



KPI Metrics Configuration Guide

An Open Source Asset for use with TIBCO® Data Virtualization

TIBCO Software empowers executives, developers, and business users with Fast Data solutions that make the right data available in real time for faster answers, better decisions, and smarter action. Over the past 15 years, thousands of businesses across the globe have relied on TIBCO technology to integrate their applications and ecosystems, analyze their data, and create real-time solutions. Learn how TIBCO turns data—big or small—into differentiation at www.tibco.com.

Project Name	AS Assets KPI Metrics
Document Location	This document is only valid on the day it was printed. The source of the document will be found in the ASAssets_KPI folder (https://github.com/TIBCOSoftware)
Purpose	Self-paced instructional



www.tibco.com

Global Headquarters
3303 Hillview Avenue
Palo Alto, CA 94304

Tel: +1 650-846-1000
+1 800-420-8450
Fax: +1 650-846-1005

Revision History

Version	Date	Author	Comments
1.0	Oct 2012	Manny Bhatia	Initial revision
1.4	Nov 2012	Matthew Lee	Minor updates to deployment language to clarify steps
1.5	Nov 2012	Matthew Lee	Revised setup and configuration steps for clarification
1.6	Sep 2013	Matthew Lee	Revised document for KPImetrics version 2.0
1.7	Jan 2014	Matthew Lee	Migrated to Cisco documentation standards
1.8	Aug 2017	Mike Tinius	Upgraded to work with Data Virtualization metrics with Postgres. Add-on to DV metrics capability.
1.9 – 2017.4	Dec 2017	Mike Tinius	Add additional capabilities and migrated to Tibco documentation standards
2018.1	March 2018	Mike Tinius	Optimization of Oracle and SQL Server views to achieve better push-down. Improved installation scripts.
2018.101	March 30 2018	Mike Tinius	Fixed an issue with the installation scripts. Added KPImetrics_worksheet.xlsx.
2018.102	April 23 2018	Mike Tinius	Added ldap and metrics history “userkey” field to allow normalization of the join key. Fixed an issue with CPUemChecker.
2018.103	May 2 2018	Mike Tinius	Modified data transfer procedure. Added column CURRENT_OPERATION to METRICS_JOB_DETAILS table. Deprecated cache_status and cache_tracking usage.
2018.3	Sep 18 2018	Mike Tinius	Modified to allow datasource names in the format of KPI_oracle_11g, KPI_oracle_12c, KPI_sqlserver_2012 or KPI_sqlserver_2014 in order to allow for adding a datasource. Modified sqlParserV1_2 to allow for parallel processing of SQL when a cluster is present. A node can process rows other than its own rows when it has no work to do. Modified tables: METRICS_ALL_RESOURCES, METRICS_SQL_REQUEST Added tables: METRICS_SQL_CONTROL, METRICS_SQL_CONTROL_LOG Modified the native SQL data transfer script to use stage tables resulting in only inserts to the history tables to improve efficiency. Modified /shared/ASAssets/KPImetrics/Business/Logical/metrics *_hist views to use where starttime > TO_TIMESTAMP('1900-01-01 00:00:00.000') to force the use of parallel queries.
2018.301	Oct 2 2018	Mike Tinius	Modified parseSqlScriptComplex and parseSqlScriptTemplate to change template select from SELECT DISTINCT to SELECT TOP 1 to make it more efficient.
2018.302	Oct 12 2018	Mike Tinius	Fixed a collection table row volume problem by adding an independent delete trigger for the collection tables: metrics_requests and metrics_resources_usage.
2018.400	Dec 1 2018	Mike Tinius	Added processing for SQL request datasource lineage thus allowing reporting on what physical data source resources are related to which published resources. Requires Utilities 2018Q4.
2018.401	Dec 3 2018	Mike Tinius	Removed all reference and processing for METRICS_RESOURCES_USAGE_UD. This is essentially a copy of metrics_resources_usage_hist where resourcekind='user defined'. This is a simplification.
2019.100	Jan 7 2019	Mike Tinius	Fixed installation script. Added “domainkey” to METRICS_LDAP_PERSON and history tables to qualify the “user”. The same user can be in multiple domains.
2019.101	Jan 30 2019	Mike Tinius	Added support for DV 8.0 as the out-of-the-box metrics tables [metrics_requests and metrics_sessions] added user, domain and

			group. KPI installation automatically checks for DV version 7.0 or 8.0 and executes the appropriate scripts to install KPImetrics.
2019.102	Mar 12 2019	Mike Tinius	Added upgrade scripts. Added "group" to metrics_requests and metrics_sessions for 7.0 to bring up to par with 8.0.

Related Documents

Name	Version
How To Use Utilities.pdf	2018Q4

Supported Versions

Name	Version
TIBCO® Data Virtualization	7.0.4 or later
AS Assets Utilities open source	2018Q4 or later

Table of Contents

1	Introduction	6
	Purpose	6
	Audience	6
	References	6
	Overview	6
2	Requirements	7
3	New Installation and Configuration	10
	Supported Database Platforms	10
	Installing KPImetrics	10
	Turn off DV metrics	10
	Download and Import the KPImetrics components to your DV instance.....	10
	Configuration Overview	12
	[1.] Configure the KPImetrics data source.....	12
	[2.] Configure DV Email	13
	[3.] Configure Common Values.....	13
	[4.] Configure LDAP Data Source	15
	[5.] Configure Metrics Job Lookup Tables	18
	[6.] Configure KPImetrics Triggers	21
	Execute Post-Installation Script.....	23
	Execute Installation Script	23
	Deploy CPU and Memory Checker shell scripts (Windows and UNIX).....	26
	Enable Triggers.....	27
	Execute ClusterSafeCache Scripts [Cluster Only]	29
	Configure DV Out-of-the-box Metrics	29
	Information Only Section.....	31
	Create the KPImetrics storage tables for Oracle	32
	Create the KPImetrics storage tables for SQL Serer	34
	Common Configuration for all Databases.....	36
4	Upgrading KPImetrics.....	39
	Introduction.....	39
	How to Upgrade KPImetrics	39
	Current Version:	39
	New Version:.....	40
5	KPImetrics Administration Scenarios	46
	Turn KPI On/Off.....	46
	Turn Data Virtualization (DV) metrics On/Off	48
	Modify Triggers	48
	Perform Oracle Database Maintenance on Collection Tables	49
	Configure Third Party Tool Access	50
	Get the Current Row Distribution for the History Tables/Partitions	50
6	KPImetrics Resources	52
	Configuration Resources.....	52
	KPI Version Overview.....	52
	Configuration Folder Overview	52

- Published Resources 55
 - KPImetrics Catalog..... 55
- Data Sources 68
 - Metadata Data Source for LDAP..... 68
 - Metadata Data Source for CPUAndMemChecker..... 68
 - Metadata Table Relationship Diagram 70
 - Metadata Data Source Tables and Procedures for KPI_<database_type> 71
 - Metadata System Triggers and Load Scripts..... 81
 - Metadata System Helpers Scripts..... 89
 - Physical Oracle Data Transfer Script 90
 - Physical SQL Server Data Transfer Script 99
- 7 Release Notes..... 107**
 - Added or Modified in this Release 107
 - Release 2019Q102 [Mar 12 2019] 107
 - Release 2019Q101 [Jan 30 2019] 108
 - Release 2019Q100 [Jan 7 2019] 109
 - Release 2018Q401 [Dec 4 2018]..... 110
 - Release 2018Q400 [Dec 1 2018]..... 111
 - Release 2018Q302 [Oct 12 2018]..... 112
 - Release 2018Q301 [Oct 1 2018]..... 112
 - Release 2018Q3 [Sep 2018]..... 112
- 8 Appendix A – Partitioning Schemes..... 115**
 - Oracle Partition Scheme 115
 - SQL Server Partition Scheme 121

1 Introduction

Purpose

The purpose of this document is to provide guidance on how install, monitor and use the AS Assets KPI Metrics.

Audience

This document is intended to provide guidance for the following users:

- Data Virtualization Administrators – provides a guide for installation.
- Architects – provides the KPImetrics architecture.
- Data professionals – provides background on the published views and usage.
- Operations users – provides insight into triggers and procedures that are executed.
- Project Managers – provides general information on KPImetrics.

References

Product references are shown below. Any references to CIS or DV refer to the current TIBCO® Data Virtualization.

- TIBCO® Data Virtualization was formerly known as
 - Cisco Data Virtualization (DV)
 - Composite Information Server (CIS)

Overview

Please review the document “**KPImetrics Overview 1.20.pdf**”.

2 Requirements

The following requirements and pre-requisites must be met:

- **Minimum of Data Virtualization (DV) 7.0.8.00 HF01** is recommended as there have been DV Metrics bug fixes.
- TIBCO Open Source **ASAssets Utilities_2018Q4.zip** is required as a baseline. Follow the instructions within the zip file for installation.
- Use **KPImetrics worksheet.xlsx** as a way to gather the required information for installation.
- Acquire the LDAP connection information and credentials that will be used to setup the KPImetrics LDAP data source.
- A DV admin user will be created at installation time as “metrics_app_id@composite” so that resource ownership can be assigned to all /shared/ASAssets/KPImetrics and /services/databases/ASAssets/KPImetrics. Required for metrics data filter deletion.
- Database Schema for each DV environment to store the Metrics and KPImetrics data for a period of time specified by the user. Suggestions: 4 months in lower-level environments (LLE) and 13 months in production (PROD) environments.
 - [Oracle 11g or 12c, SQL Server 2012 or 2014]
 - Recommended database=Oracle
 - Recommended due to better push-down capabilities
 - Straight-forward partitioning support
 - Recommended schema (user) or database name=CIS_KPI
 - Must be granted privileges for the schema user given the following:
 - SELECT/INSERT/UPDATE/DELETE
 - CREATE/DROP/ALTER TABLE
 - CREATE/DROP/ALTER INDEX
 - CREATE/DROP PROCEDURE
 - CREATE/DROP SEQUENCE
 - Oracle:

- CREATE/ALTER SESSION
- RESOURCE
- EXECUTE SYS.DBMS_STATS.GATHER_TABLE_STATS
- CREATE/DROP/ALTER PARTITION
 - ALTER TABLE ADD PARTITION
 - ALTER TABLE DROP PARTITION
- SELECT SYS.ALL_TAB_PARTITIONS
- SQL Server:
 - UPDATE STATISTICS
 - CREATE/DROP/ALTER PARTITION
 - CREATE PARTITION FUNCTION
 - CREATE PARTITION SCHEME
 - ALTER TABLE SWITCH PARTITION
 - ALTER PARTITION SCHEME [NEXT USED]
 - ALTER PARTITION FUNCTION [SPLIT RANGE | MERGE RANGE]
 - TRUNCATE TABLE
 - SELECT sys.tables, sys.indexes, sys.partition_schemes, sys.partition_functions, sys.partitions, sys.partition_range_values
- Recommended Oracle Tuning
 - Increase the archiver timeout to 2 hours for each process.
 - Increase the deletion of the archive logs to 1 hour if errors start occurring referencing rollback segments or archive space issues.
 - Increase the UNDO tablespace to 50 GB or higher if errors start occurring regarding rollback segment too old.
 - Calibrate the I/O (dbms_resource_manager.calibrate_io).
- Recommended tablespaces or filegroups
 - Tablespace/Filegroup name=METRICS_DATA_COLL
 - This will contain the DV metrics collection tables [metrics_sessions, metrics_requests, metrics_resources_usage]
 - Configured to be a NOLOGGING tablespace so that it is more efficient for providing insert/delete operations every 2 hours without impacting performance. Rows will be transferred via to the history tables which are stored in the tablespace METRICS_DATA.
 - Tablespace/Filegroup name=METRICS_DATA_HIST
 - This will contain the KPI metrics tables.

- Configured with logging and large enough to hold 100 GB and grow as needed.
- Tablespace/Filegroup name=METRICS_DATA_IDX
 - This will contain the KPImetrics indexes.
 - Configured with logging and large enough to hold 100 GB and grow as needed.

3 New Installation and Configuration

Supported Database Platforms

The majority of metrics discussed in the previous sections are generated using custom aggregation procedures. Because DV does not retain the system metrics data needed to generate KPI metrics data long enough for historical reporting, the KPI metrics module must store this cached data to a dedicated database in order to retain the generated results.

The KPI metrics module supports the following database platforms at this time as incremental caching targets.

- Oracle 11g, 12c or later (recommended due to better push-down capabilities)
- SQL Server 2012 or 2014

Support for additional platforms would require customization of the KPI metrics module by a professional services consultant. Please contact TIBCO's professional service group for details.

Please note that it is strongly recommended that the database chosen to cache KPI metrics data have case sensitivity and ignore trailing space settings that match your DV server to maximize query pushdowns in order to minimize the amount of additional load the KPI metrics module adds to your DV environment.

Installing KPI metrics

Turn off DV metrics

Turn off DV metrics if it is running.

1. Use DV Studio and open /policy/metrics under the root folder at the bottom of Studio
2. If metrics is "enabled", uncheck the box to turn it off.

Download and Import the KPI metrics components to your DV instance

Deploy the KPI metrics components to your DV instance in order to use the KPI metrics module.

3. **Download** the ASAssets Utilities and KPI metrics from the TIBCO Open Source GIT site
 - a. Utilities: [Utilities_2018Q4.zip](#)
 - i. Follow the Utilities documentation "[How To Use Utilities.pdf](#)" for installation. Do this first.
 - ii. Import Utilities_2018Q4.car
 - iii. **IMPORTANT** – Configure the environment procedure uniquely for each DV

environment: /shared/ASAssets/Utilities/environment/getEnvName

1. Example: DEV, TEST, PROD
- b. KPImetrics: [KPImetrics_2019Q102.zip](#)
 - i. KPImetrics_2019Q102_installation.car
 - ii. KPImetrics_2019Q102.car
 - iii. KPImetrics Configuration Guide v1.21.pdf
 - iv. KPImetrics Overview.pdf
 - v. KPImetrics_worksheet.xlsx
 - vi. KPImetrics_scripts directory

Complete the following steps to configure the KPImetrics components

1. **Upload CAR file:**

- a. Upload [KPImetrics_YYYYQnnn.zip](#) to the target DV server file system and record the full path for use during installation.

2. **Login to DV Studio as “admin”**

- a. All configuration operations should be completed as DV “admin”.

3. **Import KPImetrics Installation CAR file:**

- a. In the Studio left resource panel tree, right click on the root folder (/) icon and select Import.
- b. Import the file [KPImetrics_YYYYQnnn_installation.car](#) with the overwrite checkbox enabled. The folders /shared/ASAssets/KPImetrics_installation and should appear after the import completes.

4. **Execute Pre-Installation Script:**

- a. Execute “[1_Pre_Installation](#)” and provide parameters:
 - i. Location: /shared/ASAssets/KPImetrics_installation/1_Pre_Installation
 - ii. IN **metrics_app_id_password** - DV password for the user/owner of KPImetrics source code which is “metrics_app_id”.
 - iii. IN **car_file_os_full_path** - Full path to the car file archive in the OS file system. If null the import is skipped and the KPImetrics_YYYYQnnn.car will need to be imported manually.
- b. Information Only Section
 - i. This script performs the following operations:
 - ii. Create the published data source “ASAssets” if it does not already exist.
 - iii. Create the “KPImetrics” catalog for ASAssets data source.
 - iv. Create a DV user called “metrics_app_id” in the “composite” domain. This way to can determine the process id that is executing requests and

filter these requests out of the metrics history tables if you choose, using a strategy to be discussed later.

- v. Import KPImetrics_YYYYQnnn.car if path is provided.

5. **Import KPImetrics CAR File:** [Optional if not done in step 4.]
 - a. **Bypass this step if you provided the path in Step 4.** above and the car file was successfully imported. Otherwise proceed with the instructions below.
 - b. In the Studio left resource panel tree, right click on the root folder (/) icon and select Import.
 - c. Import the file **KPImetrics_YYYYQnnn.car** with the overwrite checkbox enabled. The folders /shared/ASAssets/KPImetrics and /services/databases/ASAssets/KPImetrics appear after the import completes.

Configuration Overview

The following is an overview of resources to be configured prior to post-installation execution.

Recommended approach: Use **KPImetrics worksheet.xlsx** as a way to gather the required information for this section.

1. [Configure KPImetrics Data Source](#)
2. [Configure DV Email](#)
3. [Configure Common Values Procedure](#)
4. [Configure LDAP Data Source](#)
5. [Configure Metrics Job Lookup Tables](#)
6. [Configure KPImetrics Triggers](#)

[1.] Configure the KPImetrics data source

The KPImetrics module makes use of several custom tables to store logging and metrics data. You must configure a data source connection in order to view KPImetrics data.

1. **Configure KPImetrics data source:**
 - a. Locate and configure the appropriate data source for your KPImetrics database.
 - i. **NOTE:** The recommended user is CIS_KPI but any user will work fine.
 - ii. **Oracle: Recommended schema (user)=CIS_KPI**
 1. /shared/ASAssets/KPImetrics/Physical/Metadata/KPI_oracle_11g
 2. /shared/ASAssets/KPImetrics/Physical/Metadata/KPI_oracle_12c
 - iii. **SQL Server: Recommended database name=CIS_KPI**
 1. /shared/ASAssets/KPImetrics/Physical/Metadata/KPI_sqlserver_2012

2. `/shared/ASAssets/KPImetrics/Physical/Metadata/KPI_sqlserver_2014`
 - b. Enable the data source that is required.
 - c. Test the connection to make sure it is working properly.
 - d. Disable the data sources not being used.
 - e. If you used a schema name or database name different than CIS_KPI, the post installation script will automatically take care of this.

[2.] Configure DV Email

The KPImetrics module uses sends alerts and notifications.

2. Configure DV email:

- a. Select Administration → Configuration → Server → Configuration → E-mail
 - i. From Address
 - ii. SMTP Host Name
 - iii. SMPT Port
 - iv. If required, provide user name and password

[3.] Configure Common Values

The KPImetrics module uses several constant values that are set in the procedure `/shared/ASAssets/KPImetrics/Configuration/commonValues`. You will need to update some of these constants with values for your environment to ensure that KPImetrics functions correctly. Update the following minimum values for operation.

1. Configure Common Values:

- a. Open the procedure `/shared/ASAssets/KPImetrics/Configuration/commonValues` and modify the following properties:
- b. **cisServerNickname** – This uses the Utilities: `/shared/ASAssets/Utilities/environment/getEnvName()` procedure. Configure `getEnvName()` with the DV server nickname for the instance of DV. E.g. DEV1, SIT1, UAT1, PROD1. Alternatively, a static string could be used as well. This is used in email notifications to alert you of issues. You determine what your environment nicknames are. Be consistent. *Each DV server environment must be unique*. For a DV cluster, each node in the cluster will have the same nickname as it describes the environment and not the node.
- c. **replyTo** – Provide a single email as a reply to.
- d. **sendTo** – Provide the comma separated list of email address to send notifications to. This setting is used as a fall back if there are no METRICS_EVENT_REGISTRATION rows configured and, in the event, that the database is completely down. Note: You must configure email in DV.

