IBM® Security Information Queue (ISIQ)

Version 1.0.5

Performance Tuning Guide

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

3rd Edition notice
Note: This edition applies to Version 1.0.5 of IBM Security Information Queue and to all subsequent releases and modifications until otherwise indicated in new editions.

Table of Contents

About This Publication

This edition includes sections associated with troubleshooting the IBM® Security Information Queue (ISIQ), IBM® Security Identity Manager (ISIM), and IBM® Security Identity Governance and Intelligence (IGI) tiers, best practices for configuring, tuning, and managing ISIQ, ISIM, and IGI, and regular recommended maintenance tasks for ensuring optimum ISIQ performance.

Intended Audience

This document is designed for the system administrators, security administrators, and operations teams which are responsible for an organization that uses IBM® Security Information Queue (ISIQ). In addition, readers are expected to understand Docker system administration and information systems security concepts. The readers must understand administration concepts for the following types of products:

- Docker
- IBM® Security Identity Manager (ISIM)
- IBM® Security Identity Governance and Intelligence (IGI)
- Grafana
- Kibana
- Kafka
- InfluxDB

ISIQ Library

You can obtain the product documentation from the ISIQ information center. The information center is available here:

Related Publications

- ISIQ Starter kit/installation instructions

  https://www-01.ibm.com/support/docview.wss?uid=ibm10787861
- ISIQ Troubleshooting Guide

  https://www-01.ibm.com/support/docview.wss?uid=ibm10717877
- ISIQ Deployment Guide

https://www-01.ibm.com/support/docview.wss?uid=ibm10717671

- ISIQ Users Guide

  https://www-01.ibm.com/support/docview.wss?uid=ibm10715715

- ISIQ FAQ

  https://www-01.ibm.com/support/docview.wss?uid=ibm10733903

- IBM Software Support Home Page

  http://www.ibm.com/support/entry/portal/overview/software/software_support_(general)

- IBM Publication Center

  http://www-05.ibm.com/e-business/linkweb/publications/servlet/pbi.wss.

- IBM Developer Works

  http://www.ibm.com/developerworks/

ISIQ Performance Tuning Introduction

ISIQ is a cross-product integrator that utilizes Kafka technology and a publish-subscribe model to integrate data between IBM® Security products such as IBM Security Identity Manager (ISIM) and IBM Security Identity Governance and Intelligence (IGI).

ISIQ Features

- A dashboard for viewing system status such as system notification, logs, etc.
- Analysis and diagnostics tools such as memory statistics, CPU utilization, and troubleshooting log files (Grafana).
- REST API client program
- InfluxDB to store metrics data related to Kafka consumer groups, alert message as well as sample logging stack JMX information from Kafka.

ISIQ Overview

IBM® Security Information Queue (ISIQ) helps organizations manage, automate, and track the use of shared privileged identities. It works as a general-purpose data exchange broker. This feature is obtained by several services running in a separate container in a cluster of Docker nodes.
The solution provides the following features:

- Cross Product Integration
- Stateful Services

- Rolling Update Supported

ISIQ Environmental Components

Critical Docker Containers

- Kafka
- Grafana
- NGINX
- Zookeeper
- Kibana
- InfluxDB
- ISIQ containers (Such as: metrics, connect, products, rest, etc.)

IBM® Security Information Queue Tuning Focus Areas

Docker Host Optimization

/etc/sysctl.conf

Elasticsearch uses a mmapfs directory by default to store its indices. The default operating system limits on mmap counts is likely to be too low, which may result in out of memory (OOM) exceptions. To avoid such issues, vm.max_map_count needs to be increased from its default value “65530” to “262144.”

The value can be changed by adding the following line to /etc/sysctl.conf.

```
vm.max_map_count=262144
```

/etc/docker/daemon.json

The Docker daemon does not include log rotation by default. This feature creates a risk that ISIQ could consume signficant disk space. It is recommended to perform the log rotation in the ISIQ server by creating/modifying the **daemon.json** file located in /etc/docker. This parameter controls the connect logs created by ISIQ. Total number of files and their maximum size can be specified as shown in the example of **dameon.json** below:

```
{
"log-driver": "json-file",
"log-opts": {
"max-size": "20m",
"max-file": "10"
}
```

```
}
```

ISIQ has integrated the feature into all its services. This helps lower the footprint of ISIQ. For this reason, logging is defined in each of the yaml files that comes with the ISIQ starter pack. Customers have flexibility to edit the yaml files based on server configuration and system requirements. Customers can enable log rolling either by updating the daemon.json or ISIQ yaml files. *If the logging section on the yaml files is not removed, then the logging settings defined in the yaml files will override the one defined in daemon.json.* Example of log rotation in ISIQ yaml files is as follows:

```
Logging:
  driver: “json-file”
  options:
    max-size: “10m”
    max-file: “5”
```

connect_stack.yml

There are various topic-level configs values in the connect_stack.yml which can be modified to improve the performance. There are various configurations which pertains to topics which could be both server default and optional per-topic override. If the per-topic configuration is not provided, then the server default is used. In the connect_stack.yml ISIQ provide various topic-level configs that can be modified such as: CONNECT_DEPENDENCY_WAIT_TIMEOUT, CONNECT_OFFSET_FLUSH_TIMEOUT_MS, CONNECT_OFFSET_FLUSH_INTERVAL_MS, CONNECT_REQUEST_TIMEOUT_MS, etc. In the laboratory environment, some of the values were modified which resulted in the better performance. They are explained below:

CONNECT_DEPENDENCY_WAIT_TIMEOUT

This value is related to the time ISIQ will wait before timing out while searching for a dependency. This value can be modified from **connect_stack.yml** that comes with the starter pack of the ISIQ. By default, the CONNECT_DEPENDENCY_WAIT_TIMEOUT is set to 10 minutes. It is not appropriate for an initial load but would be reasonable for the steady state afterwards.

This value can be changed based on the objects in the ISIM environment, and the type of scenario being run. For the initial roll out for the large

environment, it should be set to the higher value and the initial roll out for the smaller environment, it can be reduced to smaller value. Similarly, after the initial load is transferred into IGI, this value can be reduced and set to smaller number, because there should not be much of a processing delay. In the laboratory environment, tests were performed with the following ISIM objects:

- Env1: 301,014 Users, 226,014 ITIM Accounts, and 150,004 LDAP Accounts.
- Env2: 9k Users/Accounts, 62k Service groups, 12k Services
- Env3: 40k Users/Accounts

During the laboratory testing the value of the timeout was set to over four hours for Env1 and Env2, whereas the timeout value was set to about thirty minutes for the Env3 for the initial load scenario. The value of CONNECT_DEPENDENCY_WAIT_TIMEOUT is set in millisecond as follows:

```
CONNECT_DEPENDENCY_WAIT_TIMEOUT=14400000
```

CONNECT_OFFSET_FLUSH_INTERVAL_MS

The CONNECT_OFFSET_FLUSH_TIMEOUT_MS is the interval (milliseconds) at which to try committing offsets for tasks. The value for CONNECT_OFFSET_FLUSH_INTERVAL_MS is set to 80000 by default. Multiple tests were performed in the performance laboratory environment and observed that the best result was obtained when the value was set to 30000.

CONNECT_OFFSET_FLUSH_TIMEOUT_MS

The value defined for this is the maximum, number of milliseconds to wait for records to flush and partition offset data to be committed to offset storage before cancelling the process and restoring the offset data to be committed in a future attempt. By default, the value for CONNECT_OFFSET_FLUSH_TIMEOUT_MS is set to 65000.

broker_stack.yml

The configuration of Kafka service is defined in **broker_stack.yml** file in the ISIQ starter pack. The following values track the time that events were created in the various Kafka topics. This setting does not have any impact on ISIQ itself but can be very helpful for troubleshooting purposes if humans need to look at the topics to see what messages exist and when they were put there.

```
KAFKA_MESSAGE_TIMESTAMP_TYPE: CreateTime
KAFKA_LOG_MESSAGE_TIMESTAMP_TYPE: CreateTime
```

ISIQ Resource Allocation

Resources such as memory, CPU (sockets/cores), and storage (HDD/SSD) should be monitored as ISIQ processes new data from the topics. Allocating adequate resources to the ISIQ related systems ensures timely completion of workload and system. *It is recommended to deploy ISIQ on **bare metal** systems as opposed to virtual machines (VMs)*. In the event that VMs are utilized, it is recommended that the system administrator utilizes dedicated resources (cores, memory, storage) for the virtual appliance. It is not recommended to use shared resources such as resource pools.

Docker Host Recommendations

- *CPU*

  2 Socket, 4 Cores per socket (dedicated)

- *Storage*

  300 GB HDD/SSD

- *Memory*

  32 GB Memory

Performance Limiting Factors

- network throughput constraints
- firewall throttling
- network intrusion prevention systems
- network intrusion detection systems

Enable NTP

Network Time Protocol (NTP) is used to synchronize computer clocks on a network of information. The use of NTP server ensures that the time is observed accurately between the servers and the internet. *It is recommended that NTP is enabled for all the ISIQ nodes*.

Tuning IBM® Security Identity Manager (ISIM)

Tuning ISIM Virtual Appliance

Increase JDBC Connection Max Pool to 100

A database connection pool is a cache of connections to the database that can be modified to improve performance. The default number of connections per resource for ISIM is 30. In laboratory testing, it was observed that multiple connections were on *UOW Waiting* status, waiting for the connection resources to be free. To avoid this situation and potential connection timeout or failures, the DB connection pool should be increased. In the laboratory testing, a setting of 100 was able to reduce the connection errors and *UOT Waiting* status. To Increase the JDBC Connection Max Pool follow the following steps:

a Go to the top level of the Appliance Dashboard click Configure > Database Server Configuration.
b Select Identity data store and click either Configure or Reconfigure.
c Click the Connection Pool tab in the Edit Identity data store details window.
d Specify the values that you want to set for the connections.

Email notification

Enrole.workflow.notifyoption
The enrole.workflow.notifyoption option specifies the behavior of workflow email notifications. The following values can be used for this option.
0 (NOTIFY_NONE): Security Identity Manager does not send email notifications when the workflow process is complete.
1 (NOTIFY_REQUESTER): A process completion notification is sent to the requester when the workflow process completes. Account email notifications are then sent to the requestee for the following accounts requests:

- New Account
- New Password
- Change Account
- Deprovision Account
- Suspend Account
- Restore Account

In laboratory testing with 2000 users assigned as system administrators and the value for enrole.workflow.notifyoption set to “1”, substantial performance degradation was observed in ISIM. When the value was changed to “0”, the performance improved significantly. In order to change the value, the following steps are required:

a From the Appliance Dashboard navigate to Configure -> Update Property
b In the left window search for enRole.properties.

c On the search field on the right type in enrole.workflow.notifyoption
d Select enrole.workflow.notifyoption -> click edit, and change the value to 0

Email Notification Templates

a From the ISIM Console UI navigation tree, select Configure System > Workflow Notification Properties.
b On the Workflow Notification Properties page, for each of the E-mail Notification Templates, in the Status column of the table, click the popup menu icon, and then click Disable.
c After the value of the field changes to Disabled, click OK.
d On the Success page, click Close.

