Using IBM InfoSphere Guardium for monitoring and auditing IBM DB2 for i database activity

Ensure compliance and create a tamper-proof audit trail

Kathryn Zeidenstein
Mark J. Anderson

December 16, 2013
(First published October 09, 2012)

IBM® InfoSphere® Guardium® is an enterprise information audit and protection solution that helps enterprises to protect and audit information across a diverse set of relational and nonrelational data sources such as Oracle, Teradata, IMS, VSAM, Microsoft® SharePoint, and IBM Netezza®, and IBM DB2® for z/OS®, and DB2 for Linux, UNIX and Windows. With InfoSphere Guardium V9.0, DB2 for i can now be included as a data source, enabling you to monitor access through native interfaces and through SQL. This article provides a brief overview of the InfoSphere Guardium architecture, describes how to configure access (including best practices for performance), and describes how to access data activity reports.

Overview

InfoSphere Guardium is an enterprise information database audit and protection solution that helps enterprise protect and audit information across a diverse set of relational and nonrelational data sources such as Oracle, Teradata, IMS, VSAM, Microsoft SharePoint, IBM Netezza, and DB2 for z/OS and DB2 for Linux, UNIX and Windows. With InfoSphere Guardium V9.0, DB2 for i can now be included as a data source, enabling you to monitor accesses from native interfaces as well as through SQL.

This article provides a brief overview of the InfoSphere Guardium architecture, describes how to configure access (including best practices for performance), and shows how to access data activity reports.

A short introduction to InfoSphere Guardium

The IBM InfoSphere Guardium solution evolved to address the particular needs of organizations that need to implement more automated and auditable data security practices. InfoSphere Guardium continuously monitors database transactions through lightweight software probes (refer

to Figure 1) installed on the database server (or file share or Microsoft SharePoint). These probes (known as S-TAPs, for software taps) monitor all database transactions, including those of the privileged users. The S-TAP also does not require any changes to the database or its applications.

The probes forward transactions to a hardened collector on the network, where they are compared to the previously defined policies to detect violations. The system can respond with a variety of policy-based actions, including generating an alert and for some databases can block the transaction in real time. (Blocking is not currently available for DB2 for i database activity monitoring.)

Figure 1. InfoSphere Guardium Database Activity Monitoring

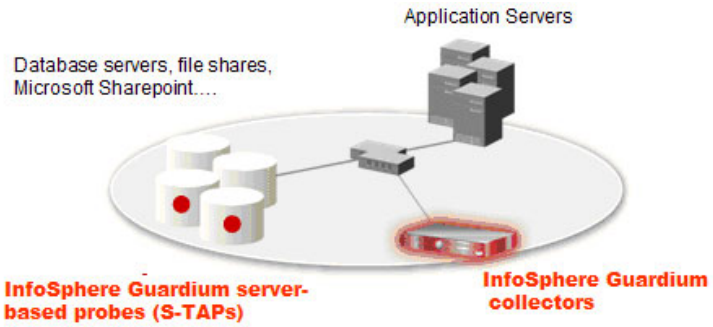

InfoSphere Guardium supports a wide variety of deployments to support very large and geographically distributed infrastructures.

As we have barely scratched the surface of what InfoSphere Guardium can do, refer to the Resources section for more information about the capabilities of InfoSphere Guardium, including data classification to help you discover sensitive data and vulnerability assessments that help you find soft spots in your infrastructure. Note that not all capabilities are available for all data sources.

What is new?

InfoSphere Guardium support for IBM i monitoring was previously available using three main methods:

- Import of audit journal entries (QSYS/QAUDJRN) and subsequent analysis and reporting
  While the audit journal support in IBM i provides a very good support of auditable events, the amount of detail in the audit entries is minimal compared to other Guardium database product support. For example, SQL statements and variable values are not audited in QAUDJRN. Also, as the support required an export and import, the support was not optimal as a real-time solution.

- Import of database monitor entries and subsequent analysis and reporting
  A database monitor (STRDBMON) can be used to capture SQL statements and write them to a database table. Subsequently, the table could be imported into the Guardium collector. While this method could capture SQL statements, variables, and more; the database monitor support was primarily designed for performance analysis. The result was that a significant

amount of data that was only interesting in a performance analysis context was captured resulting in the consumption of significant storage and processing resources. Also, as the support required an import, the support was not as optimal as a real-time solution. This method did not provide any support for native access to database objects.

- Network monitoring to capture SQL access
  While this support was able to capture SQL statements in real time that flowed over a network, any SQL statements than ran in programs, procedures, and functions on the IBM i server could not be monitored. This method did not provide any support for native access to the database objects.

The new method introduced in Version 9.0 of InfoSphere Guardium provides an integrated solution that overcomes the limitations of the previous methods.

- Any SQL access whether it is initiated on a client or the IBM i server can be captured and audited.
- Any native access that is captured in the audit journal can also be captured and sent to the InfoSphere Guardium collector.
- Both SQL access and native access are sent to the InfoSphere Guardium collector in real time.
- Much more detail than that available in the audit journal including SQL statements, variable values, client special registers, interface information, users, jobs, Transmission Control Protocol/Internet Protocol (TCP/IP) addresses, and ports is captured. However, unlike the traditional database monitoring, only the data that is interesting in a security context is captured and sent to the InfoSphere Guardium collector. This dramatically reduces the storage and resource consumption necessary.
- Filtering can be specified on the IBM i server to capture only that information which is required by auditors. For example, it is quite simple to set up auditing of any SQL or native access performed by privileged users.
- The data that is collected for InfoSphere Guardium is never written to disk on the IBM i server, providing a level of secure logging.

The new method is primarily for auditing database access. If you require auditing on a greater variety of non-database object access, the existing IBM i auditing support of exporting and importing the audit journal can still be used.

Introducing InfoSphere Guardium database activity monitoring for DB2 for i

As we mentioned in the previous section, InfoSphere Guardium Version 9.0 database activity monitoring has much more detailed auditing information for DB2 for i, including:

- Session start and end times
- Object names (tables or views, for example)
- Users

- SQLSTATEs
- Job and Job numbers
- SQL statements and variables
- Client special register values
- TCPIP address and port
- Interface information, such as ODBC, ToolboxJDBC, Native JDBC, .NET, and so on

This information can be used to create activity reports, help you meet auditing requirements, and generate alerts of unauthorized activity. Figure 2 shows you some database activity from one particular user on the system, including both a summary and more detailed version of the data. What is important to remember is that the InfoSphere Guardium reporting infrastructure is incredibly powerful with alerting capabilities and the ability to be automated into repeatable, regularly scheduled audit processes.

Figure 2. A sample SQL activity report

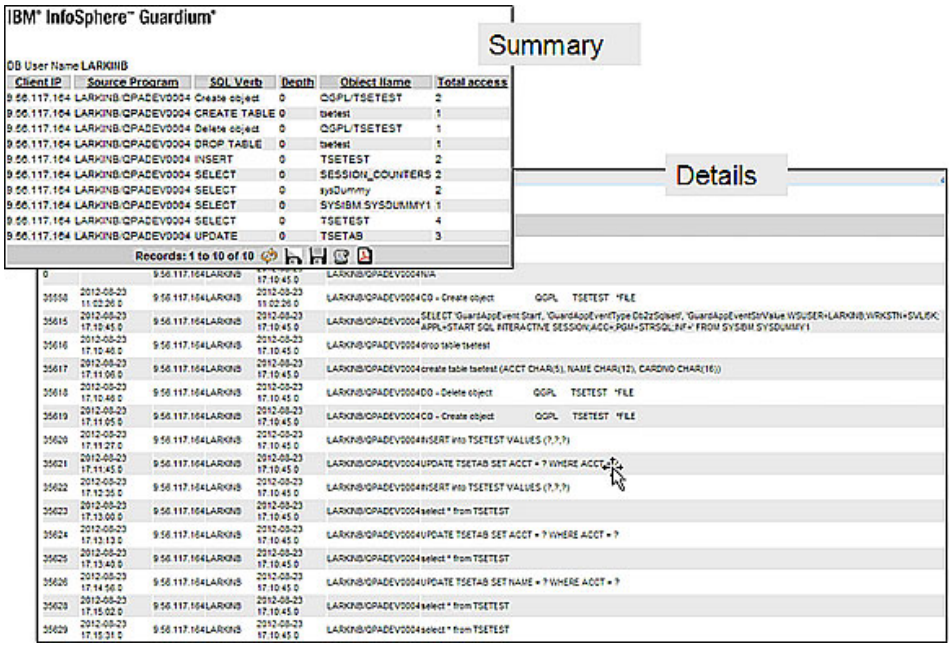

By using an InfoSphere Guardium S-TAP, you can monitor both SQL and native database application programming interface (API) traffic for DB2 for i. The configuration is similar to other database S-TAPs in that the processor usage on the database server is low, and the database events are sent to the InfoSphere Guardium collector for reports and alerting along with any other monitored data sources in your environment.