- e. **CpuCheckerCommandPath** – Determine which script format to use based on your installation environment.
 - i. The following section stores CPU and memory checker default values for WINDOWS and UNIX. The following are the command line execution statements:
 - ii. CPU Utilization - processor time percentage
 - iii. Windows 7 or 10: 'powershell.exe -file '||getCisHome()|'|
'bin\KPImetricsCpuUtilization.ps1'
 - iv. UNIX (Linux 6): getCisHome()|'|/bin/KPImetricsTopCommandGrepCpu_linux6.sh'
 - v. UNIX (Linux 7): getCisHome()|'|/bin/KPImetricsTopCommandGrepCpu_linux7.sh'
- f. **memoryCheckerCommandPath** – Determine which script format to use based on your installation environment.
 - i. Memory Utilization - Memory used and memory available
 - ii. Windows 7 or 10: 'powershell.exe -file '||getCisHome()|'|
'bin\KPImetricsMemUtilization.ps1'
 - iii. UNIX (Linux 6): getCisHome()|'|/bin/KPImetricsFreeMemCommand_linux6.sh'
 - iv. UNIX (Linux 7): getCisHome()|'|/bin/KPImetricsFreeMemCommand_linux7.sh'
- g. **dataSourceName** – Provide the name of the data source used to store KPImetrics data. Valid values are KPI_oracle_11g, KPI_oracle_12c, KPI_sqlserver_2012 or KPI_sqlserver_2014.
- h. **dataSourceCatalog** – Identifies the data source catalog name if applicable [SQL Server]. Set to null if not applicable.
- i. **dataSourceSchema** – Identifies the data source schema name [Oracle, SQL Server].
- j. **collectionTablespaceName** – Identifies the Metrics collection tablespace name which will contain the DV out-of-the-box metrics tables. To use the default tablespace or no tablespace set this value to NULL. You should have already created a tablespace (with no logging if Oracle) or a filegroup if SQL Server.
- k. **historyTablespaceName** – Identifies the Metrics history tablespace name which will contain all of the history and reporting tables. To use the default tablespace or no tablespace set this value to NULL. You should have already created a tablespace or a filegroup if SQL Server.
- l. **indexTablespaceName** – Identifies the Index tablespace name which will contain the indexes for the history tables. To use the default tablespace or no tablespace set this value to NULL. You should have already created a tablespace or a filegroup if SQL Server.
- m. **historyTableCompression** – Identifies the type of compression that is allowed for the history tables [metrics_sessions_hist, metrics_resources_usage_hist, metrics_requests_hist]. Applied on table creation only.
 - i. Oracle:
 1. NOCOMPRESS - used for an Oracle database that does not support compression

2. COMPRESS FOR QUERY HIGH - used for an Oracle Exadata platform.
 3. COMPRESS - basic compression.
 - ii. SQL Server: parameters is ignored and should be set to null
- n. **partitionNumber** – Identifies the number of table partitions to use for the metrics history tables. The default is 0 indicating no partitioning will be used. It is “highly” recommended to partition the three history tables as it will make queries more performant and make it easier to drop a partition to purge data rather than executing delete statements. Recommendation: Always add 1 additional month to your standard data retention policy to allow for dropping a partition of data.
- i. (LLE) lower-level env: policy=3. Set to 4 months of data
 - ii. (PROD) production-level env: policy=12. Set to 13 months of data.
- o. **partitionStartDate** – Identifies the starting date for the partitioning scheme in the format 'YYYY-MM-DD'. If null then no partitioning date is used. If partitionNumber is > 0 and partitionStartDate is null an error is thrown. If the partition format is not correct and it cannot be cast to a DATE then an error is thrown. This should be the first of the month in which metrics are turned on and capturing data. If metrics have been previously turned on and have been capturing data for a while then use the first of the month of the date they were turned on. Execute the following query against the current metrics table to determine the date: `select min(starttime) starttime from metrics_requests`
- p. **dataTransferBatchInsertMax** – Identifies the insert batch maximum number for the P_METRICS_ALL_TABLES data transfer procedure. This identifies the insert batch size commit when transferring data from the metrics_requests collection table into the metrics_requests_stg stage table. Note: This is currently only implemented for Oracle..
- q. **dataTransferBatchUpdateMax** – Identifies the batch maximum number for the P_METRICS_ALL_TABLES data transfer procedure. After the metrics_requests collection table is transferred into the metrics_requests_hist table, an update routine executes to correlate the resourcekind and dataservicename from the metrics_resources_usage_hist with the metrics_requests_hist table. This routine commits rows from the metrics_requests_stg table for the batch number specified here. Note: This is currently only implemented for Oracle and is set to 50000.

[4.] Configure LDAP Data Source

The KPImetrics module is designed to retrieve user data from an LDAP directory server in order to provide additional detail on which users are making use of a monitored DV environment. You may configure the LDAP data source provided with the KPImetrics module to connect to your corporate LDAP directory server. The essence of many of the queries is based on LDAP data and LDAP users who connect to DV to execute queries. If you do not

have LDAP, then consider the options below.

1. Configure the LDAP data source

- a. **Option 1** – Configure the KPImetrics LDAP data source to connect to your local ldap: /shared/ASAssets/KPImetrics/Physical/Metadata/LDAP1.
 - i. Proceed to #2 below to map LDAP_PERSON to the LDAP data source.
 - ii. NOTE: If more than one DV domain exists, then configure LDAP2. Add additional LDAP data sources for each domain. For the LDAP_PERSON view, a UNION statement will be needed between each SELECT for each LDAP data source domain that is represented.
- b. **Option 2** – Configure your own LDAP data source.
 - i. Proceed to #2 below to map LDAP_PERSON to your data source.
 - ii. NOTE: If more than one DV domain exists, then configure LDAP2. Add additional LDAP data sources for each domain. For LDAP_PERSON view, a UNION statement will be needed between each SELECT for each LDAP data source domain that is represented.
- c. **Option 3** – Configure a pseudo LDAP to satisfy the inner workings of KPImetrics.
 - i. Proceed directly to #2 below and modify the existing UNION statement with composite users. There is no mapping to LDAP in this scenario.

2. Modify the LDAP_PERSON view

/shared/ASAssets/KPImetrics/Physical/Metadata/System/LDAP_PERSON
Correctly map to your LDAP directory structure.

- For Option 1 and 2, comment out the entire default UNION section of 'user1', 'user2' and 'user3' at the bottom of the view.
 - ii. Uncomment the top section which is the LDAP implementation section.
 - iii. For multiple LDAP data source "domains", you will need a SELECT statement for each one with a UNION to tie them altogether.
 - iv. Comment out the bottom section containing the "UNION" selects.
 - v. Determine which table the LDAP data should come from. The default is "user" but it may be person, organizationPerson, inetOrgPerson or a custom table.
 - vi. **[required]** Modify the source location "FROM" clause as required to point to the correct table that was determined in the previous step:
 1. FROM /shared/ASAssets/KPImetrics/Physical/Metadata/LDAP1/"user"
 - vii. **[required]** Create the proper transformation for "**userkey**" so that the data can be joined with metrics history "userkey" data. Analyze the LDAP data and make sure the "userkey" will contain the proper data that will join with metrics history "userkey" which is derived from metrics_resources_usage_hist."user".
 - viii. **[required]** Create the proper 'value' for "**domainkey**" so that the data can be joined with metrics history "domainkey" data. The "domainkey"

will be the name of the DV domain which can be found in the DV Manager browser under Domain Configuration. There is a direct correlation between the LDAP data source connection and the domain connection. The names need to be “exactly” the same in order to properly join with metrics history “domainkey” which is derived from metrics_resources_usage_hist.”domain”.

- For Option 3, modify the existing UNION and modify ‘user1’, ‘user2’, and ‘user3’ with your own composite users.
- Note: “userkey” is required.
- Note: “domainkey” is required.
- Guidelines:
 - ix. Keep the alias column names the same.
 - x. The physical column names may be different based on your LDAP data source.
 - xi. If a column does not exist then simply do a CAST(null as VARCHAR(...)) columnName.
 - xii. Keep the number of columns the same.
 - xiii. Keep the datatype and length the same.
 - xiv. Implement a where clause as required based on your LDAP data source. E.g. WHERE userid IS NOT NULL AND employeeStatus IS NOT NULL and employeeStatus = ‘A’ – active employees only
- Key fields to map include the following:
 - xv. cn (lower case cn field)
 - xvi. uid
 - xvii. userkey (lower case cn|uid|employeeID)
 - 1. Transformation of the user key which could be cn or uid or employeeID. Use SQL CASE if needed.
 - 2. This field is required and is used to join with metrics history tables. The data needs to be normalized so that it matches with metrics history “userkey” data.
 - xviii. CAST(CASE WHEN mail IS NOT NULL AND INSTR(mail, '@') > 0 THEN SUBSTRING(mail, INSTR(mail, '@')+1) ELSE mail END AS VARCHAR(255)) domainkey,
 - 1. Transformation of mail to get the domain name being used.
 - xix. objectClass
 - xx. displayName
 - xxi. name
 - xxii. objectGUID (user id field)
 - xxiii. sn (surname)
 - xxiv. givenName
 - xxv. employeeNumber

- xxvi. employeeID
- xxvii. mail
- xxviii. baseDN
- xxix. relativeDN
- Nice to have fields to map include the following:
 - xxx. description
 - xxxi. telephoneNumber
 - xxxii. c (country)
 - xxxiii. l (city)
 - xxxiv. st (state)
 - xxxv. street (street address)
 - xxxvi. ou
 - xxxvii. title
 - xxxviii. postalAddress
 - xxxix. postalCode (zip code)
 - xl. postOfficeBox
 - xli. physicalDeliveryOfficeName
 - xlii. initials
 - xlili. employeeType
 - xliv. manager
 - xlvi. homePhone
 - xlvi. mobile
 - xlvi. pager

[5.] Configure Metrics Job Lookup Tables

The KPImetrics contains several lookup tables that need to be pre-populated prior to the installation script running. The tables include:

- METRICS_JOB_ENVIRONMENTS
- METRICS_JOB_FILTERS
- METRICS_EVENT_REGISTRATION

1. Configure 07_pqInsert_KPI_Tables_METRICS_JOB_tables

- a. Edit /shared/ASAssets/KPImetrics/Physical/Metadata/DDL/Common/07_pqInsert_KPI_Tables_METRICS_JOB_tables
 - i. Configure the **METRICS_JOB_ENVIRONMENTS**
 1. Add a unique row for each environment in your pipeline. Remove any example rows that you are not relevant.
 2. Example: DEV1, SIT1, UAT, PROD
 - ii. Configure the **METRICS_JOB_FILTERS** – The job filters are used to filter out rows by the data transfer routine when transferring from the collection

tables to the history tables. The objective is to identify potentially high-use, non-user defined requests that are of little or no value to the ultimate KPI metrics reporting. *Since this feature is not built into DV metrics, it is “highly” recommended to filter out admin and other user accounts associated with deployment or non-query type of activity in DV.*

1. Example 1 would be for each environment:

```
ENV_TYPE TABLE_NAME          USER   DOMAIN RESOURCE_KIND
'UAT',    'metrics_resources_usage', 'admin', 'composite', 'system'
```

2. Example 2 would be any of the KPI metrics functionality that produces their own events. The assumption is that the KPI metrics folder resources were configured for resource ownership by metrics_app_id during installation:

```
ENV_TYPE TABLE_NAME          USER   DOMAIN RESOURCE_KIND
'UAT', 'metrics_resources_usage', 'metrics_app_id', 'composite', 'system'
```

3. Example 3 would be a deployment process that produces an inordinately high volume of events. Assume you have a deployment user called “deploy_app_id” that runs a deployment process but you don’t want to log those events:

```
ENV_TYPE TABLE_NAME          USER   DOMAIN RESOURCE_KIND
'UAT', 'metrics_resources_usage', 'deploy_app_id', 'composite', 'system'
```

4. Add a unique row for each environment, user, domain and resourcekind combination.

- b. Post-installation maintenance

- i. This script can be executed by itself post-installation by simply executing it to reload what is configured.

2. Configure 08_pqInsert_KPI_Tables_METRICS_EVENT_REGISTRATION

- a. Edit /shared/ASAssets/KPI metrics/Physical/Metadata/DDL/Common/

08_pqInsert_KPI_Tables_METRICS_EVENT_REGISTRATION

- a. **Requirement:** Prior to execution of this script LDAP data must first be configured and loaded. The post-installation script will automatically cache LDAP_PERSON as long as it is configured and caching is turned on server-wide.

- b. Configure the METRICS_EVENT_REGISTRATION

1. This table contains the event registrations for sending emails based on the subscriptions to various events. This procedure is used for either initial load or maintenance. If a row already exists, it does not update it. It simply bypasses it. This means that you can run this procedure as many times as you want and not impact existing rows. It does not delete or unsubscribe requester events. To delete a subscription, invoke `pMetricsEventRegistrationUnsubscribe()`. A subscription in the `METRICS_EVENT_REGISTRATION` table consists of a unique record for the combination of `SUBSCRIBER_EMAIL`, `GROUP_NAME`, `ENVIRONMENT_TYPE`, `EVENT_TYPE` and `REQUESTER_EMAIL`.
2. Edit a row to provide the subscription information. At a minimum, add a group subscriber email for the DV administration group for each event. If you don't have a group alias, then choose the DV administrator's email to receive alerts. A row includes the following:
 - a. `REQUESTER_EMAIL` [PK] – Primary requester email.
 - b. `SUBSCRIBER_EMAIL` [PK] – userid email or a group email alias. Who the email alert will be sent to.
 - c. `GROUP_NAME` [PK] – Group name subscribing to. When an alert occurs for a user the groups will be checked and cross-referenced with this registered group. The group [all] is a composite group and a catch-all for any user belonging to this composite group.
 - d. `ENVIRONMENT_TYPE` [PK] – Register for all environments [ALL] or a certain environment type [DEV1, CIT1, SIT1, UAT, TT, PROD]
 - e. `EVENT_TYPE` [PK] – [LONG_RUNNING, EXCEEDED_MEMORY, INACTIVITY, WORKFLOW_FAILURE, DBMS_SCHEDULER_ERROR, PURGE_HISTORY]
 - f. `EXCLUDE_TEXT` – A comma separate list of text that when found will signal exclusion and the email will not be sent. This is a way of filtering out emails based on text.

c. Post-installation maintenance

- i. This script can be executed by itself post-installation by simply executing it and answering “Y” to the parameter to delete the rows and reload what is configured.

[6.] Configure KPImetrics Triggers

The KPImetrics module uses a series of triggers to cache various tables of information. The different triggers provide flexibility to turn on and off processing as required. If certain functionality and data is not required, the trigger can be turned off saving on database space and DV processing.

1. Configure default triggers for your use case

- a. Open/Edit the resource
/shared/ASAssets/KPImetrics/Configuration/**defaultTriggersToEnable**
- b. Only modify the ON/OFF settings for each trigger. Leave all other settings alone.
 - i. Refer to the following sections for details on each trigger:
 - 1. [Metadata System Triggers and Load Scripts](#)
- c. The current triggers defaulted to OFF are as follows:
 - i. **kpimetricsTrig_00_CheckMetricsActivityDebug** – Only turn this on if you suspect that DV metrics is not working properly and you want to debug the DV metrics every hour.
 - ii. **kpimetricsTrig_11_Cache_METRICS_SQL_REQUEST_REPROCESS** – Only turn this on if you get a code update from the Open Source site and there were changes to the SQL Parser code.
 - iii. **kpimetricsTrig_17_CheckExceedMemoryPercentRequests** – You may choose to keep this off in lower-level environments but turn it on in PROD environments.
 - iv. **kpimetricsTrig_18_CheckLongRunningRequests** – You may choose to keep this off in lower-level environments but turn it on in PROD environments.
- d. The current triggers defaulted to ON that you may wish to evaluate:
 - i. **kpimetricsTrig_10_Cache_METRICS_SQL_REQUEST** – This trigger is defaulted to ON. If you do not wish to perform SQL parsing on the request description SQL statement to parse out the table and column resources used in the SQL then turn this trigger off. There is quite a bit of overhead associated with this trigger.
 - ii. **kpimetricsTrig_23_Cache_METRICS_SQL_RESOURCE_LINEAGE** – This trigger is defaulted to ON. If the above

kpimetricsTrig_10_Cache_METRICS_SQL_REQUEST is turned off then turn this trigger off as it uses data produced by that trigger to produce lineage data.

- iii. **kpimetricsTrig_33_DeleteCollection** – This trigger is defaulted to ON. You may want to turn this trigger off if you do not have a very active system such as a lower level environment or the system is not producing 100K or more transactions in 1 hour. If `commonValues.queryRequestTemplate='Y'` it can produce 100K or more transactions itself within 1 hour and therefore it is good to have this trigger ON. This setting is used to keep the collection tables as small as possible of unwanted data based on METRICS_JOB_FILTERS rows.
 1. `metrics_resources_usage` – delete from this collection table where a row matches METRICS_JOB_FILTERS based on user, domain and resourcekind. This delete gets executed first.
 2. `metrics_requests` – delete from this collection table when a row does not exist in `metrics_resources_usage` collection table. This delete gets executed second.

e. Trigger Category – Essential

- i. Provides baseline data for the data transfer and other processes
 1. `kpimetricsTrig_01_Cache_ALL_RESOURCES`
 2. `kpimetricsTrig_02_Cache_ALL_USERS`
 3. `kpimetricsTrig_03_Cache_LDAP_PERSON`
- ii. Used to transfer data from collection tables into history tables
 1. `kpimetricsTrig_30_DBMSScheduler`
- iii. Used to check for errors with the data transfer
 1. `kpimetricsTrig_31_DBMSSchedulerError`
- iv. Used to perform partition management on the history tables
 1. `kpimetricsTrig_32_DBMSPartitionManager`
- v. Provides clean-up/purge capability based on define schedule
 1. `kpimetricsTrig_16_PurgeHistoryData`
- vi. Provides monitoring capabilities
 1. `kpimetricsTrig_14_CheckCISWorkflowStatusFail`
 2. `kpimetricsTrig_15_CheckMetricsActivity`

f. Trigger Category – 2nd level processing (lite-weight)

- i. A series of value-added metrics that are an addition to the DV out-of-the-box metrics that are lite-weight in terms of processing. DV does not incur much overhead when these triggers execute.
 1. `kpimetricsTrig_04_Cache_CIS_SYSTEM_RESOURCES`

2. kpimetricsTrig_05_Cache_CPU_MEMORY_CHECKER
3. kpimetricsTrig_06_Cache_LOG_DISK
4. kpimetricsTrig_07_Cache_LOG_IO
5. kpimetricsTrig_08_Cache_LOG_MEMORY
6. kpimetricsTrig_12_Cache_SYS_CACHES
7. kpimetricsTrig_13_Cache_SYS_DATASOURCES
8. kpimetricsTrig_17_CheckExceedMemoryPercentRequests
9. kpimetricsTrig_18_CheckLongRunningRequests
10. kpimetricsTrig_33_DeleteCollection

g. Trigger Category – 3rd level processing (heavy-weight)

- i. A series of value-added metrics that are an addition to the DV out-of-the-box metrics

1. kpimetricsTrig_10_Cache_METRICS_SQL_REQUEST
2. kpimetricsTrig_19_AllCustom_AccessByUserOverTime
3. kpimetricsTrig_20_AllCustom_ActiveResourcesOverPeriodOfTime
4. kpimetricsTrig_21_AllCustom_ResourceCount_Details
5. kpimetricsTrig_22_AllCustom_ResourceCount_Total
6. kpimetricsTrig_23_Cache_METRICS_SQL_RESOURCE_LINEAGE

Execute Post-Installation Script

The KPImetrics module provides an automated script to complete the installation. The “Information Only Section” below will describe in detail what the script is going to execute. When the installation is complete and there are no red/impacted resources then turn on the triggers to begin processing KPI metrics data.