Post Office

a From the navigation tree, select Configure System > Post Office.
b Unselect the Enable store forwarding to disable the post office.

ISIM Local Management Interface (LMI)

a From the Appliance Dashboard navigate to Configure -> Update Property
b In the left window search for ui.properties and set the following values:
    - enrole.ui.maxSearchResults: 1000 (default value)
    - enrole.ui.pageLinkMax: 10 (default value)
    - enrole.ui.pageSize: 50 (default value)
c In the left window search for enRole.properties and set the following values:
    - enrole.reconciliation.threadcount: 4
    - enrole.connectionpool.protocol: plain ssl (default value)
    - enrole.search.paging.enable=true
    - enrole.search.paging.pagesize=128 (default value)
    - enrole.userACICache.maxSize=50 (default value)
    - enrole.accesscontrollist.refreshInterval=10 (default value)
    - enrole.accesscontrollist.maxSize=1000 (default value)

Enable NTP

Network Time Protocol (NTP) is used to synchronize computer clocks on a network of information. The use of NTP server ensures that the time is observed accurately between the servers and the internet. It is

recommended that NTP is enabled in all the ISIM components ISIM VA, ISIM SDS, and ISIM DB2.

Tuning ISIM Directory Server

Enable monitoring in ISDS DB

It is recommended to enable monitoring for any DB2 servers because it allows the system administrator with the ability to tune DB2 to obtain maximum performance. Perform the following actions in order to enable monitoring in ISDS DB:

```
db2 connect to <itimdb> user <user> using
<password>
db2 update dbm cfg using DFT_MON_STMT ON
db2 update dbm cfg using DFT_MON_BUFPOOL ON
db2 update dbm cfg using DFT_MON_LOCK ON
db2 update dbm cfg using DFT_MON_SORT ON
db2 update dbm cfg using DFT_MON_TIMESTAMP
ON
db2 update dbm cfg using DFT_MON_UOW ON
db2 disconnect current
```

Move the DB2 transaction log to a separate disk

The performance of the ISIM ISDS DB can be adversely affected because of the I/O wait time. One of the reasons for high I/O wait time is a very active DB transaction log on the DB2. To mitigate this situation, the transaction log should be moved to a separate, dedicated disk. The detailed steps we followed for moving DB2 transaction log are explained here**.**

Required indexes in ISDS DB

The target attribute (MODIFY_TIMESTAMP) is not a standard attribute. It is an operational attribute. Standard attributes in ISDS have their own tables whereas MODIFY_TIMESTAMP is actually a column on the LDAP_ENTRY table. Before applying the index, directory connector lookups were taking 5-30 seconds, because they are reading all of the entries before filtering out the ones with the right timestamp. The following index must be applied in ISIM ISDS to avoid expensive queries on LDAP_ENTRY table:

```
db2 "create index ldap_entry_modify_timestamp on
ldap_entry(modify_timestamp, eid) collect detailed statistics"
```

Enable change log

ISIQ tracks changes to ISIM by monitoring change logs. To enable change log, perform the following steps:

a Stop the Instance:
/opt/< SDS Install Directory >/sbin/ibmslapd -I <instance-name> -k

b Enable the Change log:
/opt/< SDS Install Directory >/sbin/idscfgchglg -I <instance-name> -m <size limit>

Additionally, it is also possible to configure the days and hours using the -y and -h. The example below shows that the Change log have the size limit of 1,000,000 and an entry age of 25 hours:
/opt/<SDS Install Directory>/sbin/idscfgchglg -I <instance-name> -m 1000000 -y 1 -h 1

For more information about SDS feature visit the following link:
https://www.ibm.com/support/knowledgecenter/SSVJJU_6.2.0/com.ibm.IBMDS.doc_6.2/install17.htm#cfgchgl

SDS Audit log maintenance

When the ISIM product is created in ISIQ, the size of audit.log located under /home/<instance-owner>/idsslapd-idsldap/logs/ will increase up to 2GB in 1 day. The audit.log will continue to increase in size as long as the product remains connected to ISIQ. Hence, it is strongly recommended to roll-over the audit.log. Download the script roll_audit.sh from https://github.com/IBM-Security/performance/tree/master/IAM/scripts and apply it in the SDS server to roll-over the log.

Runstats

As noted in the DB2 Performance Optimization section here, it is recommended to periodically update stable statistics so that the DB2 optimizer will make use of the most efficient access plan for executing queries.

DB2_ANTIJOIN=EXTENDED

DB2_ANTIJOIN=EXTENDED causes the optimizer to equally consider subqueries of a NOT EXISTS clause. This parameter can benefit most queries with a NOT IN clause and a NOT EXISTS clause. You can identify all the queries in your environment using these clauses and validate the benefits of this variable for your specific environment.

When this variable is set to YES, the optimizer searches for opportunities to transform NOT EXISTS subqueries into anti-joins which can be processed more efficiently by DB2.
When this variable is set to EXTEND, the optimizer searches for opportunities to transform both NOT IN and NOT EXISTS subqueries into anti-joins. It is recommended to have the value of DB2_ANTIJOIN=EXTENDED.

Instance_Memory=Automatic

This parameter specifies the maximum amount of memory that can be allocated for a database partition. The DB2 database products can be set with the memory usage restrictions or to a specific value. Since the integration of ISIQ requires the change log to be enabled it is recommended that the instance_memory is set to automatic.
The AUTOMATIC setting results in a value that is computed at database partition activation. The computed value ranges between 75 percent and 95 percent of the system memory capacity on the system – the larger the system, the higher the percentage. The AUTOMATIC settings allow Instance Memory usage to grow as needed. If the Self tuning Memory Manager (STMM) is enabled and tunes the overall database memory size, STMM tunes based on available system memory. This option indirectly determines the actual instance memory usage.

Enable NTP

Network Time Protocol (NTP) is used to synchronize computer clocks on a network of information. The use of NTP server ensures that the time is observed accurately between the servers and the internet. It is recommended that NTP is enabled in all the ISIM components ISIM VA, ISIM SDS, and ISIM DB2 so the clock is synchronized.

Tuning ISIM DB Server

Enable Monitoring in ISIM DB

It is recommended to enable monitoring for any DB2 servers because it allows the system administrator with the ability to tune DB2 to obtain maximum performance. Perform the following actions in order to enable monitoring in ISIM DB:

```
db2 connect to <itimdb> user <user> using
<password>
db2 update dbm cfg using DFT_MON_STMT ON
db2 update dbm cfg using DFT_MON_BUFPOOL ON
db2 update dbm cfg using DFT_MON_LOCK ON
db2 update dbm cfg using DFT_MON_SORT ON
```

```
db2 update dbm cfg using DFT_MON_TIMESTAMP
ON
db2 update dbm cfg using DFT_MON_UOW ON
db2 disconnect current
```

Move the DB2 transaction log to the separate disk

The performance of the ISIM DB2 can be adversely affected because of the I/O wait time. One of the reasons for high I/O wait time is a very active DB transaction log on the DB2. To mitigate this situation, the transaction log should be moved to a separate, dedicated disk. The detailed steps we followed for moving DB2 transaction log are explained here**.**

Runstats

As noted in the DB2 Performance Optimization section here, it is recommended to periodically update stable statistics to that the DB2 optimizer will make use of the most efficient access plan for executing queries.

CPU

After connecting the ISIM product with ISIQ, one should expect to observe a moderate increase in average CPU utilization. This increase could average 2%-3% when no entries are being processed and could increase to 13% during times in which larger number of entries are being processed by ISIQ.

Required Indexes for ISIQ optimization

When the ISIM is configured in ISIQ, ISIQ queries ISIM DB server which holds audit event and process workflow information not contained in the ISIM SDS. This causes the CPU usage of ISIM DB server to spike up as soon as the ISIQ connection is made. The following indexes will help to reduce the CPU usage.

```
CREATE INDEX INTIMUSER.T1_TIMESTAMP_ITIM_EVENT_CATEGORY ON
ITIMUSER.AUDIT_EVENT (“ITIM_EVENT_CATEGORY” ASC, “TIMESTAMP”
ASC) ALLOW REVERSE SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX ITIMUSER.COMPLETED ON ITIMUSER.PROCESS
(“COMPLETED” ASC) ALLOW REVERSE SCANS COLLECT DETAILED
STATISTICS;
CREATE INDEX ITIMUSER.EVENTTYPE ON ITIMUSER.PROCESSLOG
(“EVENTTYPE” ASC) ALLOW REVERSE SCANS COLLECT DETAILED
STATISTICS;
```

DB2_ANTIJOIN=EXTENDED

DB2_ANTIJOIN=EXTENDED causes the optimizer to equally consider subqueries of a NOT EXISTS clause. This parameter can benefit most queries with a NOT IN clause and a NOT EXISTS clause. You can identify all the

queries in your environment using these clauses and validate the benefits of this variable for your specific environment.
When this variable is set to YES, the optimizer searches for opportunities to transform NOT EXISTS subqueries into anti-joins which can be processed more efficiently by DB2. When this variable is set to EXTEND, the optimizer searches for opportunities to transform both NOT IN and NOT EXISTS subqueries into anti-joins. It is recommended to have the value of DB2_ANTIJOIN=EXTENDED

Instance_Memory=Automatic

This parameter specifies the maximum amount of memory that can be allocated for a database partition. The DB2 database products can be set with the memory usage restrictions or to a specific value. The AUTOMATIC setting results in a value that is computed at database partition activation. The computed value ranges between 75 percent and 95 percent of the system memory capacity on the system – the larger the system, the higher the percentage. The AUTOMATIC settings allow Instance Memory usage to grow as needed. If STMM is enabled and tunes the overall database memory size, STMM tunes based on available system memory. This option indirectly determines the actual instance memory usage. To ensure that enough memory is available to process long running or large request in DB2 it is recommended to have the value of instance_memory to automatic

Enable NTP

Network Time Protocol (NTP) is used to synchronize computer clocks on a network of information. The use of NTP server ensures that the time is observed accurately between the servers and the internet. It is recommended that NTP is enabled in all the ISIM components ISIM VA, ISIM SDS, and ISIM DB2 so the clock is synchronized

Tuning IBM® Security Identity Governance and Intelligence (IGI)

Tuning IGI VA

Ib_recon_failure_threshold

Set the ib_recon_failure_threshold to 100% to ensure the SDS acts on all the changes in your target.

a From Command Line Interface (CLI), log onto the target as ‘admin’
b IGI → utilities → ib_settings → ib_recon_failure_threshold
c set (Type in 100 and hit enter, after it prompts you)
d get (To verify that it is set to 100)

Enable NTP

Network Time Protocol (NTP) is used to synchronize computer clocks on a network of information. The use of NTP server ensures that the time is observed accurately between the servers and the internet. It is recommended that NTP is enabled in all the IGI components IGIVA, and IGI DB2 so the clock is synchronized.

Worker threads for queues

There are five queues which are serviced by the Rule Engine. Four queues (IN, TARGET, OUT, and INTERNAL queues) have multiple threads to process events. The number of worker threads for each queue is now configurable by setting the value of thread number from 1 to 10 threads. By default, the value of the thread for all the queues is set to 3.
For IGI version 5.2.4.1 the queues can be accessed at Task Planner → Manage → Rule Engine → Jobs (on the middle pane) as shown in the image below.

Image 1: Rule Engine

Whereas in IGI version 5.2.5 or newer, the four queues can be found at Task Planner → Manage.

The Rule Engine must be stopped for this value to be adjusted, but it is not necessary to restart the appliance. It is recommended to monitor the CPU utilization closely when the number of Rule Engine threads is adjusted and increase this number slowly to avoid over-committing the CPU resources.

Increasing the Heap Size

Starting with IGI version 5.2.4, it provides a mechanism to increase the heap size to a max of 8GB. This step might be necessary when running large bulk loads, or launching many or particularly large hierarchies at the same time, etc. If the VA reports an error message that the heap has been exhausted, the following procedure can be used to increase the heap. From the CLI navigate to igi → jvm_heapsize → set_max_heapsize. The menu will offer an option to change the IGI application server heap or the Broker application server heap. The user may choose any size between 4096 MB and 8192 MB. A restart of the application server is necessary. Up to 5.2.6 Version, for configurations with large numbers of users, 500K+, it is advisable to set the heap to 8GB to ensure the heap is not exhausted during report generation. In 5.2.6.1 this mechanism has improved and in configurations with about 700K users, the default 4092MB has safe to run reports.

IGI Admin Console Settings

1 Synchronize the Scheduler

Log in to Admin console → Task Planner and make sure that the active tasks are synchronized.

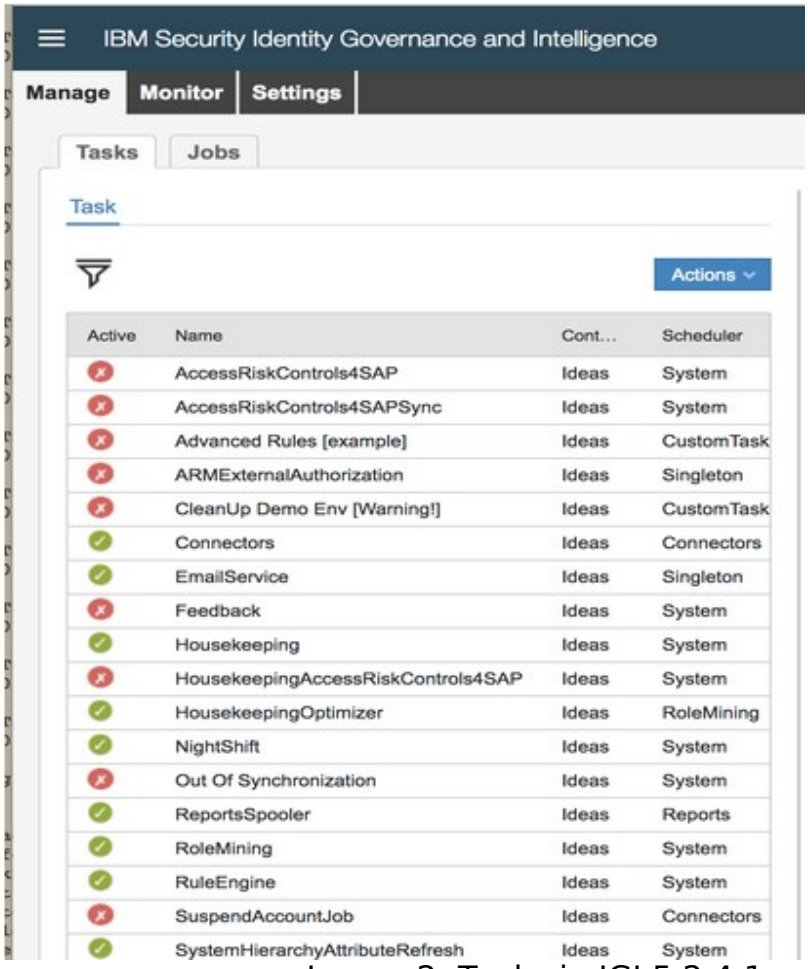

Image 2: Tasks in IGI 5.2.4.1

If any of the task that are running are not synchronized use the following steps to make sure that the tasks are synchronized:
Log in to Admin console → Task Planner → Settings → Scheduler → Select each Scheduler as shown in the image below → Action → Synchronize (for IGI 5.2.4.1) or Synchronize All (for IGI 5.2.5).

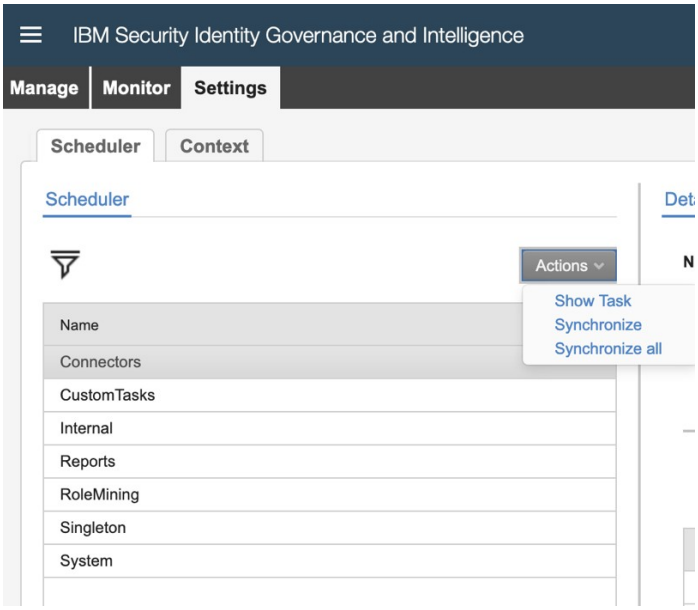

Image 3: Synchronize example

2 Enable History for each of the following Tasks:

- Connectors
- HouseKeeping
- HouseKeepingOptimizer
- RoleMining
- RoleEngine
- SystemHierarchyAttributeRefresh

Click on each Task above → Actions → Stop → Select the Enable History option → Start the Task.

3 While subscribing the ISIM product in ISIQ, the performance of IGI could be increased by stopping the task called NightShift. To stop the NightShift, go to Task Planner Settings → Click on NightShift → Actions → Stop from the drop-down menu.

4 Prior to starting the initial load by subscribing the ISIM through ISIQ, it is recommended that the data analytics be turned off (Risk Scans, Role Mining, etc.) and resumed after the operation is complete.

Tuning IGIDB

Move the DB2 Transaction log to the separate disk

The performance of the IGI DB2 can be adversely affected because of the I/O wait time. One of the reasons for high I/O wait time is a very active DB transaction log on the DB2. To mitigate this situation, the transaction log should be moved to a separate, dedicated disk. The detailed steps we followed for moving DB2 transaction log are explained here.

Increase log buffer size

This is required due to the expectation of high disk utilization and considerable read activity on the dedicated log disk. Execute the following command as the instance owner. Restart the DB2 database so that the changes will take effect.

```
db2 update db cfg using logbufsz
3072;
```

IGA_SERVICE_BLOB

Avoid running out of space by disabling a max size for table space “IGA_SERVICE_BLOB”. Run the following command on the DB2 as the instance owner, then restart the DB2 and IGI application server.

```
db2 “alter tablespace IGA_SERVICE_BLOB MAXSIZE
NONE”
```

Increase log size

Increase the log size with the following command by running it as the instance owner:

```
db2 “UPDATE DB CFG USING logprimary 200 logsecond 50
logfilsiz 16384”
```