Two sources of data can be sent to InfoSphere Guardium (refer to Figure 3):

- SQL Performance Monitor (otherwise known as **database monitor**) data for SQL applications
- Audit entries from QSYS/QAUDJRN for applications using non-SQL interfaces

The DB2 for i S-TAP requires Portable Application Solutions Environment (PASE), which is automatically started and stopped as needed when a user who has the *JOBCTL authority (or

QIBM_DB_SQLADM function usage privilege) starts and stops the DB2 for i S-TAP from the InfoSphere Guardium user interface.

Figure 3. Two sources of information for database activity monitoring

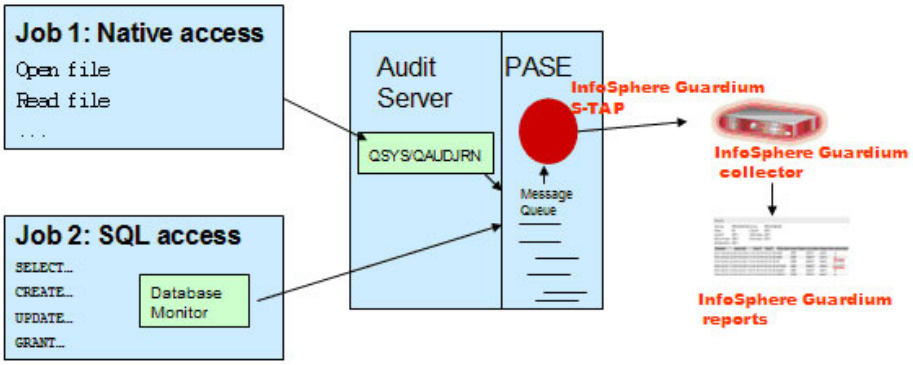

Requirements

The integration requires the following prerequisites:

- On IBM i: for the recommended minimum releases and PTFs, see IBM i Technology Updates wiki.
  - Refer to DB2 for IBM i 2012 Group PTF Schedule to subscribe to or review DB2 for IBM i PTF group schedule and availability.
  - License program 5722SS1-33 Portable App Solutions Environment (PASE) for i is a free of charge, optionally installable component of the operating system. Verify that PASE is installed on your IBM i server. If not, refer to the DB2 for i Information Center.
- IBM InfoSphere Guardium V9.0 or later appliance (configured as a collector) and the Standard Activity Monitoring for Databases software entitlement.
- For the DB2 server-side agent, you need to download the appropriate software tap (S-TAP) from IBM Fix Central. To ensure that you get the right S-TAP, filter on IBM i as shown here.

The InfoSphere Guardium appliance

The InfoSphere Guardium Data Security and Compliance solution is available as:

- A fully configured software solution delivered on physical appliances provided by IBM.

- Software images you can deploy on your own hardware either directly or as virtual appliances.

Before attempting to monitor DB2 for i, ensure that you check the IBM support site for additional patches that might be required.

This article does not provide information about the installation and configuration of the IBM InfoSphere Guardium appliance and assumes that you have at least one appliance connected to the IBM i server.

What gets collected

The information sent from the QAUDJRN and the information sent from the database monitor are not identical. The following table describes the information provided by each method.

Table 1. Database monitor vs Audit journal data that can be collected for auditing

Audit Data	SQL Monitor	Audit Journal
Job name	Yes	Yes
Job user	Yes	Yes
Job number	Yes	Yes
Start time	Yes	Yes
End time	Yes	Always the same as the Start time
SQLSTATE	Yes	08001 for invalid password (PW) and for general purpose audit records (GR) 42501 for authority failure (AF) 00000 everything else
SQLCODE	Yes	-30080 for invalid password (PW) and for general purpose audit records (GR) -551 for authority failure. (AF) 0 everything else
SQL statement	Yes – limited to 60K	No - basic journal entry description instead
SQL variables	Yes - limited to 1000 bytes	No
Interface	Yes	Always QAUDJRN
Client application name	Yes,	No
Client user ID	Yes	No
Client workstation	Yes	No
Client accounting	Yes	No
Client program	Yes	No
Current user	Yes	Yes
Thread ID	Yes	Yes
Program schema	Yes, if the statement is executed from a program or service program	Yes, if the statement is executed from a program or service program
Program name	Yes, if the statement is executed from a program or service program	Yes, if the statement is executed from a program or service program
Client IP Address	Yes	Yes

Local or server port number	Yes	Yes
RDB name	Yes	Yes
Number of rows	Yes, only for INSERT, DELETE, UPDATE, MERGE, OPEN*, VALUES INTO, CREATE TABLE AS, DECLARE GLOBAL TEMPORARY TABLE AS, and SET VARIABLE	No

*OPEN appears as SELECT in InfoSphere Guardium reports.

Note that the database monitor used for audit purposes with InfoSphere Guardium does not include events that are not security-related. For example, activities such as FREE LOCATOR or RELEASE are not audited. EXECUTE is not audited, but the SQL statement that ran is audited. PREPARE is not audited, but any authorization errors are audited.

Configure DB2 for i for QAUDJRN auditing (optional)

If auditing has already been configured on the IBM i or you are only interested in SQL auditing, you can skip this step.

On the DB2 for i server, create the QSYS/QAUDJRN journal and enable auditing if not already done. For more information on setting up security auditing, refer to the IBM i information center.

For example, on an IBM i command line:

```
CRTJRNRCV JRNRCV(QSYS/RCV1)
CRTJRN JRN(QSYS/QAUDJRN)
 JRNRCV(QSYS/RCV1) DLTRCV(*YES)
```

Next, specify the amount of auditing that you prefer to happen by setting the QAUDCTL, QAUDLVL, and QAUDLVL2 system values. For example:

```
CHGSYSVAL SYSVAL(QAUDCTL)
VALUE('*AUDLVL *OBJAUD')
CHGSYSVAL SYSVAL(QAUDLVL)
VALUE('*CREATE *DELETE *OBJMGT *SECURITY *SERVICE *SYSMGT *SAVRST');
```

If you only want to audit specific users, use the CHGUSRAUD command to change auditing for a user. For example, the following command enables a variety of auditing for user MJA, who might be one of your privileged users. For example:

```
CHGUSRAUD USRPRF(MJA)
 OBJAUD(*ALL) AUDLVL(*CREATE *DELETE *OBJMGT
*SECURITY *SERVICE *SYSMGT *SAVRST *AUTFAIL)
```

You can use the CHGOBJAUD command to change auditing for specific objects. For example, the following command enables auditing for all tables, views, indexes, and aliases (*FILE objects) in the PRODLIB schema:

```
CHGOBJAUD OBJ(PRODLIB/*ALL) OBJTYPE(*FILE) OBJAUD(*ALL)
```

Relevant QAUDJRN audit entries

The QAUDJRN audit journal can contain a wide variety of journal entries, but only a relevant subset is processed and sent to the InfoSphere Guardium collector.

