 11-49.

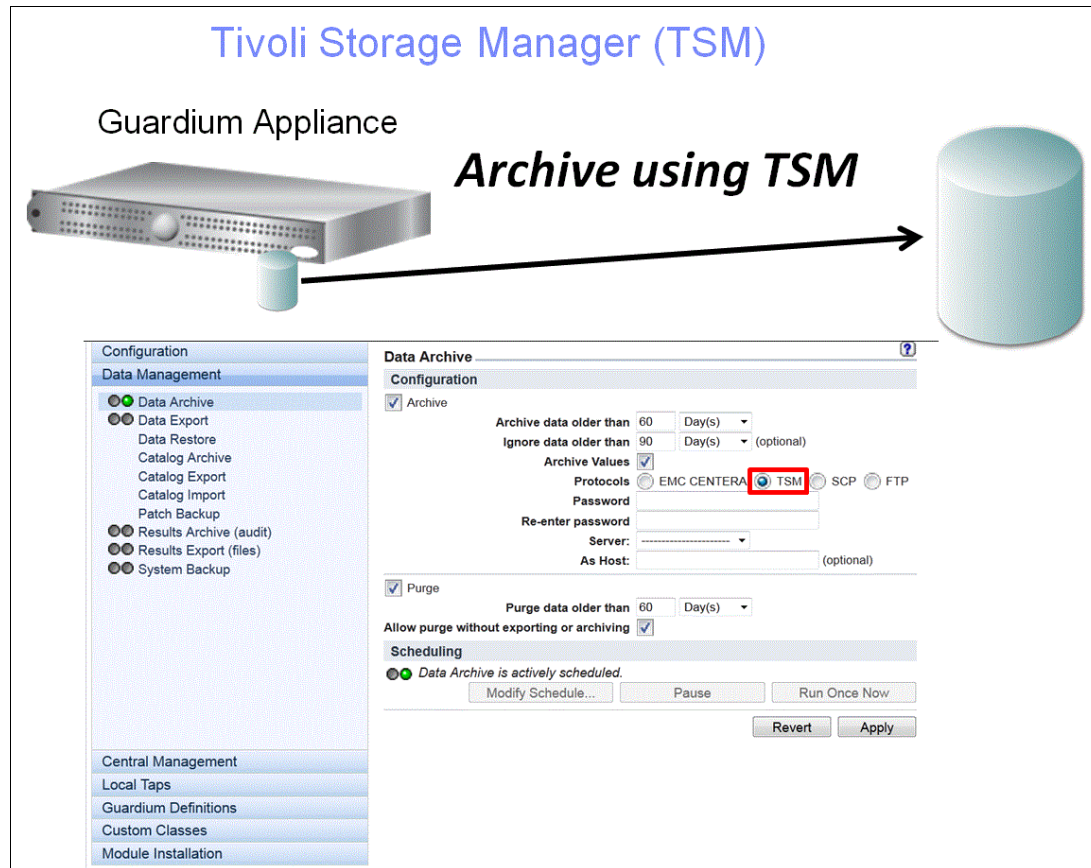


Figure 11-49 Archiving Guardium audit data with Tivoli Storage Manager

You can use the CLI commands to configure the Guardium appliance to archive information to Tivoli Storage Manager with the following archive schedule:

```
v9GA.ibm.com> com tsm
import tsm ?
import tsm config
import tsm property
ok
v9GA.ibm.com> store storage-system TSM
USAGE: store storage-system tsm backup|archive on|off
ok
```




Deployment Guide for InfoSphere Guardium

(0.5" spine)
0.475" <-> 0.875"
250 <-> 459 pages

Related publications

The publications that are listed in this section are considered particularly suitable for a more detailed discussion of the topics that are covered in this book.

Online resources

The following websites also are relevant as further information sources:

- ▶ IBM InfoSphere Guardium V9.1 Information Center:
<http://pic.dhe.ibm.com/infocenter/igsec/v1/index.jsp>
- ▶ Processor Value Unit [PVU] licensing for Distributed Software:
http://www-01.ibm.com/software/lotus/passportadvantage/pvu_licensing_for_customers.html

Help from IBM

IBM Support and downloads:

<http://www.ibm.com/support>

IBM Global Services:

<http://www.ibm.com/services>



Deployment Guide for InfoSphere Guardium



Becoming compliance and audit ready with InfoSphere Guardium

Understanding the InfoSphere Guardium solution architectures

Implementing InfoSphere Guardium solutions

IBM InfoSphere Guardium provides the simplest, most robust solution for data security and data privacy by assuring the integrity of trusted information in your data center. InfoSphere Guardium helps you reduce support costs by automating the entire compliance auditing process across heterogeneous environments. InfoSphere Guardium offers a flexible and scalable solution to support varying customer architecture requirements.

This IBM Redbooks publication provides a guide for deploying the Guardium solutions.

This book also provides a roadmap process for implementing an InfoSphere Guardium solution that is based on years of experience and best practices that were collected from various Guardium experts. We describe planning, installation, configuration, monitoring, and administrating an InfoSphere Guardium environment. We also describe use cases and how InfoSphere Guardium integrates with other IBM products.

The guidance can help you successfully deploy and manage an IBM InfoSphere Guardium system.

This book is intended for the system administrators and support staff who are responsible for deploying or supporting an InfoSphere Guardium environment.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks