



A Practical Guide to Data Auditing for Multiplatforms

Document version 1.0

Kelly Smith: kellys@mainstar.com

John Jensen: john.jensen@rocketsoftware.com

Andy Nguyen: andy.nguyen@rocketsoftware.com

Jeremy Weatherall: Jeremy.weatherall@rocketsoftware.com

Rajesh Nandwani: rajesh.nandwani@rocketsoftware.com

Mary Petras: marypetr@us.ibm.com

Thomas A. Kisielewicz: takisiel@us.ibm.com

Rita Vuong: rvuong@us.ibm.com

CONTENTS

List of Tables	vi
Revision History	vii
1 Executive Summary.....	viii
1.1 Auditing Today	ix
1.2 Why Audit	ix
1.3 Company Perspective.....	xi
1.4 Auditor’s Perspective	xii
1.5 The DBA Perspective	xii
1.6 Traditional Auditing	xiii
1.7 Achieving Integrity through Segregation of Duties.....	xiii
2 DB2 Audit Management Expert for Multiplatforms	xiv
2.1 Expertise.....	xiv
2.2 Centralization.....	xiv
2.3 Simplification.....	xiv
2.4 Segregation of Duties	xv
2.5 Internal Security.....	xv

3	Best Practices.....	xv
3.1	Preparation for an Install.....	xvi
3.1.1	Are you using 32- or 64-bit instances?	xvi
3.1.2	Port numbers	xvi
3.1.3	DB2 instance to audit.....	xvi
3.2	Overview of Install Steps	xvi
3.2.1	Install DB2 Audit Management Expert for MP	xvi
3.2.2	Configure the server	xvi
3.2.3	Configure the agent	xvi
3.2.4	Start the server, then the agent	xvii
3.2.5	Start the Administration User Interface.....	xvii
3.2.6	Create an authorization so an auditor can view the audit data	xvii
3.2.7	Start the Reporting User Interface	xvii
3.3	Installation Preparation	xvii
3.4	Tips for Installing DB2 Audit Management Expert MP.....	xviii
3.5	Repository.....	xix
3.5.1	Increase the DB2 Transaction Log	xix
3.6	Configuration Files to Use	xx
3.7	Availability of Audit Data	xx
3.8	Server Configuration Tips	xxi
3.8.1	Local Environment Settings	xxi
3.8.2	Sever Collection Settings.....	xxi
3.8.3	Server Configuration parameters and recommendations.....	xxi
3.9	Agent Configuration Tips	xxii
3.9.1	DB2 Audit Facility Message.....	xxii
3.9.2	Changing Permissions in the DB2 Audit Log File.....	xxii
3.9.3	Local Environment Settings	xxii

3.9.4	Agent Collection Settings.....	xxiii
3.9.5	Agent Configuration parameters and recommendations.....	xxiii
3.9.6	How to Check If the Server and Agent Are Up.....	xxiv
3.9.7	Data Collection Considerations	xxiv
3.9.8	Securing and Monitoring the Audit Data.....	xxv
4	Administration User Interface	xxvi
4.1	Logging in to the Administration User Interface.....	xxvi
4.2	Add Users and Groups	xxvi
4.3	Check Agent Status	xxvii
4.4	Add a Collection Profile	xxvii
4.4.1	Adding Rules to the Collection Profile	xxviii
4.4.2	Determine Collection Profile Schedule	xxix
4.4.3	Select General Audit Options	xxx
4.4.4	Include or Exclude Identities.....	xxxiii
4.4.5	Include or Exclude Applications.....	xxxiv
4.5	Collections	xxxv
4.6	Authorizations	xxxvi
4.7	Repository Tab	xxxvii
5	Reporting User Interface	xxxviii
5.1	Logging in to the Reporting User Interface.....	xxxviii
5.2	Overview report	xxxix
5.3	Filters	xl
5.4	Summary Report.....	xlii

5.5	Detailed Data	xliv
5.6	Log Analysis	xlviii

LIST OF TABLES

Table 1:	Government regulations that require auditing	ix
Table 2:	Information required during installation	xvii

1 Executive Summary

No company wants to end up in the media with unwanted publicity about any event that can affect company integrity, resulting in the loss of stockholder or customer confidence. One way this can be avoided is by obtaining a better understanding of corporate data auditing procedures, as well as the associated fraud exposures and auditing costs.

Auditing is important, not only because it is required by law, but because the integrity of your company is important. Auditing helps insure the integrity of company data. Auditing does not generate revenue, so companies want to audit with the least expense possible while remaining in compliance. The problems with using traditional methods of auditing are many: lack of segregation of duties between the auditors and those they may need to monitor; no centralization of the audit data from the many systems and sources a company may have; and no (or little) automation of auditing processes. These problems can result in fraud exposure and unnecessary costs.

Auditors lose their independence because they need to go to other, highly technical people to get the data they require – those they may need to audit – in order to acquire the necessary information. An auditor's reliance on others not only increases costs by using these highly paid people, but also increases fraud exposure. Without a segregation of duties in this area, there is a greater possibility that the data may be manipulated before reaching the auditor.

The auditing process can be complicated by the sheer size of the data center, causing critical exposures to be overlooked. When audit data is collected from many systems and sources, the data must be combined, correlated and displayed in a clear format, providing auditors with factual and easy-to-read material.

Little or no automation results in great amounts of time spent on the auditing process, which becomes error-prone, and costly. Programs can be written to automate the collection and correlation of the audited data, but those programs need to be maintained on a regular basis. Additionally, those types of programs are specific to accessing data required at that point in time, and often conflict with the required segregation of duties between auditors and database administrators, or DB2 programmers in the case where programs are written by the same people the auditors will monitor.

Without segregation of duties, fraud is always a possibility. Without centralization and automation, a more comprehensive audit results in a higher labor cost, and a less comprehensive audit runs the risk of a company being out of compliance. The right software can greatly improve the likelihood of a successful audit and give auditors the necessary insight to answer questions about accessed data: who, what, when, where and how.

IBM's DB2 Audit Management Expert for Multiplatforms (MP) pulls together disparate data sources from different systems into a central repository with a simple-to-use interface, giving auditors a complete view of the business activity collected without reliance on the technical personnel they need to monitor. Collecting data with an auditing software product enables the product repository to also be audited to provide integrity and prevent audit data tampering.

DB2 Audit Management Expert for MP is a comprehensive auditing solution that provides the three keys to auditing success: segregation of duties to ensure integrity; centralization of the data to be audited in order to eliminate the complexities of collecting data from many

systems; and automation, to achieve more thorough audits, reduce the cost of auditing and reduce the risk of being out of compliance.

Auditors now have an automated, simple method to gain the information required to determine compliance. The easy-to-use interface gives them the tools they need to audit the data they want from one central location, filter it based on their requirements, and display data of interest using standard or custom reports.

1.1 Auditing Today

Several challenges affect auditing today. It is important to accurately collect and correlate data into useful report representations that auditors can easily use. The data must adhere to regulatory compliance regardless of the size of a company's IT department. Also, many auditors depend on developers or database administrators (DBAs) to set up or gather the information they require, despite the fact that these personnel may also need to be monitored.

These challenges raise several significant questions. How do auditors ensure that the person providing the information has not updated sensitive data or excluded it from the reports? How can auditors do a thorough job without being dependent on database personnel when there are a large number of systems to monitor? How can a company ensure the external auditor has precise, accurate information to determine if they meet all applicable regulatory compliance?

This white paper focuses specifically on data auditing, which is just one aspect of regulatory compliance. There are three levels of data auditing: ensuring business controls are in place, internal audits, and external audits. This white paper is targeted to the first two levels.

1.2 Why Audit

Your data is valuable. It has always been a good practice to perform audits as a method of maintaining checks and balances. Not only does this include auditing the quality of the data, but more importantly, who has access to the data. This was thought to insure that no one person has the ability to maintain and manipulate information that could be considered highly sensitive and negatively impact the company's bottom line.

In recent years, there have been many publicized incidents where fraud has occurred, and in most cases, these incidents have had major financial ramifications. With the possibility of such occurrences, the government has had to intervene in an attempt to prevent repeated incidents by establishing several regulations that permeate several industries throughout corporate America. Not only is auditing a good practice, but now, in most industries, it's the law. Many countries have similar regulations, such as those regulations shown in the table below.

Table 1: Government regulations that require auditing

Regulation	Threat
Sarbanes-Oxley Act of 2002	<ul style="list-style-type: none">• Act passed to prevent corporate and accounting scandals• CEO and CFO certifications of annual and quarterly SEC reports• Evaluates the effectiveness of internal controls• Requires rapid disclosure of material changes in financial conditions or operations

	<ul style="list-style-type: none"> • Set up automatic controls repository to identify deficiencies • Public Company Accounting Oversight Board is an agency that regulates auditors in public companies
Gramm-Leach-Bliley Act	<ul style="list-style-type: none"> • Act passed to legalize mergers between banking and insurance companies • Financial institutions are required to have a policy to protect information from security threats and protect data integrity • Financial Privacy Rule: requires a privacy notice from financial institutions to their customers every year • Safeguards Rule: financial institutions should have a security plan to protect their consumer's non-public personal information • Pretexting Protection: financial institutions have to protect their consumer's non-public information by preventing someone without authority from accessing the information
Health Insurance Portability and Accountability Act (HIPAA)	<ul style="list-style-type: none"> • Act passed to restrict access to patient treatment and payment information to approved personnel • Protect people when they lose their jobs or change occupation
Basel II (primarily banking)	<ul style="list-style-type: none"> • Capital requirement should be more risk-sensitive • Market discipline: people who deposit money into banks can influence the way bank managers are involved in risky activities • Help financial system in the bank become more stable
Solvency II (insurance)	<ul style="list-style-type: none"> • Help protect policyholders against the risk of a company failing • Used in insurance industry to ensure a more efficient capital allocation • Provide financial stability
Japanese Financial Instruments and Exchange Law (FIEL)	<ul style="list-style-type: none"> • Intended to protect investors • Criminal penalties increased to the maximum for market frauds • Disclosure rules applies to any investment fund that invests in securities • Corporate reorganization will require securities to be registered • Companies required to have a quarterly report • Statements in annual and quarterly report are required to be certified
Japanese Protecting Personal Freedom Act	<ul style="list-style-type: none"> • Act passed to protect personal information, or any information that can identify an individual (name, date of birth) • A person's consent is needed before someone can access his/her personal information
Financial Services and Markets Act (FISMA)	<ul style="list-style-type: none"> • Act passed is intended to reduce financial crime • Ensures consumers are protected • Insurance, banking, or investment business need to be authorized before they can conduct regulated activities
Payment Card Industry (PCI)	<ul style="list-style-type: none"> • Regulation passed to protect someone in the event their credit card is stolen • Protect against unauthorized charges on a stolen credit card • Protects cardholder's information • Access to cardholder's information will be restricted on a business need-to-know • All access to cardholder's information and network resources will be tracked and monitored • Required to maintain information security
Patriot Act	<ul style="list-style-type: none"> • Act passed mandating publicly and privately held companies to assist law enforcement agencies in surveilling terror suspects

	<ul style="list-style-type: none"> • Provide private information on-demand, such as email and telephone communications, financial or medical records • Act passed after the World Trade Center attacks of 9/11/01
Various anti-money laundering (AML) regulations	<ul style="list-style-type: none"> • Requires financial institutions to monitor, investigate and report any suspicious transactions related to money laundering or currency crimes

Within each company there are several views of auditing: the corporate view, the auditor view, and the DBA view. The essential issue, however, is how companies can achieve compliance and maintain stockholder and customer confidence in the corporation, while simultaneously ensuring that as data grows, auditing costs (which do not generate revenue) are managed appropriately, and that the auditing method ensures accurate data.

1.3 Company Perspective

Such government regulations are pressuring companies to audit the viewing and updating of data. Since auditing your own company does not bring in revenue, this is now considered part of the cost of doing business. Companies try to accomplish this with as little funding and resources as possible while remaining in compliance. This adds another layer of complexity to the overall business model.

Auditing is a key component of the overall security, compliance and risk management of any company. Audit policies needs to complement the plans and policies of other business areas to reduce the risk of problems, as well as to ensure that any errors are caught as soon as possible.

A company needs to assess its organization and decide how to approach and implement the audit process. This can be done by bringing in an outside company or organizing an internal audit department to oversee the process. They will maintain control of the - *who*, *what*, *where*, *when*, and *how* of the audit controls.

The audit team needs to know who is involved in the processing of data and at what point a breakdown can occur. Most companies have these tasks covered, as specific employees are granted access to sensitive data in order to perform their job duties. The challenge arises when privileged users are involved (usually those who ultimately control the data: the systems teams, or system administrators with a high level of authority to access data) or when access controls are not well-monitored.

The data or process that will be audited will vary depending on the industry and the regulations to which it is subject. Overall, it is standard practice to have audits on any data considered personal or sensitive.

The auditing process raises many questions from a company perspective:

- What data will be audited?
- At what point in the process must checks be established?
- When will they audit? (How often is too much or too little? How much will this cost to accomplish?)
- How will they audit? Will it be done manually or in some type of automated fashion? Are there tools to help with the audit? (This is another difficult subject to address. Part of the answer lies in the industry and the regulations that pertain to that industry.)

As you can see, there is a lot to consider when it comes to auditing, an area that is sometimes overlooked since it does not generate revenue. It is nonetheless a crucial area, since thorough audits can help to prevent fraud.

1.4 Auditor's Perspective

Auditors want to know *who, what, when, where* and *how*. It can be difficult pulling together all the information required in an audit due to an auditor's dependence on developers or database administrators (DBAs). This dependency has drawbacks:

- **Collection:** Existing developer and DBA tools are not audit-oriented, nor are they designed to collect all the relevant audit information from the source.
- **Reporting:** Existing developer and DBA tools are not audit-oriented, nor are they designed to present information in a useful way for an audit.
- **Integrity:** DBAs are part of the audited population, and should therefore not be relied upon to provide key audit information. Furthermore, DBA user identifications (user IDs) have more system-level privileges than typical business users, giving them more opportunity to circumvent normal business controls.

Alternatively, auditors could collect and correlate the information themselves, but this direct approach has drawbacks:

- **Privileges:** Auditors generally are not granted the system privileges needed to collect the needed information themselves.
- **I.T. Skills:** Even if auditors were given these privileges, they need substantial knowledge of information technology to collect and correlate data at an application level. Developing such skill is costly, time-consuming, and tangential to the auditor's primary role.
- **Complexity:** Because data can be proliferated across the enterprise, it is increasingly difficult to pull information together from "all" systems.
- **Cost:** A more comprehensive audit results in a higher labor cost.
- **Repercussion:** A less comprehensive audit runs the risk of missing important events and could allow a company to operate out of compliance.

1.5 The DBA Perspective

It is important to know the business process as well as the audit process. When audits are performed, it is also important to have the right people involved. When any audit process discusses data, there will ultimately be a discussion about the database administrators (DBAs).

The amount of DBA involvement in the auditing process varies widely. Because different industries are held to different security standards, some audit requirements can result in a substantially greater workload for both the DBA and the audit team. The workload can also be affected by the auditing process itself. For obvious reasons, the less work that is

required by the DBA for the auditing process, the better for both the DBA and the auditor. However, while assisting the auditor, the DBA is aware that they too are within the scope of the audit.

Most employees in a company have pre-defined data access privileges associated with their job role. Prior to the enhanced auditing regulations, it was an accepted practice to allow DBAs and system administrator's access privileges to all data. Today, access to sensitive data is split among the DBA and system administrators. While each still has access to sensitive data, they do not have access to data that isn't within their business scope. That is, their access to sensitive data is mostly compartmentalized and each only has the appropriate access to perform their job duties. Despite how a DBA's access to sensitive data has changed with the implementation of each regulation, a DBA's role in the audit process is still required and important.

1.6 Traditional Auditing

Using traditional auditing methods, auditors require a multitude of resources: a system user ID for each system they need to collect data from; database access on each of these systems; tools to collect the data; tools to put the pieces of data together in a meaningful way, help from database administrators and DB2 system programmers, and so on. The larger the environment, the more difficult it is to coordinate these resources.

From a high-level perspective, the response is usually surprise and dismay at the cost of obtaining data during an audit, which in turn creates the motivation to reduce the cost. A company may wonder if it is necessary to spend money in order to train auditors to be nearly as experienced as database administrators when the auditor will still require the DBAs help to get the data. It all boils down to a conflict of duties between the auditor and the database administrator.

Certain concerns arise from the company's perspective. Not only are the efforts mostly manual, but how thorough can the audit be using these methods? Was something critical missed? If so, you could end up in reactive mode. The audit data is gathered after the event and could be difficult to find or unavailable.

It is understood that applications to audit data are sometimes written in-house, but that can create an exposure. It raises additional questions: Who wrote the code? Is the code maintained or even secured? Using that program, can someone manipulate audit data, and would anyone know?

Since auditing doesn't bring in revenue, companies will try to accomplish it with as little funding and resources as possible to adhere to regulations. The question is whether companies are really saving as much as they can while ensuring the integrity of the audit.