Drop all foreign key constraints

The foreign key constraint definition can cause the database deadlock in IGI. This can be avoided by dropping the constraints. To drop the foreign key constraint, copy the following and create a new file called FK-drop.ddl to /home/<instance_owner>.

```
ALTER TABLE IGAQRZ.QRZ1_BLOB_TRIGGERS DROP FOREIGN KEY
QRZ1_BLOB_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ1_CRON_TRIGGERS DROP FOREIGN KEY
QRZ1_CRON_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ1_SIMPLE_TRIGGERS DROP FOREIGN KEY
```

```
QRZ1_SIMPLE_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ1_SIMPROP_TRIGGERS DROP FOREIGN KEY
QRZ1_SIMPROP_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ1_TRIGGERS DROP FOREIGN KEY
QRZ1_TRIGGER_TO_JOBS_FK;
ALTER TABLE IGAQRZ.QRZ2_BLOB_TRIGGERS DROP FOREIGN KEY
QRZ2_BLOB_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ2_CRON_TRIGGERS DROP FOREIGN KEY
QRZ2_CRON_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ2_SIMPLE_TRIGGERS DROP FOREIGN KEY
QRZ2_SIMPLE_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ2_SIMPROP_TRIGGERS DROP FOREIGN KEY
QRZ2_SIMPROP_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ2_TRIGGERS DROP FOREIGN KEY
QRZ2_TRIGGER_TO_JOBS_FK;
ALTER TABLE IGAQRZ.QRZ3_BLOB_TRIGGERS DROP FOREIGN KEY
QRZ3_BLOB_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ3_CRON_TRIGGERS DROP FOREIGN KEY
QRZ3_CRON_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ3_SIMPLE_TRIGGERS DROP FOREIGN KEY
QRZ3_SIMPLE_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ3_SIMPROP_TRIGGERS DROP FOREIGN KEY
QRZ3_SIMPROP_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ3_TRIGGERS DROP FOREIGN KEY
QRZ3_TRIGGER_TO_JOBS_FK;
ALTER TABLE IGAQRZ.QRZ4_BLOB_TRIGGERS DROP FOREIGN KEY
QRZ4_BLOB_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ4_CRON_TRIGGERS DROP FOREIGN KEY
QRZ4_CRON_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ4_SIMPLE_TRIGGERS DROP FOREIGN KEY
QRZ4_SIMPLE_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ4_SIMPROP_TRIGGERS DROP FOREIGN KEY
QRZ4_SIMPROP_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ4_TRIGGERS DROP FOREIGN KEY
QRZ4_TRIGGER_TO_JOBS_FK;
ALTER TABLE IGAQRZ.QRZ5_BLOB_TRIGGERS DROP FOREIGN KEY
QRZ5_BLOB_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ5_CRON_TRIGGERS DROP FOREIGN KEY
QRZ5_CRON_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ5_SIMPLE_TRIGGERS DROP FOREIGN KEY
QRZ5_SIMPLE_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ5_SIMPROP_TRIGGERS DROP FOREIGN KEY
QRZ5_SIMPROP_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ5_TRIGGERS DROP FOREIGN KEY
QRZ5_TRIGGER_TO_JOBS_FK;
ALTER TABLE IGAQRZ.QRZ6_BLOB_TRIGGERS DROP FOREIGN KEY
QRZ6_BLOB_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ6_CRON_TRIGGERS DROP FOREIGN KEY
QRZ6_CRON_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ6_SIMPLE_TRIGGERS DROP FOREIGN KEY
```

```
QRZ6_SIMPLE_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ6_SIMPROP_TRIGGERS DROP FOREIGN KEY
QRZ6_SIMPROP_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ6_TRIGGERS DROP FOREIGN KEY
QRZ6_TRIGGER_TO_JOBS_FK;
ALTER TABLE IGAQRZ.QRZ7_BLOB_TRIGGERS DROP FOREIGN KEY
QRZ7_BLOB_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ7_CRON_TRIGGERS DROP FOREIGN KEY
QRZ7_CRON_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ7_SIMPLE_TRIGGERS DROP FOREIGN KEY
QRZ7_SIMPLE_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ7_SIMPROP_TRIGGERS DROP FOREIGN KEY
QRZ7_SIMPROP_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ7_TRIGGERS DROP FOREIGN KEY
QRZ7_TRIGGER_TO_JOBS_FK;

ALTER TABLE IGAQRZ.QRZ1_BLOB_TRIGGERS DROP FOREIGN KEY
QRZ1_BLOB_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ1_CRON_TRIGGERS DROP FOREIGN KEY
QRZ1_CRON_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ1_SIMPLE_TRIGGERS DROP FOREIGN KEY
QRZ1_SIMPLE_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ1_SIMPROP_TRIGGERS DROP FOREIGN KEY
QRZ1_SIMPROP_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ1_TRIGGERS DROP FOREIGN KEY
QRZ1_TRIGGER_TO_JOBS_FK;
ALTER TABLE IGAQRZ.QRZ2_BLOB_TRIGGERS DROP FOREIGN KEY
QRZ2_BLOB_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ2_CRON_TRIGGERS DROP FOREIGN KEY
QRZ2_CRON_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ2_SIMPLE_TRIGGERS DROP FOREIGN KEY
QRZ2_SIMPLE_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ2_SIMPROP_TRIGGERS DROP FOREIGN KEY
QRZ2_SIMPROP_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ2_TRIGGERS DROP FOREIGN KEY
QRZ2_TRIGGER_TO_JOBS_FK;
ALTER TABLE IGAQRZ.QRZ3_BLOB_TRIGGERS DROP FOREIGN KEY
QRZ3_BLOB_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ3_CRON_TRIGGERS DROP FOREIGN KEY
QRZ3_CRON_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ3_SIMPLE_TRIGGERS DROP FOREIGN KEY
QRZ3_SIMPLE_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ3_SIMPROP_TRIGGERS DROP FOREIGN KEY
QRZ3_SIMPROP_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ3_TRIGGERS DROP FOREIGN KEY
QRZ3_TRIGGER_TO_JOBS_FK;
ALTER TABLE IGAQRZ.QRZ4_BLOB_TRIGGERS DROP FOREIGN KEY
QRZ4_BLOB_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ4_CRON_TRIGGERS DROP FOREIGN KEY
QRZ4_CRON_TRIG_TO_TRIG_FK;
```

```
ALTER TABLE IGAQRZ.QRZ4_SIMPLE_TRIGGERS DROP FOREIGN KEY
QRZ4_SIMPLE_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ4_SIMPROP_TRIGGERS DROP FOREIGN KEY
QRZ4_SIMPROP_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ4_TRIGGERS DROP FOREIGN KEY
QRZ4_TRIGGER_TO_JOBS_FK;
ALTER TABLE IGAQRZ.QRZ5_BLOB_TRIGGERS DROP FOREIGN KEY
QRZ5_BLOB_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ5_CRON_TRIGGERS DROP FOREIGN KEY
QRZ5_CRON_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ5_SIMPLE_TRIGGERS DROP FOREIGN KEY
QRZ5_SIMPLE_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ5_SIMPROP_TRIGGERS DROP FOREIGN KEY
QRZ5_SIMPROP_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ5_TRIGGERS DROP FOREIGN KEY
QRZ5_TRIGGER_TO_JOBS_FK;
ALTER TABLE IGAQRZ.QRZ6_BLOB_TRIGGERS DROP FOREIGN KEY
QRZ6_BLOB_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ6_CRON_TRIGGERS DROP FOREIGN KEY
QRZ6_CRON_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ6_SIMPLE_TRIGGERS DROP FOREIGN KEY
QRZ6_SIMPLE_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ6_SIMPROP_TRIGGERS DROP FOREIGN KEY
QRZ6_SIMPROP_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ6_TRIGGERS DROP FOREIGN KEY
QRZ6_TRIGGER_TO_JOBS_FK;
ALTER TABLE IGAQRZ.QRZ7_BLOB_TRIGGERS DROP FOREIGN KEY
QRZ7_BLOB_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ7_CRON_TRIGGERS DROP FOREIGN KEY
QRZ7_CRON_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ7_SIMPLE_TRIGGERS DROP FOREIGN KEY
QRZ7_SIMPLE_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ7_SIMPROP_TRIGGERS DROP FOREIGN KEY
QRZ7_SIMPROP_TRIG_TO_TRIG_FK;
ALTER TABLE IGAQRZ.QRZ7_TRIGGERS DROP FOREIGN KEY
QRZ7_TRIGGER_TO_JOBS_FK;
```

Invoke it as the instance owner with the following command:

```
db2 -tvf /home/<instance_owner>/FK-drop.ddl
```

Increase Catalog cache size

Increase Catalog cache size (CATALOGCACHE_SZ) DB Configuration parameter. The catalog cache size is 300 by default. While running the ISIQ connection if the message, “ADM4000W A catalog cache overflow condition has occurred” appears which means that the catalog cache as exceed the configured maximum size. If the condition persists, it is a good idea to

increase the CATALOGCACHE_SZ. Run the following command as the DB2 instance owner to increase it:

```
db2 “update db cfg using CATALOGCACHE_SZ <new_number>”
```

Indexes

Various indexes for IGIDB were identified that improves the performance of IGI. Following were the different indexes identified:

a For IGI V 5.2.4.1: igi5241_indexes.ddl

```
CREATE INDEX DISPATCHED_WORK_KEYS_IDX ON
IGACORE.DISPATCHED_WORK_KEYS(DISPATCHER_ID) COLLECT detailed
statistics ALLOW REVERSE SCANS;
CREATE INDEX EMPLOYMENT_REVIEW_IDX ON IGACORE.EMPLOYMENT_REVIEW
(ATTESTATION ASC, ID ASC) COLLECT detailed statistics ALLOW REVERSE
SCANS;
CREATE INDEX ENTITLEMENT_IDX ON
IGACORE.ENTITLEMENT(LOWER(EXT_REF), PROFILE_TYPE) COLLECT detailed
statistics ALLOW REVERSE SCANS;
CREATE INDEX EVENT_TARGET_SOS_IDX ON IGACORE.EVENT_TARGET
(SYNC_STATUS, OPERATION, STATE) COLLECT DETAILED STATISTICS;
CREATE INDEX EVENT_OUT_SOS_IDX ON IGACORE.EVENT_OUT (SYNC_STATUS,
OPERATION, STATE) COLLECT DETAILED STATISTICS;
CREATE INDEX INTERNAL_LOCK_ID_IDX ON IGACORE.INTERNAL_LOCK(ID)
collect detailed statistics ALLOW REVERSE SCANS;
CREATE INDEX TIMER_HISTORY_REA_GRO_IDX ON IGASERV.TIMER_HISTORY
(REALM, GROUP_KEY) COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.ORG_UNIT_IDX ON
IGACORE.ORANIZATIONAL_UNIT(HIERARCHY, LAST_MOD_TIME) COLLECT
DETAILED STATISTICS;
CREATE INDEX IGAQRZ.IDX_QRZ7_T_NEXT_FIRE_TIME_ONLY ON
IGAQRZ.QRZ7_TRIGGERS(NEXT_FIRE_TIME) ALLOW REVERSE SCANS COLLECT
STATISTICS;
```

igi5241_indexes.ddl

b For IGI V 5.2.5: igi525_indexes.ddl

```
CREATE INDEX IGACORE.EMP_REV_ATT_PERSON ON
IGACORE.EMPLOYMENT_REVIEW(ATTESTATION,PERSON) COLLECT DETAILED
STATISTICS;
CREATE INDEX IGACORE.EMP_REVER_CERT_FIRST ON
IGACORE.EMPLOYMENT_REVIEWER(CERT_FIRST_OWNER,EMPLOYMENT_REVIEW
) COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.EMP_REV_ATT_SIGNOFF_REVSTATE ON
IGACORE.EMPLOYMENT_REVIEW(ATTESTATION,SIGNED_OFF,REVIEW_STATE)
```

```
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.CHANGELOG_PWDMANAGEMENT_IDX ON
IGACORE.CHANGELOG (PWDMANAGEMENT ASC) ALLOW REVERSE SCANS
COLLECT SAMPLED DETAILED STATISTICS;
CREATE INDEX IGACORE.PERSONRISK_REMEDIATION_IDX ON
IGACORE.PERSON_RISK(PERSON ASC, REMEDIATION ASC, ENVIRONMENT ASC,
LAST_MOD_TIME ASC, RISK ASC, ID ASC) ALLOW REVERSE SCANS COLLECT
SAMPLED DETAILED STATISTICS;
CREATE INDEX IGACORE.CHANGELOG_PWDMANAGEMENT_IDX ON
IGACORE.CHANGELOG (PWDMANAGEMENT ASC) ALLOW REVERSE SCANS
COLLECT SAMPLED DETAILED STATISTICS;
```