QAUDJRN journal entries that are sent for a specific object contain the object library, object name, and object type. Only journal entries associated with the following IBM i object types will be processed (irrespective of whether they are associated with an SQL object or not):

- *FILE (a table, view, index, logical file, alias, or device file)
- *SQLUDT (an SQL user-defined type)
- *SQLPKG (an SQL package)
- *PGM (a procedure, function, or program)
- *SRVPGM (a procedure, function, global variable, or service program)
- *DTAARA (an SQL sequence)
- *USRPRF (a user profile object)

QAUDJRN journal entries can contain a wide variety of audit entry types. Only the following entry types are processed because they have been identified to be of most use to auditors:

- ZR Read object
- ZC Change object
- AD Auditing change
- AF Authority failure
- CA Authority change
- CD Command string (Note: CD is not included in the default settings of filter_audit_entry_types)
- CO Create object
- CP User Profile changes
- DO Delete object
- GR General purpose audit record
- OM Object moved or renamed
- PG Primary group change
- PW Invalid password or user ID
- OW Change owner
- OR Object restored
- RA Restore authority change
- RO Restore owner change
- RZ Restore primary group change
- SV System value change

QAUDJRN journal entries do not contain the SQL statement. For journal entries that identify an object, the following information will be concatenated and be returned in place of the SQL statement:

- 30-byte-description of the operation

- 10-byte-system-schema_name
- 10-byte-system-object-name
- 8-byte-object-type

For example:

ZC - Change object MJATST T1 *FILE

For more information on the journal audit entry types, refer to Audit Journal (QAUDJRN) entry types in the IBM i Information Center.

Install the DB2 for i S-TAP

1. 1. In the PASE shell environment on the IBM i server, create a temporary directory to put the S-TAP installation script (such as /tmp). You can use a 5250 emulator software to connect to the IBM i system remotely and enter the PASE shell by entering **call qp2term.**
2. Use FTP to move the following S-TAP installation shell script to that temporary directory:`uard-itap-9.0.0_r`*`nnnnn`*`-aix-5.3-aix-powerpc.sh`
3. In the same directory, run the following command:

`guard-itap-9.0.0_r`*`nnnnn`*`-aix-5.3-aix-powerpc.sh` *`guardium_host_IP`*
where *guardium_host_IP* is the IP address of the InfoSphere Guardium collector. The installation program will install under `/usr/local/guardium`.

After the installation is complete, InfoSphere Guardium attempts to start the processes that enable activity monitoring and to locate the InfoSphere Guardium collector using the IP address specified at the installation time.

To validate the successful installation and start of the audit process, log in to the InfoSphere Guardium web console as an administrator and navigate to the System View tab and check the status of the S-TAP, which should show green as shown in Figure 4.

Figure 4. System monitor shows that configuration is successful

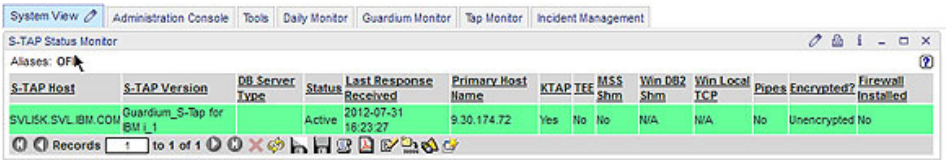

Troubleshooting

If the S-TAP monitor is not showing green in the status monitor, make sure you correctly ran step 3 in the installation instructions above. If you need to correct something, such as an incorrectly specified IP address, you can invoke the start_istap_monitor Guardium CLI command or create the Status report on the Guardium console to invoke that API from the

GUI. To create the Status report, see the instructions in the section entitled Recommended: Set up the DB2 for i Status report on the collector

The next step, configuring the S-TAP on the InfoSphere Guardium collector is strongly recommended because it enables you to view S-TAP status on the IBM i server, update the configuration as needed, and specify filtering values.

Note: The IBM InfoSphere Guardium Installation Manager (GIM) is not supported for the DB2 for i S-TAP.

Recommended: Set up DB2 for i S-TAP configuration capability on the collector

As we mentioned in the previous section, it is strongly recommended to go ahead and set up the configuration capability on the InfoSphere Guardium collector. You can do this by creating a configuration report, which enables you to invoke APIs that run on the IBM i server, which can start and stop processes and update the configuration file, QSYS2.SYSAUDIT.

You must have the *JOBCTL authority or the QIBM_DB_SQLADM function usage privilege on IBM i to configure the environment.

You must also have completed the installation steps above and the monitor process must have been started on the IBM i server (as validated by the STAP monitor by displaying green).

In this section, you'll be doing the following steps:

1. Defining DB2 for i as a recognized data source to InfoSphere Guardium and testing the connection. For this, you will need to know the database name, port, and credentials.
2. Populating the InfoSphere Guardium collector with information from the configuration file on IBM i that was created when you installed the DB2 for i S-TAP, using the Custom Table Builder process.
3. Creating a DB2 for i configuration report. It is from this report interface that you can invoke the APIs that start and stop the monitoring process, get status information, and update configuration parameters, including filtering values.

Define the DB2 for i data source to InfoSphere Guardium

In this step, you need to define your DB2 for i as a data source that the InfoSphere Guardium collector can recognize. You can do this by creating a custom domain and defining DB2 for i as the data source for that domain using the Datasource Builder.

To create a data source for the DB2 for i with the InfoSphere Guardium Datasource Builder:

1. Click **Tools>Datasource Definitions** then select **Custom Domain** from the **Application Selection** box. Click **Next.**

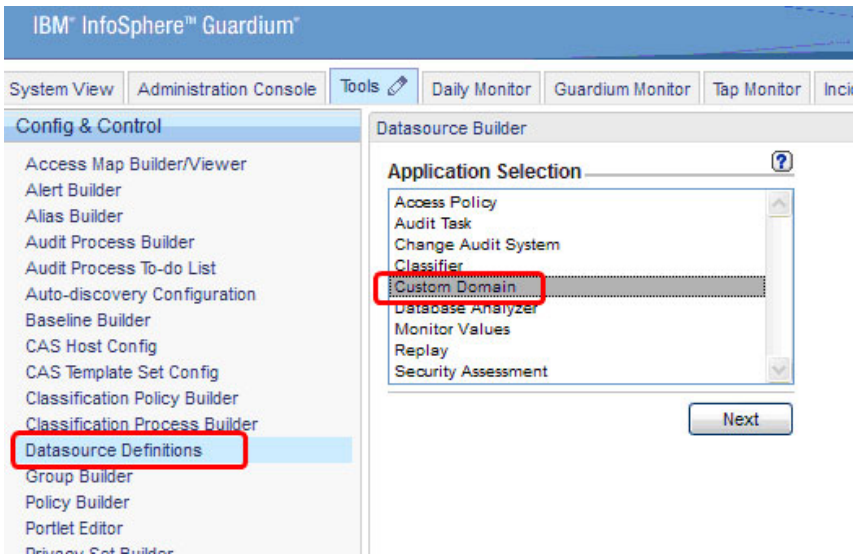

2. In the Datasource Finder, click **New**, which brings up the Datasource Builder. Select **DB2 for i** as the database type and then add the appropriate information for the port, host, service name (which is the database name), and credentials. Also, enter a meaningful name for this definition.
3. Click **Apply** and then click **Test Connection** to ensure all is configured correctly.