1.7 Achieving Integrity through Segregation of Duties

The key to gathering data with integrity, meaningful representations of the data, and maintaining a separation of the roles of auditor and DBA, is to automate the process with auditing software. Auditing software gives the auditors independence so they can adhere to published industry standards without relying on personnel who are also being monitored. The right software can help organizations audit more successfully, and less expensively, by providing an easy-to-use tool to access the required data.

2 DB2 Audit Management Expert for Multiplatforms

DB2 Audit Management Expert for Multiplatforms (MP) is an auditing solution that centralizes audit information, greatly simplifies access for auditors, and provides data integrity through segregation of duties on DB2 AIX platforms so auditors can easily find out *who, what, where, when, and how*. The key advantages DB2 Audit Management Expert for MP provides are described in the following sections.

2.1 Expertise

Auditors often must consult multiple sources because no one person has the security authorizations to access, nor knowledge about, all of the necessary data. DB2 Audit Management Expert for MP pulls together data from all of the disparate sources and collects it into a central repository with a simple-to-use graphical user interface, so auditors can analyze the data without relying on a DB2 systems programmer, DBA, or developer. From the auditor's perspective, it is like working with an expert DBA or a combination of a systems programmer and a DBA. If an auditor wants activity for a specific table from specific plans or users, DB2 Audit Management Expert collects what is needed.

2.2 Centralization

DB2 Audit Management Expert for MP uses audit data from the DB2 Audit Facility and log analysis and stores the audit data in a single repository to produce a complete view of this business activity for auditors.

There are several types of database events that can be tracked and audited. Some of these events include instances of denied access attempts without proper authorization, explicit grant and revoke statements, and the assignment and change of authorization IDs to access DB2. In addition, all selects (reads), all changes, and all create, alter, and drops are recorded.

A centralized repository creates a single source for reporting, institutional controls, summarization of the data including high-level trending of audit anomalies and drill-down capability (one layer at a time), as well as a robust level of reporting events controlled by the auditor without DBA involvement.

As data proliferates across the enterprise, centralization is integral to reducing auditing costs and increasing productivity, creating easier and more thorough audits, thereby reducing the risk of being out of compliance.

2.3 Simplification

DB2 Audit Management Expert for MP reduces manual auditing and empowers non-technical users to easily audit the data without requiring logins to each system.

In a traditional environment, auditors require logins to all the systems and require authorization to access each of the DB2 instances. In large sites, setting up and keeping track of all of these logins can be an administrative nightmare.

Auditors using DB2 Audit Management Expert for MP do not need to go to a large number of sources to access data and they do not need user IDs for DB2 or the operating system. They log into one place, DB2 Audit Management Expert for MP, to gain complete visibility of all auditable objects. An auditor can display collected data for all DB2 instances, or just the DB2 instances of interest, all from the central repository. The administration user interface, usually managed by the lead auditor, provides the ability to assign auditor's access to the tool which in turn allows them access to the repository data. For these reasons, DB2 Audit Management Expert for MP makes auditing data much more manageable.

2.4 Segregation of Duties

Segregation of duties has always been a challenge to the auditing process. In general, auditors usually depend on developers or database administrators (DBAs) to collect and report information. As described in the Auditors Perspective section, the most critical drawback with this approach pertains to the integrity of the data provided to the auditor.

DB2 Audit Management Expert for MP maintains the segregation of duties, resulting in assurance of data integrity, which results in more accurate reports. This allows DBAs to perform their own job duties and allows auditors to run audit reports independently of the DBAs, which results in easier, more accurate audits. Auditors now have the ability to adhere to published industry standards and external auditing without relying on the personnel being monitored.

The DB2 Audit Management Expert for MP administrator can specify how much visibility each auditor has to the auditable objects.

2.5 Internal Security

DB2 Audit Management Expert for MP is well-suited to enforce controls that govern DBAs, as well as to report on their activity. DBAs are trusted with sensitive data in order to do their jobs. They need to be able to maintain, copy, and recover sensitive data, as well as load and reorganize it, to name a few of their responsibilities. The continuous, automated auditing provided by DB2 Audit Management Expert for MP removes the opportunity to alter or even omit important data from the audit reports. Thus, an independent audit mechanism in place of personnel involvement provides assurance that reported data has not been modified. Consequently, the accuracy of data and reports is more reliable.

3 Best Practices

IBM's DB2 Audit Management Expert for MP software is an auditing solution that enables companies to easily segregate duties, while providing essential centralization of the data and automation of the auditing processes to reduce fraud exposures, and the costs associated with manual auditing methods.

For ease of use, and utmost value, this section focuses on employing best practices with DB2 Audit Management Expert for MP.

3.1 Preparation for an Install

You will need to have some information ready to be able to install, configure and run DB2 Audit Management Expert for MP. In addition to the questions in this section, review section 3.3, "Installation Preparation", and fill in the blanks in Table 2.

3.1.1 Are you using 32- or 64-bit instances?

Install the version of DB2 Audit Management Expert for MP that matches the instance size. Also, the size of the instance containing the repository must match the size of any instance that is being audited.

3.1.2 Port numbers

DB2 Audit Management Expert for MP needs two UNIX port numbers to configure the agent and server, and some companies require that they be assigned by the system administrators.

3.1.3 DB2 instance to audit

DB2 Audit Management Expert for MP can be run using only the instance that holds the repository. Therefore, a separate instance to audit is needed only if you want to see audit data from a specified instance.

3.2 Overview of Install Steps

3.2.1 Install DB2 Audit Management Expert for MP

Install DB2 Audit Management Expert for MP using InstallShield (see User's Guide).

3.2.2 Configure the server

Configure the server (see User's Guide).

- Configure the server environment
- Create the DB2 Audit Management Expert repository
- Enable other user IDs, if different users will run the server
- Bind the repository packages
- Configure and run the adhuap utility to create an initial DB2 Audit Management Expert for MP administrative user
- Modify the server configuration file
 - specify location specific parameters (port numbers, instance)
 - reduce the summarizer interval if audit data is needed sooner
 - discussed in section 3.8

3.2.3 Configure the agent

Configure the agent (see User's Guide).

- Configure the agent environment
- Start the DB2 Audit Facility, db2audit, and adjust permissions to access the audit log file
- Enable other user IDs, if different users will run the agent

- Bind all databases to be monitored
- Modify the agent configuration file
 - specify location specific parameters (port number, instance)
 - reduce collection parameters
 - reduce collection parameters if audit data is needed more frequently
 - write audit data to the repository more frequently

3.2.4 Start the server, then the agent

Start the server first, and then start the agent (see User's Guide).

3.2.5 Start the Administration User Interface

Start the Administration User Interface and create a collection profile (see User's Guide and also instructions below in section 4.1 and 4.4).

- Create a collection profile
- Create a collection using the collection profile and activate the profile

3.2.6 Create an authorization so an auditor can view the audit data

Create an authorization so an auditor can view the audit data (see instructions below in section 4.6).

3.2.7 Start the Reporting User Interface

Start the Reporting User Interface (see User's Guide and also instructions below in section 5.1).

- If no data is available, wait a few minutes and refresh the display
- View details of collected data

3.3 Installation Preparation

Having all the necessary information at hand before beginning the DB2 Audit Management Expert for MP installation will help the installation proceed quickly and smoothly. You can use the following table to gather the necessary information:

Table 2: Information required during installation

	Item	Purpose	Default	My Value
For Server				
	Server host name or IP Address	Identifies the system on which the server is running	none	
	agent-listener-port	Same as server-port	52521	
	client-listener-port	Same as Admin Client's Settings: Server port	52522	
For Agent				
	Agent host name or IP Address	Identifies the system on which	none	

		the agent is running		
	Server-address	Server host name or IP Address	none	
	server-port	Same as agent-listener-port in ADHCFGS	52521	
For Reporting User Interface				
	Settings: Server host	Same as server-host	none	
	Settings: Server port	Same as client-listener-port	52522	
For Admin User Interface				
	Settings: Server host	Same as server-host	none	
	Settings: Server port	Same as client-listener-port	52522	
	Repository tab: Host Name	Host name or IP Address of repository's DB2 database	none	
	Repository tab: Location	Location of repository's DB2 database	parameter to adhDbSetup.sh	
	Repository tab: Port	DB2 port for reporting-to-repository communication via JDBC	run: grep <i>instancename</i> /etc/services grep db2c	

3.4 Tips for Installing DB2 Audit Management Expert MP

When installing and configuring DB2 Audit Management Expert for MP, be sure to refer to the User's Guide to ensure no steps are skipped. Configuration errors often occur because a step was inadvertently left out. This white paper is only intended to provide tips on specific steps.

Some sites have experienced a problem with InstallShield that is caused by an incompatibility with system runtime libraries. The incompatibility causes InstallShield install to fail when it is started. If the error occurs, you'll see messages like:

```

Initializing Wizard.....
  Launching InstallShield Wizard.....

```

```
/adh11serveraix[20]: 684180 Memory fault(coredump)
JVMDG217: Dump Handler is Processing Signal 11 - Please Wait.
JVMDG303: JVM Requesting Java core file
JVMDG304: Java core file written to
    db2backup/db2audittool/javacore684180.1197489368.txt
JVMDG215: Dump Handler has Processed Exception Signal 11.
```

There is a segmentation fault in libaixppk during the install process. This is the result of an incompatibility between the third party installer that the product uses and the AIX installation tracking software, ISMP.

IBM has encountered this problem with their WebSphere product as described in the link below. In the support article, IBM recommends updating the runtime libraries on the target machine following the procedure documented in the WebSphere support article:

<http://www-1.ibm.com/support/docview.wss?uid=swg21202151>

3.5 Repository

For evaluation purposes, DB2 Audit Management Expert for MP can be run using the instance that holds the repository objects, so a separate instance to audit is only needed if you want to see audit data from a specific instance if it is different than the instance that contains the repository objects.

The repository table spaces should have regular runstats, reorgs, and backups run like any production data. Repository data can grow quickly especially if more than data of interest is captured so it is important to establish a plan for archiving the data that needs to be kept to satisfy the regulatory compliance rules.

3.5.1 Increase the DB2 Transaction Log

On systems that have a large number of DB2 transactions, the agent can fail if the DB2 transaction log fills up. The DB2 Audit Management Expert for MP installation increases the transaction log's size, but we suggest increasing it further with the following DB2 commands.

The following statements can be found in the file titled `crtfoundation.sql` that is installed with DB2 Audit Management Expert for MP.

```
UPDATE DB CFG FOR #ADHDATABASE USING DBHEAP 10000;

UPDATE DB CFG FOR #ADHDATABASE USING LOGBUFSZ 4096;
```

Where the value, `#ADHDATABASE`, is the repository database name that was created during the installation of DB2 Audit Management Expert for MP. Edit these statements to use the increased limits before creating the repository.

To increase the size of the transaction log after the repository objects have been created, you may run the commands from the command line. Then terminate the connection to the repository database using the following command to have the increased log file size take effect:

```
DB2 TERMINATE
```

3.6 Configuration Files to Use

There are four configuration files within the DB2 Audit Management Expert for MP that can be used to configure the server and the agent. We recommend using the files `fullServerconfig.xml` and `fullAgentconfig.xml` which contain the most robust option list with which to configure the server and agent.

3.7 Availability of Audit Data

In general, audit data does not need to be available in real time. Depending on how often you want to view the audit data, DB2 Audit Management Expert for MP parameters can be used to define how often to update the DB2 Audit Management Expert for MP repository.

The raw DB2 audit data is loaded into a temporary staging area where events will accumulate before being loaded into the DB2 Audit Management Expert for MP repository. These events are periodically loaded into the normalized repository tables. The frequencies with which the events are loaded into the repository are controlled by two parameters in the agent configuration file: **refresh-count** and **refresh-interval**. Data is loaded from the staging area to the repository when there are more than the specified **refresh-count** events in the staging area, or when the data is older than the specified seconds in the **refresh-interval** parameter.

The agent configuration parameter **refresh-count** specifies the maximum number of audit events that will be accumulated in the staging area before they are loaded into the repository tables. The value of the refresh count can vary from 1 to 5000, with a default value of 1800. For efficiency, this number should be set as high as is compatible with the corporate auditing policy unless you are evaluating the product and want the data to be available sooner. Recommendations for evaluation settings are shown in the, "Agent Configuration Tips", section 3.9, of this white paper.

The agent configuration parameter **refresh-interval** controls the maximum amount of time (in seconds) that audit events will be stored in the staging area before they are loaded into the normalized audit data repository tables. The values for the refresh interval can range from 120 seconds to 1800 seconds with a default value of 1800 seconds. As with the refresh count, larger values allow the agent to operate more efficiently. For a product evaluation, data may need to be available sooner. Recommendations for evaluation settings are shown in the, "Agent Configuration Tips", section 3.9, of this white paper.

At periodic intervals the server configuration parameter, **summarizer-refresh-interval**, is used by the server to read selected audit data from the repository and condense it into a summary table which is also stored in the audit repository. This way, the reporting user interface can display high level statistics without having to read the entire set of audit data. There are some metrics that are commonly reported, so putting them into the summary table makes them available to the user without having to explicitly request them.

The summary table is summarized by Access Attempts, Read, Change, Create Alter and Drop, Explicit Grant and Revoke, and Authorization Failures. This data is grouped by hour, day, week, month, AUTHID, application, and End User ID. For evaluations, summary data needs to be available sooner. Recommendations for evaluation settings are shown in the, "Server Configuration Tips", section 3.8 of this white paper.

3.8 Server Configuration Tips

This section contains recommendations for the server configuration settings. These settings primarily define how often data is written to the repository. These thresholds can be set much lower during evaluation situations to allow the data to propagate to the reporting user interface much faster.

3.8.1 Local Environment Settings

The server configuration file needs to be updated to reflect the machine's environment. The agent and client-listener ports may need to be changed to use available TCP/IP port numbers. The object-qualifier (schema) and server-repository names are the values defined when the repository was created using the script `adhDbSetup.sh`.

```
<agent-listener-port>52521</agent-listener-port>
<client-listener-port>52522</client-listener-port>
<object-qualifier>AMESHEMA</object-qualifier>
<server-repository>AMEDB2</server-repository>
```

3.8.2 Sever Collection Settings

The default settings for the agent and server assume a "real" workload, and it may take over 30 minutes for collected data to show up in the summary table. If you are evaluating the product and want to see the data sooner, the time interval for refreshing the summary table should be decreased to 300 seconds to allow data to display in the reporting user interface much faster.

```
<summarizer-refresh-interval>300</summarizer-refresh-
interval>
```

3.8.3 Server Configuration parameters and recommendations

The following settings are recommended for the server configuration for evaluations.

```
<server-config>
  <agent-listener-port>52521</agent-listener-port>
  <bind-retry-delay>10</bind-retry-delay>
  <bind-retry-max>30</bind-retry-max>
  <client-listener-port>52522</client-listener-port>
  <community-string>mystring</community-string>
  <log-level>1</log-level>
  <multicast-address>236.1.2.4</multicast-address>
  <multicast-delay>5</multicast-delay>
  <multicast-interface></multicast-interface>
  <multicast-port>52523</multicast-port>
  <multicast-ttl>5</multicast-ttl>
  <object-collection>ADHMPV11</object-collection>
  <object-qualifier>AMESHEMA</object-qualifier>
  <server-con-alias>server</server-con-alias>
  <server-pwd>xxxxxxx</server-pwd>
  <server-repository>AMEDB2</server-repository>
  <server-usr>xxxxxxx</server-usr>
  <summarizer-refresh-interval>300</summarizer-refresh-interval> (Recommended)
  <trace-config>true</trace-config>
  <trace-events>true</trace-events>
```

```
<trace-network>true</trace-network>
```

3.9 Agent Configuration Tips

This section contains recommendations for the Agent configuration settings. These settings primarily apply to how often data is written to the repository. If you are evaluating the product and want to see the data sooner, these settings need to be set much lower to allow the data to propagate to the reporting user interface much faster.

3.9.1 DB2 Audit Facility Message

Use the following command to start the DB2 Audit Facility,

```
$ db2audit start
```

If the audit facility is already running, the following message is displayed and is for your information only.

```
AUD0026I  A request to start the DB2 audit facility has been
processed. Note that audit may have already been started on the
instance.
```

3.9.2 Changing Permissions in the DB2 Audit Log File

It is necessary to start the DB2 audit facility manually before starting the agent for the first time so that it will create the directory that holds the DB2 audit log. This directory is created with permissions that allow the instance owner to read and modify the log file. It is necessary to adjust the permissions in this directory to allow the DB2 Audit Management Expert for MP agent user to read the log and prune it to remove audit records that have been processed.

Follow the instructions in the User's Guide to modify the permissions for the "sqllib" directory of the DB2 instance. Changing the permissions at the directory level allows the DB2 Audit Management Expert for MP agent user to access the directory and update the log file. You should run the "chmod -R 775 security" command as the instance owner or as the "root" user.

3.9.3 Local Environment Settings

The agent configuration file needs to be updated to reflect the machine's environment.