Igi525_indexes.ddl

c create_indexes_igacore.ddl:

```
CREATE INDEX IGACORE.IDX_AR_REV_PWDCFG_FK_FPIDX ON
IGACORE.ACCOUNT_REVIEW ("PWDCFG") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_AR_REV_REVIEWED_BY_FK ON
IGACORE.ACCOUNT_REVIEW ("REVIEWED_BY") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_AR_REVIEWER_FIRSTOWNER_FK ON
IGACORE.ACCOUNT_REVIEWER ("CERT_FIRST_OWNER") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_AR_REVIEWER_SECONDOWNER_FK ON
IGACORE.ACCOUNT_REVIEWER ("CERT_SECOND_OWNER") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_ACCOUNT_REV_HIS_PERSON_FK ON
IGACORE.ACCOUNT_REVIEW_H ("PERSON") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_ACCOUNT_REV_HIS_PWDCFG_FK ON
IGACORE.ACCOUNT_REVIEW_H ("PWDCFG") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_ACCOUNT_REV_HIS_REVIEWED_BY_FK ON
IGACORE.ACCOUNT_REVIEW_H ("CERT_FIRST_OWNER") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_ARNOTE_ACCOUNT_REVIEWER_FK ON
IGACORE.ACCOUNT_REVIEW_NOTE ("ACCOUNT_REVIEWER") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_ARNOTE_PERSON_FROM_FK ON
IGACORE.ACCOUNT_REVIEW_NOTE ("PERSON_FROM") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_ARNOTE_TO_FK ON
IGACORE.ACCOUNT_REVIEW_NOTE ("PERSON_TO") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_ADDITIONAL_DATA_REQUEST_FK ON
```

```
IGACORE.ADDITIONAL_DATA ("REQUEST") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_APPLICATION_ENT_EXT_TYPE_FK ON
IGACORE.APPLICATION ("SOD_EXT_LEVEL") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_APPLICATION_FLOW_FK ON
IGACORE.APPLICATION ("FLOW") ALLOW REVERSE SCANS COLLECT DETAILED
STATISTICS;
CREATE INDEX IGACORE.IDX_APPLICATIONS_REQUEST_FK ON
IGACORE.APPLICATIONS ("REQUEST") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_ARM_AUTHORIZATIONS_REQUEST_FK ON
IGACORE.ARM_AUTHORIZATIONS ("REQUEST") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_ARM_AUTH_SU_RED_FK ON
IGACORE.ARM_AUTH_S_USER ("REDIRECTED_BY") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_ARM_AUTH_SU_RED_REP_FK ON
IGACORE.ARM_AUTH_S_USER ("RED_REPRESENTED_BY") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_ARM_AUTH_SU_REP_FK ON
IGACORE.ARM_AUTH_S_USER ("REPRESENTED_BY") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_ARM_I18N_MESSAGE_LOC_CODE_FK ON
IGACORE.ARM_I18N_MESSAGE ("LOCALIZATION_CODE") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_ATT_DEF_REVIEWER_FK ON
IGACORE.ATTESTATION ("DEF_REVIEWER") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_ATT_SUPERVISOR_FK ON IGACORE.ATTESTATION
("SUPERVISOR") ALLOW REVERSE SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_ATTBL_USER_ATTESTATION_FK ON
IGACORE.ATTESTATION_BLACKLIST ("ATTESTATION") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_ATTBL_USER_PERSON_FK ON
IGACORE.ATTESTATION_BLACKLIST ("PERSON") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_ATTPROP_ATTESTATION_FK ON
IGACORE.ATTESTATION_PROP ("ATTESTATION") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_ATTVIEW_ATTESTATION_FK ON
IGACORE.ATTESTATION_VIEW ("ATTESTATION") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_ATTR_PERM_MAPPING_TARGET_ID_FK ON
IGACORE.ATTR_PERMISSION_MAPPING ("TARGET_ID") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_CFG_FIX_APP_SCOPE_FK ON
IGACORE.CFG_FIXED_APPLICATION_SCOPE ("ACTIVITY_ID") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_FIXED_REPRSCOPE_FK ON
```

```
IGACORE.CFG_FIXED_REPRESENTED_SCOPE ("ACTIVITY_ID") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_USR_SCOPE_ACTID_FK ON
IGACORE.CFG_FIXED_USER_SCOPE ("ACTIVITY_ID") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_CFG_NOTIF_EMS_TEMPLATE_FK ON
IGACORE.CFG_NOTIFICATION ("EMS_TEMPLATE") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_CFG_PRIORITY_NOTIF_EMAIL_FK ON
IGACORE.CFG_PRIORITY ("CFG_NOTIFICATION_EMAIL") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_CFG_PRIORITY_NOTIF_SMS_FK ON
IGACORE.CFG_PRIORITY ("CFG_NOTIFICATION_SMS") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_CFG_PROCESS_EXPIRATION_WF_FK ON
IGACORE.CFG_PROCESS ("EXPIRATION_WF") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_CFG_PROCESS_NOTIFICATION_FK ON
IGACORE.CFG_PROCESS ("CFG_NOTIFICATION") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_CFG_PROCESS_NOTIF_SMS_FK ON
IGACORE.CFG_PROCESS ("CFG_NOTIFICATION_SMS") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_CFG_PROCESS_PRIORITY_FK ON
IGACORE.CFG_PROCESS ("CFG_PRIORITY") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_CFG_PROCACT_ESC_PROCESS_FK ON
IGACORE.CFG_PROCESSACTIVITY ("ESCALATION_PROCESS") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_CFG_PROCACT_LINK_PROCESS_FK ON
IGACORE.CFG_PROCESSACTIVITY ("LINKED_PROCESS") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_CFG_ACTION_FK ON
IGACORE.CFG_PROCESSACTIVITY_ACTION ("CFG_PROCESSACTIVITY") ALLOW
REVERSE SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_CHANGELOG_PWDMANAGEMENT_FK ON
IGACORE.CHANGELOG ("PWDMANAGEMENT") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_CLASS_NAME_TYPE_FK ON IGACORE.CLASS_NAME
("TYPE") ALLOW REVERSE SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_COMPARE_LOG_PARENT_ID_FK ON
IGACORE.COMPARE_LOG ("PARENT_ID") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_CONNECTOR_HISTORY_CONNECTOR_FK ON
IGACORE.CONNECTOR_HISTORY ("CONNECTOR") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_CONN_PROPERTY_CONNECTOR_FK ON
IGACORE.CONNECTOR_PROPERTY ("CONNECTOR") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_COUNTER_LOG_EXECUTION_LOG_FK ON
```

```
IGACORE.COUNTER_LOG ("EXECUTION_LOG") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_DOMAIN_PROFILE_DOMAIN_FK ON
IGACORE.DOMAIN_PROFILE ("ENVIRONMENT", "DOMAIN") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_DRIVER_CLASS_NAME_FK ON IGACORE.DRIVER
("CLASS_NAME") ALLOW REVERSE SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_DRIVER_HIER_CHILD_FK ON
IGACORE.DRIVER_HIER ("CHILD") ALLOW REVERSE SCANS COLLECT DETAILED
STATISTICS;
CREATE INDEX IGACORE.IDX_DRIVER_HIER_PARENT_FK ON
IGACORE.DRIVER_HIER ("PARENT") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_DRIVER_PROPERTY_DRIVER_FK ON
IGACORE.DRIVER_PROPERTY ("DRIVER") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_EMPLOYMENT_HIERARCHY_FK ON
IGACORE.EMPLOYMENT ("HIERARCHY") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_EMP_REV_ENTITLEMENT_FK ON
IGACORE.EMPLOYMENT_REVIEW ("ENTITLEMENT") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_EMP_REV_GROUP_FK ON
IGACORE.EMPLOYMENT_REVIEW ("ORGANIZATIONAL_UNIT") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_EMP_REV_REVIEWED_BY_FK ON
IGACORE.EMPLOYMENT_REVIEW ("REVIEWED_BY") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_EMP_REVIEWER_FIRSTOWNER_FK ON
IGACORE.EMPLOYMENT_REVIEWER ("CERT_FIRST_OWNER") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_EMP_REVIEWER_SECONDOWNER_FK ON
IGACORE.EMPLOYMENT_REVIEWER ("CERT_SECOND_OWNER") ALLOW
REVERSE SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_EMP_REV_HIS_ENTITLEMENT_FK ON
IGACORE.EMPLOYMENT_REVIEW_H ("ENTITLEMENT") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_EMP_REV_HIS_PERSON_FK ON
IGACORE.EMPLOYMENT_REVIEW_H ("PERSON") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_EMP_REV_HIS_REVIEWED_BY_FK ON
IGACORE.EMPLOYMENT_REVIEW_H ("CERT_FIRST_OWNER") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_ERNOTE_EMPLOYMENT_REVIEWER_FK ON
IGACORE.EMPLOYMENT_REVIEW_NOTE ("EMPLOYMENT_REVIEWER") ALLOW
REVERSE SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_ERNOTE_PERSON_FROM_FK ON
IGACORE.EMPLOYMENT_REVIEW_NOTE ("PERSON_FROM") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_ERNOTE_TO_FK ON
```

```
IGACORE.EMPLOYMENT_REVIEW_NOTE ("PERSON_TO") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_EMPREVR_EMPREV_FK ON
IGACORE.EMPLOYMENT_REVIEW_RIGHT ("EMPLOYMENT_REVIEW") ALLOW
REVERSE SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_EMPREVR_PERMISSION_FK ON
IGACORE.EMPLOYMENT_REVIEW_RIGHT ("PERMISSION") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_EMPREVR_PERSON_SERVICE_FK ON
IGACORE.EMPLOYMENT_REVIEW_RIGHT ("PERSON_SERVICE") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_EMPREVR_REVIEWED_BY_FK ON
IGACORE.EMPLOYMENT_REVIEW_RIGHT ("REVIEWED_BY") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_EMP_REV_R_H_ENTITLEMENT_FK ON
IGACORE.EMPLOYMENT_REVIEW_RIGHT_H ("ENTITLEMENT") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_EMP_REV_R_H_PERSON_FK ON
IGACORE.EMPLOYMENT_REVIEW_RIGHT_H ("PERSON") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_EMP_REV_R_H_PER_SERV_FK ON
IGACORE.EMPLOYMENT_REVIEW_RIGHT_H ("PERSON_SERVICE") ALLOW
REVERSE SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_EMP_REV_R_H_REVIEWED_BY_FK ON
IGACORE.EMPLOYMENT_REVIEW_RIGHT_H ("CERT_FIRST_OWNER") ALLOW
REVERSE SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_EMS_I18N_LANGUAGE_FK ON
IGACORE.EMS_I18N_MESSAGE ("LANGUAGE") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_EMS_TEMPLATE_ALTERNATIVE_FK ON
IGACORE.EMS_TEMPLATE ("ALTERNATIVE") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_EMS_TEMPLATE_BODY_FK ON
IGACORE.EMS_TEMPLATE ("BODY") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_EMS_TEMPLATE_SUBJECT_FK ON
IGACORE.EMS_TEMPLATE ("SUBJECT") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_EMS_TEMPLATE_TEMPLATE_TYPE_FK ON
IGACORE.EMS_TEMPLATE ("TEMPLATE_TYPE") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_ENTITLEMENT_APPLICATION_FK ON
IGACORE.ENTITLEMENT ("APPLICATION") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_ENTITLEMENT_ENT_FAMILY_FK ON
IGACORE.ENTITLEMENT ("FAMILY") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_ENTITLEMENT_FLOW_AUTH_FK ON
IGACORE.ENTITLEMENT ("FLOW_AUTH") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
```

```
CREATE INDEX IGACORE.IDX_ENTITLEMENT_FLOW_CHECK_FK ON
IGACORE.ENTITLEMENT ("FLOW_CHECK") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_ENTITLEMENT_OWNER_FK ON
IGACORE.ENTITLEMENT ("OWNER") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_ENTITLEMENT_PROFILE_TYPE_FK ON
IGACORE.ENTITLEMENT ("PROFILE_TYPE") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_ENT_F_HIER_CHILD_APP_FK ON
IGACORE.ENTITLEMENT_FLAT_HIER ("CHILD_APPLICATION") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_ENT_F_HIER_CHILD_PROF_TYPE_FK ON
IGACORE.ENTITLEMENT_FLAT_HIER ("CHILD_PROFILE_TYPE") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_ENT_F_HIER_PARENT_APP_FK ON
IGACORE.ENTITLEMENT_FLAT_HIER ("PARENT_APPLICATION") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_ENT_F_HIER_PARENT_PROF_TYPE_FK ON
IGACORE.ENTITLEMENT_FLAT_HIER ("PARENT_PROFILE_TYPE") ALLOW
REVERSE SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_ENTITLEMENT_INCOMP_ENT_FK ON
IGACORE.ENTITLEMENT_INCOMP ("ENTITLEMENT") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_ENT_REV_REVIEWED_BY_FK ON
IGACORE.ENTITLEMENT_REVIEW ("REVIEWED_BY") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_ENT_REVIEWER_FIRSTOWNER_FK ON
IGACORE.ENTITLEMENT_REVIEWER ("CERT_FIRST_OWNER") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_ENT_REVIEWER_SECONDOWNER_FK ON
IGACORE.ENTITLEMENT_REVIEWER ("CERT_SECOND_OWNER") ALLOW
REVERSE SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_ENT_REV_HIS_ENTITLEMENT_FK ON
IGACORE.ENTITLEMENT_REVIEW_H ("ENTITLEMENT") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_ENT_REV_HIS_REVIEWED_BY_FK ON
IGACORE.ENTITLEMENT_REVIEW_H ("CERT_FIRST_OWNER") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_ENTNOTE_ENT_REVIEWER_FK ON
IGACORE.ENTITLEMENT_REVIEW_NOTE ("ENTITLEMENT_REVIEWER") ALLOW
REVERSE SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_ENTNOTE_PERSON_FROM_FK ON
IGACORE.ENTITLEMENT_REVIEW_NOTE ("PERSON_FROM") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_ENTNOTE_TO_FK ON
IGACORE.ENTITLEMENT_REVIEW_NOTE ("PERSON_TO") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_ENTITLEMENT_SERV_SERV_ATTR_FK ON
IGACORE.ENTITLEMENT_SERVICE ("SERVICE_ATTRIBUTE") ALLOW REVERSE
```

```
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_EXTMAPPING_LOCALIZATION_FK ON
IGACORE.EXTERNALMAPPING ("LOCALIZATION") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_FLOW_IDEAS_MODULE_FK ON IGACORE.FLOW
("IDEAS_MODULE") ALLOW REVERSE SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_FLOW_ENTITY_IDEAS_MODULE_FK ON
IGACORE.FLOW_ENTITY ("IDEAS_MODULE") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_FLOW_EV_TYPE_IDEAS_MODULE_FK ON
IGACORE.FLOW_EVENT_TYPE ("IDEAS_MODULE") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_FLOW_RULE_IDEAS_MODULE_FK ON
IGACORE.FLOW_RULE ("IDEAS_MODULE") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_FLOW_STD_CONF_IDEAS_MODULE_FK ON
IGACORE.FLOW_STD_CONF ("IDEAS_MODULE") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_GROUP_PERSON_HIERARCHY_FK ON
IGACORE.GROUP_PERSON ("HIERARCHY") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_HIERARCHY_FLOW_FK ON IGACORE.HIERARCHY