Execute Installation Script

Perform the post-installation configuration.

1. Execute Post-Installation Script:

- a. **WARNING:** If you are executing an installation into an existing KPImetrics environment, this script **WILL DROP** existing tables and data. Do not use this script for upgrading an environment. We recommend using TIBCO Professional Services for doing upgrades.
- b. Execute “[2_Post_Installation](#)” and provide parameters:
 - i. Location: /shared/ASAssets/KPImetrics_installation/2_Post_Installation
 - ii. IN inDebugSecondary – Debug Secondary provides a deep debug with detailed debug output. Y=deep debug. N=cursory, high-level debug. N

is recommended unless there are issues.

- iii. IN **performInstallationAction** – Y=perform the installation which will drop and recreate KPI metrics tables/sequences/procedures. N=Do nothing.
- iv. IN **destroyCIS_metrics_collection_Tables** – Y=destroy and recreate DV metrics collection tables including metrics_requests, metrics_resources_usage, and metrics_sessions. N=do not destroy if they exist.
- v. IN **forceOverwrite** – Force overwriting the CPU and Memory checker Windows/UNIX scripts. Y=force overwrite. N=do not force overwrite.
- vi. IN **userTransformation** – Leave null if no transformation is required. Contains a native database SQL-based transformation that is valid within the context of SELECT and uses "user" as the column name to transform. It must be a valid transformation for Oracle or SQL Server. If null, the default is simply "user". The column name that is "user" is transformed into the "userkey". The same transformation is applied to metrics_sessions_hist, metrics_requests_hist and metrics_resources_usage_hist. The objective is for the "userkey" field data to match the METRICS_LDAP_PERSON."userkey" field. This will only be applicable when the actual "user" requires a transformation to match with data found in LDAP.
 - 1. Always use double quotes around the "user" field.
 - 2. Never include an alias in front of the "user" field. Only reference the "user" field itself within the transformation.
 - 3. The user transformation is actually applied on the fly during the creation of the data transfer script for either Oracle or SQL Server. When these scripts are instantiated in the database, they will contain the transformation which will take the "user" field and normalize it into the "userkey" field:
 - a. 06_pqCreateDrop_KPI_Plsql_oracle_data_xfer_script
 - b. 06_pqCreateDrop_KPI_Plsql_sqlserver_data_xfer_script
 - 4. Example of an extreme case.
 - a. Given a metrics_resources_usage_hist."user" = u12345678.
 - b. Given METRICS_LDAP_PERSON."userkey" = 12345678 where the u is dropped or assumed.
 - c. Input Oracle Transformation would be:


```

CASE WHEN INSTR('abcdefghijklmnopqrstuvwxyz',
SUBSTR(LOWER("user"),1,1)) > 0 and INSTR('0123456789',
SUBSTR("user",2,1)) > 0

THEN SUBSTR(LOWER("user"),2)

ELSE LOWER("user")

END

```

d. Input SQL Server Transformation would be:

```

CASE WHEN patindex(SUBSTRING(LOWER("user"),1,1),
'abcdefghijklmnopqrstuvwxyz') > 0 and patindex(SUBSTRING("user",2,1),
'0123456789') > 0

THEN SUBSTRING(LOWER("user"),2,len("user"))

ELSE LOWER("user")

END

```

- vii. **IN domainTransformation** – Leave null if no transformation is required. Contains a native database SQL-based transformation that is valid within the context of SELECT and uses "domain" as the column name to transform. It must be a valid transformation for Oracle or SQL Server. If null, the default is simply "domain". The column name that is "domain" is transformed into the "domainkey". The same transformation is applied to metrics_sessions_hist, metrics_requests_hist and metrics_resources_usage_hist. The objective is for the "domainkey" field data to match the METRICS_LDAP_PERSON."domainkey" field. This will only be applicable when the actual "domain" requires a transformation to match with data found in LDAP.

1. Always use double quotes around the "domain" field.
2. Never include an alias in front of the "domain" field. Only reference the "domain" field itself within the transformation.
3. The domain transformation is actually applied on the fly during the creation of the data transfer script for either Oracle or SQL Server. When these scripts are instantiated in the database, they will contain the transformation which will take the "domain" field and normalize it into the "domainkey" field:
 - a. 06_pqCreateDrop_KPI_Plsql_oracle_data_xfer_script
 - b. 06_pqCreateDrop_KPI_Plsql_sqlserver_data_xfer_script
4. Example of an extreme case.
 - a. Most likely a transformation will not be needed unless the DV domain does not match the LDAP email domain. If that is the case, then a CASE statement will be required to map the LDAP email domain to the DV domain.

- b. Given a metrics_resources_usage_hist."domain" = "tib1".
- c. Given METRICS_LDAP_PERSON."domainkey" = "tibco1.com".

- d. Input Oracle Transformation would be:

```
CASE WHEN LOWER("domain") = "tib1" THEN LOWER("tibco1.com")
      WHEN LOWER("domain") = "tib2" THEN LOWER("tibco2.com")
      ELSE LOWER("domain")
END
```

- e. Input SQL Server Transformation would be:

```
CASE WHEN LOWER("domain") = "tib1" THEN LOWER("tibco1.com")
      WHEN LOWER("domain") = "tib2" THEN LOWER("tibco2.com")
      ELSE LOWER("domain")
END
```

Deploy CPU and Memory Checker shell scripts (Windows and UNIX)

IF ERROR ONLY

During post-installation, it was attempted to write the scripts to the \$CIS_HOME/bin directory on either Windows or UNIX depending on your operating system.

- a. If no error was returned then all scripts were installed correctly into \$CIS_HOME/bin and are able to be executed. Continue to the next major step/section.
- b. If the error message "23. INSTALL KPImetrics CpuAndMemoryChecker scripts manually." was received during post-installation, then you will need to install the scripts manually and set the file permissions accordingly especially for UNIX such as [rwxr-xr-x]. The recommend location to deploy the scripts is \$CIS_HOME/bin so that they can be executed by the user account that DV is running under. Take note of where the scripts have been deployed, you will need to provide the path to the scripts when configuring the KPImetrics "commonValues" script. The data source "/shared/ASAssets/KPImetrics/Physical/Metadata/CPUAndMemChecker" is used to execute the scripts. The following details the scripts:
 - a. Windows Powershell
 - i. KPImetricsCpuUtilization.ps1
 - ii. KPImetricsMemUtilization.ps1
 - b. Linux6_scripts
 - i. KPImetricsFreeMemCommand_linux6.sh
 - ii. KPImetricTopCommandGrepCpu_linux6.sh
 - c. Linux7_scripts
 - i. KPImetricsCpuFormat_linux7
 - ii. KPImetricsFreeMemCommand_linux7.sh

iii. `KPImetricTopCommandGrepCpu_linux7.sh`**Enable Triggers**

Enabling triggers starts the processing of KPI metrics data. The KPImetrics module makes use of pure database tables in order to retain DV metrics for a longer period than supported by the base DV logging functionality.

Please note that executing this procedure will cache the minimum required tables: METRICS_ALL_RESOURCES, METRICS_ALL_USERS and METRICS_LDAP_PERSON. Those 3 tables must be completed before any of the other triggers are allowed to run.

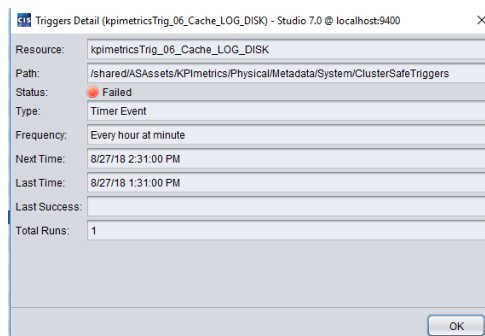
2. Enable Triggers:

- a. Execute the **updateTriggers** procedure
 - i. Location: `/shared/ASAssets/KPImetrics/Configuration/updateTriggers`
 - ii. `enable` – 0=disable, 1=enable, 2=display trigger list
 1. Select 1 to enable all configured triggers.
 - iii. `includeList` – Comma-separated list of trigger numbers to include in the (enable/disable) action. Leave null if the “defaultTriggersToEnable” are configured as desired.
 - iv. `excludeList` – Comma-separated list of trigger numbers to exclude in the (enable/disable) action. The excludeList overrides includeList. Leave null if the “defaultTriggersToEnable” are configured as desired.
- b. Review the Studio Manager → Triggers tab and look for any “Failed” triggers
 - i. Keep checking this screen until you have verified that at least one of the triggers has executed successfully. If the trigger has fired there will be a timestamp in the “Last Time” column. If the trigger is red “Failed” then go then proceed to point ii. below. If the trigger is green “Ready” then it is executing without issues. The majority of the triggers will fire within an hour of turning them on.

kpimetricsTrig_00_CheckMetricsActivityDebug	Disabled
kpimetricsTrig_01_Cache_ALL_RESOURCES	Ready
kpimetricsTrig_02_Cache_ALL_USERS	Ready
kpimetricsTrig_03_Cache_LDAP_PERSON	Ready
kpimetricsTrig_04_Cache_CIS_SYSTEM_RESOURCES	Ready
kpimetricsTrig_05_Cache_CPU_MEMORY_CHECKER	Ready
kpimetricsTrig_06_Cache_LOG_DISK	Failed
kpimetricsTrig_07_Cache_LOG_IO	Ready
kpimetricsTrig_08_Cache_LOG_MEMORY	Ready
kpimetricsTrig_10_Cache_METRICS_SQL_REQUEST	Ready
kpimetricsTrig_11_Cache_METRICS_SQL_REQUEST_REPROCE...	Disabled
kpimetricsTrig_12_Cache_SYS_CACHES	Ready
kpimetricsTrig_13_Cache_SYS_DATASOURCES	Ready
kpimetricsTrig_14_CheckCISWorkflowStatusFail	Failed
kpimetricsTrig_15_CheckMetricsActivity	Ready
kpimetricsTrig_16_PurgeHistoryData	Ready
kpimetricsTrig_17_CheckExceedMemoryPercentRequests	Disabled
kpimetricsTrig_18_CheckLongRunningRequests	Disabled
kpimetricsTrig_19_AllCustom_AccessByUserOverTime	Ready
kpimetricsTrig_20_AllCustom_ActiveResourcesOverPeriodOffTime	Ready
kpimetricsTrig_21_AllCustom_ResourceCount_Details	Ready
kpimetricsTrig_22_AllCustom_ResourceCount_Total	Ready
kpimetricsTrig_30_DBMSScheduler	Ready
kpimetricsTrig_31_DBMSSchedulerError	Failed
kpimetricsTrig_32_DBMSPartitionManager	Ready

ii. If a trigger has failed

1. Unfortunately, double clicking on the trigger does not yield any additional insights as shown in the screen shot



2. Trigger failures will most likely show up in the “**cs_server_events.log**”. It also may be necessary to check “**cs_server.log**”.

iii. The most common solution to a trigger failure will be permission problems on resources.

1. Make sure the privileges are set correctly on /shared/ASAssets/Utilities folder. It is recommended that /shared/ASAssets/Utilities be set to the composite group “all” with Read Execute Select.
2. Specifically, check privileges for /shared/ASAssets/Utilities/environment/getEnvName().

- iv. To verify the above issue and recommendation, check for the following in the “**cs_server_events.log**”. There would be an entry in that log that looks similar to this:

```

2018-08-27 13:31:00.002 -0400 INFO START trigger
name=/shared/ASAssets/KPImetrics/Physical/Metadata/System/ClusterSafeTriggers/kpim
etricsTrig_06_Cache_LOG_DISK type=TIMER action=PROCEDURE 1
    metrics_app_id composite      20400  6141564240009  6141564240009
    /shared/ASAssets/KPImetrics/Physical/Metadata/System/ClusterSafeTriggers/kpim
etricsTrig_06_Cache_LOG_DISKTIMER PROCEDURE
2018-08-27 13:31:00.112 -0400 ERROR FAIL trigger
name=/shared/ASAssets/KPImetrics/Physical/Metadata/System/ClusterSafeTriggers/kpim
etricsTrig_06_Cache_LOG_DISK type=TIMER action=PROCEDURE, reason failed: User
"metrics_app_id/composite" has insufficient privileges to execute
"/shared/ASAssets/Utilities/environment/getEnvName", on line 55, column 65. User
has READ privileges for that resource. [repository-1900321] 1
    metrics_app_id composite      20402  6141564240449  6141564240009
    /shared/ASAssets/KPImetrics/Physical/Metadata/System/ClusterSafeTriggers/kpim
etricsTrig_06_Cache_LOG_DISKTIMER PROCEDURE User
"metrics_app_id/composite" has insufficient privileges to execute
"/shared/ASAssets/Utilities/environment/getEnvName", on line 55, column 65. User has
READ privileges for that resource. [repository-1900321]

```

v. Verify the trigger issue is resolved. The screen shot indicates that the previously failed trigger now has successfully executed identified by a timestamp in the “Last Success” field and a status of green “Ready”.

Resource ^	Status	Type	Frequency	Next Time	Last Time	Last Success	Total Runs
kpimetricsTrig_10_Cache_METRICS_SQL_REQUEST	Ready	Timer Event	Every 30 minutes	8/27/18 2:45:00 PM	8/27/18 2:15:00 PM	8/27/18 2:15:00 PM	2

Execute ClusterSafeCache Scripts [Cluster Only]

For a cluster only, manually execute the following scripts on each node where **updateTriggers** was not executed. This is required so that METRICS_ALL_USERS and METRICS_ALL_RESOURCES gets populated for each node.

1. /shared/ASAssets/KPImetrics/Physical/Metadata/System/ClusterSafeCache/Cache_ALL_USERS
2. /shared/ASAssets/KPImetrics/Physical/Metadata/System/ClusterSafeCache/Cache_ALL_RESOURCES

Configure DV Out-of-the-box Metrics

To enable DV metrics:

1. Locate and open /policy/metrics
2. Since the tables were created earlier, choose the option to configure the path to the tables without re-creating them.
3. Data Source – browse and set the data source to the database you have configured.
4. Select “Edit Tables”

Status

Enable Buffering Status: ? DISABLED | Truncate Status: ? DISABLED

Storage

Data Source: /shared/ASAssets/KPImetrics/Physical/Metadata/KPI_oracle

Metrics Tables

metrics_sessions
metrics_requests
metrics_resources_usage

5. Browse to the schema path
 - a. Browse to the Sessions table: metrics_sessions
 - b. Browse to the Requests table: metrics_requests
 - c. Browse to the Usage table: metrics_resources_usage
 - d. DO NOT execute DDL
 - e. Click OK to finish

Create or Bind Tables - Studio 7.0 @ localhost:9400

Select Location: /shared/ASAssets/KPImetrics/Physical/Metadata/KPI_oracle/CIS_KPI

Sessions: metrics_sessions

Requests: metrics_requests

Usage: metrics_resources_usage

6. Enable metrics and save

Status

Enable Buffering Status: ✓ | Truncate Status: ✓

7. If the Buffer Status shows **RED** then review the following:
 - a. Make sure the three collection tables exist in the database and metrics are properly configured to point to them.
 - b. Make sure the tables have the correct permissions for writing to them.
 - c. If Oracle, make sure the tablespace has the correct permissions for writing to it.

KPImetrics Configuration is COMPLETE!

Information Only Section

This section provides a background on what gets executed by the post-installation script.

Script Requirements:

1. This script must be executed by the user admin or someone who has admin privileges
2. The KPImetrics CAR file has been imported either manually or with 1_Pre_Installation.

Steps to be performed after the KPImetrics CAR file is imported:

1. Update impacted resources.
2. Set resource privileges
3. Enable and test the KPImetrics data source
4. Rebind physical database type tables to the data source schema
5. Rebind the physical abstraction folder "/Physical/Abstraction" to the correct physical data source folder "/Physical/KPI_[oracle|sqlserver]"
6. Remove CIS_KPI folder if not needed
7. Update impacted resources
8. Validate and Create KPImetrics collection tables exist
9. Create KPImetrics cache tables, sequences and procedures [This will drop existing]
10. Introspect / Reintrospect the KPImetrics data source
11. Change resource ownership
12. Refresh the LDAP_PERSON view
13. Load the METRICS_JOBS table
14. Load the METRICS_EVENT_REGISTRATION table
15. Install the CpuAndMemCheckerCjp scripts into the file system for either Windows or UNIX.
16. Update the CPUAndMemChecker data source url and reintrospect it
17. Test the /System/CPU_MEMORY_CHECKER view
18. Installation and configuration are complete

How the scripts work:

The KPImetrics module requires several tables in the KPImetrics data source database in order to store metrics data for reporting. You must create these storage tables using the provided DDL in order for the KPImetrics module to function correctly.

If you choose to create the tables from within Studio, execute the procedures under /shared/ASAssets/KPImetrics/Physical/Metadata/DDL for your data source type. Proceed to the following section associated for your database type.

Procedure Parameters:

IN displayDDL – Y=Display the DDL, N=Do not display the DDL.

IN executeDDL – Y=execute the DDL, N=Do not execute the DDL. If you choose to execute the DDL externally, you will need to execute each of the 01-07 DDL procedures with the variable set as executeDDL=N and displayDDL=Y. This will output the DDL that you need to execute externally without actually performing the DDL operations. Once you have executed each procedure 01-07, you are now ready to run the DDL externally. Once you have executed the DDL proceed to the next section “*Common Configuration for all Databases*”.