- The server-port needs to match the port number specified by the parameter agent-listener-port in the server configuration file, and the parameter server-address is the system name or IP address for the machine where the server is located.
- The parameter agent-monitor specifies the name of the instance that is being monitored, and the parameters agent-monitor-path and nodes-config-path indicate the pathnames for the instance.
- The parameters object-qualifier (schema) and server-repository names are those used when the repository was created using the script adhDbSetup.sh.

```
<agent-monitor>adhv9t</agent-monitor>
<agent-monitor-path>/db2home/adhv9t/sqllib/security</agent-
monitor-path>
<nodes-config-path>/db2home/adhv9t/sqllib</nodes-config-path>
```

```
<object-qualifier>AMESHEMA</object-qualifier>
<server-address>txaix01</server-address>
<server-port>52521</server-port>
<server-repository>AMEDB2</server-repository>
```

3.9.4 Agent Collection Settings

The default settings for the agent and server assume a "real" workload, and it may take over 30 minutes for collected data to show up in the repository. If you are evaluating the product and want the data to show up sooner, the time intervals for the agent-interval and refresh-interval should be decreased to 300 seconds, and the refresh-count, decreased to 200 seconds to allow data to display in the reporting user interface much faster.

```
<agent-interval>300</agent-interval>
<refresh-count>200</refresh-count>
<refresh-interval>300</refresh-interval>
```

3.9.5 Agent Configuration parameters and recommendations

The following settings are recommended for the agent configuration for evaluations.

```
<agent-config>
  <agent-interval>300</agent-interval>                                (Recommended)
  <agent-monitor>adhv9t</agent-monitor>
  <agent-monitor-con-alias>monitor</agent-monitor-con-alias>
  <agent-monitor-path>/db2home/adhv9t/sqllib/security</agent-monitor-path>
  <agent-monitor-pwd>xxxxxxx</agent-monitor-pwd>
  <agent-monitor-usr>xxxxxxx</agent-monitor-usr>
  <community-string>mystring</community-string>
  <load-event-text>true</load-event-text>
  <log-level>I</log-level>
  <log-size-max>5</log-size-max>
  <multicast-address>236.1.2.4</multicast-address>
  <multicast-port>52523multicast-port>
  <nodes-config-path>/db2home/adhv9t/sqllib</nodes-config-path>
  <nodes-file-name>db2nodes.cfg</nodes-file-name>
  <object-collection>ADHMPV11</object-collection>
  <object-qualifier>AMESHEMA</object-qualifier>
  <process-cancel-timeout>5</process-cancel-timeout>
  <process-poll-rate>5</process-poll-rate>
  <record-delimiter>0xff</record-delimiter>
  <refresh-count>200</refresh-count>                                (Recommended)
  <refresh-interval>300</refresh-interval>                            (Recommended)
  <request-thread-timeout>100</request-thread-timeout>
  <server-address>txaix01</server-address>
  <server-con-alias>server</server-con-alias>
  <server-connect-retry-delay>10</server-connect-retry-delay>
  <server-connect-retry-max>30</server-connect-retry-max>
  <server-port>52521</server-port>
  <server-pwd>xxxxxxx</server-pwd>
  <server-repository>AMEDB2</repository>
  <server-usr>xxxxxxx</server-usr>
  <trace-audit-events>true</trace-audit-events>
  <trace-config>true</trace-config>
  <trace-db2-attachment>true</trace-db2-attachment>
  <trace-events>true</trace-events>
```

```
<trace-filters>true</trace-filters>
<trace-network>true</trace-network>
<trace-sql>true</trace-sql>
<trace-log-analysis>true</trace-log-analysis>
</agent-config>
```

3.9.6 How to Check If the Server and Agent Are Up

Issue the following command from the command line to check that the server is running:

```
$ ps -ef | grep adhae
ameuser 2064470      1  36 14:08:07 pts/11   0:00 ./adhaea -c
/home/adhmpptest/demoinstall/adh64/agent/conf/fullAgentconfig.xml -m
/home/adhmpptest/dem
oinstall/adh64/agent/aix64/messages/enu
ameuser 2129992      1   0 14:04:26  pts/6    0:00 ./adhaes -c
/home/adhmpptest/demoinstall/adh64/server/conf/fullServerconfig.xml
-m /home/adhmpptest/d
emoinstall/adh64/server/aix64/messages/enu
$
```

Information will be returned displaying the server and agent processes that are running.

3.9.7 Data Collection Considerations

A DB2 instance can process a huge amount of data. If DB2 Audit Management Expert for MP were configured to capture all of that activity, it would incur unnecessary overhead, require a huge repository, and most likely, not all of the activity will be useful.

It is essential to audit just the data of interest. Data of interest is any data that is sensitive in nature and requires auditing -- the activity that is truly useful to an auditor.

Filtering capability is available on both the collection side (before the data has been written to the repository), and on the reporting side (after the data has been collected and stored in the repository). It is wise to filter on the collection side instead of the reporting side so unnecessary data it is not written to the repository.

The data collection capabilities enable you to audit individual tables in a database. You can audit different users, applications and/or plans.

Two levels of filtering are available to reduce the audit data to a useful subset

Collection

The DB2 Audit Management Expert for MP audit administrator controls the amount of data collected and stored in the audit repository using a collection profile. With this collection profile you can collect a subset of the audit activity by specifying particular table names, DB2 users, plan names, schedule, and other criteria. A collection profile can be created in a test environment for refinement then promoted to the production environment for auditing purposes.

Use of Includes and/or Excludes

A major performance advantage of DB2 Audit Management Expert for MP is its ability to include and exclude data in a collection profile. For example, if we are sure that package A accesses table B securely, we may want to exclude that plan from the collection profile. The input/output (I/O) to the repository will be correspondingly reduced, and the overall performance of DB2 Audit Management Expert for MP improved. Consider when there are a million accesses and a large number of includes and/or excludes – in this case, saving the I/O to the repository is extremely beneficial.

Including and excluding data will increase CPU usage slightly, but from initial performance tests, the CPU usage of the DB2 Audit Management Expert for MP agent was a small percentage of the total CPU processing. It is always more efficient to use data filtering to exclude an unwanted event instead of collecting it and inserting it into the repository.

Reporting

The auditor can view a subset of the collected data by specifying time, instance, database, table, DB2 user, and other criteria.

3.9.8 Securing and Monitoring the Audit Data

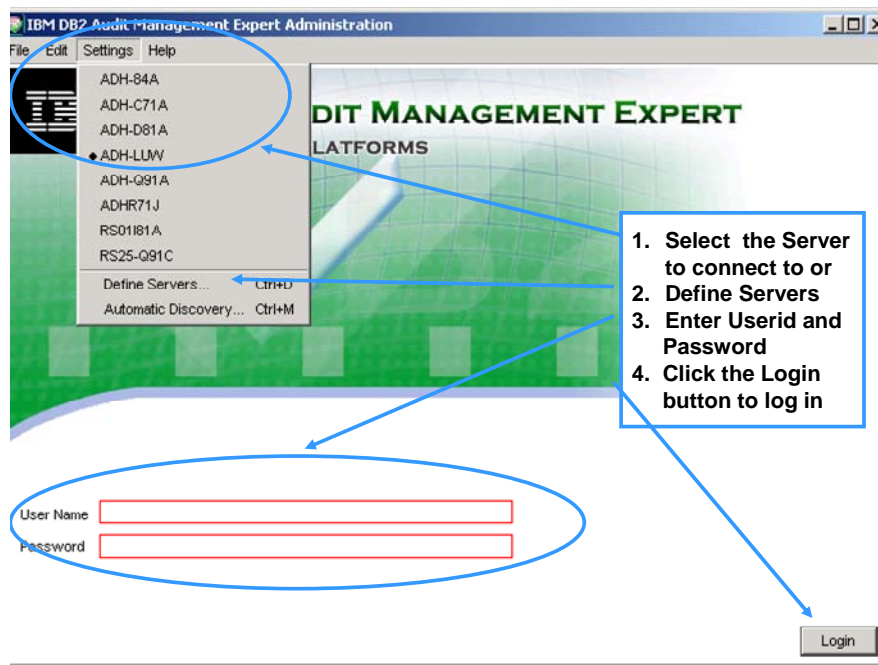
The auditor must have confidence that the audit data in the repository accurately reflects the audited instances and has not been tampered with. Some guidelines to follow are listed below:

- The server, agents, and JDBC need different authorities. Use a separate user ID for each. Give the minimum authority possible to each necessary user ID for the server, the agent and the reporting user interface.
 - Server: See script grantServer.sh
 - Agents: See script grantAgent.sh
 - Reporting user interface (JDBC): See script grantReport.sh
- The DB2 Audit Management Expert repository should be audited to ensure no one with authority has manipulated any audit data.
- Create a collection profile that monitors repository table activity by any user ID other than these three. Caution: monitoring these three User IDs is recursive and will cause the repository to grow exponentially.
- Segregation of duties: Ideally, production DBAs should not have access to the DB2 Audit Management Expert server and repository objects.

4 Administration User Interface

There are many tasks that you can accomplish with the DB2 Audit Management Expert Administration for MP User Interface. This section walks you through common tasks and includes instructions and screen captures.

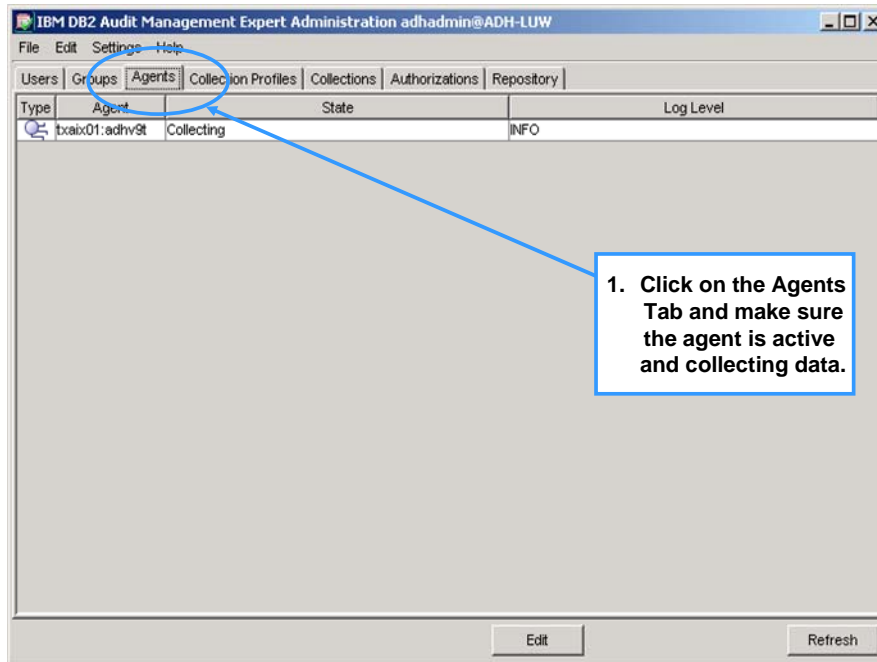
4.1 Logging in to the Administration User Interface



4.2 Add Users and Groups

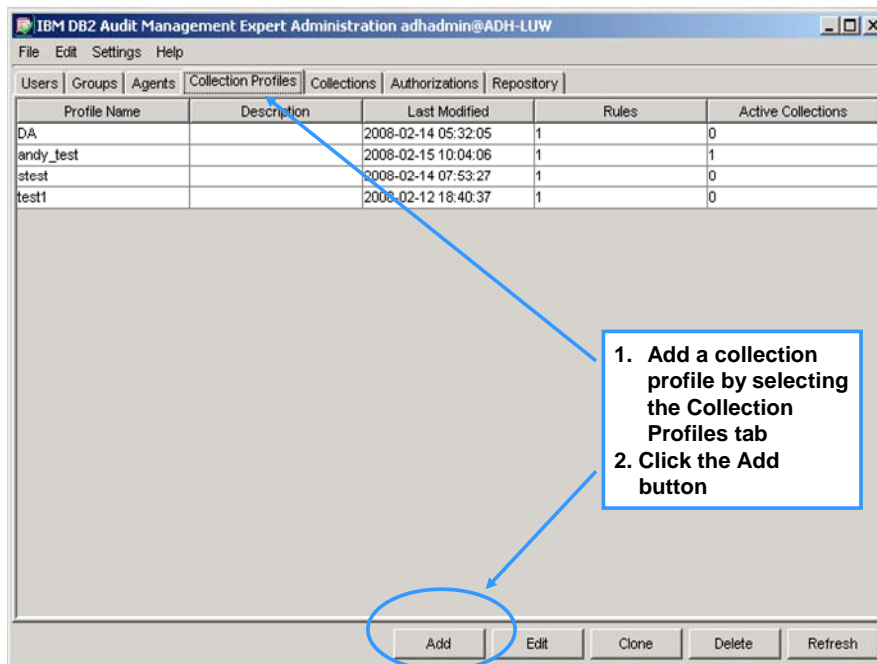
Adding Users and Groups are not demonstrated here. See the User's Guide for more information.

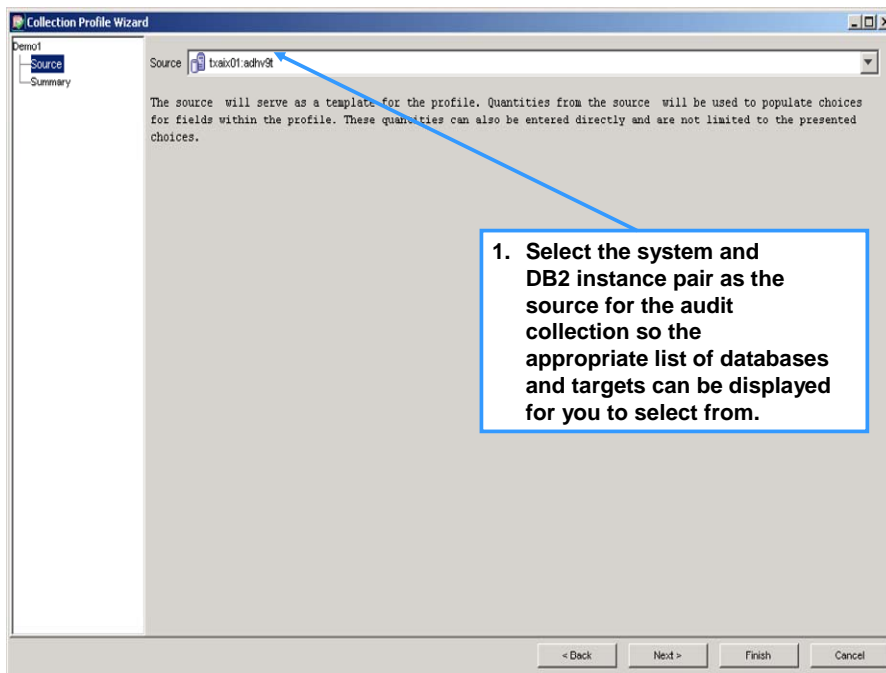
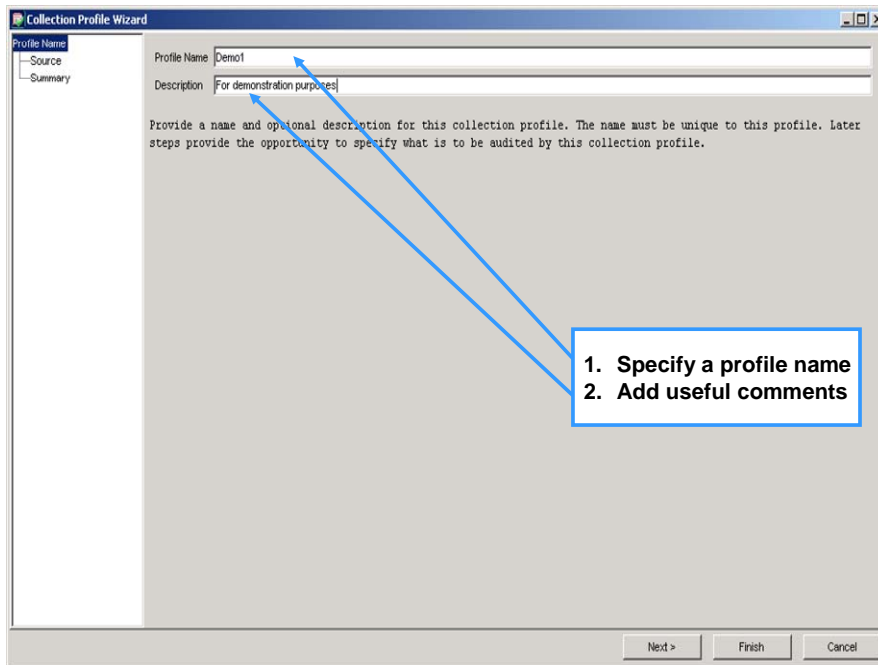
4.3 Check Agent Status



If the agent status is not in the collecting state, make sure that the agent configuration file is properly set up and that the agent has been started. Furthermore, check the agent log file to make sure that no errors have occurred after startup. For more information on how to properly configure the agent configuration file and on how to diagnose the agent log file, please refer to the DB2 Audit Management Expert for Multiplatform User's Guide.