("FLOW") ALLOW REVERSE SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_HAPH_PWDCFG_ATTR_KEY_FK ON
IGACORE.HR_ACCOUNT_PHOLDER ("PWDCFG_ATTR_KEY") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_HAPH_PWDCFG_FK ON
IGACORE.HR_ACCOUNT_PHOLDER ("PWDCFG") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_HPC_PERSON_FK ON
IGACORE.HR_PHOLDER_CHANGES ("PERSON") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_HPV_PERSON_FK ON
IGACORE.HR_PHOLDER_VALUES ("PERSON") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_INT_RES_REQUEST_FK ON
IGACORE.INT_RESOURCES ("REQUEST") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_INT_RES_AUTH_RISK_FK ON
IGACORE.INT_RES_AUTH ("ENVIRONMENT", "RISK") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_INT_RES_AUTH_TASK_FK ON
IGACORE.INT_RES_AUTH ("ENVIRONMENT", "TASK") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_JOB_IDEAS_MODULE_FK ON IGACORE.JOB
("IDEAS_MODULE") ALLOW REVERSE SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_JOB_PROPERTY_IDEAS_MODULE_FK ON
IGACORE.JOB_PROPERTY ("IDEAS_MODULE") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_JOB_PROPERTY_EXT_JOB_PROP_FK ON
```

```
IGACORE.JOB_PROPERTY_EXT ("JOB_PROPERTY") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_JOB_PROP_EXT_IDEAS_MODULE_FK ON
IGACORE.JOB_PROPERTY_EXT ("IDEAS_MODULE") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_JOB_UNIT_HIERARCHY_FK ON IGACORE.JOB_UNIT
("HIERARCHY") ALLOW REVERSE SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_JU_REV_ATTESTATION_FK ON
IGACORE.JOB_UNIT_REVIEW ("ATTESTATION") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_JU_REV_REVIEWED_BY_FK ON
IGACORE.JOB_UNIT_REVIEW ("REVIEWED_BY") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_JU_REVIEWER_FIRSTOWNER_FK ON
IGACORE.JOB_UNIT_REVIEWER ("CERT_FIRST_OWNER") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_JU_REVIEWER_SECONDOWNER_FK ON
IGACORE.JOB_UNIT_REVIEWER ("CERT_SECOND_OWNER") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_JU_REV_HIS_ENTITLEMENT_FK ON
IGACORE.JOB_UNIT_REVIEW_H ("ENTITLEMENT") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_JU_REV_HIS_ORG_UNIT_FK ON
IGACORE.JOB_UNIT_REVIEW_H ("ORGANIZATIONAL_UNIT") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_JU_REV_HIS_REVIEWED_BY_FK ON
IGACORE.JOB_UNIT_REVIEW_H ("CERT_FIRST_OWNER") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_JURNOTE_JOB_UNIT_REVIEWER_FK ON
IGACORE.JOB_UNIT_REVIEW_NOTE ("JOB_UNIT_REVIEWER") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_JURNOTE_PERSON_FROM_FK ON
IGACORE.JOB_UNIT_REVIEW_NOTE ("PERSON_FROM") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_JURNOTE_TO_FK ON
IGACORE.JOB_UNIT_REVIEW_NOTE ("PERSON_TO") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_MAPPING_FUNCTION_FK ON IGACORE.MAPPING
("FUNCTION") ALLOW REVERSE SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_MAPPING_MAPPING_OBJ_CLASS_FK ON
IGACORE.MAPPING ("MAPPING_OBJ_CLASS") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_MPS_FIELD_FK ON
IGACORE.MAPPING_SOURCE_FIELD ("SOURCE_FIELD") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_MPS_MAPPING_FK ON
IGACORE.MAPPING_SOURCE_FIELD ("MAPPING") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_CFG_PROACTIVITY_FK ON IGACORE.MN_LINK
("CFG_PROCESSACTIVITY") ALLOW REVERSE SCANS COLLECT DETAILED
```

```
STATISTICS;
CREATE INDEX IGACORE.IDX_MN_LINK_CFG_ACTIVITY_FK ON IGACORE.MN_LINK
("CFG_ACTIVITY") ALLOW REVERSE SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_MN_LINK_MN_ROLEMAPPING_FK ON
IGACORE.MN_LINK ("MN_ROLE_MAPPING") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_MN_ROLE_MAP_LOCALIZ_CODE_FK ON
IGACORE.MN_ROLE_MAPPING ("LOCALIZATION_CODE") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_MN_TREE_LOCALIZ_CODE_FK ON
IGACORE.MN_TREE ("LOCALIZATION_CODE") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_MN_TREE_MN_LINK_FK ON IGACORE.MN_TREE
("MN_LINK") ALLOW REVERSE SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_OBJECT_CLASS_PARENT_FK ON
IGACORE.OBJ_CLASS ("PARENT") ALLOW REVERSE SCANS COLLECT DETAILED
STATISTICS;
CREATE INDEX IGACORE.IDX_OCF_PARENT_FK ON IGACORE.OBJ_CLASS_FIELD
("PARENT") ALLOW REVERSE SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_ORG_UNIT_OWNER_FK ON
IGACORE.ORGANIZATIONAL_UNIT ("OWNER") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_ORG_UNIT_PARENT_FK ON
IGACORE.ORGANIZATIONAL_UNIT ("PARENT") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_OU_INCOMP_HIERARCHY_FK ON
IGACORE.ORGANIZATIONAL_UNIT_INCOMP ("HIERARCHY") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_OU_INCOMP_ORG_UNIT_FK ON
IGACORE.ORGANIZATIONAL_UNIT_INCOMP ("ORGANIZATIONAL_UNIT") ALLOW
REVERSE SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_ORG_UNIT_PROPERTY_OU_FK ON
IGACORE.ORG_UNIT_PROPERTY ("ORGANIZATIONAL_UNIT") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_PERSON_INCOMP_ENT_FK ON
IGACORE.PERSON_INCOMP ("PERSON") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_PER_PRO_SERV_SA_FK ON
IGACORE.PERSON_PROFILE_SERVICE ("SERVICE_ATTRIBUTE") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_PERSON_PWDSYNC_PERSON_FK ON
IGACORE.PERSON_PWDSYNC ("PERSON") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_PERSON_PWDSYNC_PWDPOLICY_FK ON
IGACORE.PERSON_PWDSYNC ("PWDPOLICY") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_UR_REV_PERSONREM_FK ON
IGACORE.PERSON_REM_REVIEW ("PERSON_REMEDIATION") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_UR_REV_PERSON_FK ON
```

```
IGACORE.PERSON_REM_REVIEW ("PERSON") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_UR_REV_REMEDIATION_FK ON
IGACORE.PERSON_REM_REVIEW ("REMEDIATION") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_UR_REV_REVIEWED_BY_FK ON
IGACORE.PERSON_REM_REVIEW ("REVIEWED_BY") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_UR_REV_RISK_FK ON
IGACORE.PERSON_REM_REVIEW ("ENVIRONMENT", "RISK") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_PR_REVIEWER_FIRSTOWNER_FK ON
IGACORE.PERSON_REM_REVIEWER ("CERT_FIRST_OWNER") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_PR_REVIEWER_SECONDOWNER_FK ON
IGACORE.PERSON_REM_REVIEWER ("CERT_SECOND_OWNER") ALLOW
REVERSE SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_UR_REV_HIS_PERSON_FK ON
IGACORE.PERSON_REM_REVIEW_H ("PERSON") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_UR_REV_HIS_REMEDIATION_FK ON
IGACORE.PERSON_REM_REVIEW_H ("REMEDIATION") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_UR_REV_HIS_REVIEWED_BY_FK ON
IGACORE.PERSON_REM_REVIEW_H ("CERT_FIRST_OWNER") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_PRRNOTE_PERSON_FROM_FK ON
IGACORE.PERSON_REM_REVIEW_NOTE ("PERSON_FROM") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_PRRNOTE_PERSON_TO_FK ON
IGACORE.PERSON_REM_REVIEW_NOTE ("PERSON_TO") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_PERSON_RISK_PERSON_FK ON
IGACORE.PERSON_RISK ("PERSON") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_PERSON_RISK_RISK_FK ON
IGACORE.PERSON_RISK ("ENVIRONMENT", "RISK") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_PERSON_SERVICE_PWDCFG_FK ON
IGACORE.PERSON_SERVICE ("PWDCFG") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_PERSON_SERVICE_PWDMGM_FK ON
IGACORE.PERSON_SERVICE ("PWDMANAGEMENT") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_PERSON_SERVICE_SERV_ATTR_FK ON
IGACORE.PERSON_SERVICE ("SERVICE_ATTRIBUTE") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_PSERVICEPENDING_PROFILE_FK ON
IGACORE.PERSON_SERVICE_PENDING ("PROFILE") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
```

```
CREATE INDEX IGACORE.IDX_PWDCFG_PWDPOLICY_FK ON IGACORE.PWDCFG
("PWDPOLICY") ALLOW REVERSE SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_PWDCFG_AK_LOCALIZATION_FK ON
IGACORE.PWDCFG_ATTR_KEY ("LOCALIZATION") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_PWDMANAGEMENT_PWDMAN_TYPE_FK ON
IGACORE.PWDMANAGEMENT ("PWDMANAGEMENT_TYPE") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_PWDPOLICY_PROPS_PWDPOLICY_FK ON
IGACORE.PWDPOLICY_PROPS ("PWDPOLICY") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_PWDPOLICY_WORDS_PWDPOLICY_FK ON
IGACORE.PWDPOLICY_WORDS ("PWDPOLICY") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_RECORD_LOG_EXECUTION_LOG_FK ON
IGACORE.RECORD_LOG ("EXECUTION_LOG") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_REMEDIATION_PROPERTY_REM_FK ON
IGACORE.REMEDIATION_PROPERTY ("REMEDIATION") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_REQUEST_CFG_PRIORITY_FK ON
IGACORE.REQUEST ("CFG_PRIORITY") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_REQUEST_ESC_PROCESSACTIVITY_FK ON
IGACORE.REQUEST ("ESCALATION_PROCESSACTIVITY") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_REQUEST_PROCESSACTIVITY_FK ON
IGACORE.REQUEST ("PROCESSACTIVITY") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_REQUEST_SUBPROCESSACTIVITY_FK ON
IGACORE.REQUEST ("SUB_PROCESSACTIVITY") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_REQUEST_NOTIFICATION_PA_N_FK ON
IGACORE.REQUEST_NOTIFICATION ("PA_NOTIFICATION") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_REQUEST_NOTIFICATION_PA_W_FK ON
IGACORE.REQUEST_NOTIFICATION ("PA_WORKING") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_RESOURCES_REQUEST_FK ON
IGACORE.RESOURCES ("REQUEST") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_RISK_OWNER_FK ON IGACORE.RISK ("OWNER")
ALLOW REVERSE SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_RISK_RISK_TYPE_FK ON IGACORE.RISK
("ENVIRONMENT", "RISK_TYPE") ALLOW REVERSE SCANS COLLECT DETAILED
STATISTICS;
CREATE INDEX IGACORE.IDX_SERVICE_ROLES_FK ON IGACORE.SERVICE
("ROLES") ALLOW REVERSE SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_SATTR_PROFILE_FK ON
IGACORE.SERVICE_ATTRIBUTE ("PROFILE") ALLOW REVERSE SCANS COLLECT
```

```
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_SATTR_PWDCFG_FK ON
IGACORE.SERVICE_ATTRIBUTE ("PWDCFG") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_SUPERVISOR_ATTESTATION_FK ON
IGACORE.SUPERVISOR ("ATTESTATION") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_SUPERVISOR_PERSON_FK ON
IGACORE.SUPERVISOR ("PERSON") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_SURVEYCFGEXT_SURVEY_CONFIG_FK ON
IGACORE.SURVEY_CONFIG_EXT ("SURVEY_ID") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_SURVEY_PAGES_SURVEY_CONFIG_FK ON
IGACORE.SURVEY_PAGES ("SURVEY_ID") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_SURVEY_USER_SURVEY_CONFIG_FK ON
IGACORE.SURVEY_USER ("SURVEY_ID") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_TARGET_PWDCFG_FK ON IGACORE.TARGET
("PWDCFG") ALLOW REVERSE SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_TARGET_AUTH_CONNECTOR_FK ON
IGACORE.TARGET_AUTH ("CONNECTOR") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_TARGET_RES_SCHEMA_DEF_TRSID_FK ON
IGACORE.TARGET_RESOURCE_SCHEMA_DEF ("TR_SCHEMA_ID") ALLOW
REVERSE SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_TAR_UID_LOCK_USER_EV_ERC_FK ON
IGACORE.TARGET_UID_LOCK ("LAST_EVENT_ID") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_TASK_OWNER_FK ON IGACORE.TASK ("OWNER")
ALLOW REVERSE SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_TASK_PARENT_FK ON IGACORE.TASK
("ENVIRONMENT", "PARENT") ALLOW REVERSE SCANS COLLECT DETAILED
STATISTICS;
CREATE INDEX IGACORE.IDX_TASK_ENTITLEMENT_ENTITL_FK ON
IGACORE.TASK_ENTITLEMENT ("ENTITLEMENT") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_TASK_GROUP_ORG_UNIT_FK ON
IGACORE.TASK_GROUP ("ORGANIZATIONAL_UNIT") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_TASK_PERSON_PERSON_FK ON
IGACORE.TASK_PERSON ("PERSON") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_TEMPLATE_ENTITY_APPLICATION_FK ON
IGACORE.TEMPLATE_ENTITY ("APPLICATION") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_TEMPLATE_ENTITY_ENTITLEMENT_FK ON
IGACORE.TEMPLATE_ENTITY ("ENTITLEMENT") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
```

```
CREATE INDEX IGACORE.IDX_TEMPLATE_ENTITY_OU_FK ON
IGACORE.TEMPLATE_ENTITY ("ORGANIZATIONAL_UNIT") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_TEMPLATE_ENTITY_PERSON_FK ON
IGACORE.TEMPLATE_ENTITY ("PERSON") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_TEMPLATE_ENTITY_PWDCFG_FK ON
IGACORE.TEMPLATE_ENTITY ("PWDCFG") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_TEMPLATE_ENTITY_PWDM_FK ON
IGACORE.TEMPLATE_ENTITY ("PWDMANAGEMENT") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_TEMPLATE_ENTITY_RISK_FK ON
IGACORE.TEMPLATE_ENTITY ("ENVIRONMENT", "RISK") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_TEMPLATE_ENTITY_TEMPLATE_FK ON
IGACORE.TEMPLATE_ENTITY ("TEMPLATE") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_TEMPLATE_PROPERTY_FLOW_FK ON
IGACORE.TEMPLATE_PROPERTY ("FLOW") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGACORE.IDX_TEMPLATE_PROPERTY_TEMPLATE_FK ON
IGACORE.TEMPLATE_PROPERTY ("TEMPLATE") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;

commit work;
```