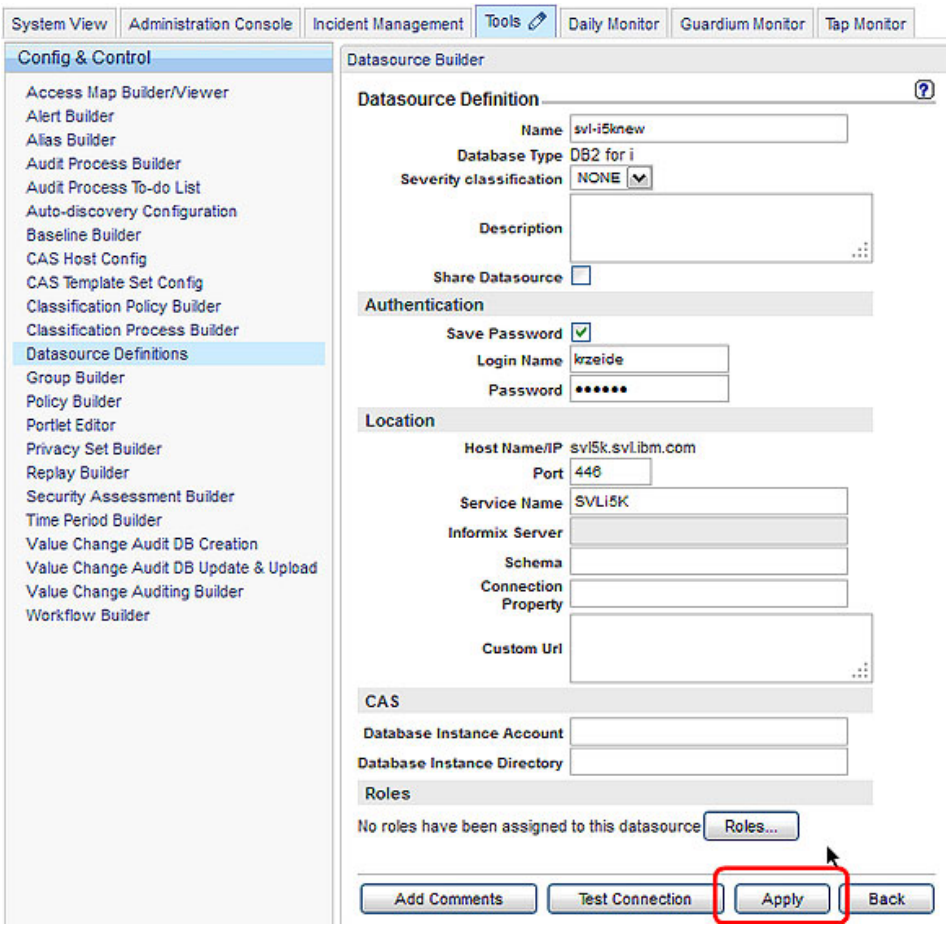

Upload the DB2 for i configuration settings to the InfoSphere Guardium collector using a custom table builder

In this step, you use the InfoSphere Guardium interface to import the configuration information from the IBM i system. You do this by performing the following steps.

1. Invoke the report building interface.
2. Create a custom table on the local InfoSphere Guardium to hold the configuration data from the DB2 for i data source.
3. Import the configuration data from DB2 for i to that custom table.

Here are the detailed steps:

1. Click **Tools**>**Report Building**. (Hint: You might need to scroll down to find the Report Building option on the left.)
2. Click **Custom Table Builder**, and select **DB2 for i S-TAP Configuration** and then click **Upload Data**.

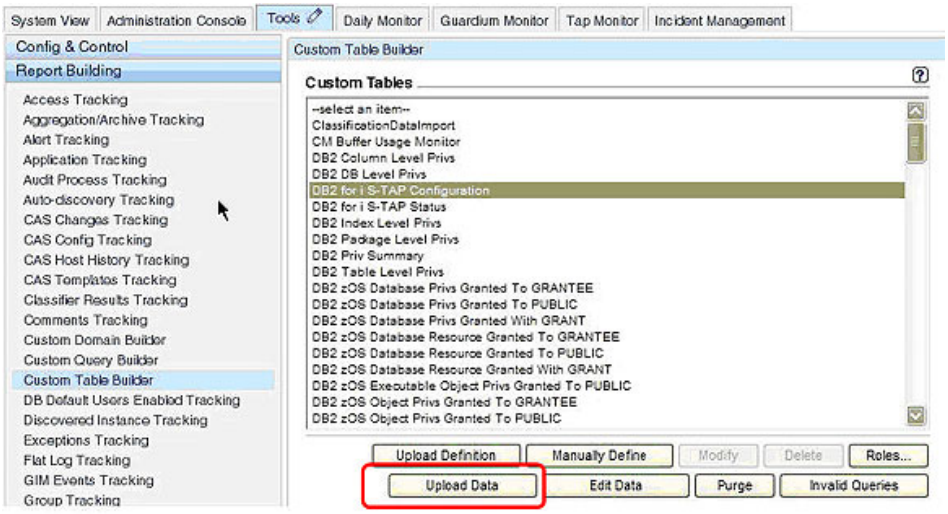

3. Click **Add Datasource** .

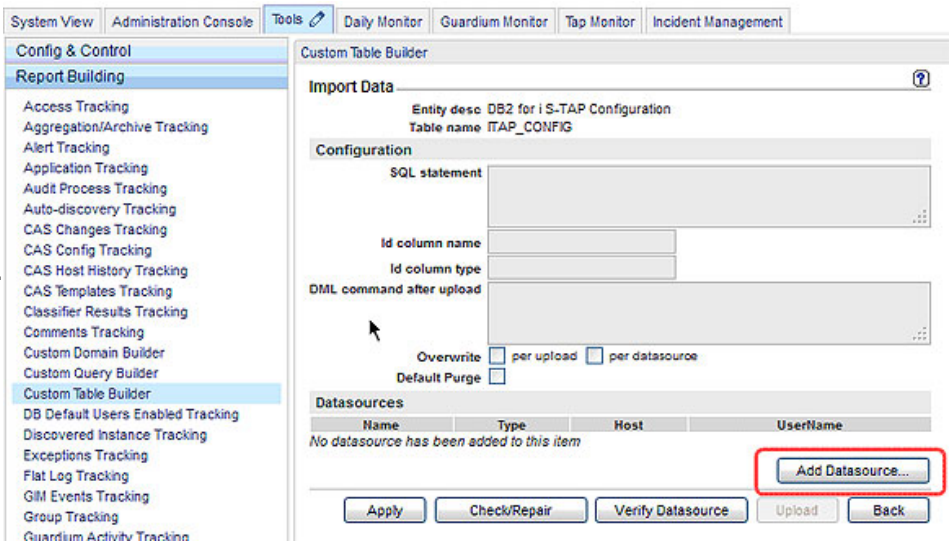

4. On the Datasource Finder, locate your DB2 for i data source on the list and then click **Add.**

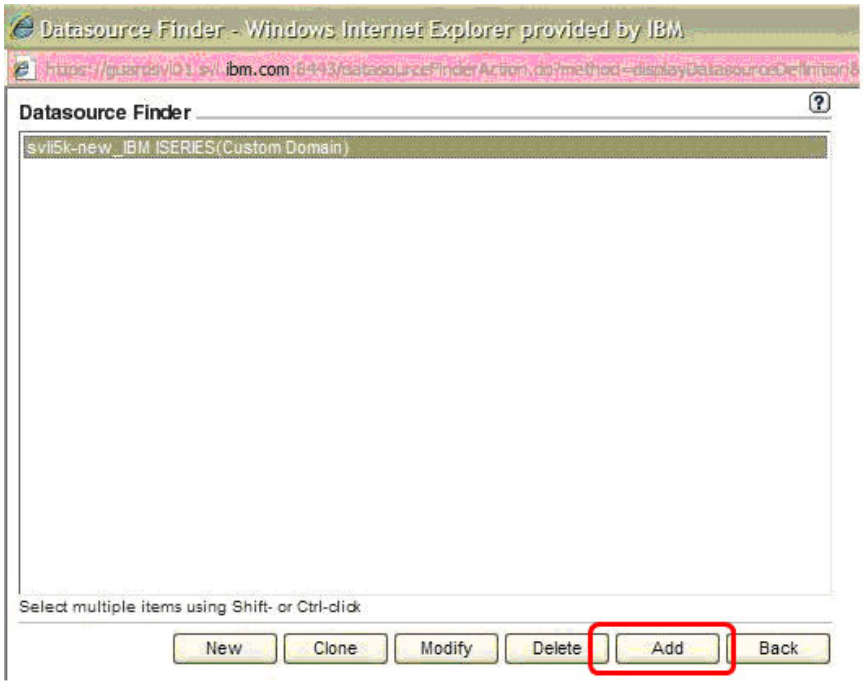

5. On the Import Data screen, ensure the DB2 for i data source appears. Click **Apply** and then click **Run Once Now**. You should see a message that the operation ended successfully.