IN dropIndexes– Y=drop the indexes before creating the first. N=do not drop the indexes.

IN dropTables – Y=drop the tables before creating the tables. N=do not drop the tables.

IN createTables – Y=execute the table creation DDL, N=display the table creation DDL in the console window only.

IN createIndexes– Y=execute index creation DDL. N=display the index creation DDL in the console window only.

OUT cursCombinedResult – Provides a status on each SQL statement executed.

OUT sqlScript – Generates an output of the entire script which can be used for external execution.

The common 08-09 DML procedures to populate the database must be run from within DV as there is no external SQL generation for those.

Create the KPImetrics storage tables for Oracle

1. Do the DV metrics collection tables exist?
 - a. NO
 - i. Create a tablespace called “METRICS_DATA_COLL” with NOLOGGING.
 1. Example:
 2. create tablespace METRICS_DATA_COLL nologging datafile
'C:/DV/oracle/metrics_data_coll01.dbf' size 500m autoextend on next 50m extent management local;
 - ii. Execute
02_pqCreateDrop_KPI_Tables_oracle_metrics_collection_stage_tables_[700|800]
 1. SchemaName – derived from commonValues.dataSourceSchema
 2. TablespaceName – derived from
commonValues.collectionTablespaceName

iii. Execute

02_pqCreateDrop_KPI_Tables_oracle_metrics_collection_tables_[700|800]

1. SchemaName – derived from commonValues.dataSourceSchema
2. TablespaceName – derived from commonValues.collectionTablespaceName

b. YES

- i. Consider exporting the data and turning off metrics in order to create these tables with NO LOGGING option enabled in their own tablespace called “METRICS_DATA_COLL”.
- ii. If this is not desired, then bypass this step.

2. Create the metrics history tables and indexes

a. Create a tablespace called “METRICS_DATA_HIST” and “METRICS_DATA_IDX” with logging turned on.

i. Example:

- ii. create tablespace METRICS_DATA_HIST logging datafile 'C:/DV/oracle/metrics_data_hist01.dbf' size 500m autoextend on next 50m extent management local;
- iii. create tablespace METRICS_DATA_IDX logging datafile 'C:/DV/oracle/metrics_data_idx01.dbf' size 500m autoextend on next 50m extent management local;

b. Drop and Create the metrics history tables and indexes: metrics_requests_hist, metrics_resources_usage_hist and metrics_sessions

c. Execute

03_pqCreateDrop_KPI_Tables_oracle_metrics_history_tables_[700|800] with

input:

- i. partitionNumber– derived from commonValues.partitionNumber
- ii. partitionStartDate– derived from commonValues.partitionStartDate
- iii. SchemaName – derived from commonValues.dataSourceSchema
- iv. TablespaceName – derived from commonValues.historyTablespaceName and commonValues.indexTablespaceName.

3. Create the metrics KPI reporting tables and indexes

a. Execute **04_pqCreateDrop_KPI_Tables_oracle_kpi_tables** with input:

- i. SchemaName – derived from commonValues.dataSourceSchema

- ii. `TablespaceName` – derived from `commonValues.historyTablespaceName` and `commonValues.indexTablespaceName`.
- 4. Create the metrics KPI reporting sequence
 - a. Execute ***05_pqCreateDrop_KPI_Tables_oracle_kpi_sequence*** with input:
 - i. `SchemaName` – derived from `commonValues.dataSourceSchema`
- 5. Create the metrics KPI “native” PLSQL data transfer script
 - a. Execute ***06_pqCreateDrop_KPI_Plsql_oracle_data_xfer_script_[700|800]*** with no input:
 - i. `userTransformation` – Contains a SQL-based transformation that is valid within the context of SELECT and uses "user" as the column name to transform. It must be a valid transformation for Oracle. If null, the default is simply "user". The column name that is "user" is transformed into the "userkey". The column "userkey" is joined with `METRICS_LDAP_PERSON`.
 - ii. `SchemaName` – derived from `commonValues.dataSourceSchema`

Create the KPI metrics storage tables for SQL Server

1. Do the DV metrics collection tables exist?
 - a. NO
 - i. Create a tablespace (filegroup) called “METRICS_DATA_COLL”.
 1. Example:
 2. alter database CIS_KPI add filegroup METRICS_DATA_COLL;
 3. alter database CIS_KPI add FILE (NAME='METRICS_DATA_COLL1', FILENAME='C:\SQLServer_FileGroup\METRICS_DATA_COLL1') TO FILEGROUP METRICS_DATA_COLL;
 - ii. Execute ***02_pqCreateDrop_KPI_Tables_sqlserver_metrics_collection_stage_tables_[700|800]***
 1. `SchemaName` – derived from `commonValues.dataSourceSchema`
 2. `TablespaceName` – derived from `commonValues.collectionTablespaceName`
 - iii. Execute ***02_pqCreateDrop_KPI_Tables_sqlserver_metrics_collection_tables_[700|800]***

1. SchemaName – derived from commonValues.dataSourceSchema
 2. TablespaceName – derived from
commonValues.collectionTablespaceName
- b. YES
- i. Consider exporting the data and turning off metrics in order to create these tables with NO LOGGING option enabled in their own tablespace called “METRICS_DATA_COLL”.
 - ii. If this is not desired, then bypass this step.
2. Create the metrics history tables and indexes
 - a. Create a tablespace (filegroup) called “METRICS_DATA_HIST” and “METRICS_DATA_IDX” with logging turned.
 - i. Example:
 - ii. alter database CIS_KPI add filegroup METRICS_DATA_HIST;
 - iii. alter database CIS_KPI add FILE (NAME='METRICS_DATA_HIST_1', FILENAME='C:\SQLServer_FileGroup\METRICS_DATA_HIST_1') TO FILEGROUP METRICS_DATA_HIST;
 - iv. alter database CIS_KPI add filegroup METRICS_DATA_IDX;
 - v. alter database CIS_KPI add FILE (NAME='METRICS_DATA_IDX_1', FILENAME='C:\SQLServer_FileGroup\METRICS_DATA_IDX_1') TO FILEGROUP METRICS_DATA_IDX;
 - b. Drop and Create the metrics history tables and indexes: metrics_requests_hist, metrics_resources_usage_hist and metrics_sessions
 - c. Execute **03_pqCreateDrop_KPI_Tables_sqlserver_metrics_history_tables** with input:
 - i. partitionNumber– derived from commonValues.partitionNumber
 - ii. partitionStartDate– derived from commonValues.partitionStartDate
 - iii. SchemaName – derived from commonValues.dataSourceSchema
 - iv. TablespaceName – derived from commonValues.historyTablespaceName and commonValues.indexTablespaceName.
 3. Create the metrics KPI reporting tables and indexes
 - a. Execute **04_pqCreateDrop_KPI_Tables_sqlserver_kpi_tables** with input:
 - i. SchemaName – derived from commonValues.dataSourceSchema
 - ii. TablespaceName – derived from commonValues.historyTablespaceName and commonValues.indexTablespaceName.

4. Create the metrics KPI reporting sequence
 - a. Execute **05_pqCreateDrop_KPI_Tables_sqlserver_kpi_sequence** with input:
 - i. SchemaName – derived from commonValues.dataSourceSchema
5. Create the metrics KPI “native” PLSQL data transfer script
 - a. Execute **06_pqCreateDrop_KPI_Plsql_sqlserver_data_xfer_script** with no input:
 - i. userTransformation – Contains a SQL-based transformation that is valid within the context of SELECT and uses "user" as the column name to transform. It must be a valid transformation for SQL Server. If null, the default is simply "user". The column name that is "user" is transformed into the "userkey". The column "userkey" is joined with METRICS_LDAP_PERSON.
 - ii. SchemaName – derived from commonValues.dataSourceSchema

Common Configuration for all Databases

1. Execute the procedure
/shared/ASAssets/KPImetrics/Configuration/**updateEnableDatasource**
 - a. This procedure will enable the proper KPImetrics data source [KPI_oracle or KPI_sqlserver] based on commonValues configuration.
2. Reintrospect or Introspect the data source
 - a. Was the default CIS_KPI schema or CIS_KPI/dbo catalog/schema used?
 - i. YES: Reintrospect the KPImetrics data source to confirm that the tables are visible.
 1. Right-click on your chosen data source “KPI_<database_type>” and select “Re-Introspect Now” and wait for it to complete.
 2. Click OK when completed.
 - ii. NO: A new schema or catalog name was used. Introspect the new schema or catalog/schema.
 1. Right-click on your chosen data source “KPI_<database_type>” and select “Add/Remove Tables”.
 2. Select your schema or catalog/schema
 3. Select all tables/procedures:
 - a. P_METRICS_ALL_TABLES
 - b. “cache_status” and “cache_tracking”

- c. Starting with “METRICS_...”
 - d. Starting with “metrics_...”
 - e. Click Next. Click Finish. Wait for the introspection to complete.
Review the list of tables in this section:
 - f. [“Metadata Data Source for KPI_<database_type>”](#)
 4. Click OK when completed.
3. Execute the rebind of the physical database type folder resources
`/shared/ASAssets/KPImetrics/Configuration/rebindPhysicalDatabaseType`
 - a. This procedure rebinds the
`/shared/ASAssets/KPImetrics/Physical/Physical/KPI_<database_type>` folder to the configured schema and catalog path found in `commonValues`.
 - b. No input is required. It uses `commonValues` to determine the target data source to rebind to. *Note:* `/shared/ASAssets/KPImetrics/Configuration/commonValues` must be configured prior to executing this procedure.
4. Execute the rebind procedure
`/shared/ASAssets/KPImetrics/Configuration/rebindPhysicalAbstraction` to rebind all KPImetrics abstraction layer views to the appropriate data source.
 - a. No input is required. It uses `commonValues` to determine the target data source to rebind to. *Note:* `/shared/ASAssets/KPImetrics/Configuration/commonValues` must be configured prior to executing this procedure.
 - b. Based on the `commonValues.dataSourceName`, it will rebind the `/Physical/Abstraction` views to either `/Physical/KPI_oracle` or `/Physical/KPI_sqlserver`.
 - c. Additionally, there is a list of views “`commonValues.viewOptimizationPathList`” used by `rebindPhysicalAbstraction`, that is used to modify the view SQL Script to optimize for Oracle or SQL Server in order to achieve push-down of the query.
5. Remove default CIS_KPI schema/catalog if not used.
 - a. If a different schema/catalog was chosen other than CIS_KPI then remove the old CIS_KPI name from your chosen data source “KPI_<database_type>”.
 - i. Right-click on CIS_KPI and select delete
 1. Update Impacted Resources

4 Upgrading KPImetrics

Introduction

This section describes how to upgrade KPImetrics from a previous version to the current version without data loss. The upgrade supports upgrading on the same TDV version such as 7.x or upgrading the KPImetrics from TDV 7.x to 8.x. When upgrading from 7.x to 8.x, the collection tables will change. Specifically, metrics_requests and metrics_sessions add “user”, “domain” and “group” columns.

How to Upgrade KPImetrics

Start with a few house-keeping activities described below such as turning off metrics and determining your current version.

Current Version:

1. Turn off all triggers using /shared/ASAssets/KPImetrics/Configuration/updateTriggers(0)
2. Turn off /policy/metrics so no data is inserted into the metrics collection tables.
3. Export a backup of /shared/ASAssets/KPImetrics and /services/databases/ASAssets/KPImetrics.
4. Create the folder /shared/ASAssets/KPImetrics_backup

The resources identified below contains all of the specific customizations for a given environment.

- a. Copy /shared/ASAssets/KPImetrics/Physical/Metadata and paste it into KPImetrics_backup.

Key resources to save from /Metadata folder:

- The active datasource starting with one of [KPI_oracle..., KPI_sqlserver...]
- The LDAP datasources which may be LDAP or LDAP1, LDAP2
- /DDL/Common/07_pqInsert_KPI_Tables_METRICS_JOB_tables
- /DDL/Common/08_pqInsert_KPI_Tables_METRICS_EVENT_REGISTRATION
- /System/LDAP_PERSON

- b. Copy /shared/ASAssets/KPImetrics/Configuration

Key resources to save from /Configuration folder:

- commonValues
- defaultTriggersToEnable

5. Determine the current version KPImetrics by executing getKPIVersion().
 - a. Use this knowledge to know what upgrade script to start executing with.
e.g. 2018.103

6. NOTE: Any version less than 2018.103 is not currently supported for upgrade.
 Contact Tibco Data Virtualization PSG if this use case comes up.

New Version:

7. If upgrading versions of TDV and creating a new TDV environment, create the user "metrics_app_id" in the new environment.
- a. This applies to either a new 7.x or 8.x environment.
 - b. This must be done prior to import of the KPImetrics car file.

8. Download the new KPImetrics package from the open source site

**** STOP. READ. ****

- DO NOT RUN THE CONFIGURATION "2_Post_Installation" WHICH DROPS AND RECREATES TABLES OR DATA WILL BE LOST.
- PERFORM A MANUAL IMPORT OF THE .CAR FILE.

- a. Import KPImetrics_YYYYQnnn.car at the root "hostname (/)" folder with overwrite checked. ----- <-- REQUIRED

- b. Import options checked:

X Caching	X Data Source Connections	Include Users Jars
X Include Cache Policies	X Overwrite	Merge Users Folder
Create Caching Tables	Override Locks	X Privileges
X Custom Jars	Discovery Files	

9. Set the proper privileges

- a. Right click on /services/databases/ASAssets/KPImetrics and select "Privileges".
 - i. Click on the radio button "Hide users and groups without explicit privileges".
 - ii. Set composite Groups all to "Read Execute Select"
 - iii. Set composite Users metrics_app_id to "Read Write Execute Select Update Insert Delete"
 - iv. Make sure all other users and groups are unchecked. The group all takes care of the access.
 - v. Recursively apply privileges
 - 1. Select the check box "Apply recursively to child resources and folders"
 - 2. Select the check box "Make child resources look like this resource".

3. Click OK
 4. Click Yes to apply
 - b. Right click on /shared/ASAssets/KPImetrics and select "Privileges".
 - i. Click on the radio button "Hide users and groups without explicit privileges".
 - ii. Set composite Groups all to "Read Execute Select"
 - iii. Set composite Users metrics_app_id to "Read Write Execute Select Update Insert Delete"
 - iv. Make sure all other users and groups are unchecked. The group all takes care of the access.
 1. Select the check box "Apply recursively to child resources and folders"
 2. Select the check box "Make child resources look like this resource".
 3. Click OK
 4. Click Yes to apply
10. Set the resource ownership to metrics_app_id
 - a. Right click on /services/databases/ASAssets/KPImetrics and select "Change Resource Owner"
 - i. Change New owner to: composite metrics_app_id
 - ii. Check the box "Apply the change recursively"
 - iii. Click OK
 - iv. If a box pops up that says "No resources will be changed", then you are done. Click Cancel.
 - v. If a box pops up that lists resources to change, then click OK to change them.
 - b. Right click on /shared/ASAssets/KPImetrics and select "Change Resource Owner"
 - i. Change New owner to: composite metrics_app_id
 - ii. Check the box "Apply the change recursively"
 - iii. Click OK
 - iv. If a box pops up that says "No resources will be changed", then you are done. Click Cancel.
 - v. If a box pops up that lists resources to change, then click OK to change them.
11. Re-Configure the resources that were backed up in step 4.

- a. Configure "commonValues" - due to changes in the parameters, the values from the backed up "commonValues" need to be updated one-by-one into the new "commonValues". Don't just copy and paste everything as there may be new variable entries. Copy the values. Do not copy and overwrite the entire resource.
- b. Configure "defaultTriggersToEnable" - due to changes in triggers, the on/off values from the backed up "defaultTriggersToEnable" need to be updated one-by-one in the new "defaultTriggersToEnable". Copy the values. Do not copy and overwrite the entire resource.
- c. Configure "LDAP_PERSON" - due to changes in the view, copy the configuration from /shared/ASAssets/KPImetrics/Physical/Metadata/System/LDAP_PERSON but be sure to incorporate the new "userkey" and "domainkey" fields. Copy the values. Do not copy and overwrite the entire resource.
- d. Copy the active datasource from the backed-up folder back to /shared/ASAssets/KPImetrics/Physical/Metadata. In the case where the datasource name is KPI_oracle or KPI_sqlserver, you will need to reconfigure one of the following instead of copying the old resource: [KPI_oracle_11g, KPI_oracle_12c, KPI_sqlserver_2012, KPI_sqlserver_2014].
- e. Copy all of the backed up /DDL/Common procedures to /shared/ASAssets/KPImetrics/Physical/Metadata/DDL/Common.
- f. Copy the LDAP datasource from the backed-up folder. In the case where the original datasource was call LDAP do the following:
 - i. Copy the old LDAP back to /Metadata folder
 - ii. Delete the current LDAP1.
 - iii. Rename the old LDAP datasource to LDAP1.

12. Rebind the resources.

- a. Execute /shared/ASAssets/KPImetrics/Configuration/rebindPhysicalDatabaseType to synchronize the code with the datasource schema.
- b. Execute /shared/ASAssets/KPImetrics/Configuration/rebindPhysicalAbstraction() to rebind the folder /shared/ASAssets/KPImetrics/Physical/Physical/Abstraction to the proper sibling folder [KPI_oracle or KPI_sqlserver].

13. There may be impacted resources.

- a. If there are impacted resource, update impacted resource.
- b. Execute /shared/ASAssets/KPImetrics/Configuration/updateImpactedResources to fix impacted resources.

14. Perform the database upgrade.

- a. Navigate to the correct DDL database folder [Oracle or SqlServer] based on your implementation.
- b. Using the version # from the old installation, determine which upgrade script to start executing from the upgrade folder:

/shared/ASAssets/KPImetrics/Physical/Metadata/DDL/Oracle/Upgrade

For example: if your old version was Oracle and version 2018.400 or 2018.4 then start by executing

```
"upgrade_2018Q400_to_2018Q401_pqCreateDrop_Tables_oracle"
```

- c. Execute all of the subsequent versions following your version in sequence.