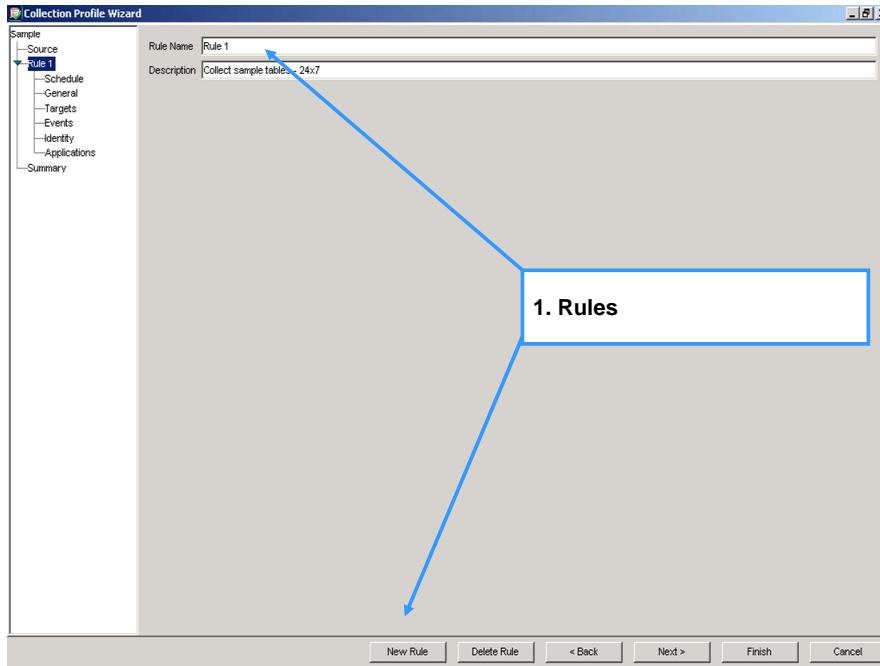
4.4 Add a Collection Profile



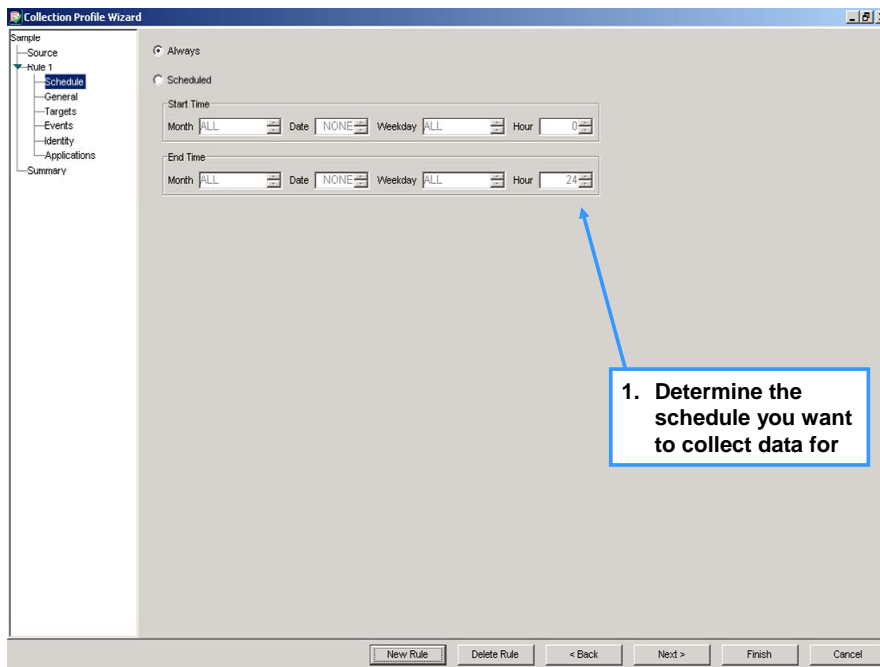


4.4.1 Adding Rules to the Collection Profile

The rules determine how DB2 Audit Management Expert for MP performs auditing. Rules are simply criteria on which the DB2 data will be collected. One or more rules comprise a collection profile. To add multiple rules, click the "New Rule" button at the bottom of the screen. Notice the indentation under, Rule1, on the left hand side of the screen. Those are the screens that contain parameters that need to be defined. The Next> button takes you through each screen.



4.4.2 Determine Collection Profile Schedule

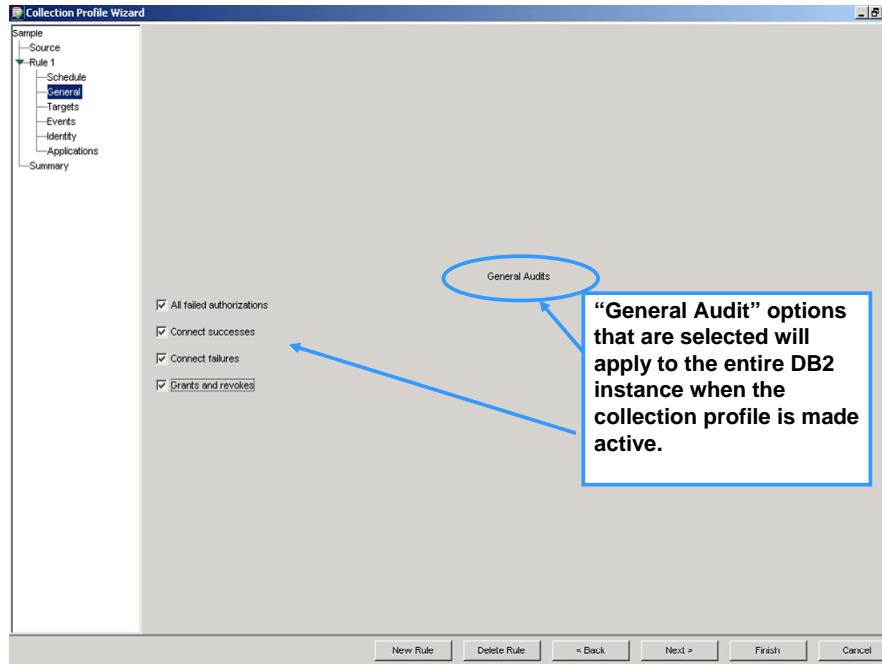


A schedule within a collection profile is simply a time-frame within which audit data should be collected for the monitored DB2 activity.

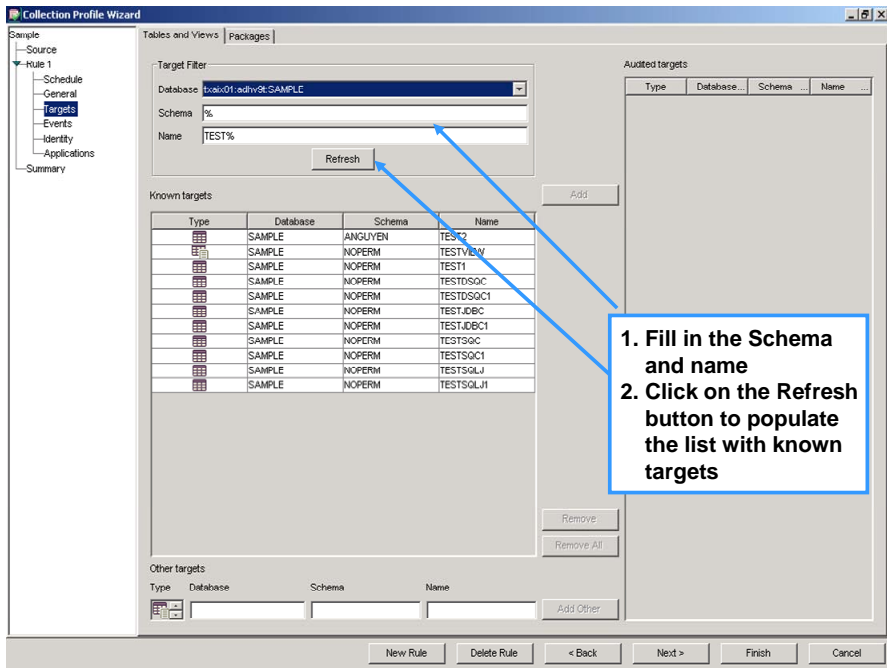
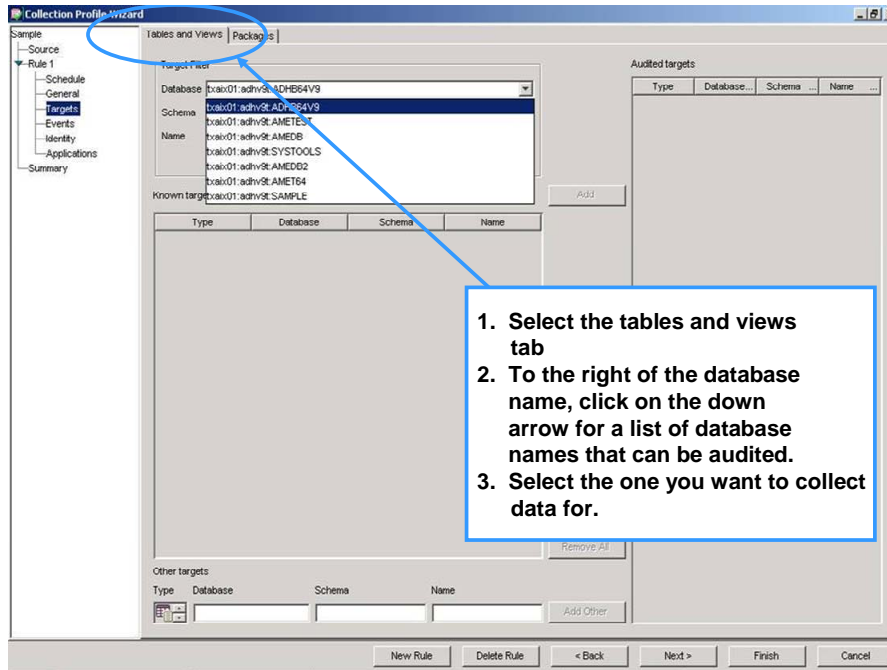
4.4.3 Select General Audit Options

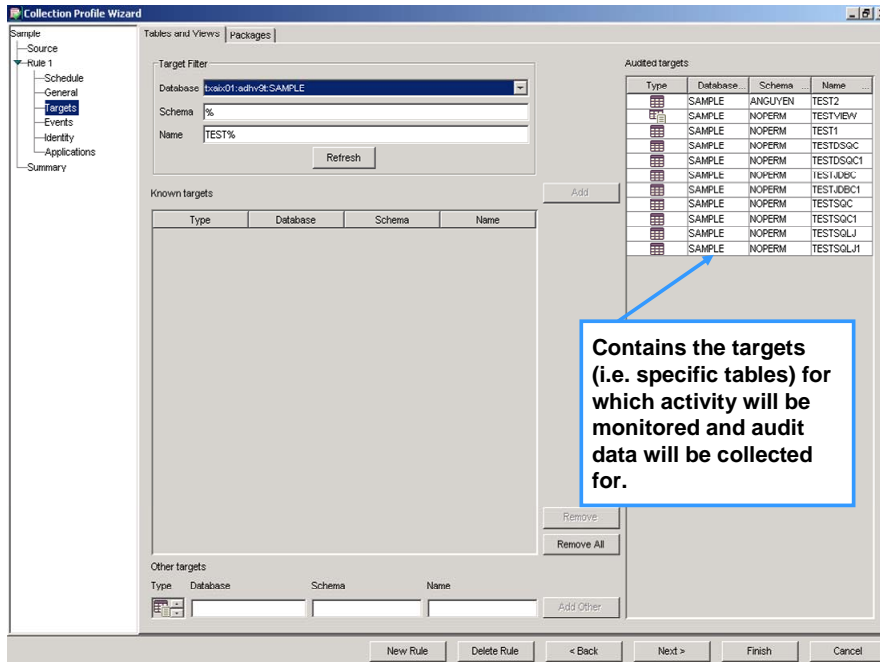
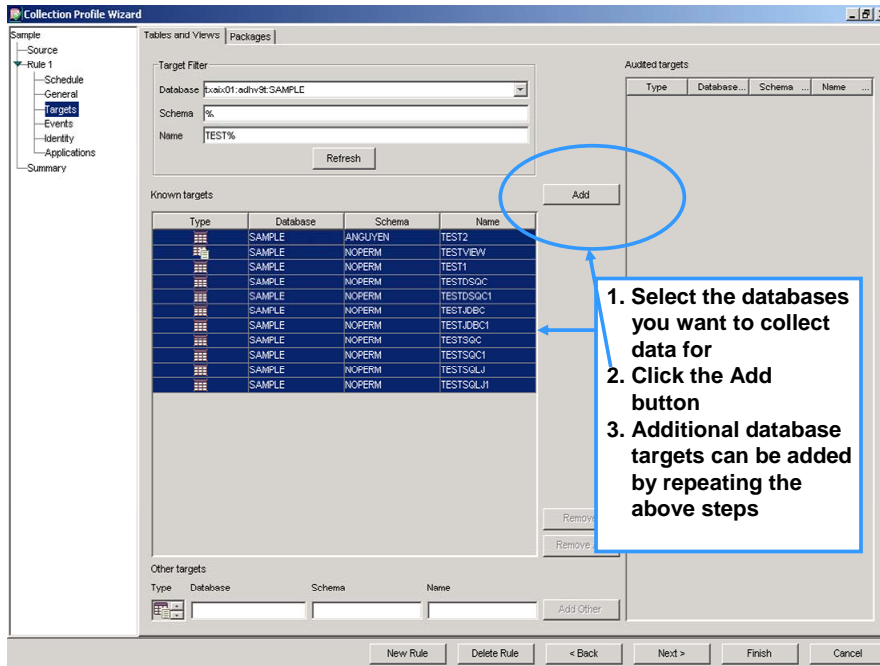
Select General Audit options. Later screens allow you to specify the data to be collected based on specific targets, events, identities and applications criteria.

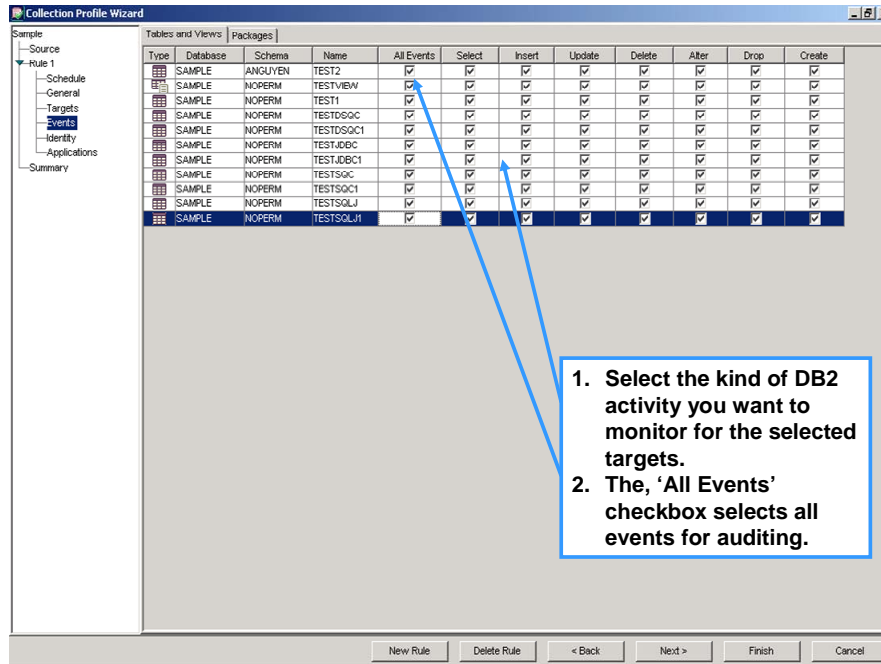
To see the capabilities of DB2 Audit Management Expert for MP, select all options.



DB2 Audit Management Expert for MP allows you to select audited targets from a list of available databases within the source DB2 instance. Targets are simply the individual tables, views, and packages within a database for which audit data is to be collected. Targets from multiple databases can be added to a rule by selecting a database, choosing the desired targets from that database, adding them to the list of Audited Targets, and repeating the process for another database. However, it is more typical to add targets from different databases into separate rules. The screenshots below will show how targets can be selected for auditing.