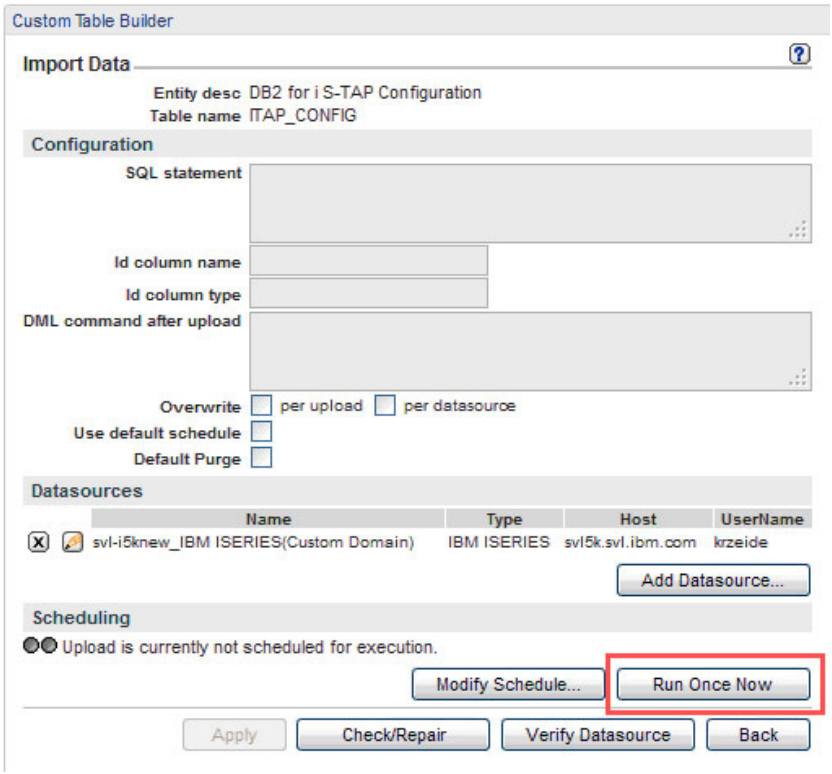

Creating the configuration report to invoke InfoSphere Guardium APIs

This section explains the following major tasks:

- An optional task to customize the InfoSphere Guardium interface to create a space (that is, a pane) where you can put the new configuration report for DB2 for i. We will use the name *My New Reports* for this pane. (If you are logged in as a user rather than an administrator, the My New Reports pane will already exist.)
- A task to search for and add the DB2 for i S-TAP configuration report to the pane. After that configuration report exists, you can invoke the APIs to change the configurations for DB2 for i.

For more information about creating reports, refer to the InfoSphere Guardium Information Center topic on How to build a report and customize parameters.

To create a report pane (only required if one does not exist):

1. To create the My New Reports pane, from the upper right corner of the Guardium UI, click **Customize** then click the **Add Pane** button, as shown below.

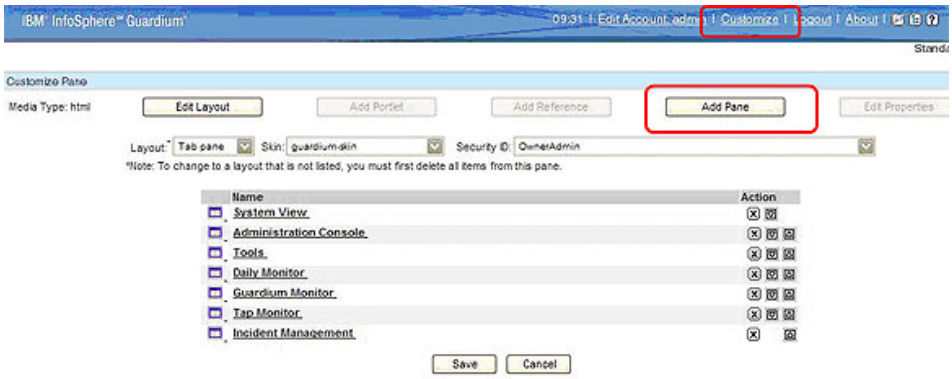

2. Give the pane a new name, My New Reports (spelled *exactly*). Click **Apply** and then **Save**.

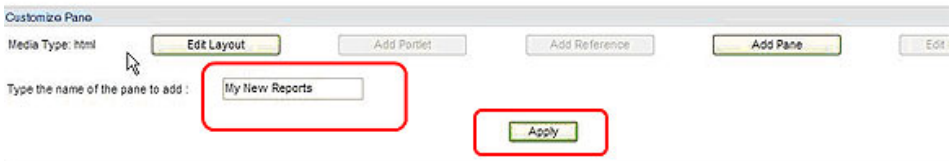

3. My New Reports will appear in the Customize Pane. Click on the icon to the left of that item. On the Layout pulldown, choose **Menu Pane**, and then **Save**. Your new pane will appear as a tab.

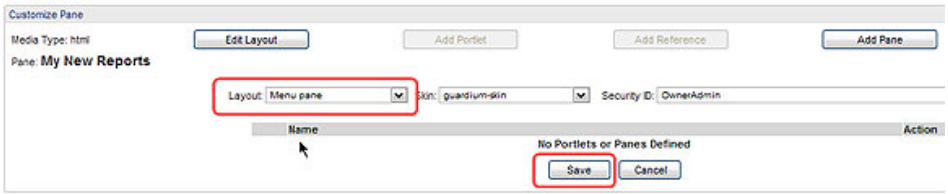

Create the configuration report and add to the report pane:

1. Now you are ready to create the configuration report to add to the new report pane. To do this, click on **Report Builder** in the left navigation pane. In the right pane, from the Query list, select DB2 for i S-TAP configuration, and then click **Search**.

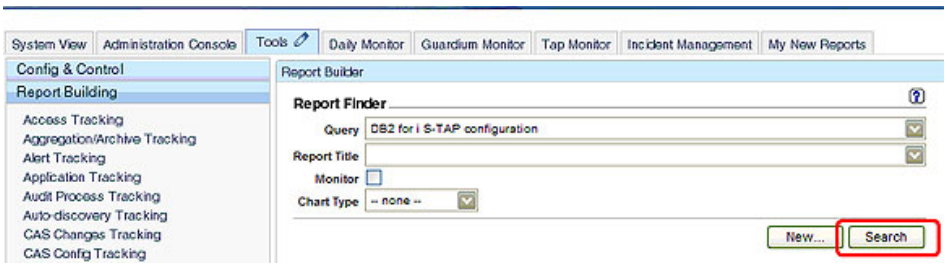

2. Select DB2 for i S-TAP configuration and then click the **Add to My New Reports** button, as shown below (or add the report to an existing pane by clicking **Add to Pane…**)