For example: if your old version was Oracle and version 2018.400 or 2018.4 then execute the following in order.

```
"upgrade_2018Q400_to_2018Q401_pqCreateDrop_Tables_oracle"
```

```
"upgrade_2018Q401_to_2019Q100_pqCreateDrop_Tables_oracle"
```

```
"upgrade_2019Q100_to_2019Q101_pqCreateDrop_Tables_oracle"
```

```
"upgrade_2019Q101_to_2019Q102_pqCreateDrop_Tables_oracle"
```

... additional scripts not displayed here ...

```
"upgrade_8000Q000_pqCreateDrop_Final_oracle" - always perform this step
```

```
"upgrade_9000Q000_pqUpdate_METRICS_ALL_RESOURCES" - always perform this step
```

For example: if your old version was SqlServer and version 2018.400 or 2018.4 then execute the following in order.

```
"upgrade_2018Q400_to_2018Q401_pqCreateDrop_Tables_sqlserver"
```

```
"upgrade_2018Q401_to_2019Q100_pqCreateDrop_Tables_sqlserver"
```

```
"upgrade_2019Q100_to_2019Q101_pqCreateDrop_Tables_sqlserver"
```

```
"upgrade_2019Q101_to_2019Q102_pqCreateDrop_Tables_sqlserver"
```

... additional scripts not displayed here ...

```
"upgrade_8000Q000_pqCreateDrop_Final_sqlserver" - always perform this step
```

```
"upgrade_9000Q000_pqUpdate_METRICS_ALL_RESOURCES" - always perform this step
```

15. Re-introspect your active datasource and determine if all changes are aligned with the current resources.

16. Update impacted resource - Execute

/shared/ASAssets/KPImetrics/Configuration/updateImpactedResources to fix impacted resources.

- a. Close Studio and re-open it. Studio may become sluggish after executing `updateImpactedResources`.
- b. If the impacted resources persist, open each one, put a space in the comment section, save and close.

The impacted resource should resolve.

17. Turn on `/policy/metrics` to start inserting rows into the metrics collection tables.
 - a. If the prior version used `KPI_oracle` or `KPI_sqlserver` then you will have to reconfigure it to use one of [`KPI_oracle_11g`, `KPI_oracle_12c`, `KPI_sqlserver_2012`, `KPI_sqlserver_2014`]
 - b. Make sure both radio button indicators are green.
 - c. If they are red then there is a problem. Try disabling metrics and then enabling again.
 - d. Also try reconfiguring the database, schema and tables from scratch.
 - e. If the problem persists, then check `cs_server.log` for errors.
18. Test each of the following scripts manually and in order shown to determine that there were no issues.

- a. After each script, look at the view to the right of the script to determine if it was successful or failed.
- b. Scripts are located here:
`/shared/ASAssets/KPImetrics/Physical/Metadata/System/ClusterSafeCache`
- c. Validation views are located here:
 - i. `/shared/ASAssets/KPImetrics/Application/workflow/vCISWorkflowStatus`
 - ii. `/shared/ASAssets/KPImetrics/Application/workflow/vJobDetails`

SCRIPT

Validation View

-- Must do these three first to establish base line resources used by other scripts.

- | | |
|-------------------------------------|---------------------------------|
| 1. <code>Cache_ALL_RESOURCES</code> | <code>vCISWorkflowStatus</code> |
| 2. <code>Cache_ALL_USERS</code> | <code>vCISWorkflowStatus</code> |
| 3. <code>Cache_LDAP_PERSON</code> | <code>vCISWorkflowStatus</code> |

-- Quick tests

- | | |
|--|---------------------------------|
| 4. <code>Cache_CPU_MEMORY_CHECKER</code> | <code>vCISWorkflowStatus</code> |
| 5. <code>Cache_LOG_DISK</code> | <code>vCISWorkflowStatus</code> |
| 6. <code>Cache_LOG_IO</code> | <code>vCISWorkflowStatus</code> |
| 7. <code>Cache_LOG_MEMORY</code> | <code>vCISWorkflowStatus</code> |
| 8. <code>Cache_CIS_SYSTEM_RESOURCES</code> | <code>vCISWorkflowStatus</code> |
| 9. <code>Cache_SYS_CACHES</code> | <code>vCISWorkflowStatus</code> |
| 10. <code>Cache_SYS_DATASOURCES</code> | <code>vCISWorkflowStatus</code> |

11. pCheckExceedMemoryPercentRequests	vCISWorkflowStatus
12. pCheckLongRunningRequests	vCISWorkflowStatus
13. pPurgeData	vCISWorkflowStatus

-- Execute some external test queries to generate the proper requests.

-- To force /policy/metrics to write data to the collection tables, disable it and then enable it.

-- These scripts rely on data from collection tables

14. pMETRICS_ALL_TABLES_exec	vJobDetails
15. pMETRICS_DELETE_COLLECTION_exec	vJobDetails
16. pPARTITION_MANAGER_exec	vJobDetails
17. pCheckDBMSSchedulerError	vJobDetails
18. Cache_AllCustom_AccessByUserOverTime	vCISWorkflowStatus
19. Cache_AllCustom_ActiveResourcesOverPeriodOfTime	vCISWorkflowStatus
20. Cache_AllCustom_ResourceCount_Details	vCISWorkflowStatus
21. Cache_AllCustom_ResourceCount_Total	vCISWorkflowStatus
22. Cache_METRICS_SQL_REQUEST_EXEC	vCISWorkflowStatus
23. pCheckCISWorkflowStatusFail	vCISWorkflowStatus

19. Turn on all triggers using /shared/ASAssets/KPImetrics/Configuration/updateTriggers(1)





























20. Upgrade is complete.

5 KPImetrics Administration Scenarios

Turn KPI On/Off

This section describes how to turn KPImetrics on and off by simply turning on/off the triggers.

1. Turn OFF KPImetrics triggers
 - a. Execute this procedure
/shared/ASAssets/KPImetrics/Configuration/updateTriggers
 - i. Enable=0
 - ii. includeList=null
 - iii. excludeList=null
 - b. Refresh Studio
 - c. Review the triggers status on the Manager tab / Triggers:

Resource 	Status
kpimetricsTrig_00_CheckMetricsActivityDebug	 Disabled
kpimetricsTrig_01_Cache_ALL_RESOURCES	 Disabled
kpimetricsTrig_02_Cache_ALL_USERS	 Disabled
kpimetricsTrig_03_Cache_LDAP_PERSON	 Disabled
kpimetricsTrig_04_Cache_CIS_SYSTEM_RESOURCES	 Disabled
kpimetricsTrig_05_Cache_CPU_MEMORY_CHECKER	 Disabled
kpimetricsTrig_06_Cache_LOG_DISK	 Disabled
kpimetricsTrig_07_Cache_LOG_IO	 Disabled
kpimetricsTrig_08_Cache_LOG_MEMORY	 Disabled
kpimetricsTrig_10_Cache_METRICS_SQL_REQUEST	 Disabled
kpimetricsTrig_11_Cache_METRICS_SQL_REQUEST_REPROCESS	 Disabled
kpimetricsTrig_12_Cache_SYS_CACHES	 Disabled
kpimetricsTrig_13_Cache_SYS_DATASOURCES	 Disabled
kpimetricsTrig_14_CheckCISWorkflowStatusFail	 Disabled
kpimetricsTrig_15_CheckMetricsActivity	 Disabled
kpimetricsTrig_16_PurgeHistoryData	 Disabled
kpimetricsTrig_17_CheckExceedMemoryPercentRequests	 Disabled
kpimetricsTrig_18_CheckLongRunningRequests	 Disabled
kpimetricsTrig_19_AllCustom_AccessByUserOverTime	 Disabled
kpimetricsTrig_20_AllCustom_ActiveResourcesOverPeriodOfTime	 Disabled
kpimetricsTrig_21_AllCustom_ResourceCount_Details	 Disabled
kpimetricsTrig_22_AllCustom_ResourceCount_Total	 Disabled
kpimetricsTrig_23_Cache_METRICS_SQL_RESOURCE_LINEAGE	 Disabled
kpimetricsTrig_30_DBMScheduler	 Disabled
kpimetricsTrig_31_DBMSchedulerError	 Disabled
kpimetricsTrig_32_DBMSPartitionManager	 Disabled
kpimetricsTrig_33_DeleteCollection	 Disabled

2. Turn ON KPImetrics triggers
 - a. Execute this procedure
/shared/ASAssets/KPImetrics/Configuration/updateTriggers
 - i. Enable=1
 - ii. includeList=null

- iii. excludeList=null
- b. Refresh Studio
- c. Review the triggers status on the Manager tab / Triggers:

Resource ▲	Status
kpimetricsTrig_00_CheckMetricsActivityDebug	ⓘ Disabled
kpimetricsTrig_01_Cache_ALL_RESOURCES	✔ Ready
kpimetricsTrig_02_Cache_ALL_USERS	✔ Ready
kpimetricsTrig_03_Cache_LDAP_PERSON	✔ Ready
kpimetricsTrig_04_Cache_CIS_SYSTEM_RESOURCES	✔ Ready
kpimetricsTrig_05_Cache_CPU_MEMORY_CHECKER	✔ Ready
kpimetricsTrig_06_Cache_LOG_DISK	✔ Ready
kpimetricsTrig_07_Cache_LOG_IO	✔ Ready
kpimetricsTrig_08_Cache_LOG_MEMORY	✔ Ready
kpimetricsTrig_10_Cache_METRICS_SQL_REQUEST	✔ Ready
kpimetricsTrig_11_Cache_METRICS_SQL_REQUEST_REPROCESS	ⓘ Disabled
kpimetricsTrig_12_Cache_SYS_CACHES	✔ Ready
kpimetricsTrig_13_Cache_SYS_DATASOURCES	✔ Ready
kpimetricsTrig_14_CheckCISWorkflowStatusFail	✔ Ready
kpimetricsTrig_15_CheckMetricsActivity	✔ Ready
kpimetricsTrig_16_PurgeHistoryData	✔ Ready
kpimetricsTrig_17_CheckExceedMemoryPercentRequests	ⓘ Disabled
kpimetricsTrig_18_CheckLongRunningRequests	ⓘ Disabled
kpimetricsTrig_19_AllCustom_AccessByUserOverTime	✔ Ready
kpimetricsTrig_20_AllCustom_ActiveResourcesOverPeriodOfTime	✔ Ready
kpimetricsTrig_21_AllCustom_ResourceCount_Details	✔ Ready
kpimetricsTrig_22_AllCustom_ResourceCount_Total	✔ Ready
kpimetricsTrig_23_Cache_METRICS_SQL_RESOURCE_LINEAGE	✔ Ready
kpimetricsTrig_30_DBMSScheduler	✔ Ready
kpimetricsTrig_31_DBMSSchedulerError	✔ Ready
kpimetricsTrig_32_DBMSPartitionManager	✔ Ready
kpimetricsTrig_33_DeleteCollection	✔ Ready

3. Turn ON KPImetrics triggers [TRICK]

- a. Usage: In this scenario you have executed the triggers and they are up and running. Now you need to stop all the triggers for some maintenance. When the maintenance is over, you want to restart the triggers but you don't need to or want to wait for triggers 1-3 to perform their operation since the data in their tables is current given that you had a short maintenance window within the same day. The trick is to turn on all of the triggers except 1-3 which you will do manually.
- b. Execute this procedure
/shared/ASAssets/KPImetrics/Configuration/updateTriggers.
 - i. The procedure will run very fast since it does not have to load the data for triggers 1-3 as you will put them in the excludeList as shown below.
 - ii. Enable=1
 - iii. includeList=null
 - iv. excludeList=1,2,3**
- c. In Studio, go to the Manager (tab) → Triggers (screen)
 - i. Locate the following three triggers and highlight them

kpimetricsTrig_01_Cache_ALL_RESOURCES	Disabled
kpimetricsTrig_02_Cache_ALL_USERS	Disabled
kpimetricsTrig_03_Cache_LDAP_PERSON	Disabled

- ii. Click on “Change Enabling” button

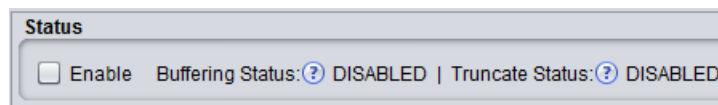
kpimetricsTrig_01_Cache_ALL_RESOURCES	Ready
kpimetricsTrig_02_Cache_ALL_USERS	Ready
kpimetricsTrig_03_Cache_LDAP_PERSON	Ready

- iii. Now all of the triggers are operational and will execute based on their next time slot.

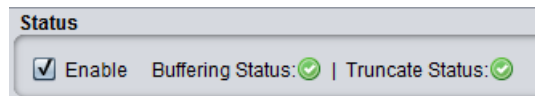
Turn Data Virtualization (DV) metrics On/Off

This section describes how to turn DV metrics on/off.

1. To Stop DV metrics, you must log in as the “admin” user
 - a. Open /policy/metrics
 - b. Uncheck the “Enable” button
 - c. The metrics will indicate they are DISABLED



2. To Start DV metrics, you must log in as the “admin” user
 - a. Open /policy/metrics
 - b. Check the “Enable” button
 - c. Both Buffering Status and Truncate Status should be green.



- d. If they are not, then there is a problem with the metrics configuration. Check the log entries for errors.

Modify Triggers

This section describes how to modify the triggers once they are installed.

1. Open/Edit the resource /shared/ASAssets/KPImetrics/Configuration/defaultTriggersToEnable
 - a. Only modify the ON/OFF settings for each trigger. Leave all other settings alone.
 - i. Refer to the following sections for details on each trigger:
 1. [Metadata System Triggers and Load Scripts](#)
 - b. The current triggers defaulted to OFF are as follows:
 - i. **kpimetricsTrig_00_CheckMetricsActivityDebug** – Only turn this on if you suspect that DV metrics is not working properly and you want to debug the DV metrics every hour.
 - ii. **kpimetricsTrig_11_Cache_METRICS_SQL_REQUEST_REPROCESS** – Only

turn this on if you get a code update from the Open Source site and there were changes to the SQL Parser code.

- iii. **kpimetricsTrig_17_CheckExceedMemoryPercentRequests** – You may choose to keep this off in lower-level environments but turn it on in PROD environments.
 - iv. **kpimetricsTrig_18_CheckLongRunningRequests** – You may choose to keep this off in lower-level environments but turn it on in PROD environments.
- c. The current triggers defaulted to ON that you may wish to evaluate:
- i. **kpimetricsTrig_10_Cache_METRICS_SQL_REQUEST** – This trigger is defaulted to ON. If you do not wish to perform SQL parsing on the request description SQL statement to parse out the table and column resources used in the SQL then turn this trigger off. There is quite a bit of overhead associated with this trigger.
 - ii. **kpimetricsTrig_33_DeleteCollection** – This trigger is defaulted to ON. This trigger deletes unwanted rows in the metrics_resources_usage based on applicable rows found in METRICS_JOB_FILTERS. It also deletes rows in metrics_requests where a corresponding row does not exist in metrics_resources_usage. As part of the normal KPI metrics operation, the trigger kpimetricsTrig_30_DBMSScheduler runs a procedure that will transfer metrics collection data and delete rows from the metrics_resources_usage and metrics_requests collection tables. If this trigger/procedure is operating without issues and keeping up with deletes, then it is possible to turn off the trigger: **kpimetricsTrig_33_DeleteCollection**.

2. Execute this procedure /shared/ASAssets/KPImetrics/Configuration/updateTriggers
 - a. Enable=1
 - b. includeList=null
 - c. excludeList=null
3. Refresh Studio
4. Review the triggers status on the Manager tab / Triggers to ensure your trigger changes were enforced.

Perform Oracle Database Maintenance on Collection Tables

This section outlines how to perform maintenance on the Oracle collection tables in order to regain the tablespace.

1. Stop Data Virtualization metrics
2. Execute the
/shared/ASAssets/KPImetrics/Physical/Physical/KPI_oracle/P_METRICS_ALL_TABLES_exec to process the existing rows in the collection tables
3. Execute the following commands directly against Oracle

```
ALTER TABLE metrics_sessions ENABLE ROW MOVEMENT;  
ALTER TABLE metrics_sessions SHRINK SPACE CASCADE;  
ALTER TABLE metrics_sessions DEALLOCATE UNUSED KEEP 50;  
ALTER TABLE metrics_requests ENABLE ROW MOVEMENT;  
ALTER TABLE metrics_requests SHRINK SPACE CASCADE;  
ALTER TABLE metrics_requests DEALLOCATE UNUSED KEEP 50;  
ALTER TABLE metrics_resources_usage ENABLE ROW MOVEMENT;  
ALTER TABLE metrics_resources_usage SHRINK SPACE CASCADE;  
ALTER TABLE metrics_resources_usage DEALLOCATE UNUSED KEEP 50;
```

Configure Third Party Tool Access

This section outlines how to configure third party tool access for reporting tools such as Cognos, Spotfire, MicroStrategy or others.

1. Download the Data Virtualization (DV) ODBC 7 Drivers and install on the client host machine
2. Configure an ODBC data source
 - a. DSN Name: Provide different connections to different DV instances
 - i. DV_KPIMETRICS_DEV
 - ii. DV_KPIMETRICS_UAT
 - iii. DV_KPIMETRICS_PROD
 - b. Host – hostname of the DV target instance
 - c. Port – port number of the DV target instance (e.g. 9401)
 - d. User Name – the user name or service account to use
 - e. Password – the password for the user name or service account
 - f. Domain – the domain name such as “composite” or “ldap”
 - g. Data source – the data source will be “ASAssets”
 - h. Catalog – can leave this blank
 - i. Local/Code Page – can leave this blank

Get the Current Row Distribution for the History Tables/Partitions

This section describes how to get the row distribution for the three history tables and there partitions.

1. Oracle – Execute the following procedure with no input:
/shared/ASAssets/KPImetrics/Physical/Metadata/DDL/Oracle/
03_pqCreateDrop_KPI_Tables_oracle_metrics_history_tables_ROW_DISTRIBUTION
2. SQL Server – Execute the following procedure with no input:
/shared/ASAssets/KPImetrics/Physical/Metadata/DDL/Sqlserver/
03_pqCreateDrop_KPI_Tables_sqlserver_metrics_history_tables_ROW_DISTRIBUTION

6 KPImetrics Resources

Configuration Resources

This section outlines the resources that are used for configuration of KPImetrics.

KPI Version Overview

Location: /shared/ASAssets/KPImetrics

This section lists the version and “How to Install” procedures.

Configuration Name	Description
getKPIVersion	Returns the version of KPI in the format: YYYY.Q[1-4] e.g. 2017.4 If a patch is release within the Quarter then it would be 2017.401
README	Provides a brief description of the published views.
RELEASE_NOTES	Provides a description of what changes per release.

Configuration Folder Overview

Location: /shared/ASAssets/KPImetrics/Configuration

This section lists all of the configuration procedures that have been defined for the KPImetrics module. These scripts provide functionality for setup/configuration of the KPImetrics module.