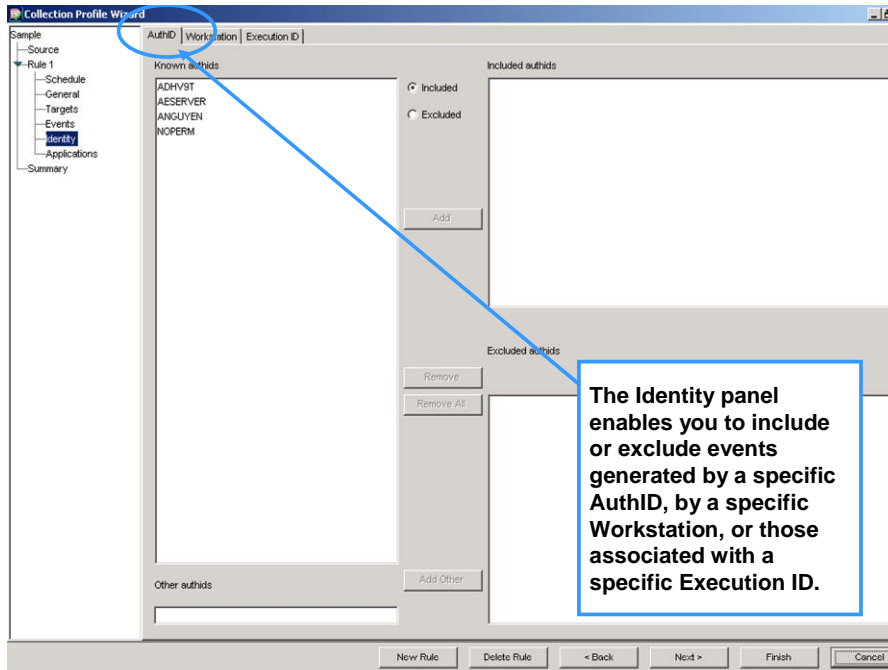


4.4.4 Include or Exclude Identities

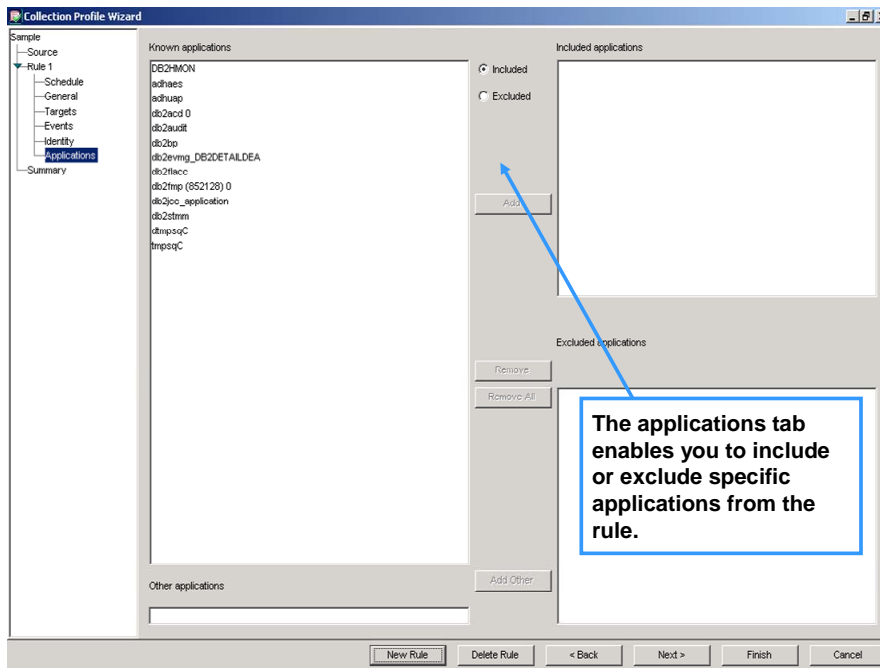
A collection profile can be set up to monitor events for specific AuthIDs, workstations, execution IDs and applications. But if you don't select to include or exclude any values, the default is to collect for all AuthIDs, workstations, execution IDs and applications. However, it should be noted that if, for example, you include one specific AuthID, then only data for that AuthID is collected and any others are implicitly excluded.

Warning: If you audit the repository tables, you should not monitor transactions for the user ids running the agent and server.

The data used to populate the lists of AuthIDs, workstations, execution IDs and applications comes from audit data collected by DB2 Audit Management Expert. The first time you create a collection profile there will be no data in the repository, so these lists will be empty. If you wish to include or exclude a specific name, type it in the "other" window at the bottom of the screen and click "add other".



4.4.5 Include or Exclude Applications

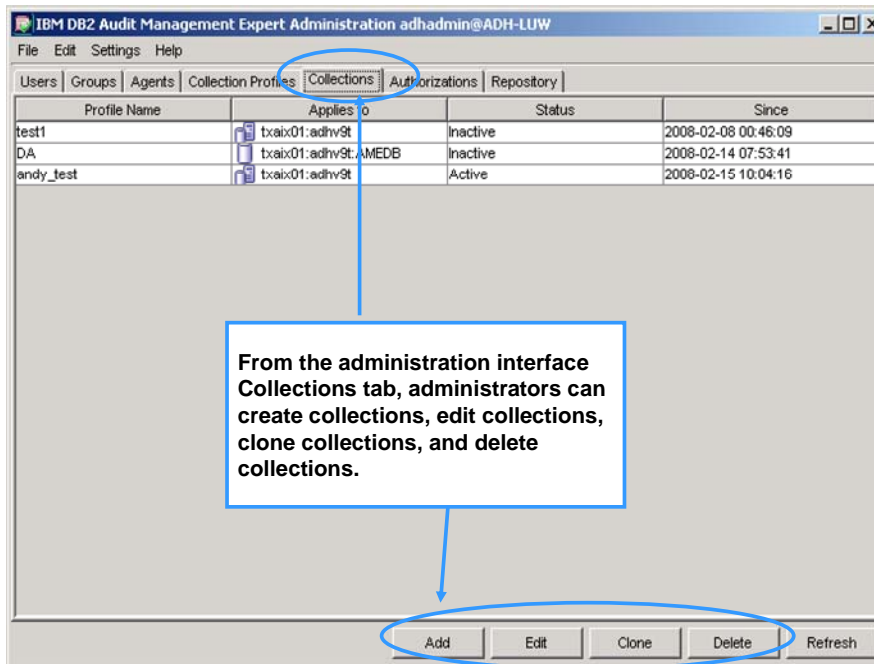


When you have completed adding a new Collection Profile, click Finish to save your changes and continue.

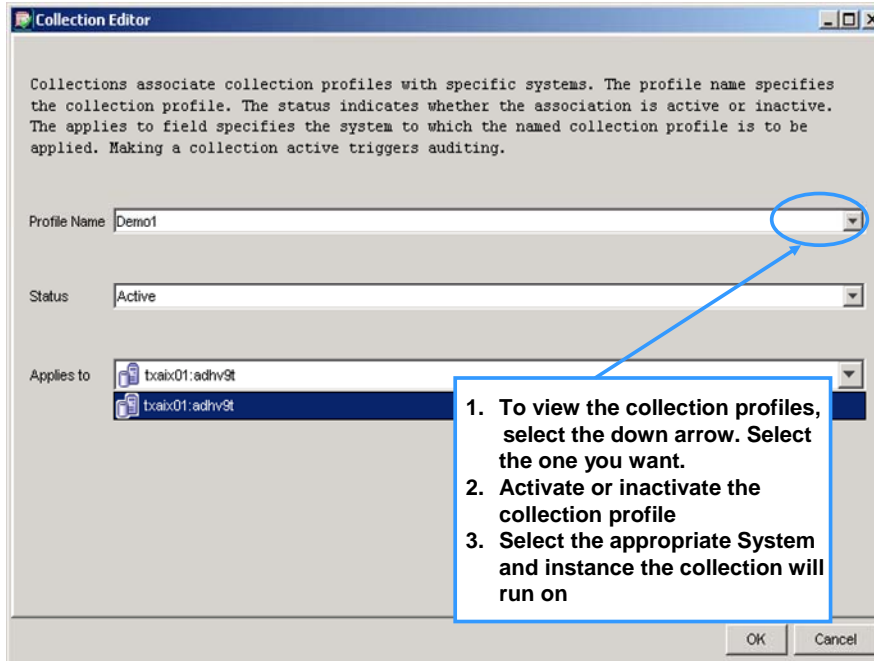
4.5 Collections

A *collection* associates a specific collection profile with a specific DB2 instance. Administrators can specify a collection as *active* (actively collecting audit data) or *inactive* (not actively collecting audit data). From the Collections tab, administrators can create collections, edit collections, clone collections, and delete collections.

For each collection, the Collections tab displays the name of the collection profile, the agent and instance to which the collection profile is applied, whether or not the collection is active, and the date of the last status change. If the collection profile to be applied to a specific DB2 instance does not appear as part of an active collection, you must create a Collection.

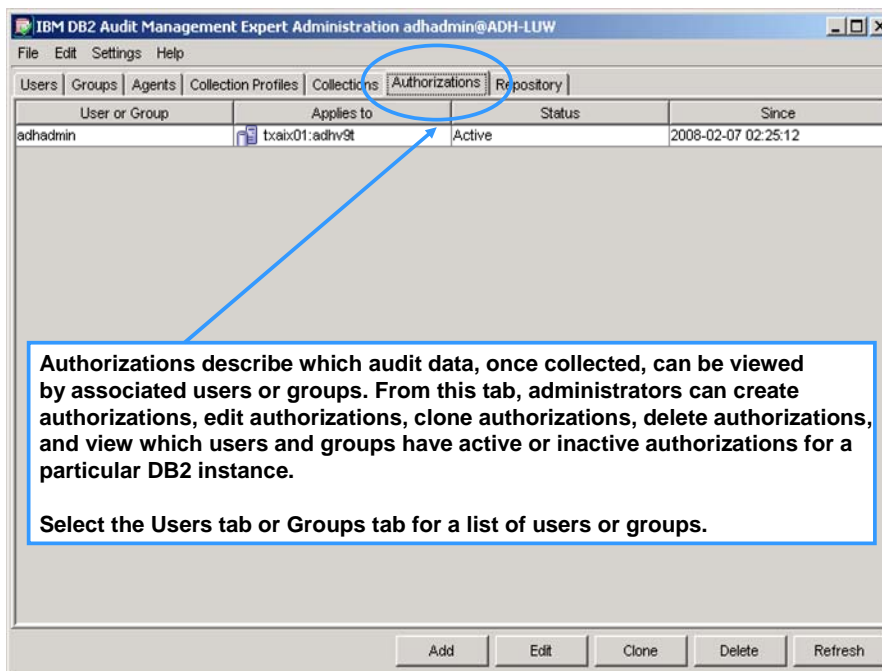


To create a collection, click the Add button. On the Collection Editor window, click the “Profile Name” list box and select the desired collection profile to be applied. Verify that the “Status” is set to be “Active”. Then click the “Applies to” list box and select the instance to which the collection profile is to be applied. Finally, click the Ok button.

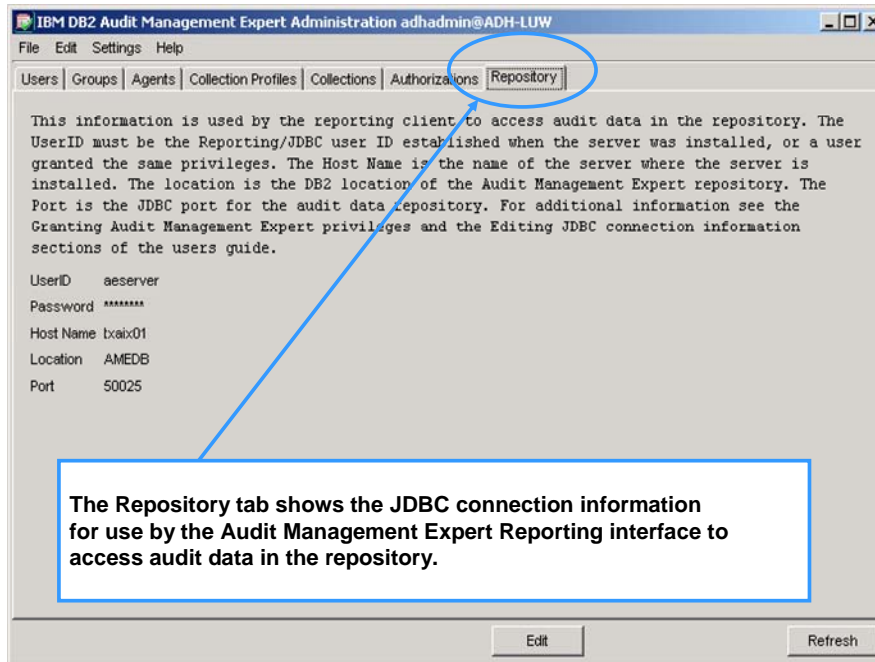


When pressing the OK button on the screen above, if a message dialog appears that says that an active collection is already targeting this location, it is because you can only have one collection profile active at a time against a single agent. However, since a collection profile can have multiple rules, a single active collection is capable of monitoring multiple audit events on targets within different databases.

4.6 Authorizations



4.7 Repository Tab

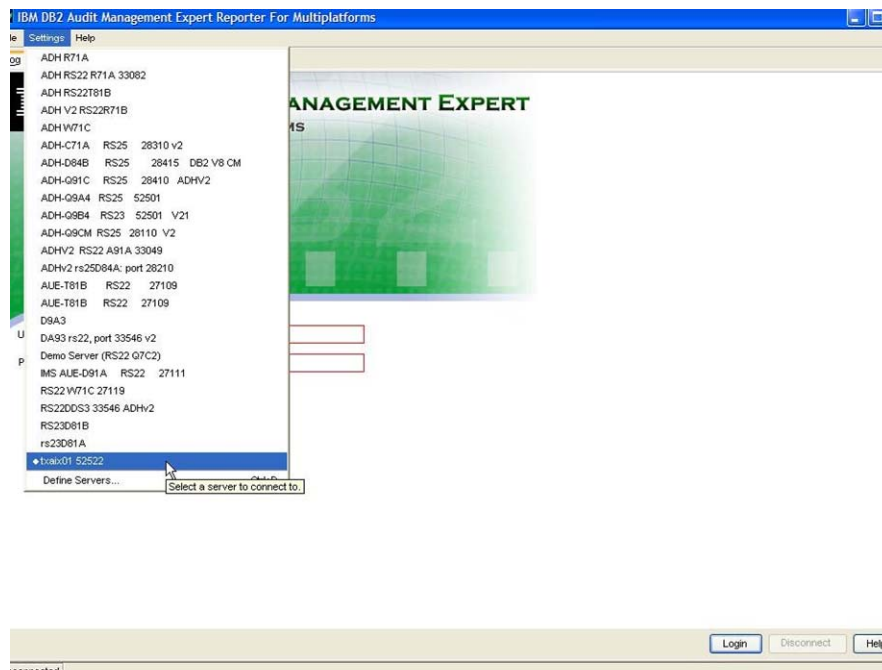


5 Reporting User Interface

There are many tasks that you can accomplish with the DB2 Audit Management Expert for MP Reporting User Interface. This section walks you through the common tasks with instructions and screen captures.

5.1 Logging in to the Reporting User Interface

Select the server to connect to or define servers. Then specify a userid and password and click on Login.



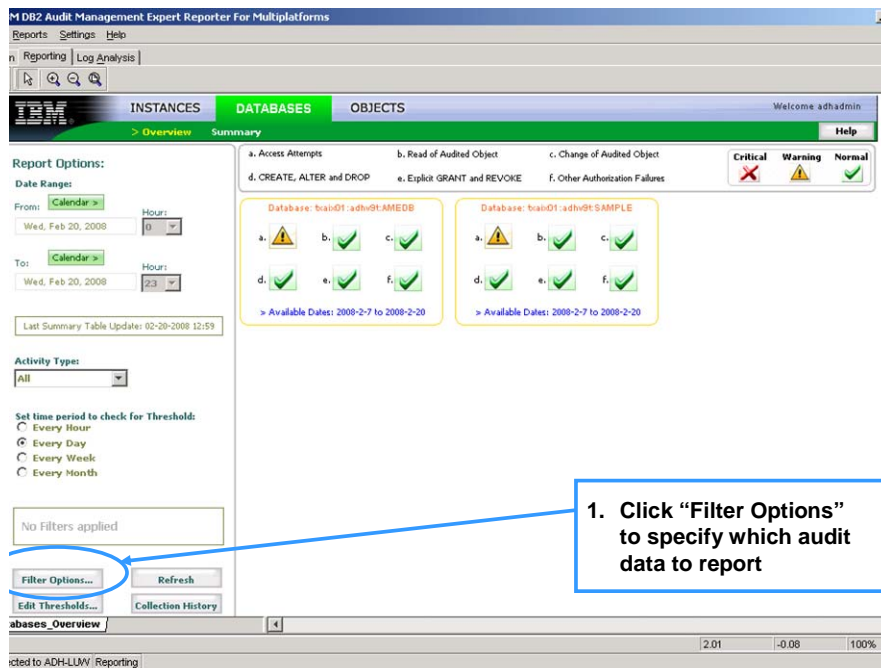
When login is complete, the status in the lower left-hand corner of the Reporting User Interface will display the text Connected to <Location>, indicating you have successfully logged into the Reporting User Interface, as shown below.