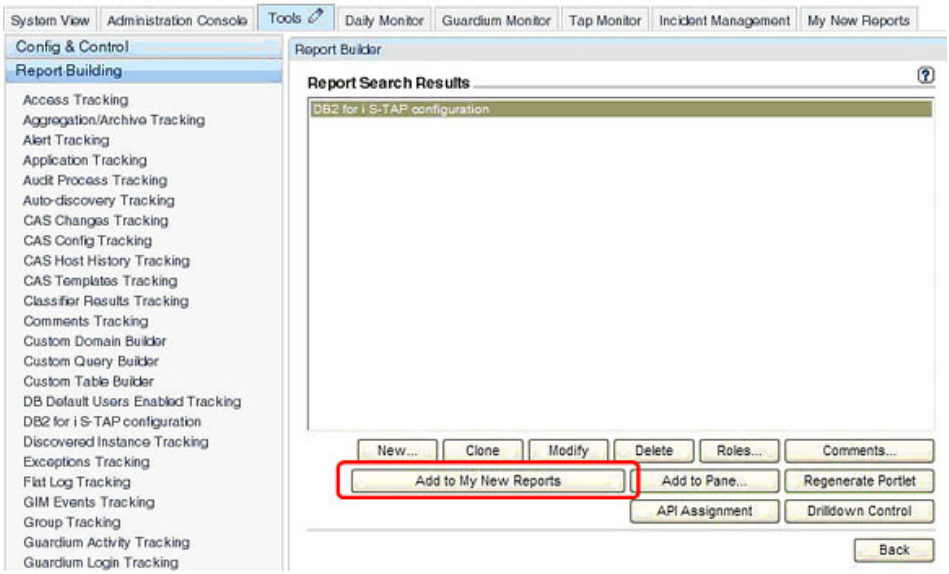

3. Click on the **My New Reports** tab which now will be displaying the IBM i report row. Double-click a row in the report and then click **Invoke**.

4. Now you can see the InfoSphere Guardium APIs. Click **update_istap_config**.

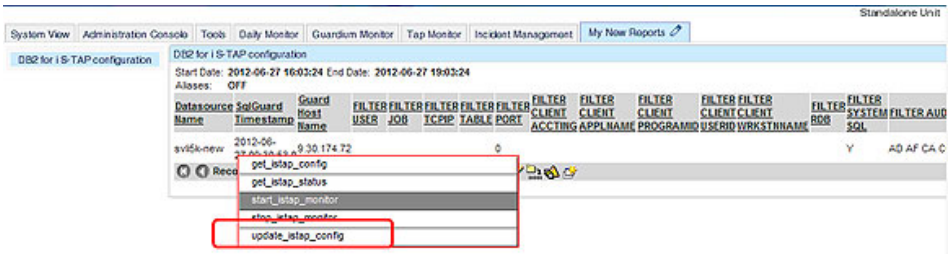

The section Overview of DB2 for i S-TAP APIs includes more information about the configuration API parameters available from that report and a brief overview of the other APIs.

Recommended: Set up the DB2 for i status report on the collector

Another very useful report you need to have handy on your console is the status report. The steps to add this report are almost exactly the same as those described in the previous section, Recommended: Set up DB2 for i S-TAP configuration capability on the Collector. In summary:

1. As an administrator, navigate to **Tools**>**Report Building**>**Customer Table Builder**.
2. Select **DB2 for i Status**, and then click **Upload Data**.
3. Add the Datasource, then on the Import Data screen, click **Apply** and then **Run Once Now**.
4. Navigate to the Report Builder, search on the report title of DB2 for i Status and then **Add to My New Reports**.

Here is an example of the Status report:

Figure 5.

You can double click on the report to start and stop the monitor on the server as well as refresh the status.

Figure 6.

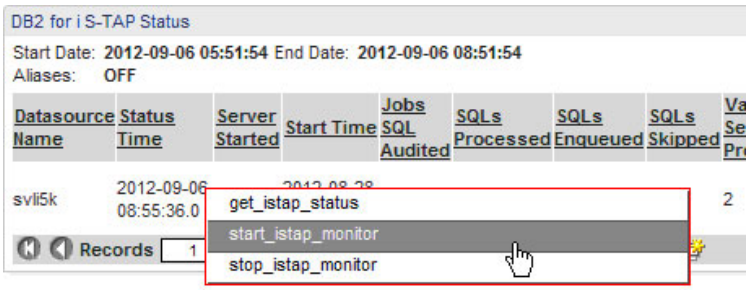

Overview of DB2 for i S-TAP APIs

There are several ways to invoke InfoSphere Guardium APIs. In this article, we just show how to invoke them from the DB2 for i S-TAP Configuration and Status Reports. As you learned in the previous section, you can double click in the report to invoke APIs Selecting an API brings up the allowable parameters for that particular API. In this case, as shown in Figure 7, you can use the update_istap_config API to send information, such as the IP address of the InfoSphere Guardium collector, and to start or restart the auditing processes on the i system. (If the value is "unchange" then that parameter will not be updated.) You'll learn more about the value of the filtering options in the section Filtering audit data on the IBM i server.

Figure 7. Options to update the IBM i S-TAP configuration using

When the value for start_monitor is set to 1 (default), the auditing process will start (or restart) on the i server after the configuration table is updated. When the auditing process is started, stored procedures on DB2 for i are invoked that will:

- Create the message queue that will be used to send entries to the InfoSphere Guardium collector and starts a global database monitor using a view with an INSTEAD OF trigger (which sends the entries to the message queue)
- Start PASE and S-TAP.
- Receive journal entries from QAUDJRN and add them to the message queue.

Invocable S-TAP APIs for IBM i

To provide for scripting and automation, the S-TAP APIs can be invoked from the command-line interface (CLI) in InfoSphere Guardium.

Here is an overview of the APIs for IBM i monitoring:

- **start_istap_monitor** will start the audit process on IBM i.
- **stop_istap_monitor** will stop the audit processes on IBM i.
- **get_istap_status** can be used to check whether the audit server is running and it includes other information (such as the number of messages on the queue, the size of the message queue, and so on) that can be useful for troubleshooting and performance tuning.
- **get_istap_config** can be used to view configuration parameters, including the current filtering options.
- **update_istap_config** can be used to update the configuration settings on IBM i.

When the S-TAP connects to the collector, a row similar to the one shown in Figure 4 appears in the System View tab.

View monitoring reports

After the system is configured and auditing is underway, you can start taking advantage of the real power of InfoSphere Guardium to run reports, set alerts, and much more. InfoSphere Guardium has a rich reporting interface, which is beyond the scope of this article.

When creating reports, depending on whether you have logged in as an administrator or as a user, the navigation paths might be different. Therefore, make sure to read the How to build a report and customize parameters and How to create custom reports from stored data topics in the InfoSphere Guardium Information Center, or by clicking on the question mark icon an the upper-right corner of the user interface to access the help book.

To use reports that show the database activity details, you need to be logged in as a user and that is what this section of the article assumes. Note that the InfoSphere Guardium user interface is highly customizable, so the screen captures and the navigation paths shown here might not work exactly as shown at your site.

This procedure assumes that the S-TAP configuration is successful and that the database activity is occurring on DB2 for i.

1. Click the **View** tab. (Optional: Rename this tab to**Standard Reports** by clicking the pencil icon on the tab and then clicking **Edit Properties.**
2. You should see some reports as shown here.