Configuration Name	Description
commonValues	Script to store default purge time values and data source information.
getQueryResponseTime	This procedure executes the passed in query and calculates the average time taken to retrieve the first row of data.
pMaintenanceUpdateHostnamePort	This procedure is used to convert the values for NODE_HOST/nodehost and NODE_PORT/nodeport from one value to another. It is “highly” unlikely that this procedure would ever be used but it is provided nonetheless for unforeseen circumstances. For each view found in /shared/ASAssets/KPImetrics/Physical/Physical/Abstraction perform the updated operation if they contain both NODE_HOST/nodehost and NODE_PORT/nodeport. This is a dynamic procedure so it will pick up whatever is in that path. Note: If you (1) migrate from one host to another host or (2) migrate from one DV instance to another DV instance on the same host, you should <u>not modify the data</u> . The nodehost and nodeport along with requestid are needed to insure a unique row. When moving to a new DV instance, the requestid may be repeated.
pMetricsEventRegistrationList	This procedure is used to "LIST" to a metrics event registration. A subscription in the METRICS_EVENT_REGISTRATION table that consists of a unique record for the combination of SUBSCRIBER_EMAIL, GROUP_NAME, ENVIRONMENT_TYPE, EVENT_TYPE and REQUESTER_EMAIL.
pMetricsEventRegistrationSubscribe	This procedure is used to "SUBSCRIBE" to a metrics event registration. A subscription in the METRICS_EVENT_REGISTRATION table consists of a

	<p>unique record for the combination of SUBSCRIBER_EMAIL, GROUP_NAME, ENVIRONMENT_TYPE, EVENT_TYPE and REQUESTER_EMAIL. The column EVENT_TYPE is one of: [LONG_RUNNING EXCEEDED_MEMORY INACTIVITY WORKFLOW_FAILURE DBMS_SCHEDULER_ERROR]</p>
pMetricsEventRegistrationUnsubscribe	<p>This procedure is used to "UNSUBSCRIBE" to a metrics event registration. A subscription in the METRICS_EVENT_REGISTRATION table that consists of a unique record for the combination of SUBSCRIBER_EMAIL, GROUP_NAME, ENVIRONMENT_TYPE, EVENT_TYPE and REQUESTER_EMAIL.</p>
rebindPhysicalAbstraction	<p>The 1st step of this procedure is used to rebind all of the resources (Views) in /shared/ASAssets/KPImetrics/Physical/Physical/Abstraction to the KPI_<database_type> folder as configured in commonValues. All views above the /Abstraction layer will be redirected to use the correct data source. This is a one-time configuration done during setup. The 2nd step is used to optimize various views and procedures as defined by commonValues.viewOptimizationPathList.</p> <p>Some examples of optimization:</p> <p><u>SQL SERVER:</u></p> <p>ex 1. {OPTION IGNORE_TRAILING_SPACES="TRUE", CASE_SENSITIVE="FALSE"}</p> <p>ex 2. CAST(DATEDIFF('DAY', requestdate, CURRENT_DATE)/31.00 AS DECIMAL(19,2)) requestdatemonths, CAST(DATEDIFF('DAY', requestdate, CURRENT_DATE) AS DECIMAL(19,0)) requestdatedays,</p> <p>ex 3. FROM /shared/ASAssets/KPImetrics/AllCustomReports/ ResourceCount_DetailsRT_sqlserver</p> <p>ex 4. FROM /shared/ASAssets/KPImetrics/AllCustomReports/ ResourceCount_TotalRT_sqlserver</p> <p><u>ORACLE:</u></p> <p>ex 1. {OPTION IGNORE_TRAILING_SPACES="FALSE", CASE_SENSITIVE="TRUE"}</p> <p>ex 2. CAST(ROUND(MONTHS_BETWEEN(CURRENT_DATE, requestdate),2) AS DECIMAL(19,2)) requestdatemonths, CAST(DAYS_BETWEEN(CURRENT_DATE, requestdate) AS DECIMAL(19,0)) requestdatedays,</p> <p>ex 3. FROM /shared/ASAssets/KPImetrics/AllCustomReports/ ResourceCount_DetailsRT_oracle</p> <p>ex 4. FROM /shared/ASAssets/KPImetrics/AllCustomReports/ ResourceCount_TotalRT_oracle</p>
rebindPhysicalDatabaseType	<p>Rebind the folder /shared/ASAssets/KPImetrics/Physical/Physical/KPI_<database_type> [KPI_oracle, KPI_sqlserver]. Rebind from whatever the current folder is currently pointing to and modify to the new catalog/schema path for tables and procedures. This is required when the catalog and/or the schema name are changed from the original setting. This may be required on the initial setup in the development environment. This may be required if upper environments have a different catalog or schema from the DEV environment although this is not recommended. This procedure is used when the schema and catalog for Oracle</p>

	is not 'CIS_KPI' or SQLServer is not 'CIS_KPI/dbo'. Configure the following /shared/ASAssets/KPImetrics/Configuration/commonValues.dataSourceName PRIOR to running this procedure.
updateEnableDatasource	This procedure is used to enable the KPImetrics data source [KPI_oracle or KPI_sqlserver].
updateImpactedResources	Provides a way to iterate through /shared/ASAssets/KPImetrics and attempt to fix any impacted resources due to an anomaly in the DV repository.
updateTriggers	This procedure is used to enable and disable the triggers based on the stored in defaultTriggersToEnable().
defaultTriggersToEnable	This view contains a series of UNION statements for each trigger that exists in the KPImetrics. It provides information on the following: triggerName, triggerNumber, recommendation, executeImmediate, isCache and cacheTableName. Example row: 'kpimetricsTrig_01_Cache_ALL_RESOURCES' triggerName, 1 triggerNumber, 'ON' recommendation, 1 executeImmediate, 0 isCache, " cacheTableName

Published Resources

This section outlines the resources that are published under the ASAssets virtual database to expose metrics data. Resources are organized under catalogs and schemas based upon their functionality.

See the section titled KPImetrics Metrics Resources for descriptions of result sets returned by each resource. There are 124 out-of-the-box views described below that can be used to build reports from.

KPImetrics Catalog

AllCustomReports Schema

Resource (Resource Type)	Description
AccessByUserOvertime (TABLE) – cached AccessByUserOvertimeRT – real time	<p>lineage {pm}=physical metadata table</p> <p>Cached report of the most active resources by a user over time. Sorted resource count, user and date. The algorithm for this report is shown below. The group by is the key to this report. The key grouping is on the user first and then the resource.</p> <pre> SELECT "user", "domain", fullname, email, requestdate, resourcepath, resourcename, resourcetype, parentpath, resourcekind, datasourcetype, dataservicename, categoryname, nodehost, nodeport, COUNT(resourceid) countname, CAST(ROUND(MONTHS_BETWEEN(CURRENT_DATE, requestdat),2) AS DECIMAL(19,2)) requestdatemonths, CAST(DAYS_BETWEEN(CURRENT_DATE, requestdate) AS DECIMAL(19,0)) requestdatedays FROM /shared/ASAssets/KPImetrics/Business/Logical/resourceUs age/vResourceUsage GROUP BY "user", "domain", fullname, email, requestdate, resourceid, resourcepath, resourcename, resourcetype, parentpath, resourcekind, datasourcetype, dataservicename, categoryname, nodehost, nodeport </pre> <p>Derived from the following tables: vResourceCountUsersDate (group by clause) → vResourceUsage → metrics_resources_usage_hist {pm}</p>
ActiveResourcesOverPeriodOfTime (TABLE) – cached ActiveResourcesOverPeriodOfTimeRT – real time	<p>Cached report of the most active resources over time. Sorted by resource count and date.</p> <pre> SELECT requestdate, resourcepath, resourcename, resourcetype, parentpath, resourcekind, datasourcetype, dataservicename, categoryname, nodehost, nodeport, COUNT(resourceid) countname, CAST(ROUND(MONTHS_BETWEEN(CURRENT_DATE, requestdate),2) AS DECIMAL(19,2)) requestdatemonths, </pre>

	<p>CAST(DAYS_BETWEEN(CURRENT_DATE, requestdate) AS DECIMAL(19,0)) requestdatedays FROM /shared/ASAssets/KPImetrics/Business/Logical/resourceUsage/vResourceUsage GROUP BY requestdate, resourceid, resourcepath, resourcename, resourcectype, parentpath, resourcekind, datasourcetype, dataservicename, categoryname, nodehost, nodeport</p> <p>Derived from the following tables: vResourceCountDate (group by clause) → vResourceUsage → metrics_resources_usage_hist {pm}</p>
ResourceAccessByUsers (TABLE) – cached ResourceAccessByUsersRT – real time	Cached report of the most used resources by a user with no time period. Derived from the following tables: vResourceCountUsers (group by clause) → vResourceUsage → metrics_resources_usage_hist {pm}
ResourceCount_Details (TABLE) – cached ResourceCount_DetailsRT – real time	Cached detail report of resource count by date. Derived from the following tables: vResourceUsage → metrics_resources_usage_hist {pm}
ResourceCount_Total (TABLE) – cached ResourceCount_TotalRT – real time	Cached roll-up report of resource count by month. Derived from the following tables: vResourceUsage → metrics_resources_usage_hist {pm}
SystemCPUandMemoryStatus (TABLE)	Report of system CPU and memory utilization and DV memory over time. Derived from the following tables: vSystemResources → METRICS_CIS_SYSTEM_RESOURCES{pm}
vEventRequestSqlColumns (TABLE)	Report of the columns accessed by a SQL request query. Derived from the following tables: [METRICS_SQL_COLUMNS{pm}, METRICS_SQL_RESOURCE{pm}, RequestExpanded --> [metrics_requests_hist{pm}, metrics_sessions_hist{pm}]]
vEventRequestSqlResources (TABLE)	Report of the resources accessed by a SQL request query. Derived from the following tables: [METRICS_SQL_REQUEST{pm}, METRICS_SQL_RESOURCE{pm}, RequestExpanded --> [metrics_requests_hist{pm}, metrics_sessions_hist{pm}]]
vResourceUsage (TABLE)	A raw report of all user defined resources. Derived from the following tables: metrics_resources_usage_hist {pm}
vResourceUsagePublished (TABLE)	A raw report of published user defined resources accessed over time where the resourcekind='user defined' and resourcectype='LINK'. Derived from the following tables: vResourceUsage → metrics_resources_usage_hist {pm}

cache Schema

Resource (Resource Type)	Description
--------------------------	-------------

vCache (TABLE)	Report of the KPImetrics cache_status table. Displays all records in the table. Derived from METRICS_SYS_CACHES{pm}.
vCacheActive (TABLE)	Report of the KPImetrics cache_status table. Displays all “ active ” records in the table no matter if the cache is up or down or has a configuration error. Derived from vCache →METRICS_SYS_CACHES{pm}.
vCacheDisabled (TABLE)	Report of the KPImetrics cache_status table. Displays all “ disabled ” records in the table. Derived from vCache →METRICS_SYS_CACHES{pm}.
vCacheIssues (TABLE)	Report of the KPImetrics cache_status table. Displays all records with “ issues ” in the table that have an error state such as DOWN, CONFIG ERROR and NOT LOADED. Derived from vCache →METRICS_SYS_CACHES{pm}.
vCacheSchedule (TABLE)	Report of the KPImetrics cache_status table. Displays all active records in the table with a cache schedule and potential cache schedule dependency. It is ordered by their next schedule refresh time and dependencies upon other cached resources. Derived from vCache →METRICS_SYS_CACHES{pm}.

configuration Schema

Resource (Resource Type)	Description
pMetricsEventRegistrationList (PROCEDURE)	Provides the ability to list the subscriptions for the metrics event registration programmatically.
pMetricsEventRegistrationSubscribe (PROCEDURE)	Provides the ability to subscribe to a metrics event registration programmatically.
pMetricsEventRegistrationUnsubscribe (PROCEDURE)	Provides the ability to unsubscribe to a metrics event registration programmatically.
updateTriggers (PROCEDURE)	Provides the ability to turn on/off triggers programmatically.

metrics Schema

Resource (Resource Type)	Description
Metrics Collection Tables	
metrics_requests (TABLE)	Native DV out-of-the-box metrics requests collection table. Derived from metrics_requests.
metrics_resources_usage (TABLE)	Native DV out-of-the-box metrics resources usage collection table. Derived from metrics_resources_usage.
metrics_sessions (TABLE)	Native DV out-of-the-box metrics sessions collection table. Derived from metrics_sessions.

Metrics Data Mart History Tables	
metrics_requests_hist (TABLE)	Historical metrics requests table. Derived from metrics_requests_hist and metrics_resources_usage_hist. Expanded with user information and resourcekind and dataservicename.
metrics_resources_usage_hist (TABLE)	Historical metrics resources usage table. Derived from metrics_resources_usage_hist.
metrics_sessions_hist (TABLE)	Historical metrics sessions table. Derived from metrics_sessions_hist and metrics_resources_usage_hist. Expanded with user information if found.
Metrics Roll-up Counts	
metrics_all_kpimetrics_table_counts (TABLE)	Provides a rollup of all counts by nodehost and nodeport for all metrics tables.
metrics_all_min_max_starttime_count (TABLE)	Provides a rollup of the min/max starttime/logintime, min/max requestid/sessionid and the total count of rows for each of the 6 metrics collection and historical tables.
Metrics Requests History Roll-up	
metrics_requests_hist_groupby_date (TABLE)	Provides a rollup row count of the metrics_requests_hist table grouped by date.
metrics_requests_hist_groupby_date_month (TABLE)	Provides a rollup row count of the metrics_requests_hist table grouped by month.
metrics_requests_hist_groupby_date_month_nodehost_nodeport (TABLE)	Provides a rollup row count of the metrics_requests_hist table grouped by month, nodehost and nodeport.
metrics_requests_hist_groupby_date_nodehost_nodeport (TABLE)	Provides a rollup row count of the metrics_requests_hist table grouped by date, nodehost and nodeport.
metrics_requests_hist_groupby_date_user_domain (TABLE)	Provides a rollup row count of the metrics_requests_hist table grouped by date, user and domain.
metrics_requests_hist_groupby_date_user_domain_resourcekind (TABLE)	Provides a rollup row count of the metrics_requests_hist table grouped by date, user, domain and resourcekind [system, user defined].
metrics_requests_hist_groupby_date_user_domain_resourcekind_dataservicename (TABLE)	Provides a rollup row count of the metrics_requests_hist table grouped by date, user, domain and resourcekind [system, user defined] and dataservicename.
metrics_requests_hist_groupby_nodehost_nodeport (TABLE)	Provides a rollup row count of the metrics_requests_hist table grouped by nodehost and nodeport.
metrics_requests_hist_groupby_user_domain_nodehost_nodeport (TABLE)	Provides a rollup row count of the metrics_requests_hist table grouped by user, domain, nodehost and nodeport.

metrics_requests_hist_groupby_user_domain_resourcekind_dataservicename (TABLE)	Provides a rollup row count of the metrics_requests_hist table grouped by user, domain, resourcekind [system, user defined] and dataservicename.
metrics_requests_hist_groupby_user_domain_resourcekind_dataservicename_nodehost_nodeport (TABLE)	Provides a rollup row count of the metrics_requests_hist table grouped by user, domain, resourcekind [system, user defined], dataservicename, nodehost and nodeport.
Metrics Resources Usage History Roll-Up	
metrics_resources_usage_hist_groupby_date (TABLE)	Provides a rollup row count of the metrics_resources_usage_hist table grouped by date.
metrics_resources_usage_hist_groupby_date_month (TABLE)	Provides a rollup row count of the metrics_resources_usage_hist table grouped by month.
metrics_resources_usage_hist_groupby_date_month_nodehost_nodeport (TABLE)	Provides a rollup row count of the metrics_resources_usage_hist table grouped by month, nodehost and nodeport.
metrics_resources_usage_hist_groupby_date_nodehost_nodeport (TABLE)	Provides a rollup row count of the metrics_resources_usage_hist table grouped by date, nodehost and nodeport.
metrics_resources_usage_hist_groupby_date_user_domain_resourcekind (TABLE)	Provides a rollup row count of the metrics_resources_usage_hist table grouped by date, user, domain and resourcekind [system, user defined].
metrics_resources_usage_hist_groupby_nodehost_nodeport (TABLE)	Provides a rollup row count of the metrics_resources_usage_hist table grouped by nodehost and nodeport.
metrics_resources_usage_hist_groupby_user_domain_resourcekind (TABLE)	Provides a rollup row count of the metrics_resources_usage_hist table grouped by user, domain and resourcekind [system, user defined].
metrics_resources_usage_hist_groupby_user_domain_resourcekind_dataservicename	Provides a rollup row count of the metrics_resources_usage_hist table grouped by user, domain, resourcekind [system, user defined] and dataservicename.
Metrics Sessions History Roll-up	
metrics_sessions_hist_groupby_clienthost_nodehost_nodeport (TABLE)	Provides a rollup row count of the metrics_sessions_hist table grouped by clienthost, nodehost and nodeport.
metrics_sessions_hist_groupby_date (TABLE)	Provides a rollup row count of the metrics_sessions_hist table grouped by date.
metrics_sessions_hist_groupby_date_nodehost_nodeport (TABLE)	Provides a rollup row count of the metrics_sessions_hist table grouped by date, nodehost and nodeport.
metrics_sessions_hist_groupby_date_type (TABLE)	Provides a rollup row count of the metrics_sessions_hist table grouped by date and type='end'
metrics_sessions_hist_groupby_month (TABLE)	Provides a rollup row count of the metrics_sessions_hist table grouped by month.

metrics_sessions_hist_groupby_month_node_host_nodeport (TABLE)	Provides a rollup row count of the metrics_sessions_hist table grouped by month, nodehost and nodeport.
metrics_sessions_hist_groupby_month_type (TABLE)	Provides a rollup row count of the metrics_sessions_hist table grouped by month and type='end'
metrics_sessions_hist_groupby_nodehost_nodeport (TABLE)	Provides a rollup row count of the metrics_sessions_hist table grouped by nodehost and nodeport.
metrics_sessions_hist_groupby_sessiontype (TABLE)	Provides a rollup row count of the metrics_sessions_hist table grouped by session type [TASK,INTERNAL,JDBC, etc].

requests Schema

Resource (Resource Type)	Description
vEventRegLog (TABLE)	Derived from METRICS_EVENT_REG_LOG. Details about the events that have occurred and the emails that have been sent out. An event is only logged if an email is sent. Events are registered in the METRICS_EVENT_REGISTRATION table and include: LONG_RUNNING, EXCEEDED_MEMORY, INACTIVITY, WORKFLOW_FAILURE, DBMS_SCHEDULER_ERROR, and PURGE_HISTORY.
vEventRegLogLineage (TABLE)	Derived from METRICS_EVENT_REG_LOG_LINEAGE. Details about the data source lineage for a SQL description (request) that results from an event registered by METRICS_EVENT_REGISTRATION. Events that log data source lineage include: LONG_RUNNING and EXCEEDED_MEMORY because they have an associated SQL description that gets parsed in order to determine the lineage.
vEventRequestSqlColumns (TABLE)	Details about the column projection list from the metrics_request description (SQL). These columns were parsed. Derived from the following tables: METRICS_SQL_REQUEST, METRICS_SQL_RESOURCE, METRICS_SQL_COLUMNS, metrics_requests_hist and metrics_sessions_hist
vEventRequestSqlResources (TABLE)	Details about the resources used from the metrics_request description (SQL). These resources were parsed. Derived from the following tables: METRICS_SQL_REQUEST,