Table: create_indexes_igacore.ddl

d create_indexes_igaserv.ddl:

```
CREATE INDEX IGASERV.IDX_CFG_CONFIG_PROP_CFG_APP_FK ON
IGASERV.CFG_CONFIG_PROP ("CFG_APPLICATION") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGASERV.IDX_TIMER_FAMILY_FK ON IGASERV.TIMER ("REALM",
"FAMILY") ALLOW REVERSE SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGASERV.IDX_TIMER_HISTORY_TIMER_WORK_FK ON
IGASERV.TIMER_HISTORY ("REALM", "TIMER_WORK") ALLOW REVERSE SCANS
COLLECT DETAILED STATISTICS;
CREATE INDEX IGASERV.IDX_TIMER_WORK_PARENT_FK ON
IGASERV.TIMER_WORK ("REALM", "PARENT") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGASERV.IDX_TIMER_WORK_TIMER_FK ON
IGASERV.TIMER_WORK ("REALM", "TIMER") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGASERV.IDX_TIMER_WORK_WORK_FK ON
IGASERV.TIMER_WORK ("REALM", "WORK") ALLOW REVERSE SCANS COLLECT
DETAILED STATISTICS;
CREATE INDEX IGASERV.IDX_WORK_CLASS_NAME_FK ON IGASERV.WORK
("REALM", "CLASS_NAME") ALLOW REVERSE SCANS COLLECT DETAILED
```

```
STATISTICS;
CREATE INDEX IGASERV.IDX_CFG_APPLICATION_PROP_APP_FK ON
IGASERV.CFG_APPLICATION_PROP ("CFG_APPLICATION") ALLOW REVERSE
SCANS COLLECT DETAILED STATISTICS;
CREATE INDEX IGASERV.IDX_TIMER_SHEDULER_FK ON IGASERV.TIMER
("SCHEDULER") ALLOW REVERSE SCANS COLLECT DETAILED STATISTICS;

commit work;
```

create_indexes_igaserv.ddl

e latest_findings.ddl:

The following indexes were identified based on the analysis from the performance runs, as well as from the snapshots collected from the customer related issues.

```
CREATE INDEX IGACORE.CHANGELOG_PWDMANAGEMENT_IDX ON
IGACORE.CHANGELOG (PWDMANAGEMENT ASC) ALLOW REVERSE SCANS
COLLECT SAMPLED DETAILED STATISTICS;

CREATE INDEX IGAQRZ.IDX_QRZ7_T_NEXT_FIRE_TIME_ONLY ON
IGAQRZ.QRZ7_TRIGGERS(NEXT_FIRE_TIME) ALLOW REVERSE SCANS COLLECT
STATISTICS;

CREATE INDEX IGACORE.PERSON_PROFILE_LF1_IDX ON "IGACORE
"."PERSON_PROFILE" (LOWER(ACCOUNT_CODE) ASC,
LOWER(PERMISSION_CODE) ASC, LOWER(PERMISSION_TYPE) ASC,
LOWER(TARGET) ASC) COMPRESS NO INCLUDE NULL KEYS ALLOW REVERSE
SCANS;

CREATE INDEX IGACORE.PERSON_LOWER_DN ON IGACORE.PERSON
(LOWER(DN) ASC) ALLOW REVERSE SCANS COLLECT SAMPLED DETAILED
STATISTICS;

CREATE INDEX IGACORE.PERSON_UPPER_CODE ON IGACORE.PERSON
(UPPER(CODE) ASC) ALLOW REVERSE SCANS COLLECT SAMPLED DETAILED
STATISTICS;

CREATE INDEX IGACORE.USERERC_UPPER_PMCODE ON IGACORE.USER_ERC
(UPPER(PM_CODE) ASC) ALLOW REVERSE SCANS COLLECT SAMPLED DETAILED
STATISTICS;

CREATE INDEX IGACORE.PWDCFG_LOWER_VALUE ON IGACORE.PWDCFG
(LOWER(VALUE) ASC) ALLOW REVERSE SCANS COLLECT SAMPLED DETAILED
STATISTICS;

CREATE INDEX IGACORE.PWDMANAGEMENT_CODE ON
IGACORE.PWDMANAGEMENT ("CODE" ASC) ALLOW REVERSE SCANS COLLECT
SAMPLED DETAILED STATISTICS;
```

```
CREATE INDEX EVENT_TARGET_SOS_IDX ON
IGACORE.EVENT_TARGET(SYNC_STATUS, OPERATION, STATE) COLLECT
DETAILED STATISTICS;

CREATE INDEX EVENT_TARGET_SYST_IDX ON IGACORE
EVENT_TARGET(SYNC_STATUS ASC, STATE DESC) ALLOW REVERSE SCANS
COLLECT SAMPLED DETAILED STATISTICS;

CREATE INDEX EVENT_TARGET_LTLC_IDX ON
IGACORE.EVENT_TARGET(LOWER(TARGET), LOWER(CODE)) COLLECT DETAILED
STATISTICS;

CREATE INDEX DID_IDX ON IGACORE.EVENT_OUT (DISPATCHER_ID ASC) ALLOW
REVERSE SCANS COLLECT SAMPLED DETAILED STATISTICS;
```

Apply the indexes by running the following commands as the database Instance owner:

```
db2 connect to <db_name>
db2 -tvf <filename>.ddl
db2commit
db2 disconnect current
db2stop
db2start
```

Enable Monitoring

It is recommended to enable monitoring for any DB2 servers because it allows the system administrator with the ability to tune DB2 to obtain maximum performance. Perform the following actions as the DB2 instance owner in order to enable monitoring in ISDS DB:

```
db2 connect to <itimdb> user <user> using <password>
db2 update dbm cfg using DFT_MON_STMT ON
db2 update dbm cfg using DFT_MON_BUFPOOL ON
db2 update dbm cfg using DFT_MON_LOCK ON
db2 update dbm cfg using DFT_MON_SORT ON
db2 update dbm cfg using DFT_MON_TIMESTAMP ON
db2 update dbm cfg using DFT_MON_UOW ON
db2 disconnect current
```

DB2_ANTIJOIN=EXTENDED

DB2_ANTIJOIN=EXTENDED causes the optimizer to equally consider subqueries of a NOT EXISTS clause. This parameter can benefit most queries with a NOT IN clause and a NOT EXISTS clause. You can identify all the queries in your environment using these clauses and validate the benefits of this variable for your specific environment.
When this variable is set to YES, the optimizer searches for opportunities to transform NOT EXISTS subqueries into anti-joins which can be processed more efficiently by DB2.
When this variable is set to EXTEND, the optimizer searches for opportunities to transform both NOT IN and NOT EXISTS subqueries into anti-joins. It is recommended to have the value of DB2_ANTIJOIN=EXTENDED.