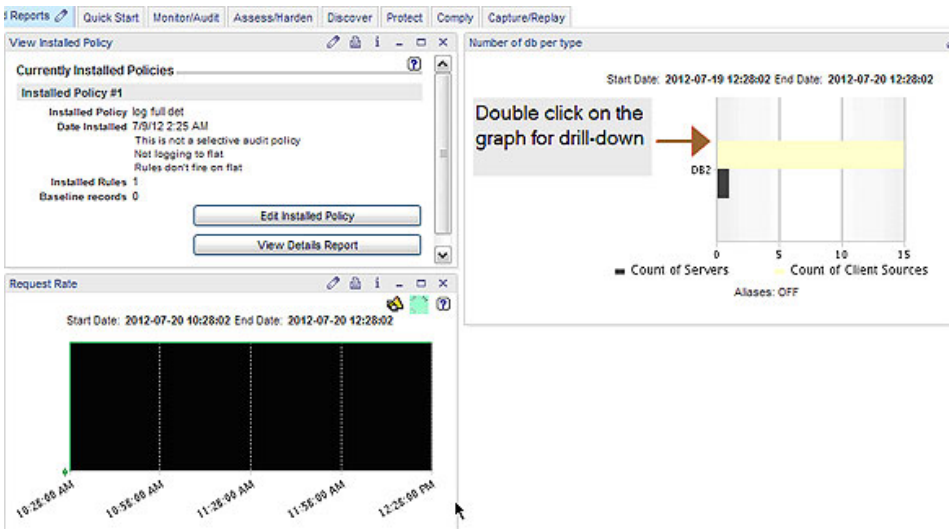

3. Double click on the graph, which brings up a tabular view that you can use to start drilldown by double clicking on subsequent report tables.

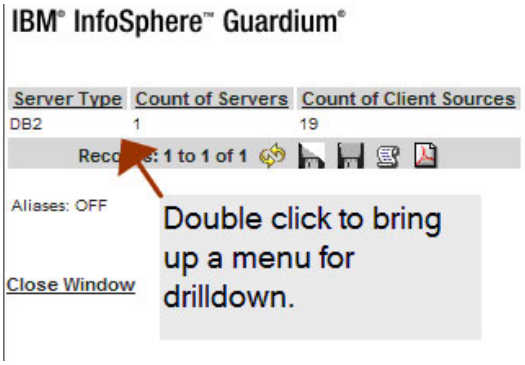

4. You can drill down through the data such as
    a. Sessions by server IP, then double click on a row
    b. Sessions by user, then double click on a row
5. Continue exploring the reports.

Reporting on string variables

The monitoring reports include the SQL statement text, and you might notice the appearance of question marks within the text. From an SQL perspective, these are known as *parameter markers* and they act as a placeholder within the SQL statement for specific values. For those cases where it is critical to understand what change was made, you can add a field to an SQL activity report called **Bind Variables Values**, as shown in the following figure on the right side of the report.

Important: If you see hexadecimal representations of the string values, that means the users are running their SQL statements with the default CCSID of 65535. Users would need to change their profiles to use a human-readable CCSID, such as the following code (for US English):

```
CHGUSRPRF JOED CCSID(37)
```

For more information about CCSIDs in IBM i, see the IBM i Information Center.

Performance recommendations

The InfoSphere Guardium integration is designed to minimize the processing required for capturing the database activity and sending it to the InfoSphere Guardium collector. However, the amount of data that is sent to the collector plays a big role in how much processing is required on the production system. Thus, it is prudent to capture only the required database activities to satisfy your audit policies. For example, perhaps only certain users, certain interfaces, or certain objects need to be audited.

For optimal performance and reduced processing, you need to consider two levels of filtering:

- The filtering performed by DB2 for i can reduce the overhead on the production system by reducing the amount of information that flows from the audit server to the S-TAP and from the IBM i to the InfoSphere Guardium collector. To do this filtering, you can specify the filtering values on the configuration pages in the InfoSphere Guardium UI.
- Filtering on the InfoSphere Guardium collector side to restrict what information sent to the collector is actually required to be analyzed and stored in the repository. This can reduce storage on the collector and the processing time for auditors. You can do this by using the policy configuration on the InfoSphere Guardium user interface.

This section also briefly describes setting the priority of the audit server job.

Filtering audit data on the IBM i server

Table 2 shows a complete list of the filtering fields that are available when you configure the DB2 for i S-TAP, which were shown in Figure 5. To change the filtering values, invoke the update configuration API (as described in the previous section) and change the values in the appropriate fields. Many of these fields map to the filtering values that are documented in the STRDBMON command for IBM i. For more information about the database monitor and monitor entries, refer to the IBM i information center.

Important: The filter criteria on DB2 for i can be combined; however, the criteria are effectively ANDed together.

Table 2. Filtering Options for DB2 for i S-TAP configuration

Filtering option	Description
FILTER_USER	The specified user or group user profile filter, if any. Only one user name or generic user name can be specified.
FILTER_JOB	The specified job filter, if any. Only one job name or generic job name can be specified.
FILTER_TCPIP	The specified TCP/IP filter, if any. Only one TCP/IP address can be specified.
FILTER_TABLE	The specified table filter, if any. Up to ten file names or generic file names can be specified. The specified library name must be the system schema name (10 character name). The file name can be either the system table name or table name (long or short name).
FILTER_PORT	The specified port filter, if any. Only one port filter can be specified. Filtering by port is only supported in release 7.1 and later.
FILTER_CLIENT_ACCTNG	The specified client accounting filter, if any. Only one client accounting filter can be specified. Filtering by client accounting is only supported in release 7.1 and later
FILTER_CLIENT_APPLNAME	The specified client application filter, if any. Only one client application filter can be specified. Filtering by client application is only supported in release 7.1 and later.
FILTER_CLIENT_PROGRAMID	The specified client program filter, if any. Only one client program filter can be specified. Filtering by client program is only supported in release 7.1 and later.
FILTER_CLIENT_USERID	The specified client user filter, if any.

	Only one client user filter can be specified. Filtering by client user is only supported in release 7.1 and later.
FILTER_CLIENT_WRKSTNNAME	The specified client workstation filter, if any. Only one client workstation filter can be specified. Filtering by client workstation is only supported in release 7.1 and later.
FILTER_RDB	The specified relational database filter, if any. Up to 10 relational database names can be specified.
FILTER_SYSTEM_SQL	The specified system SQL statement filter. Specifies whether system SQL statements should be audited (Y or N) . The default is Y.
FILTER_AUDIT_ENTRY_TYPES	The specified QAUDJRN audit entry filter, if any. Specifies which audit journal entry types should be processed. The default is 'AD AF CA CO DO GD OM OR OW PG PW RA RO RZ SV ZC ZR'

The following examples will describe some common best practices for reducing the overhead on the production system by using filtering fields.

Filtering by user (FILTER_USER)

If your business only requires auditing of certain users (perhaps those users who have a high level of authority), the most efficient way to capture only entries for those users is to create a new group profile (CRTUSRPRF) and then assign each of the users to be part of the new group (CHGUSRPRF). You can then specify a filter based on the new group profile. For example, from IBM i, create a new group user profile called GROUPGD and assign it as a primary or supplemental group to users ADMIN1, ADMIN2, and ADMIN3:

```
CRTUSRPRF USRPRF(GROUPGD) PASSWORD(*NONE) STATUS(*DISABLED) GID(*GEN)
CHGUSRPRF USRPRF(admin1) SUPGRPPRF(groupgd)
CHGUSRPRF USRPRF(admin2) SUPGRPPRF(groupgd)
CHGUSRPRF USRPRF(admin3) SUPGRPPRF(groupgd)
```

Then, you can specify a FILTER_USER value of GROUPGD which will audit users who are members of that group. Or you can use a wild card and set FILTER_USER to GROUP*, which will filter users of any group or username that begins with GROUP.