	METRICS_SQL_RESOURCE, metrics_requests_hist and metrics_sessions_hist
vEventRequestSqlResourcesAllErrors (TABLE)	Details about errors produced during parsing of the metrics_request description (SQL). This can be used to improve the SQL Parser code implemented by KPImetrics. These resources were parsed. Derived from the following tables: METRICS_SQL_REQUEST, METRICS_SQL_RESOURCE, metrics_requests_hist and metrics_sessions_hist
vEventRequestSqlResourcesCount (TABLE)	Returns a count of the number of unique requestid rows parsed for each metrics_request description (SQL). These resources were parsed. Derived from the following tables: METRICS_SQL_REQUEST, METRICS_SQL_RESOURCE, metrics_requests_hist and metrics_sessions_hist
vExceededMemoryPercentRequests (TABLE)	Details on exceeded memory queries that are occurring at the time this resource is executed. Derived from the system catalog table SYS_REQUESTS.
vGetSystemInformation (TABLE)	Returns nodehost and nodeport for a DV server. Derived from /System/Helpers/pGetSystemInformation() which in turn invokes /lib/util/GetProperty('SERVER_HOSTNAME') and /lib/util/GetProperty('SERVER_JDBC_PORT')
vLongRunningRequests (TABLE)	Details on long running queries that are occurring at the time this resource is executed. Derived from the system catalog table SYS_REQUESTS.
vMetricsSqlColumns (TABLE)	Details of the parsed SQL columns for a user defined request. Derived from METRICS_SQL_COLUMNS with no other join.
vMetricsSqlRequest (TABLE)	Details of the parsed SQL for a user defined request. Derived from METRICS_SQL_REQUEST with no other join. Note the SQL Template is either stored in KPI_DESCRIPTION_TEMPLATE VARCHAR(4000) when the SQL statement is <= 4000 characters or KPI_DESCRIPTION_TEMPLATE_CLOB [CLOB/TEXT] when it is greater than 4000. The vast majority of SQL will be less than 4000. Since the odds are in favor of smaller SQL statements, an index can be placed on KPI_DESCRIPTION_TEMPLATE allowing

	<p>faster access. However, if it is determined that the application has both then the user interface must account for the CLOB field in its queries.</p>
vMetricsSqlRequestLineage (TABLE)	<p>Details of the parsed SQL for a user defined request with the addition of the data source resource lineage. A request may contain many published resources. Each published resource may result in many data source resources. Therefore, the result may contain 0 to many resources and 0 to many data source resources.</p> <p>REQUEST → RESOURCES → DATASOURCE LINEAGE</p>
vMetricsSqlRequestUniqueSqlTemplates (TABLE)	<p>A unique grouping of parsed SQL templates derived from parsing the request “description”. Since the description is stored as a VARCHAR(4000) field it allows push down on queries. Data is only stored in this field when <= 4000 characters. Derived from METRICS_SQL_REQUEST. KPI_DESCRIPTION_TEMPLATE.</p>
vMetricsSqlRequestUniqueSqlTemplatesByUser (TABLE)	<p>A unique grouping of parsed SQL templates sorted by user and derived from parsing the request “description”. Since the description is stored as a VARCHAR(4000) field it allows push down on queries. Data is only stored in this field when <= 4000 characters. Derived from METRICS_SQL_REQUEST. KPI_DESCRIPTION_TEMPLATE and metrics_resources_usage_hist.</p>
vMetricsSqlRequestUniqueSqlTemplatesByUserByDate (TABLE)	<p>A unique grouping of parsed SQL templates sorted by user and request date and derived from parsing the request “description”. Since the description is stored as a VARCHAR(4000) field it allows push down on queries. Data is only stored in this field when <= 4000 characters. Derived from METRICS_SQL_REQUEST. KPI_DESCRIPTION_TEMPLATE and metrics_resources_usage_hist.</p>
vMetricsSqlRequestUniqueSqlTemplatesClob (TABLE)	<p>A unique grouping of parsed SQL templates derived from parsing the request “description”. Warning. Since the description is stored as a CLOB field and push down on CLOB is not supported, this may cause high memory usage in CIS. Derived from METRICS_SQL_REQUEST. KPI_DESCRIPTION_TEMPLATE_CLOB.</p>

vMetricsSqlRequestUniqueSqlTemplatesClobByUser (TABLE)	A unique grouping of parsed SQL templates sorted by user and derived from parsing the request “description”. Warning. Since the description is stored as a CLOB field and push down on CLOB is not supported, this may cause high memory usage in CIS. Derived from METRICS_SQL_REQUEST, KPI_DESCRIPTION_TEMPLATE_CLOB and metrics_resources_usage_hist.
vMetricsSqlRequestUniqueSqlTemplatesClobByUserByDate (TABLE)	A unique grouping of parsed SQL templates sorted by user and request date and derived from parsing the request “description”. Warning. Since the description is stored as a CLOB field and push down on CLOB is not supported, this may cause high memory usage in CIS. Derived from METRICS_SQL_REQUEST, KPI_DESCRIPTION_TEMPLATE_CLOB and metrics_resources_usage_hist.
vMetricsSqlResource (TABLE)	Details of the parsed SQL resources for a user defined request. Derived from METRICS_SQL_RESOURCE with no other joins.
vMetricsSqlResourceLineage (TABLE)	Details of the parsed SQL resources data source lineage for a user defined request. Derived from METRICS_SQL_RESOURCE_LINEAGE with no other joins.
vMetricsSqlResourceLineageCountReport (TABLE)	Provides a count of METRICS_SQL_RESOURCE and METRICS_SQL_RESOURCE_LINEAGE tables under different scenarios: TOTAL METRICS_SQL_RESOURCE COUNT TOTAL METRICS_SQL_RESOURCE COUNT AVAILABLE TO PROCESS TOTAL METRICS_SQL_RESOURCE NOT PROCESSED TOTAL METRICS_SQL_RESOURCE NOT PROCESSED BY NODE TOTAL METRICS_SQL_RESOURCE_LINEAGE PROCESSED TOTAL METRICS_SQL_RESOURCE_LINEAGE PROCESSED BY nodehost TOTAL METRICS_SQL_RESOURCE_LINEAGE PROCESSED BY STATUS=C=[C]opied lineage TOTAL METRICS_SQL_RESOURCE_LINEAGE PROCESSED BY STATUS=L=[L]ineage generated

	<p>TOTAL METRICS_SQL_RESOURCE_LINEAGE PROCESSED BY nodehost, status=C=[C]opied lineage</p> <p>TOTAL METRICS_SQL_RESOURCE_LINEAGE PROCESSED BY nodehost, status=L=[L]ineage generated</p> <p>TOTAL METRICS_SQL_RESOURCE_LINEAGE PROCESSED BY nodehost, starttime, status=C=[C]opied lineage</p> <p>TOTAL METRICS_SQL_RESOURCE_LINEAGE PROCESSED BY nodehost, starttime, status=L=[L]ineage generated</p>
vPublishedResourcePerRequest (TABLE)	Details on published requests correlated with user information. Derived from metrics_resources_usage_hist, metrics_sessions_hist, metrics_requests_hist.
vRequestExpandedAll (TABLE)	Details of a request expanded with user information. All records are displayed. Derived from metrics_requests_hist and metrics_sessions_hist.
vRequestExpandedUD (TABLE)	Details of a request expanded with user information. Only records of resourcekind='user defined' are displayed. Derived from metrics_requests_hist and metrics_sessions_hist.
vRequestsCountsByUser (TABLE)	A count of requests by user and date along with the following per request: avg rows, avg bytes from client, avg bytes to client, min total duration, max total duration, min server duration and max server duration.
vSessions (TABLE)	Details for user sessions. Equivalent to metrics_sessions_hist.
vSessionsUserRequests (TABLE)	Details on requests generated by each user session

resource Schema

Resource (Resource Type)	Description
METRICS_ALL_RESOURCES_groupby_nodehost_nodeport (TABLE)	Report that provides a count by nodehost and nodeport for the resources that exist in METRICS_ALL_RESOURCES
vAllResourcesHist (TABLE)	Report of all resources historically in the DV repository. Derived from METRICS_ALL_RESOURCES and is loaded only with new resources based on a trigger.
vAllResourcesMax (TABLE)	Report of the latest resources in the DV repository. Derived from METRICS_ALL_RESOURCES.

vResourceCount (TABLE)	Report of the count of resources where resourcekind='user defined'. Derived from metrics_resources_usage_hist.
vResourceCountDate (TABLE)	Report of the count of resources by date where resourcekind='user defined'. Derived from metrics_resources_usage_hist.
vResourceCountUsers (TABLE)	Report of the count of resources by user where resourcekind='user defined'. Derived from metrics_resources_usage_hist.
vResourceCountUsersDate (TABLE)	Report of the count of resources by user and date where resourcekind='user defined'. Derived from metrics_resources_usage_hist.
vResourceDistinctPublishedDatabases (TABLE)	This table returns a list of all resources published under a database on the DV server derived from metrics_resources_usage_hist.
vResourceDistinctPublishedResources (TABLE)	This table returns a list of all distinct resources published under a database on the DV server derived from metrics_resources_usage_hist.
vResourceDistinctPublishedWebServices (TABLE)	This table returns a list of all resources published as a web service operation on the DV server derived from metrics_resources_usage_hist.
vResourceDistinctResources (TABLE) →	This table returns a list of all distinct resources on the DV server derived from metrics_resources_usage_hist.
vResourcesPublishedNotUsed (TABLE) →	Report of all published resources present on the DV server that have not been used derived from METRICS_ALL_RESOURCES, metrics_resources_usage_hist.
vResourceUsageAll (TABLE) →	Report of all resources historically where resourcekind='user defined' and 'system'. Derived from METRICS_ALL_RESOURCES, metrics_resources_usage_hist.
vResourceUsageUD (TABLE) →	Report of all resources where resourcekind='user defined'. Derived from metrics_resources_usage_hist.

resourceDataCount Schema

Resource (Resource Type)	Description
getResourceDataCount (PROCEDURE)	This procedure returns a list of the top N most frequently accessed resources for the specified data range. Each row includes a count of the number of rows of data each resource contains

resourceMetadata Schema

Resource (Resource Type)	Description
vResourceListAllPublishedResources (TABLE) →	This table returns a list of all published resources present on the DV server metadata catalog derived from the cached METRICS_ALL_RESOURCES.

systemUsage Schema

Resource (Resource Type)	Description
vCpuMemUtilization (TABLE)	Details on system CPU and memory utilization. Derived from METRICS_CPU_MEMORY_CHECKER.
vDatasourceConnectionChanges (TABLE)	Details on data sources connection changes. Derived from METRICS_SYS_DATASOURCES which was cached from /services/databases/system/SYS_DATASOURCES.
vDatasourceStatusChanges (TABLE)	Details on data source status changes. Derived from METRICS_SYS_DATASOURCES which was cached from /services/databases/system/SYS_DATASOURCES.
vDatasourceUsage (TABLE)	Details on data sources usage. Derived from METRICS_SYS_DATASOURCES which was cached from /services/databases/system/SYS_DATASOURCES.
vDatasourceUsageCurrent (TABLE)	Details on current data sources usage. Derived from METRICS_SYS_DATASOURCES which was cached from /services/databases/system/SYS_DATASOURCES
vLogDisk (TABLE)	This table returns DV server disk events derived from the cached METRICS_LOG_DISK which was cached from /services/databases/system/LOG_DISK. The information includes configured disk size/used, temporary disk size/used and log disk size/used.
vLogIO (TABLE)	This table returns DV server disk events derived from the cached METRICS_LOG_IO which was cached from /services/databases/system/LOG_IO. The information includes bytes from clients, bytes to clients, bytes from data sources and bytes to data sources.
vLogMemory (TABLE)	This table returns DV server disk events derived from the cached METRICS_LOG_MEMORY which was cached from /services/databases/system/LOG_MEMORY. The

	information memory bytes, memory max, managed bytes and managed max.
vSystemResources (TABLE)	Details on system resource usage. Derived from METRICS_CIS_SYSTEM_RESOURCES.

users Schema

Resource (Resource Type)	Description
vAllUsersHist (TABLE)	Details on all users over time derived from METRICS_ALL_USERS.
vAllUsersMax (TABLE)	Details on the latest user derived from METRICS_ALL_USERS.
vLdapPerson (TABLE)	Details regarding an LDAP user derived from METRICS_LDAP_PERSON.

workflow Schema

Resource (Resource Type)	Description
vCISWorkflow (TABLE)	Report to view the status of the KPI metrics triggers derived from METRICS_CIS_WORKFLOW.
vCISWorkflowStatus (TABLE)	Report to view the status of the KPI metrics triggers sort by workflowstatus ASC and derived from METRICS_CIS_WORKFLOW. Failure message come first. Workflowstatus=F, I, S. F=Fail, I=In process, S=Success.
vEventRegistration (TABLE)	Report to view the list of email event registration subscriptions that are currently configured. Derived from METRICS_EVENT_REGISTRATION.
vJobDetails (TABLE)	Report to view the status of the METRICS_JOB_DETAILS DBMS Scheduler table. Provides a sorted status by most recent rows first. The information informs the user of the status of the data transfer from the collection tables to the history tables for each of the 3 metrics tables.
vJobDetailsReport (TABLE)	Report to summarize the results of METRICS_JOB_DETAILS grouped by the job table name, node host and node port. It provides various averages, min and max of duration and number of rows. It provides summation of rows inserted, updated, deleted and not inserted.
vJobEnvironments (TABLE)	Report to view the list of valid environments derived from METRICS_JOB_ENVIRONMENTS
vJobFilters (TABLE)	Report to view the list of job filters that are currently configured derived METRICS_JOB_FILTERS.

vSqlControl	Derived from METRICS_SQL_CONTROL. Normally has no rows unless Cache_METRICS_SQL_REQUEST_EXEC is executing. If it contains rows then they indicate that a particular node is performing the pre-processing for the METRICS_SQL_REQUEST table. Only one node at a time can perform this processing due to the fact that any single node can process another nodes SQL requests in order to achieve parallel processing. This table acts as the control table to allow each node to select their requestids that they will process. This is why this part of the process is single-threaded and must be controlled.
vSqlControlLog	Derived from METRICS_SQL_CONTROL_LOG. Contains a log of the Gatekeeper code block within Cache_METRICS_SQL_REQUEST_EXEC when debugGatekeeper = '1'. During normal operation, this will be turned off so that no rows are produced. It is only useful for debugging purposes to insure that in a clustered environment, the nodes are taking their turn initializing their own set of rows when doing parallel processing on one of the node's data.

Data Sources

This section outlines the data sources created, populated and used by KPImetrics project.

Metadata Data Source for LDAP

Location:

/shared/ASAssets/KPImetrics/Physical/Metadata/LDAP1

/shared/ASAssets/KPImetrics/Physical/Metadata/LDAP2

The data source LDAP1 and LDAP2 are LDAP data sources that connects to a client's corporate LDAP directory to lookup user information and their relation to client hierarchy. The data source's URL property should be modified to allow the data source to successfully connect to and query the target LDAP directory server. The LDAP structure may vary and could use LDAP person, organizationalPerson, inetOrgPerson or some other customer LDAP table. Determine the correct table to find person data and introspected it under this data source for LDAP integration to function successfully.

There are two data sources provided out of the box in the event there are two DV domains configured. Add additional data sources if there are more than two domains. Each data source directly corresponds to a DV domain.

It is strongly recommended that this data source should not use the same LDAP account as the DV server uses to authenticate LDAP users. This may result in the LDAP account being locked if the data source's credentials are not updated when the account's password is changed.

Metadata Data Source for CPUAndMemChecker

Location: /shared/ASAssets/KPImetrics/Physical/Metadata/CPUAndMemChecker

CPUAndMemChecker custom java procedure is used to capture system level CPU and Memory usage at the operating system level. On a linux server installation, CPUAndMemChecker invokes two shell scripts (KPImetricsTopCommandGrepCpu_linux7.sh and KPImetricsFreeMemCommand_linux7.sh) to execute 'top' and 'free' commands to returns CPU percentage, used memory and available memory. Windows installations use a couple of powershell scripts (KPImetricsCpuUtilization.ps1 and KPImetricsMemUtilization.ps1) to perform the same capability.

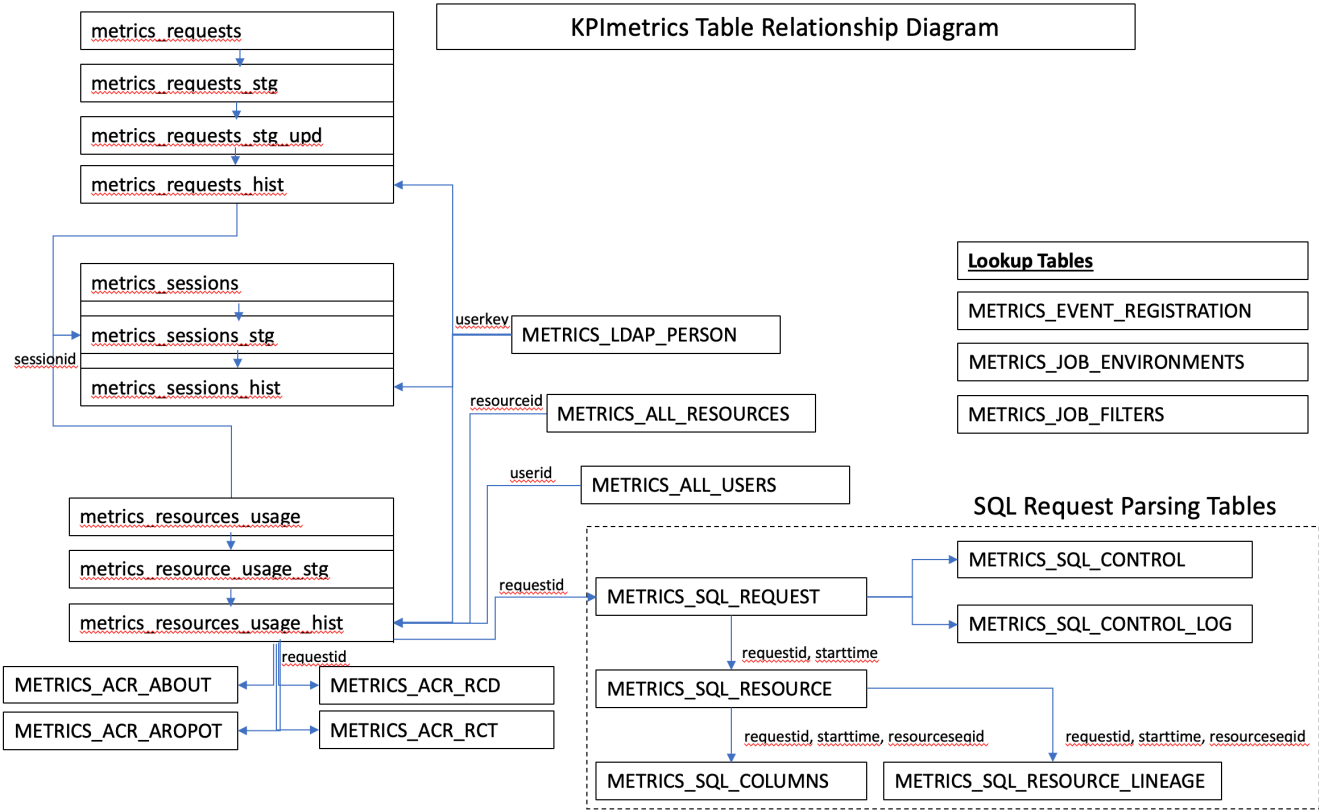
The CPUAndMemChecker procedure is invoked by Cache_CIS_SYSTEM_RESOURCES script and inserts the results in METRICS_CIS_SYSTEM_RESOURCES table of the KPImetrics data source.