Instance_Memory=Automatic

This parameter specifies the maximum amount of memory that can be allocated for a database partition. The DB2 database products can be set with the memory usage restrictions or to a specific value. The AUTOMATIC setting results in a value that is computed at database partition activation. The computed value ranges between 75 percent and 95 percent of the system memory capacity on the system – the larger the system, the higher the percentage. The AUTOMATIC settings allow Instance Memory usage to grow as needed. If STMM is enabled and tunes the overall database memory size, STMM tunes based on available system memory. This option indirectly determines the actual instance memory usage. To ensure that enough memory is available to process long running or large request in DB2 it is recommended to have the value of instance_memory to automatic.

Enable NTP

Network Time Protocol (NTP) is used to synchronize computer clocks on a network of information. The use of NTP server ensures that the time is observed accurately between the servers and the internet. It is recommended that NTP is enabled in all the IGI components IGIVA, and IGI DB2 so the clock is synchronized.

Steps to Move DB2 Transaction Log to Separate Disk

We used the following steps to move the transaction log to the separate disk. Perform the following steps in the Virtual Machine (VM) environment.

For DB2 Virtual Machine (VM)

- Stop IGI Application Server

- Power Down DB2 VM
- Edit settings, add hard drive
    - 100 GB size and “specify a datastore” that does not have the DB VM on it
- Power on DB2 MV and open console.

After completing the above steps follow the following steps based on the operating system on the DB2 server.

SUSE Linux Enterprise Server (SLES):

- yast2 → system → partitioner → ‘Yes’ to warning
- Highlight /dev/sdb (100GB), right click → add partition
- Add New Primary Partition
- Click Max Size: 100GB
- Format File System: Ext3
- Mount /igi/tranlog, ignore FSTAB options
- Finish
- Next → Finish on yast2 window, close yast2 window

REDHAT 7

Before adding the disk datastore in VM, or in the bare metal box the command ls /dev/sd* gives the following output:

After the extra dis is added the command ls /dev/sd* gives the following output:

```
/dev/sda   /dev/sda1   /dev/sda2
```

```
/dev/sda   /dev/sda1   /dev/sda2     /dev/sdb
```

Use the following steps after the state of the extra disk is verified:

- Run: fdisk /dev/sdb
    - Command (m for help): c
      DOS Compatibility flag is not set
    - Command (m for help): u
      Changing display/entry units to sectors
    - Command (m for help): n
    - Command action
      e extended

After completing the above steps run ls /dev/sd* to verify the changes were made. The output shows show as:

```
/dev/sda   /dev/sda1   /dev/sda2     /dev/sdb       /dev/sdb1 (new)
```

After verifying the changes create the file system by running the following command:

```
/sbin/mkfs.ext4 -L /backup /dev/sdb1
```

After completing the above step, use the following steps to create the directory /igi/tranlog on the newly created disk:

- mkdir -p /igi/tranlog
- mount /dev/sdb1 /igi/tranlog
- Mount shows:
  - .....
  - /dev/sdb1 on /igi/tranlog type ext4 (rw)
- Open /etc/fstab and add:
  - /dev/sdb1 /igi/tranlog ext4 acl,user_xattr 1 2
- chmod -R 775 /igi
- chown -R igiinst:igiigrp /igi
  - ls -altr /igi

    total 12

    drwxrwxr-x 3 igiinst igiigrp 4096 Mar 14 09:03 tranlog

    dr-xr-xr-x. 30 root root 4096 Mar 14 10:44 ..

    drwxrwxr-x 3 igiinst igiigrp 4096 Mar 14 10:44

This will create a new disk with the directory in REDHAT systems called /igi/tranlog. For more information about disk partition visit: https://www.techotopia.com/index.php/Adding_a_New_Disk_Drive_to_an_RHEL_6_System

DB2 commands to move the Transaction Log

After completing the steps explained in the previous section to add the additional disk based on the operating system SLES or Red hat, follow the follow the following steps to change the transaction log to the newly created disk directory.

- In terminal window as root, chmod -R 775 /igi, chown -R db2inst1.<group> /igi
- su - db2inst1
- . ~/sqllib/db2profile
- db2start
- db2 connect to igidb
- db2 get db cfg | grep "Path to log files"
- db2 update db cfg using newlogpath /igi/tranlog
- db2 get db cfg | grep "Path to log files" # to verify change took place
- sync
- db2stop force
- db2start (this will stop and start the DB to make the change effective)
- <take snapshot>
- Restart VA app server

DB2 Performance Optimization

Updating Table and Index Statistics

In order to make use of the most efficient access plan for executing queries, DB2 requires accurate statistics on the exact number of rows in the tables and available indexes. Current DB2 versions can update the statistics automatically, and we recommend manually updating the statistics in certain situations after substantial changes have been made to the system.
These situations include:

- The capturing of significant number of multi-session recordings
- Creation, modification, deletion of users, resources, credentials, etc.
- After an extended period of DB2 operations without updating table statistics

It is recommended to execute runstats/reorg utilities on an idle or lightly used database because it requires update locking on the system statistics

table. The system acquires locks on the tables that are used by the database optimizer to fulfill queries. The locks might cause transaction rollbacks on a database with a heavy load. Also, it might be necessary to stop the directory server in order to complete runstats/reorg.

Disk I/O Performance Recommendations

Performance variables can be set in DB2 in order to improve DB2 related processes such as:

- Operating Resource Policies
- Memory Tuning
- Access Plan Optimization

DB2 Registry Variables

Environment variables can be set in DB2 in order to modify configuration values of particular applications.

DB2 System Environment Variables

Due to the nature of PIM and ISIM related transactions (latency sensitive, I/O bound), IBM Security Performance recommends the Data Tier located on physical machines configured for speed and fault tolerance (RAID 5/RAID 10). DB2 registry variables can be modified to improve performance on the Data Tier related systems. For all systems, enable DB2_USE_ALTERNATE_PAGE_CLEANING. This variable specifies whether a DB2 database uses the alternate method of page cleaning algorithms or the default method of page cleaning. When this variable is set to ON, the DB2 system writes changed pages to disk, keeping ahead of LSN_GAP and pro-actively finding victims. Doing this allows the page cleaners to better utilize available disk I/O bandwidth. When this variable is set to ON, the chngpgs_thresh database configuration parameter is no longer relevant because it does not control page cleaner activity.
As the instance owner:
db2set DB2_USE_ALTERNATE_PAGE_CLEANING=ON
(instance restart required)

For SAN, RAID, and other advanced disk subsystems set the system environment variable to DB2_PARRALLEL_IO to *. This registry variable is used to change the way DB2 calculates the I/O parallelism of a table space. When I/O parallelism is enabled (either implicitly, by the use of multiple containers, or explicitly, by setting DB2_PARALLEL_IO), it is achieved by issuing the correct number of prefetch requests. Each prefetch request is a request for an extent of pages. For example, a table space has two containers and the prefetch size is four times the extent size. If the registry variable is set, a prefetch request for this table space will be broken into four requests (one extent per request) with a possibility of four prefetchers

servicing the requests in parallel. You can replace TablespaceID with an asterisk (*) to specify all table spaces. For example, if DB2_PARALLEL_IO=*, all table spaces use six as the number of disks per container.
As the instance owner:
*db2set DB2_PARALLEL_IO=**
(instance restart required)

DB2_ANTIJOIN=EXTENDED

DB2_ANTIJOIN=EXTENDED causes the optimizer to equally consider subqueries of a NOT EXISTS clause. This parameter can benefit most queries with a NOT IN clause and a NOT EXISTS clause. You can identify all the queries in your environment using these clauses and validate the benefits of this variable for your specific environment.
When this variable is set to YES, the optimizer searches for opportunities to transform NOT EXISTS subqueries into anti-joins which can be processed more efficiently by DB2.
When this variable is set to EXTEND, the optimizer searches for opportunities to transform both NOT IN and NOT EXISTS subqueries into anti-joins. It is recommended to have the value of DB2_ANTIJOIN=EXTENDED.

As the instance owner:
db2set DB2_ANTIJOIN=EXTENDED
(instance restart required)

Notices

This information was developed for products and services offered in the U.S.A. IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.
IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785 U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:
Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
1623-14, Shimotsuruma, Yamato-shi
Kanagawa 242-8502 Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:
INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION
"AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.
Some states do not allow disclaimer of express or implied warranties in certain transactions; therefore, this statement might not apply to you. This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.
Any references in this information to non-IBM 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 use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you. Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:
IBM Corporation
11501 Burnet Road
Austin, TX 78758 U.S.A.
Such information may be available, subject to appropriate terms and conditions, including in some cases payment of a fee. The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.
Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating

environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.
Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.
All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice and represent goals and objectives only.
All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.
This information is for planning purposes only. The information herein is subject to change before the products described become available.
This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

Copyright License

If you are viewing this information in softcopy form, the photographs and color illustrations might not be displayed.

Trademarks

- IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at Copyright and trademark information at www.ibm.com/legal/copytrade.shtml.
- Adobe, Acrobat, PostScript and all Adobe-based trademarks are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, other countries, or both.
- Linux is a trademark of Linus Torvalds in the United States, other countries, or both.
- Microsoft, Windows, Windows Server 2008, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
- UNIX is a registered trademark of The Open Group in the United States and other countries.
- Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and its affiliates.
- Advanced Visualization Powered by IBM ILOG Elixir Enterprise.
- Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
- ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce and is registered in the U.S. Patent and Trademark Office. IT Infrastructure Library (ITIL) is a registered trademark of the Central Computer and Telecommunications Agency which is now part of the Office of Government Commerce.
- Cell Broadband Engine and Cell/B.E. are trademarks of Sony Computer Entertainment, Inc., in the United States, other countries, or both and is used under license therefrom.

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