5.2 Overview report

The overview report is displayed under the green Databases tab and shows an overview of the activity that occurred on the databases for which an auditor is authorized to view data, by machine, instance, and then database. By default, the reporting user interface displays the initial overview report with no filtering applied and data is displayed for the current date only. Date ranges for the collected data can be altered to see data collected in a different date range. If there are any reports that have been previously saved, they can be loaded and viewed from the central repository by selecting the menu option Reports --> Report Open

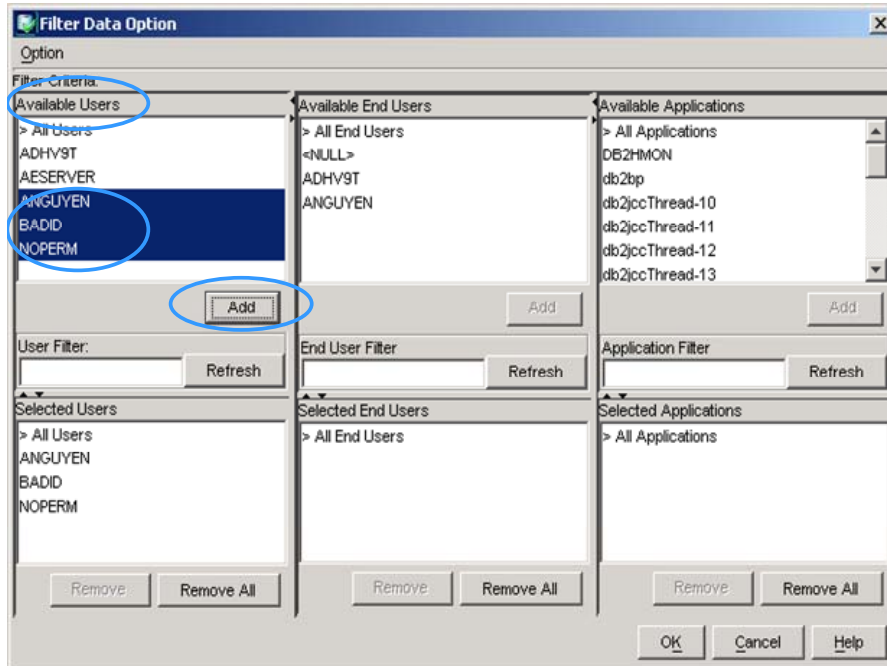
In the example below, item A in the overview represents access attempts and displays a yellow warning icon, meaning that the login attempts for this database exceeded the specified threshold of 500 attempts. Threshold values can be amended and saved with a report. Loading a saved report will then set the threshold values to the saved threshold values.



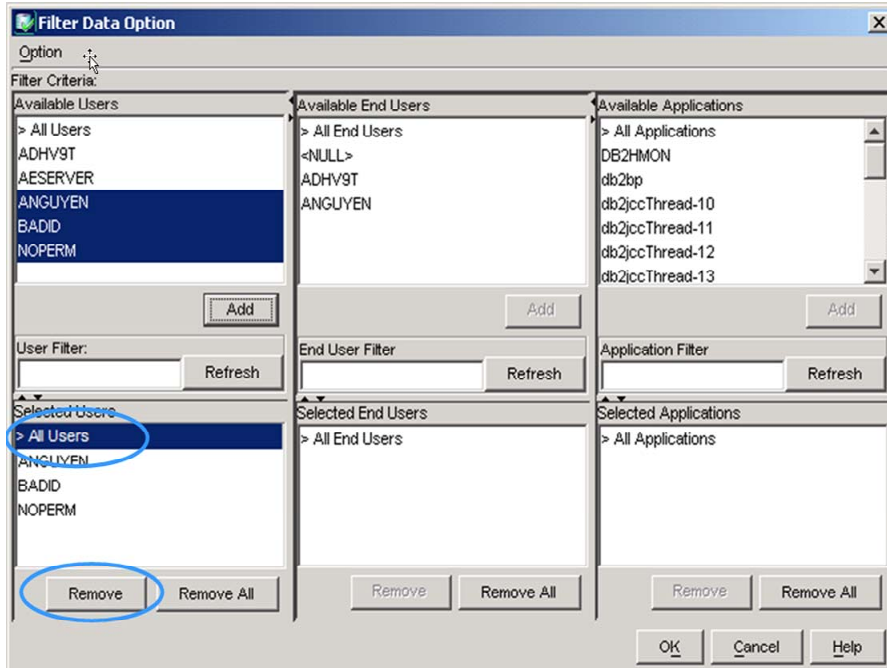
1. Click "Filter Options" to specify which audit data to report

5.3 Filters

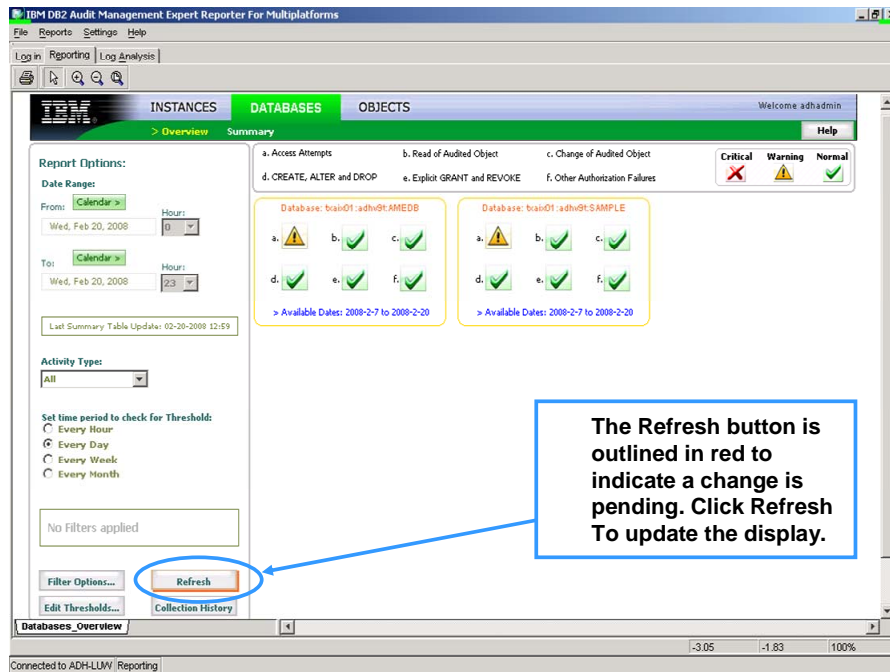
By default, the collected audit data is not filtered in reports, but you can create filters to display specific information only. In our example below, we can exclude ADHV9T data (the instance owner), and AESERVER (the Audit Management Expert server's user id) by highlighting all users in the "Available Users" window except the ones to exclude, and then clicking "Add".



In the “Selected Users” window, highlight “>All Users” and click “Remove” to report on specific users only.

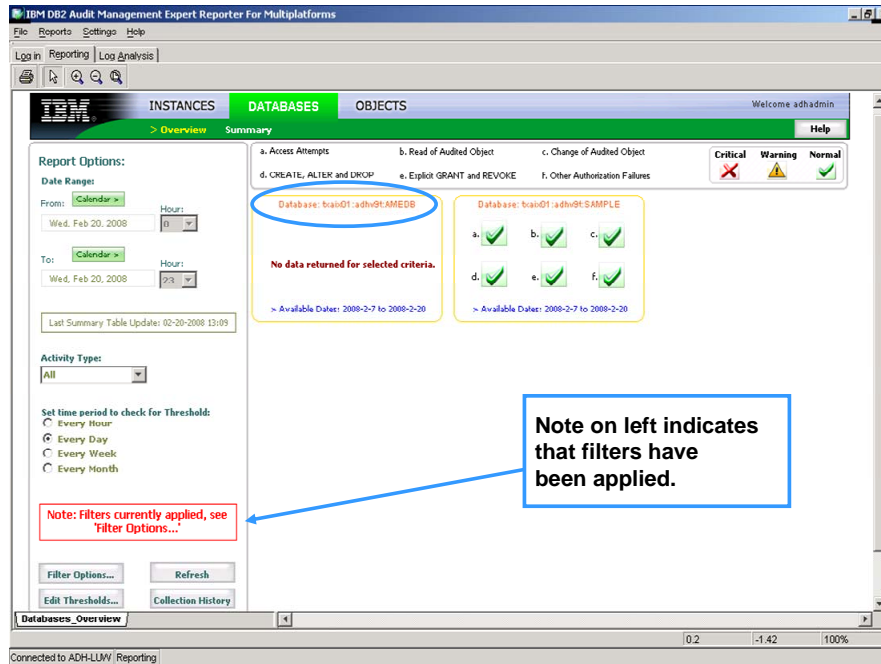


After adding filters, the overview report is shown again. The “Refresh” button is highlighted in red to indicate that a change in the report is pending. After making changes to the filter, click “Refresh” to update the summary report.



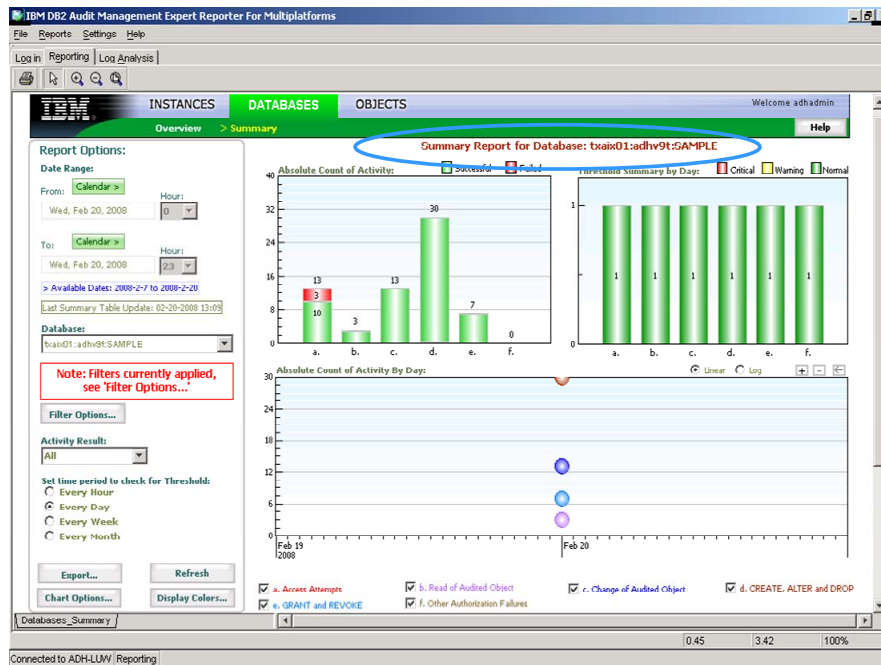
The red note on the left indicates that the filters have been applied. In our example below, there is no data for database AMEDB because the users specified in our filter did not

connect to it. Data is available, however, for these users for the database SAMPLE. Click anywhere in SAMPLE database window to drill down into the summary report.

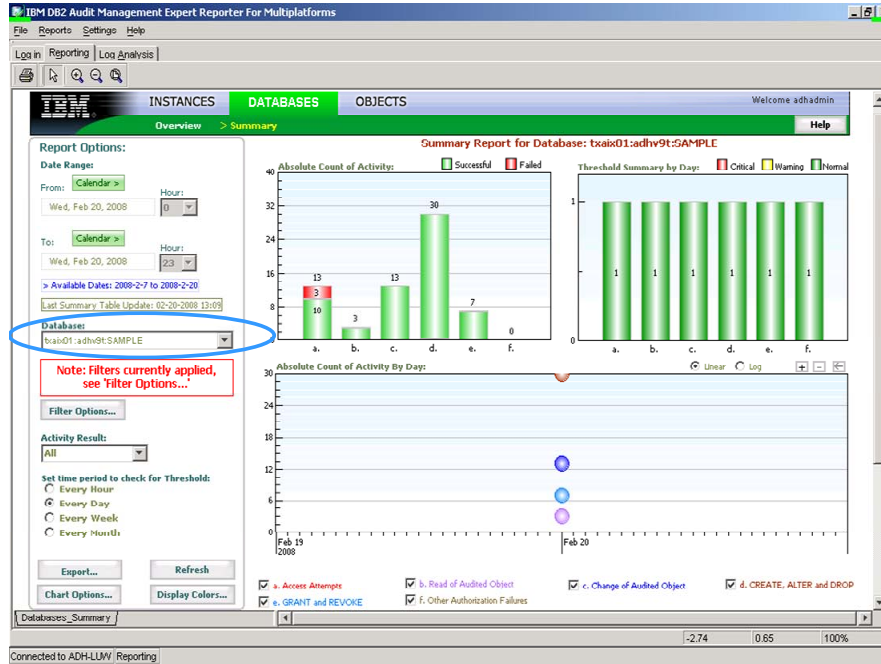


5.4 Summary Report

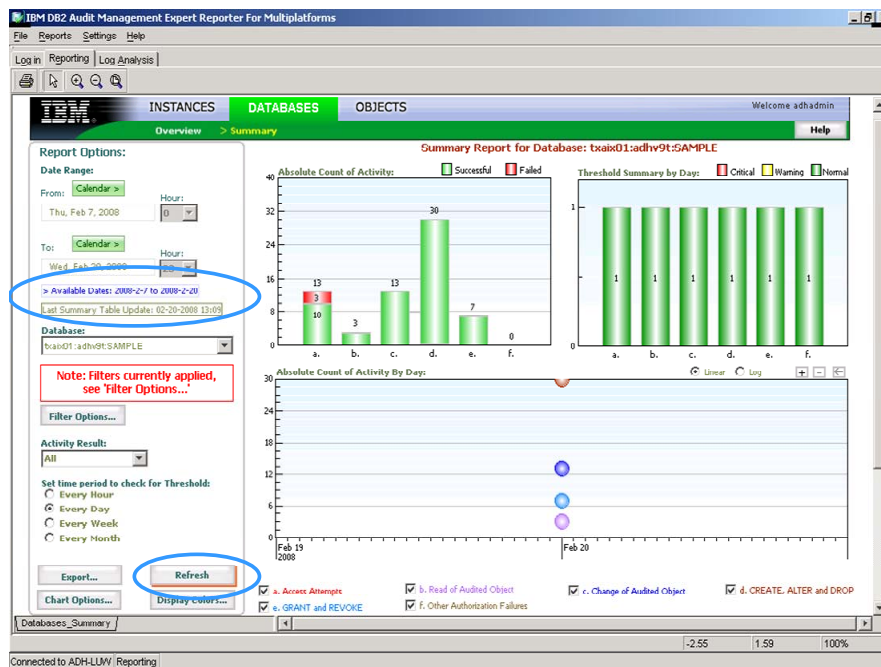
By clicking in the box for the database SAMPLE, the Summary Report "Summary report" information is displayed.



A different database can be selected using the “Database” drop-down box on the left. (Not shown in this example.)



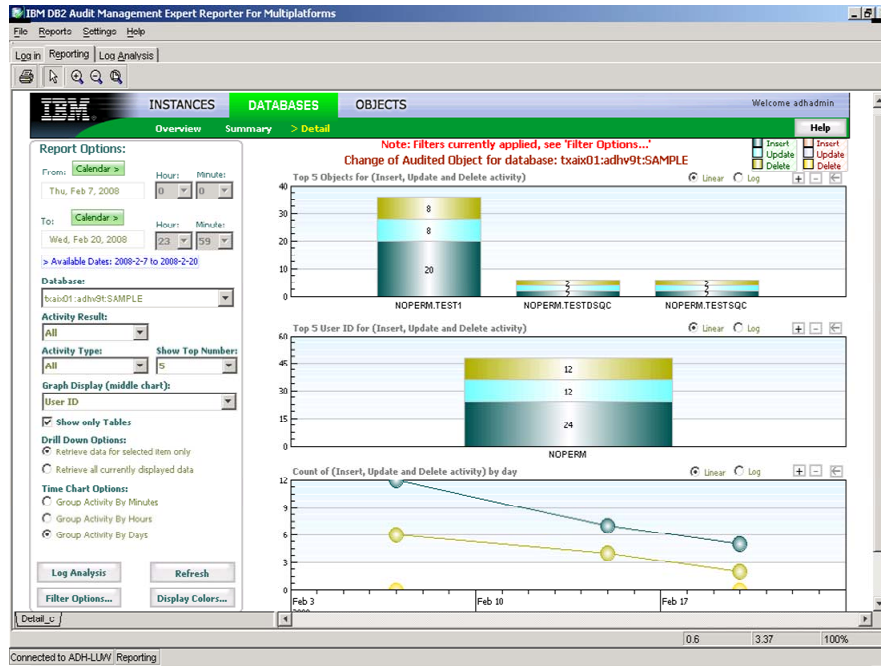
You can change the date range for which data is displayed by changing the dates on the calendar to the left. The “Refresh” button highlighted in red indicates the display is not up to date. Click “Refresh” to apply the changes.



This example shows the updated summary report for the database named SAMPLE. Note that more events are reported when the date range is changed. The colors of the bars indicate threshold levels. Drill down to see specific data by clicking the bar titled C to see changes of audited objects for the entire date range. Click the timeline dots shown in red and blue to see the data for a specific day. This results in displaying one of the many available detail reports.



As indicated at the top of the window in the detail report, the filter criteria and the existing date range is still in effect. Use the legend in the upper right-hand corner of the window to understand the data displayed in this detail report. This panel shows the changes that have occurred to this object by type: insert, update, and delete. There are no failed changes on this display. Click the table titled "NOPERM.TESTSQC" to drill down and see the collected data for that table. Note that all the displayed activity for the 3 tables was performed by user ID NOPERM, and that the table's schema is also titled NOPERM.



5.5 Detailed Data

This window shows detailed data for the changes that were made by user NOPERM for table TESTSQC. There are six records for this user/table. Note that all of the displayed activity for the table was performed by user ID NOPERM, and that the table's schema is also titled NOPERM.