Filtering by schema or table (FILTER_TABLE)

If your business requires auditing only for objects in a set of schemas (libraries), filtering can be performed on up to 10 schema names or generic schema names. If you filter by schema or table, only data manipulation (DML) statements will be captured. All other types of statements will not be captured. For example, assume you want to capture data manipulation SQL statements that are associated only with schemas PROD1, NEWPROD1, and PROD2, you might specify a FILTER_TABLE value of:

```
PROD1/*ALL NEWPROD1/*ALL PROD2/*ALL
```

Or, if no other schemas start with PROD, you can use a generic name:

```
PROD*/*ALL NEWPROD1/*ALL
```

This type of filtering can also be performed at the individual table level. For example, if you only want to capture entries for tables that start with the letters PERSONAL in library NEWPROD, you can specify

```
NEWPROD1/PERSONAL*
```

Filtering by JOB (FILTER_JOB)

If your business only requires auditing of traffic related to JDBC, ODBC, and .NET requests that come in from the IBM access drivers, you can filter by job since these requests are processed by QZDASOINIT jobs. For example, you can specify

```
*ALL/*ALL/QZDASOINIT
```

If you use the IBM DB2# Connect™ drivers you can specify:

```
*ALL/*ALL/QRWTSRVR
```

Filtering by relational database (FILTER_RDB)

Multiple relational databases can exist on an IBM i through the use of independent auxiliary storage pools (IASPs). Each IASP and *SYSBAS represents a separate relational database and each has a unique relational database name. If your business only requires auditing of two IASPs with relational database names of RDB1 and RDB2 (omitting activity against *SYSBAS and any other IASPs), you can specify a FILTER_RDB value of:

```
RDB1 RDB2
```

Filtering QAUDJRN entry types (FILTER_AUDIT_ENTRY_TYPES)

The audit journal contains a wide variety of audit entry types. By default, only a subset of the AUDJRN entries that you configured to capture in QAUDJRN will be sent to the InfoSphere Guardium collector, as described in Relevant QAUDJRN Audit entries. So, for example, if you have no need to audit object creates or object restores, you can eliminate the CO and OR entry types and specify a FILTER_AUDIT_ENTRY_TYPES value of:

```
'AD AF CA DO GD OM OW PG PW RA RO RZ ZC ZR'
```

Filtering SYSTEM SQL statements (FILTER_SYSTEM_SQL)

DB2 for i executes a variety of SQL statements to perform certain functions. These statements are flagged in the monitor as System SQL statements. However, a user-created procedure, function, program, or service program can also indicate that its SQL statements should also be flagged as a System statement through the use of the SET OPTION statement. Therefore, the best practice is to leave this value set to 'Y'.

Filtering activity through database interactive IBM i command line interfaces (FILTER_CLIENT_PROGRAMID)

DB2 for i includes the capability for users to run interactive SQL statements with commands such as Start SQL Interactive Session (STRSQL). Compliance in some organizations requires that access to those interfaces be limited to provide a second layer of defense against access to the database objects. If you want to specifically monitor access through that interface, you can set the FILTER_CLIENT_PROGRAMID to 'STRSQL' to collect detail only on SQL activity coming through the Start SQL (STRSQL) command. See the articles in Resources for more information on auditing these interactive commands.

Any of the following commands can be included as a filtering criterion:

- Run SQL Statements (RUNSQLSTM)
- Start Query (STRQRY)
- Work With Queries (WRKQRY)
- Run Query (RUNQRY)
- Run SQL (RUNSQL)

Filtering using the client information fields

The client information fields consist of 5 different fields that can be set for a particular database connection, such as by using the SQLESETI API, as described as described in the IBM i Information Center. These can be useful in an auditing context to:

- Identify different programs with greater granularity to filter out (such a trusted application) or monitor more closely.
- Help InfoSphere Guardium track down individual users if the application uses a "pooled" database connection where there is only a single "DB user". (This requires that the application set the CLIENT_USERID special register.)

The fields that make up client information are:

- FILTER_CLIENT_ACCT
- FILTER_CLIENT_APPLNAME
- FILTER_CLIENT_PROGRAMID
- FILTER_CLIENT_USERID
- FILTER_CLIENT_WRKSTNNAME

Reducing traffic sent to the InfoSphere Guardium collector

We have just described how to use the S-TAP configuration options to filter the events that are processed by S-TAP. You can use InfoSphere Guardium security policies to include IGNORE rules that will control the quantity of information that gets sent from S-TAP to the InfoSphere Guardium collector, which can reduce the amount of information to be parsed. For example, you can have rules that ignore sessions from scheduled jobs or other trusted applications.

There are also options to determine the amount of information that must be retained in the InfoSphere Guardium database after the criteria of the policy rules have been met. This is a complex topic which is beyond the scope of this article. You can read more about the policies in the InfoSphere Guardium help book (that can be accessed by clicking the question mark icon in the upper-right corner of the InfoSphere Guardium web interface).

Tuning the performance of the audit server

The audit server is responsible for receiving entries from QAUDJRN and placing them on the Guardium message queue. The audit server also contains the code that runs in PASE, removes entries from the Guardium queue, and sends them to the InfoSphere Guardium collector.

Typically, the audit server runs in a batch job that is started by InfoSphere Guardium using the SBMJOB command. On the first start of the audit server, the SBMJOB command is run with USER(*CURRENT). Subsequent starts of the audit server use the last user profile that started the audit server (USER(user-name)).

If you want to start the audit server under a different user profile name (GDUSER), update the configuration file as follows:

```
UPDATE qsys2.sysaudit SET start_user = 'GDUSER'
```

Then restart the audit server using the start_istap_monitor API from the InfoSphere Guardium UI.

As starting the audit server in this way uses the SBMJOB command, the normal rules apply to the job description that should be used to run the job (from the user profile). The job description determines how the audit server is run. For example, it specifies the initial priority given to the job.

A best practice is to create a unique job description that controls the attributes of the job, a unique job queue, and a unique user profile for running the audit server. For example:

```
CRTUSRPRF GDUSER PWDEXP(*YES) STATUS(*ENABLED) SPCAUT(*ALLOBJ *JOBCTL)
   TEXT('Guardium user profile')
CRTJOBQ GDJOBQ TEXT('Guardium job queue')
CRTJOBD GDAUDIT JOBQ(GDJOBQ) JOBPTY(2) USER(GDUSER) JOBMSGQFL(*WRAP)
   TEXT('Guardium job description')
CHGUSRPRF USRPRF(GDUSER) JOBD(QGPL/GDAUDIT)
ADDJOBQE SBSD(QBATCH) JOBQ(QGPL/GDJOBQ) MAXACT(2) SEQNBR(40)
```

Note that MAXACT must be at least 2.

Another best practice is to run the audit server at priority 1 with a larger time slice. This allows the audit server to efficiently process SQL statements, especially when the number of SQL statements sent to the collector is large. Since most users will run the audit server in the QBATCH subsystem, a new class and routing entry can be used to specify the priority and time slice. For example:

```
CRTCLS CLS(GDCLS) RUNPTY(1) TIMESLICE(10000) TEXT('Guardium class')
ADDRTGE SBSD(QBATCH) SEQNBR(800) CMPVAL(GUARDIUM) PGM(QSYS/QCMD) CLS(GDCLS)
```

For more information on SBMJOB, classes and routing entries, and controlling the job description that is used for the batch job, refer to the IBM i information center

Summary

We hope this article has given you a good starting point for implementing InfoSphere Guardium in your organization and to audit not just DB2 for i, as described in this article, but many other databases and file systems that you probably have. The auditing architecture is scalable and can be rolled out across large organizations and across geographies. In an age when every day seems to bring new news of data breaches, including those committed by privileged users, it is critical for organizations everywhere to create additional lines of defense that do not rely solely upon the native database security.

Acknowledgements

The authors would like to thank Tania Butovsky, Joe DiPietro, Scott Forstie, Rui Yu, and Larry Burroughs for their review and technical assistance.

Resources

- Recommended IBM i service level detail and pointers to useful education resources
- Guardium Database Activity Monitor and DB2 for i Serviceability Guide
- YouTube video on InfoSphere Guardium monitoring for DB2 for i
- IBM InfoSphere Guardium product page
- IBM InfoSphere Guardium Information Center
- IBM i Information Center
- IBM InfoSphere Guardium forum
- IBM InfoSphere Guardium community on developerWorks

- SOX Auditing of STRSQL and RUNSQLSTM Commands, IBM Systems Magazine, April 2009
- STRDBMON pre-filtering of QUERY-400 command use, IBM i Technology Update, June 2001
- IBM i Technology Update wiki page for Guardium integration