The screenshot shows the Audit Management Expert Data for detail_c window. The record count is 6. The table below shows the detailed data for user NOPERM on table TESTSQC.

ROW	TIME	RESULT	SYSTEM	DATABASE	APPROVAL...	APPROVAL...	ACCESS_A...
1	2008-02-07 0...	0	bcaix01:adhv9t	SAMPLE	0x00000000...	OBJECT PRI...	0x00000000...
2	2008-02-07 0...	0	bcaix01:adhv9t	SAMPLE	0x00000000...	OBJECT PRI...	0x00000000...
3	2008-02-07 0...	0	bcaix01:adhv9t	SAMPLE	0x00000000...	OBJECT PRI...	0x00000000...
4	2008-02-15 1...	0	bcaix01:adhv9t	SAMPLE	0x00000000...	OBJECT PRI...	0x00000000...
5	2008-02-15 1...	0	bcaix01:adhv9t	SAMPLE	0x00000000...	OBJECT PRI...	0x00000000...
6	2008-02-15 1...	0	bcaix01:adhv9t	SAMPLE	0x00000000...	OBJECT PRI...	0x00000000...

Scrolling to the right, the following example shows the column ACCESS_ATTEMPTED_VALUE which breaks down the access by insert, update, and delete.

Option

Record Count: 6

ACCESS_A...	ACCESS_ATTEMPTED_VALUE	CORR_ID	CATEGORY	EVENT_TYPE	CONTEXT_T...
0x00000000...	INSERT	7	CHECKING	CHECKING_...	N
0x00000000...	UPDATE	12	CHECKING	CHECKING_...	N
0x00000000...	DELETE	14	CHECKING	CHECKING_...	N
0x00000000...	INSERT	7	CHECKING	CHECKING_...	N
0x00000000...	UPDATE	12	CHECKING	CHECKING_...	N
0x00000000...	DELETE	14	CHECKING	CHECKING_...	N

Copy Export Cancel Drill Down Close

Scroll to the right to see more data. In this example, the column CONTAINER displays the schema name for this object, and the column NAME is the table's name on which the activity occurred.

Option

Record Count: 6

Y	EVENT_TYPE	CONTEXT_T...	CONTAINER	XSHEMA	NAME	TYPE	USER_ID
	CHECKING_...		NOPERM		TESTSQC	TABLE	noperm
	CHECKING_...		NOPERM		TESTSQC	TABLE	noperm
	CHECKING_...		NOPERM		TESTSQC	TABLE	noperm
	CHECKING_...		NOPERM		TESTSQC	TABLE	noperm
	CHECKING_...		NOPERM		TESTSQC	TABLE	noperm
	CHECKING_...		NOPERM		TESTSQC	TABLE	noperm

Copy Export Cancel Drill Down Close

The column AUTHORIZATION ID is the ID that connected to the repository. The column END_USER_ID is the user's login ID that accessed the object.

The screenshot shows a window titled "Audit Management Expert Data for detail_c". Below the title bar is an "Option" section. Below that, it says "Record Count: 6". The main area contains a table with the following columns: USER_ID, AUTHORIZATION ID, END_USER_ID, END_USR_T..., END_USR_..., XSHEMA2, NAME2, and SECTIC. The table contains six rows of data, all with "noperm" in the USER_ID column, "NOPERM" in the AUTHORIZATION ID column, and "anguyen" in the END_USER_ID column. A blue circle is drawn around the first three columns of the table. At the bottom of the window are buttons for "Copy", "Export", "Cancel", "Drill Down", and "Close".

USER_ID	AUTHORIZATION ID	END_USER_ID	END_USR_T...	END_USR_...	XSHEMA2	NAME2	SECTIC
noperm	NOPERM	anguyen		*LOCAL	NOPERM	TMPSQC	
noperm	NOPERM	anguyen		*LOCAL	NOPERM	TMPSQC	
noperm	NOPERM	anguyen		*LOCAL	NOPERM	TMPSQC	
noperm	NOPERM	anguyen		*LOCAL	NOPERM	TMPSQC	
noperm	NOPERM	anguyen		*LOCAL	NOPERM	TMPSQC	
noperm	NOPERM	anguyen		*LOCAL	NOPERM	TMPSQC	

Continue scrolling to the far right and expand the column STATEMENT_TXT. You will see the dynamic SQL statement that was executed for non-SYSADM and DBADM users. DB2AUDIT does not collect static SQL statements, but does record binds to a database. You can also highlight a row and click "Drill down" for specific information, but in this example, the events in this table don't provide any interesting detail.

Audit Management Expert Data for detail_c

Option

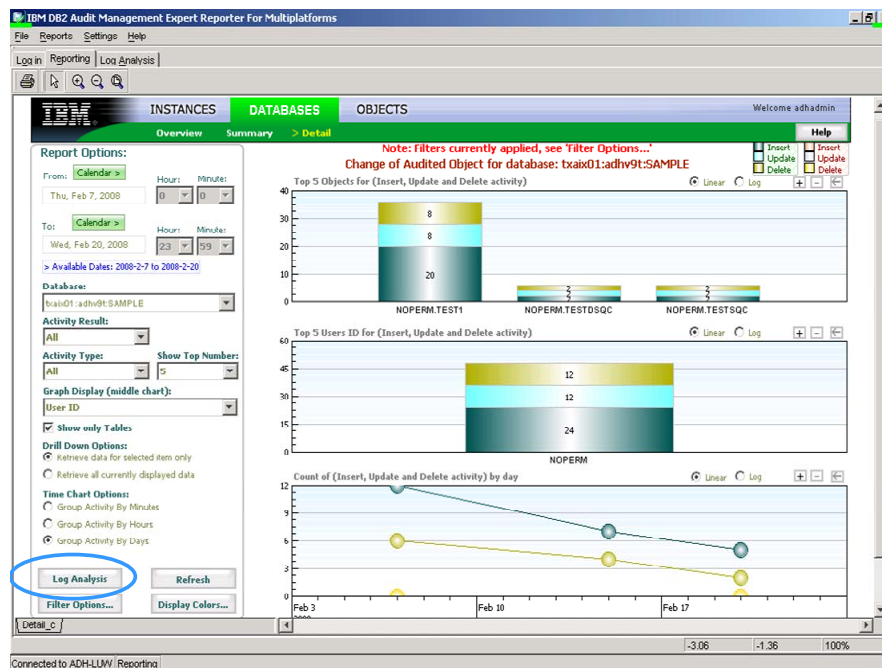
Record Count: 6

GRANTOR	GRANTEE	STATEMENT_TXT
		DELETE FROM noperm.TESTSQC WHERE 1 = ?
		DELETE FROM noperm.TESTSQC WHERE 1 = ?
		INSERT INTO noperm.TESTSQC (ID, NAME) VALUES (123, "nametest")
		INSERT INTO noperm.TESTSQC (ID, NAME) VALUES (123, "nametest")
		UPDATE noperm.TESTSQC SET ID = 999 WHERE 1 = ?
		UPDATE noperm.TESTSQC SET ID = 999 WHERE 1 = ?

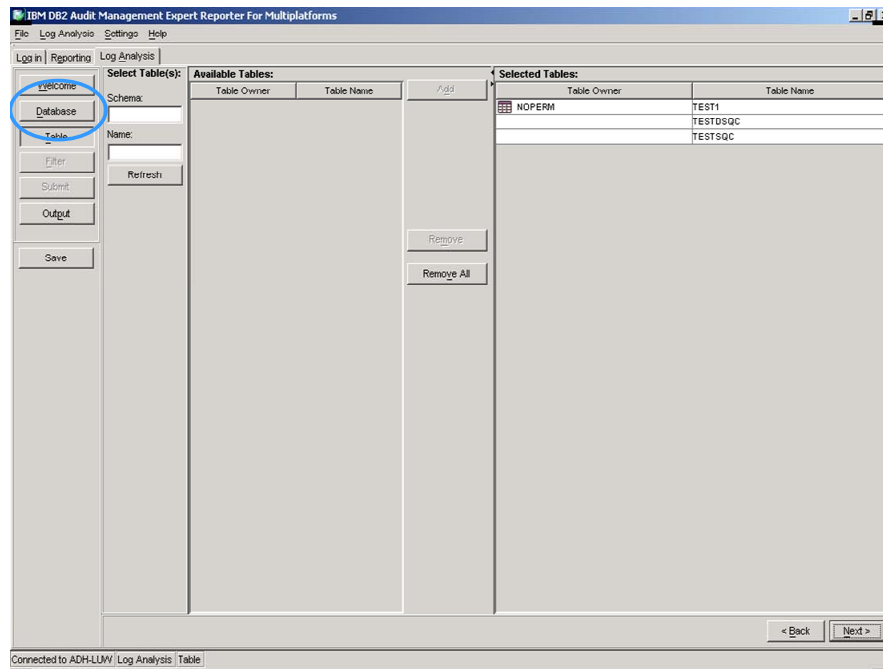
Copy Export Cancel Drill Down Close

5.6 Log Analysis


Go back to the summary report for the database SAMPLE. Click the “Log Analysis” button on the left. Click “yes” in the pop-up window to bring up the Log Analysis wizard.



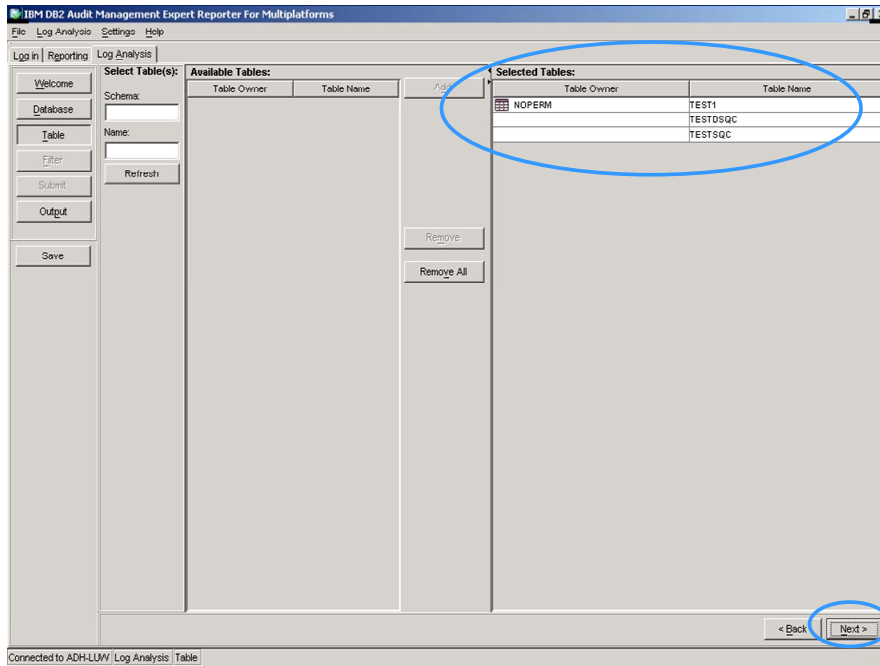
The Log Analysis wizard appears and selects the SAMPLE database by default because it was the object being viewed in the reports. Confirm this by clicking on the “Database” button.



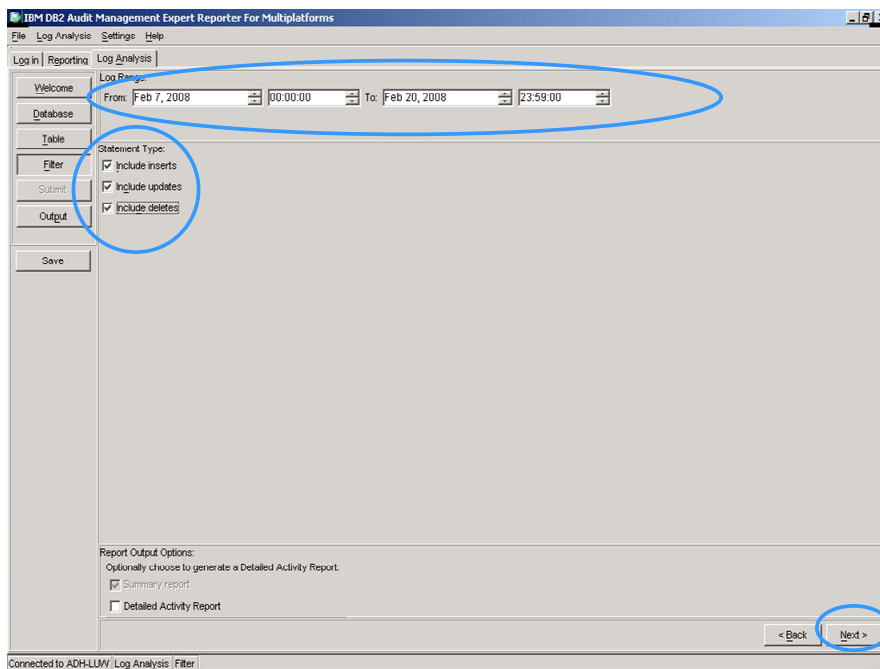
This is the Log Analysis wizard database window. The database SAMPLE is highlighted or selected. You can change the database for Log Analysis if you wish but in our example, we will use SAMPLE. Click “Next” or “Table” to move to the next window.




Note that all tables displayed in the detailed report for the database SAMPLE are pre-selected. Click “Next” to open the “Filter” window.



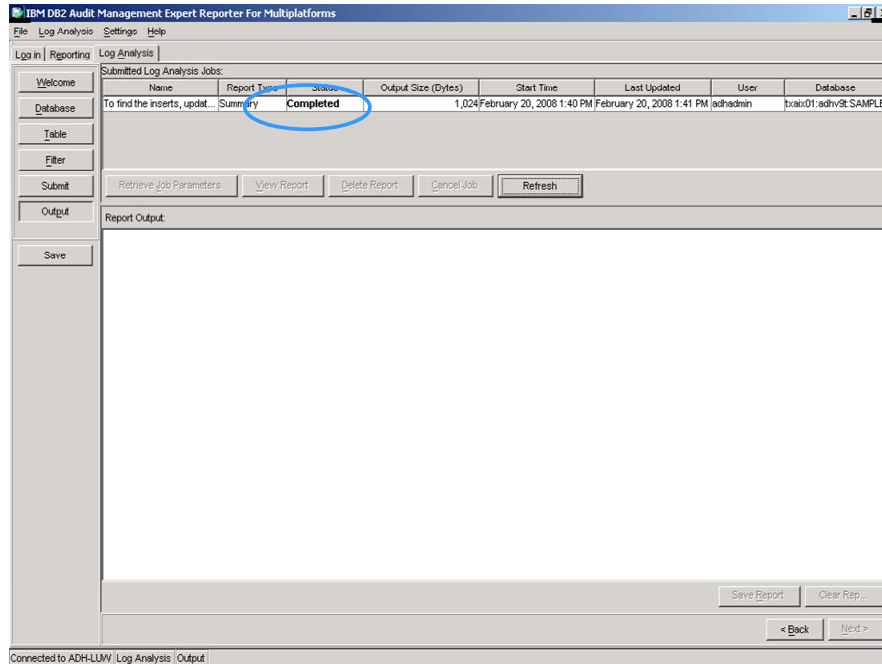
The Log Analysis Wizard filter window uses the date range of the detail report. Select what actions you are interested in viewing: inserts, updates, deletes, or all of them. Click “Next” to open the “Submit” window.



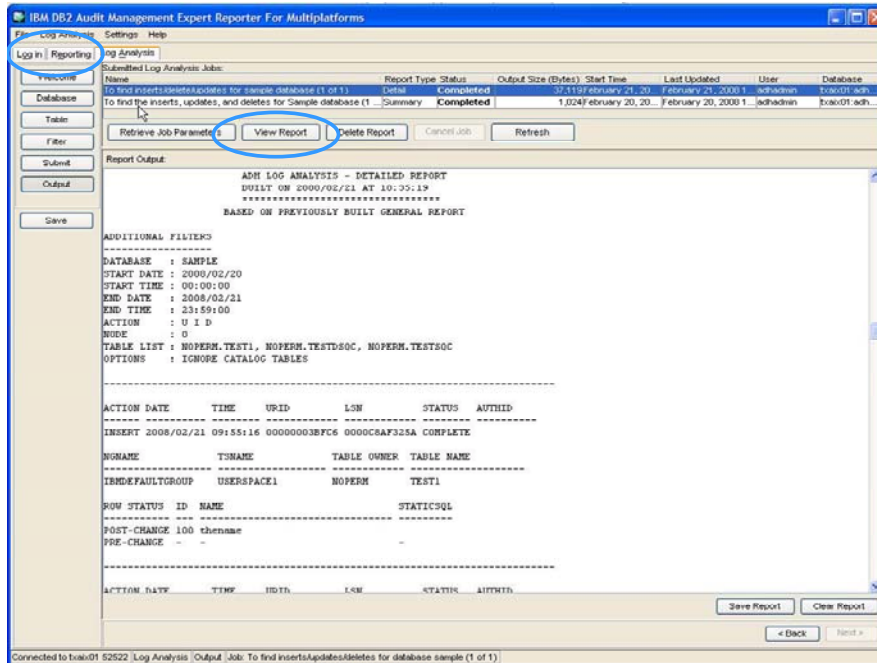
Enter a “Log Analysis Job Description” at the bottom of the window and click “Submit” on the bottom right. Click “Next” to open the “Output” window.



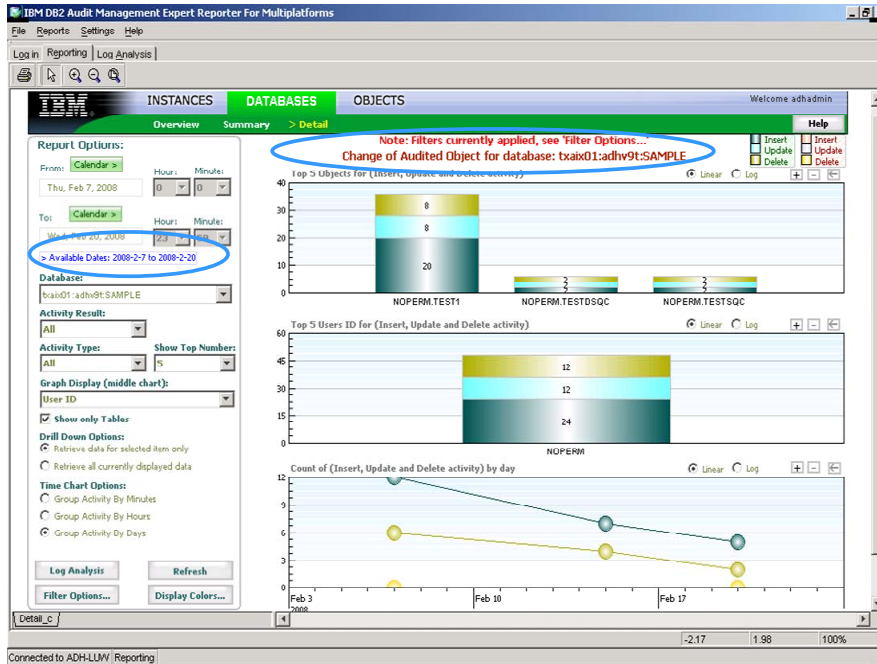
The Log Analysis Output window shows output for Log Analysis jobs. If your job does not have a status of “Completed”, click “Refresh” periodically until it finishes.



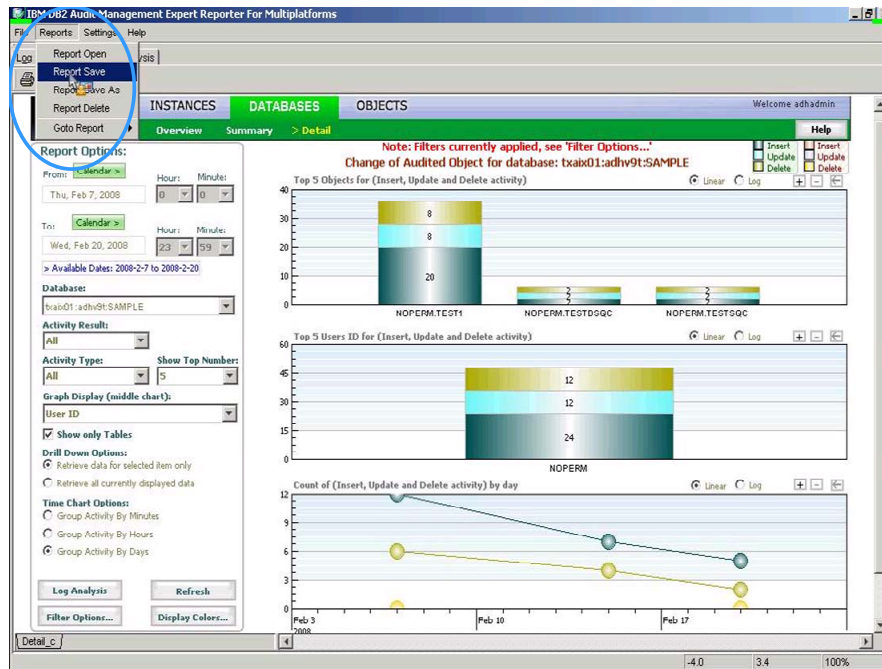
Once your Log Analysis job shows “Completed”, click “View Report” to view the Log Analysis report. Click on the “Reporting” tab on the top left to return to the summary report.



The summary report is displayed for database SAMPLE. Note that the current filters and date ranges previously selected are still in effect.



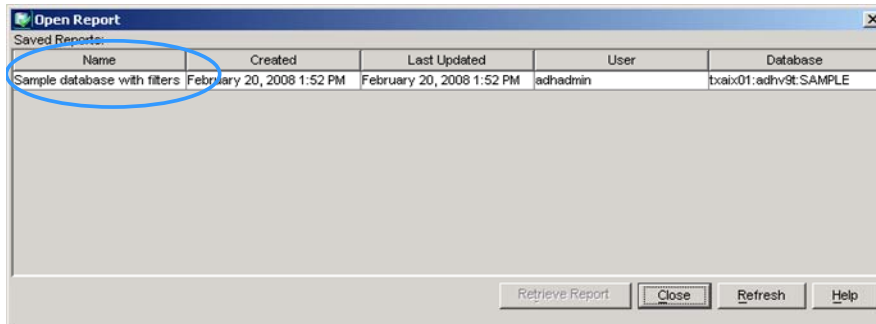
If you want to save a report, click on the “Reports” tab, and then click “Report Save”. This report can be retrieved utilizing these specific filters at any time. The saved reports are stored in a central repository so they are available from any workstation.



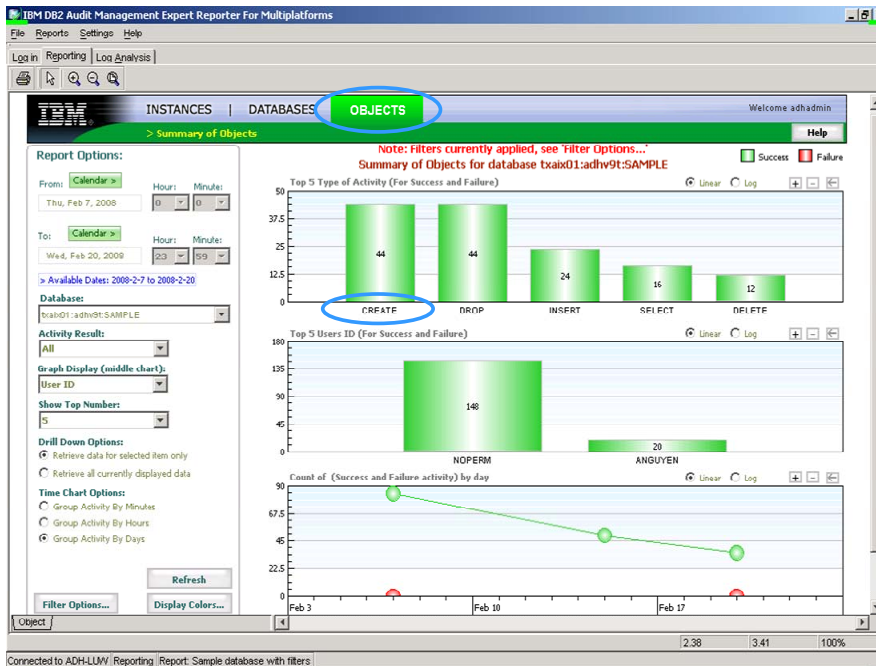
The screenshot shows the 'Save Report' dialog box. It contains a table of 'Saved Reports' with columns for Name, Created, Last Updated, User, and Database. Below the table, the 'Name' field is populated with 'Sample database with filters'. The 'Save' button is highlighted with a blue circle.

Name	Created	Last Updated	User	Database
Sample database with filters				

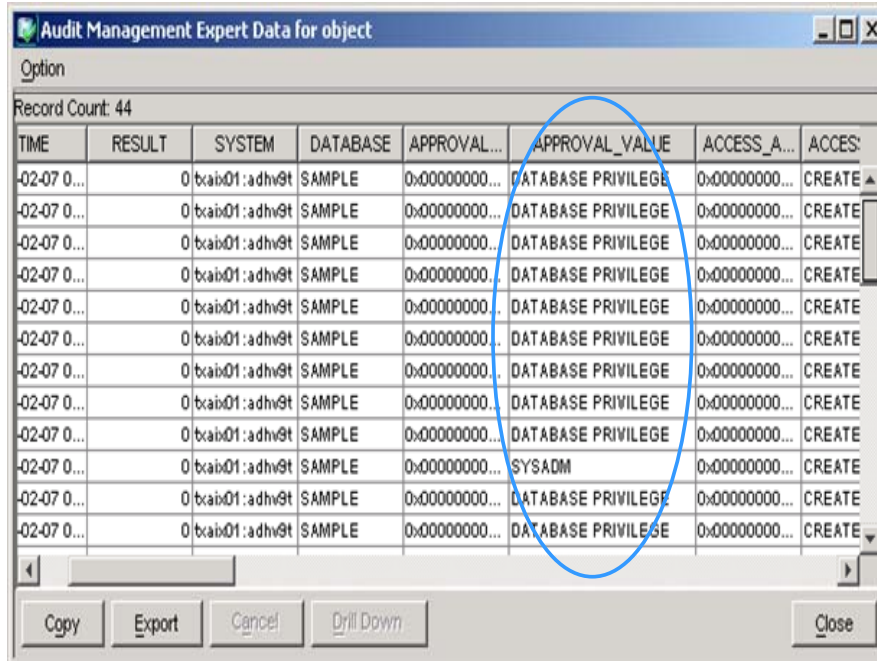
To retrieve a saved report, open the report window and click “Reports” --> “Report Open”. Highlight the report you want to view and click “Retrieve Report”. The saved filter criteria, date range, and threshold values are restored.



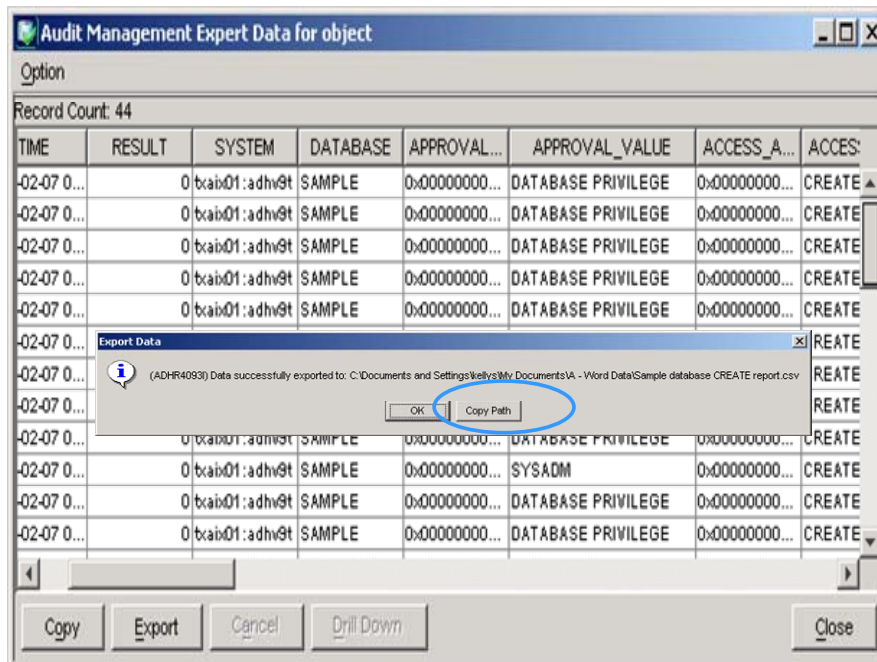
Click the “Objects” tab to see activity reports for the objects associated with the database. Still using the database SAMPLE in our example, the same filter criteria and date range is displayed, and all activity for views and tables are shown. Click on the “CREATE” bar to drill down to the details of the Create activity.



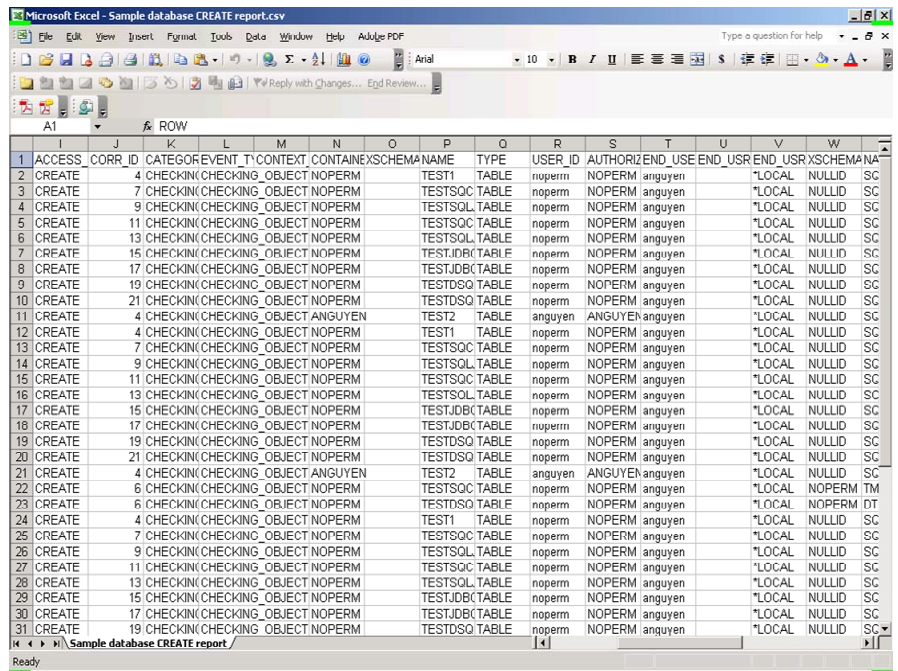
Detailed data for the objects are displayed. This data can be exported to a CSV file by clicking the “Export” button and supplying a name for the data you want to save. In this example, the privilege of the user is displayed.



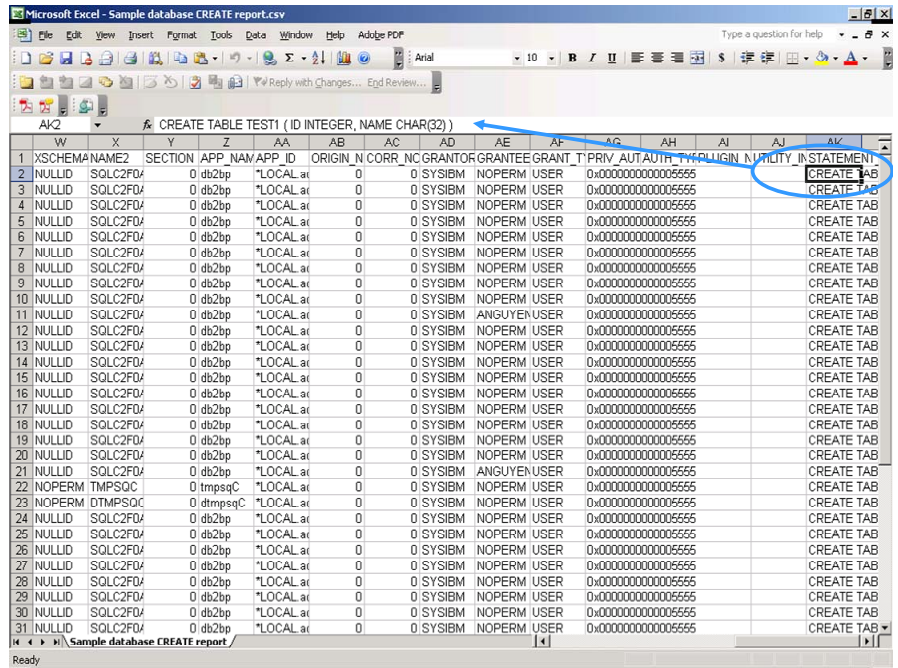
Once the data has been successfully exported, a pop-up window displays the pathname to the saved data. The "Copy Path" button copies the pathname to your PC's clipboard.



Using your favorite spreadsheet program, open your exported CSV data file.



Scroll to the right to see the column STATEMENT_TXT and left click on a cell. The full SQL statement text can be seen in the editing buffer above the spreadsheet.





© Rocket Software Inc. 2008
© Copyright IBM Corporation 2008

IBM United States of America
Produced in the United States of America
All Rights Reserved

The e-business logo, the eServer logo, IBM, the IBM logo, OS/390, zSeries, SecureWay, S/390, Tivoli, DB2, Lotus and WebSphere are trademarks of International Business Machines Corporation in the United States, other countries or both.

Lotus, Lotus Discovery Server, Lotus QuickPlace, Lotus Notes, Domino, and Sametime are trademarks of Lotus Development Corporation and/or IBM Corporation.

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

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

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PAPER "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.

Information in this paper as to the availability of products (including portlets) was believed accurate as of the time of publication. IBM cannot guarantee that identified products (including portlets) will continue to be made available by their suppliers.

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

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

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

IBM Director of Licensing
IBM Corporation
4205 South Miami Boulevard
Research Triangle Park, NC 27709 U.S.A.