The CPUAndMemChecker procedure exposes one procedure that has following parameters:

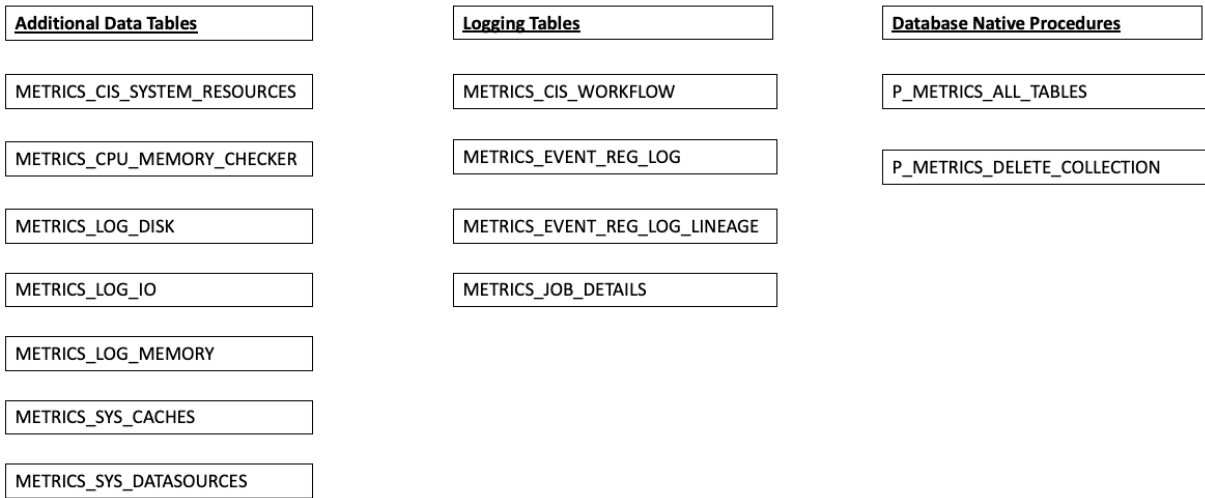
Parameter Name	Direction	Description
debug	IN	Y=debug values written to cs_server.log. N=no debug.
cpuScriptNameOrCommand	IN	Path to script to return CPU utilization. Windows example: powershell.exe -file C:\CIS7\bin\KPImetricsCpuUtilization.ps1 UNIX example: /CIS7/bin/KPImetricsTopCommandGrepCpu_linux7.sh
memScriptNameOrCommand	IN	Path to script to return memory utilization. Windows example: powershell.exe -file C:\CIS7\bin\KPImetricsMemUtilization.ps1 UNIX example: /CIS7/bin/KPImetricsFreeMemCommand_linux7.sh
cpuUsedPercent	OUT	Average CPU utilization percentage reported by the server's operating system
memoryUsedMb	OUT	Used memory in Megabytes reported by the server's operating system
memoryAvailMb	OUT	Available memory in Megabytes reported by the server's operating system

Metadata Table Relationship Diagram

The following provides a table relationship diagram for the database tables and procedures used by KPImetrics.



KPImetrics Table Relationship Diagram



Metadata Data Source Tables and Procedures for KPI_<database_type>

Location: /shared/ASAssets/KPImetrics/Physical/Metadata/KPI_<database_type>

The KPImetrics module provides data source for all currently supported storage database platforms under /shared/ASAssets/KPImetrics/Physical/Metadata.

Currently the KPImetrics module includes the following KPImetrics data sources

- /shared/ASAssets/KPImetrics/Physical/Metadata/KPI_oracle
- /shared/ASAssets/KPImetrics/Physical/Metadata/KPI_sqlserver

During deployment, the KPImetrics module must be configured to use the data source appropriate for the target KPImetrics database platform. The following instructions refer to this data source at the KPImetrics data source for simplicity

The KPImetrics data source is used to capture

- Historical server metrics captured using incremental caching
- Some pre-processed data for final reporting

The following tables have been created in CIS_KPI schema to capture the required data.

Table Name	Description
P_METRICS_ALL_TABLES	<p>This procedure is used by DV to execute the native database data transfer procedure which formats and moves data from the metrics collection tables to the metrics history tables on a scheduled basis.</p> <p>TRIGGER: /KPImetrics/Physical/Physical/KPI_oracle/kpimetricsTrig_30_DBMSScheduler_KPI → /KPImetrics/Physical/Metadata/System/ClusterSafeCache/pMETRICS_ALL_TABLES_exec → /Physical/Physical/Abstraction/P_METRICS_ALL_TABLES →</p> <p>Oracle Lineage: PROCEDURES: /KPImetrics/Physical/Physical/KPI_oracle/P_METRICS_ALL_TABLES → /KPImetrics/Physical/Metadata/KPI_oracle_11g/<schema>/P_METRICS_ALL_TABLES</p> <p>SQL Server Lineage: PROCEDURES: /KPImetrics/Physical/Physical/KPI_sqlserver/P_METRICS_ALL_TABLES → /KPImetrics/Physical/Metadata/KPI_sqlserver_2012/<catalog>/<schema>/P_METRICS_ALL_TABLES</p>
P_METRICS_DELETE_COLLECTION	<p>This procedure I sused by DV to execute native database SQL for deleting the collection table data.</p> <p>TRIGGER: /Physical/Physical/KPI_oracle/kpimetricsTrig_33_DeleteCollection → /Physical/Metadata/System/ClusterSafeCache/pMETRICS_DELETE_COLLECTION_exec → /Physical/Physical/Abstraction/P_METRICS_DELETE_COLLECTION →</p> <p>Oracle Lineage: PROCEDURES: /Physical/Physical/KPI_oracle/P_METRICS_DELETE_COLLECTION → /Physical/Metadata/KPI_oracle_11g/<schema>/P_METRICS_DELETE_COLLECTION</p> <p>SQL Server Lineage: PROCEDURES: /Physical/Physical/KPI_sqlserver/P_METRICS_DELETE_COLLECTION → /Physical/Metadata/KPI_sqlserver_2012/<catalog>/<schema>/P_METRICS_DELETE_COLLECTION</p>
METRICS_ACR_ABUOT	<p>The cache table acronym “METRICS_ACR_ABUOT” must be short and thus it stands for “All Custom Reports Access By User Over Time”. It is cached once a day from the real-time query “ACR_AccessByUserOvertime”. It is access from AllCustomReports.AccessByUserOvertime.</p> <p>TRIGGER: /KPImetrics/Physical/Metadata/System/ClusterSafeTriggers/ kpimetricsTrig_19_AllCustom_AccessByUserOverTime →</p> <p>PROCEDURE: /KPImetrics/Physical/Metadata/System/ClusterSafeCache/ Cache_AllCustom_AccessByUserOverTime →</p> <p>READ: /KPImetrics/Physical/Metadata/System/ClusterSafeCache/ACR_AccessByUserOverTime →</p> <p>INSERT:</p>

	/KPImetrics/Physical/Physical/Abstraction/METRICS_ACR_ABUOT
METRICS_ACR_AROPOT	<p>The cache table acronym “METRICS_ACR_AROPOT” must be short and thus it stands for “All Custom Reports Active Resources Over Period Of Time”. It is cached once a day from the real-time query “ACR_ActiveResourcesOverPeriodOfTime”. It is access from AllCustomReports.ActiveResourcesOverPeriodOfTime.</p> <p><u>TRIGGER:</u> /KPImetrics/Physical/Metadata/System/ClusterSafeTriggers/ kpimetricsTrig_20_AllCustom_ActiveResourcesOverPeriodOfTime →</p> <p><u>PROCEDURE:</u> /KPImetrics/Physical/Metadata/System/ClusterSafeCache/ Cache_AllCustom_ActiveResourcesOverPeriodOfTime →</p> <p><u>READ:</u> /KPImetrics/Physical/Metadata/System/ClusterSafeCache/ ACR_ActiveResourcesOverPeriodOfTime →</p> <p><u>INSERT:</u> /KPImetrics/Physical/Physical/Abstraction/METRICS_ACR_AROPOT</p>
METRICS_ACR_RCD	<p>The cache table acronym “METRICS_ACR_RCD” must be short and thus it stands for “All Custom Reports Resource Count Details”. It is cached once a day from the real-time query “ACR_ResourceCount_Details”. It is access from AllCustomReports.ResourceCount_Details.</p> <p><u>TRIGGER:</u> /KPImetrics/Physical/Metadata/System/ClusterSafeTriggers/ kpimetricsTrig_21_AllCustom_ResourceCount_Details →</p> <p><u>PROCEDURE:</u> /KPImetrics/Physical/Metadata/System/ClusterSafeCache/ Cache_AllCustom_ResourceCount_Details →</p> <p><u>READ:</u> /KPImetrics/Physical/Metadata/System/ClusterSafeCache/ACR_ResourceCountDetails →</p> <p><u>INSERT:</u> /KPImetrics/Physical/Physical/Abstraction/METRICS_ACR_RCD</p>
METRICS_ACR_RCT	<p>The cache table acronym “METRICS_ACR_RCT” must be short and thus it stands for “All Custom Reports Resource Count Total”. It is cached once a day from the real-time query “ACR_ResourceCount_Total”. It is access from AllCustomReports.ResourceCount_Total.</p> <p><u>TRIGGER:</u> /KPImetrics/Physical/Metadata/System/ClusterSafeTriggers/ kpimetricsTrig_22_AllCustom_ResourceCount_Total →</p> <p><u>PROCEDURE:</u> /KPImetrics/Physical/Metadata/System/ClusterSafeCache/ Cache_AllCustom_ResourceCount_Total→</p> <p><u>READ:</u> /KPImetrics/Physical/Metadata/System/ClusterSafeCache/ACR_ResourceCount_Total →</p> <p><u>INSERT:</u> /KPImetrics/Physical/Physical/Abstraction/METRICS_ACR_RCT</p>
METRICS_ALL_RESOURCES	<p>This table stores a cached representation of the DV system ALL_RESOURCES table because it is more efficient to query. It is processed using the system interface lineage:</p> <p><u>TRIGGER:</u> /KPImetrics/Physical/Metadata/System/ClusterSafeTriggers/ kpimetricsTrig_01_Cache_ALL_RESOURCES →</p> <p><u>PROCEDURE:</u> /KPImetrics/Physical/Metadata/System/ClusterSafeCache/CACHE_ALL_RESOURCES →</p>

	<p><u>READ:</u> /KPImetrics/Physical/Metadata/System/ALL_RESOURCES →</p> <p><u>INSERT:</u> /KPImetrics/Physical/Physical/Abstraction/METRICS_ALL_RESOURCES_STG → /KPImetrics/Physical/Physical/Abstraction/METRICS_ALL_RESOURCES</p>
METRICS_ALL_RESOURCES_STG	<p>This table is used to stage the data being inserted into METRICS_ALL_RESOURCES.</p>
METRICS_ALL_USERS	<p>This table stores user information for all user accounts that have executed queries against the DV instance. This is an incremental cache target table. It is processed using the system interface lineage:</p> <p><u>TRIGGER:</u> /KPImetrics/Physical/Metadata/System/ClusterSafeTriggers/ kpimetricsTrig_02_Cache_ALL_USERS →</p> <p><u>PROCEDURE:</u> /KPImetrics/Physical/Metadata/System/ClusterSafeCache/Cache_ALL_USERS →</p> <p><u>READ:</u> /KPImetrics/Physical/Metadata/System/ALL_USERS →</p> <p><u>INSERT:</u> /KPImetrics/Physical/Physical/Abstraction/METRICS_ALL_USERS</p>
METRICS_CIS_SYSTEM_RESOURCES	<p>This tables stores hourly snapshots of memory, disk and I/O usage captured from SYS_MEMORY, SYS_STORAGE and SYS_IO tables. It is processed using the system interface lineage:</p> <p><u>TRIGGER:</u> /KPImetrics/Physical/Metadata/System/ClusterSafeTriggers/ kpimetricsTrig_04_Cache_CIS_SYSTEM_RESOURCES→</p> <p><u>PROCEDURE:</u> /KPImetrics/Physical/Metadata/System/ClusterSafeCache/ Cache_CIS_SYSTEM_RESOURCES→</p> <p><u>READ:</u> /KPImetrics/Physical/Physical/Abstraction/METRICS_CIS_SYSTEM_RESOURCES + /KPImetrics/Physical/Metadata/System/Helpers/p15MinutesIncrements</p> <p><u>INSERT:</u> /KPImetrics/Physical/Physical/Abstraction/METRICS_CIS_SYSTEM_RESOURCES</p>
METRICS_CIS_WORKFLOW	<p>The CIS_WORKFLOW table manages workflows for processing data.</p> <p><u>PROCEDURE:</u> /KPImetrics/Physical/Metadata/System/Helpers/pStartWorkflow /KPImetrics/Physical/Metadata/System/Helpers/pEndWorkflow</p> <p><u>INSERT:</u> /KPImetrics/Physical/Physical/Abstraction/METRICS_CIS_WORKFLOW</p>
METRICS_CPU_MEMORY_CHECKER	<p>This table stores the results of executing the CPUAndMemChecker procedure over time. This is an incremental cache target table. It is processed using the system interface lineage:</p> <p><u>TRIGGER:</u> /KPImetrics/Physical/Metadata/System/ClusterSafeTriggers/ kpimetricsTrig_05_Cache_CPU_MEMORY_CHECKER→</p> <p><u>PROCEDURE:</u> /KPImetrics/Physical/Metadata/System/ClusterSafeCache/ Cache_CPU_MEMORY_CHECKER→</p> <p><u>READ:</u></p>

	<p>/KPImetrics/Physical/Metadata/System/CPU_MEMORY_CHECKER → /KPImetrics/Physical/Metadata/CPUAndMemChecker/CpuAndMemCheckerCjp → <u>CPU Utilization</u> KPImetricsCpuUtilization.ps1 KPImetricsTopCommandGrepCpu_linux6.sh KPImetricsTopCommandGrepCpu_linux7.sh <u>Memory Utilization</u> KPImetricsMemUtilization.ps1 KPImetricsFreeMemCommand_linux6.sh KPImetricsFreeMemCommand_linux7.sh <u>INSERT:</u> /KPImetrics/Physical/Physical/Abstraction/METRICS_CPU_MEMORY_CHECKER</p>
<p>METRICS_EVENT_REG_LOG</p>	<p>This table stores the details about the events that have occurred and the emails that have been sent out. An event is only logged if an email is sent. Events are registered in the METRICS_EVENT_REGISTRATION table and include: LONG_RUNNING, EXCEEDED_MEMORY, INACTIVITY and PURGE_HISTORY. The event time, the user, actual email along with the SQL description when applicable is stored. <u>TRIGGER:</u> /KPImetrics/Physical/Metadata/System/ClusterSafeTriggers/ 1. kpimetricsTrig_15_CheckMetricsActivity → Event=INACTIVITY 2. kpimetricsTrig_16_PurgeHistoryData → Event=PURGE_HISTORY 3. kpimetricsTrig_17_CheckExceedMemoryPercentRequests → Event=EXCEEDED_MEMORY 4. kpimetricsTrig_18_CheckLongRunningRequests →Event=LONG_RUNNING <u>PROCEDURE:</u> /KPImetrics/Physical/Metadata/System/ClusterSafeCache/ 1. pCheckMetricsActivity → ../Helpers/pUpdateEventRegLog 2. pPurgeData → ../Helpers/pUpdateEventRegLog 3. pCheckExceededMemoryPercentRequests → ../Helpers/pUpdateEventRegLog 4. pCheckLongRunningRequests → ../Helpers/pUpdateEventRegLog <u>READ:</u> /shared/ASAssets/KPImetrics/Physical/Metadata/System/ 1. METRICS_JOB_DETAILS 2. N/A 3. /KPImetrics/Business/Business/requests/vExceededMemoryPercentRequests 4. /KPImetrics/Business/Business/requests/vLongRunningRequests <u>INSERT:</u> /KPImetrics/Physical/Physical/Abstraction/METRICS_EVENT_REG_LOG and METRICS_EVENT_REG_LOG_LINEAGE</p>
<p>METRICS_EVENT_REG_LOG_LINEAGE</p>	<p>This table is a child table to METRICS_EVENT_REG_LOG and is used to store the data source lineage for the SQL request that was logged as a result of either LONG_RUNNING or EXCEEDED_MEMORY. The other events do not produce a SQL description so are not applicable. A SQL description is parsed using the same procedures that produce the METRICS_SQL_REQUEST data and store one or more rows associated with the requests data source and connection information. The reporting on this can look for events and quickly determine which data source is being accessed which can assist the viewer on how to take action with the data owners if necessary.</p>
<p>METRICS_EVENT_REGISTRATION</p>	<p>This table is used to register a metrics event registration. A subscription in the METRICS_EVENT_REGISTRATION table consists of a unique record for the combination of SUBSCRIBER_EMAIL, GROUP_NAME, ENVIRONMENT_TYPE,</p>

	EVENT_TYPE and REQUESTER_EMAIL. The EVENT_TYPE can be one of [LONG_RUNNING EXCEEDED_MEMORY INACTIVITY WORKFLOW_FAILURE DBMS_SCHEDULER_ERROR].
METRICS_JOB_DETAILS	This table is used to hold the data transfer job details when rows are moved from the metrics collection tables to the history tables. The native database procedure “P_METRICS_ALL_TABLES” performs the following data transfer capabilities: metrics_sessions → metrics_sessions_hist metrics_resources_usage → metrics_resources_usage_hist metrics_requests → metrics_requests_hist
METRICS_JOB_ENVIRONMENTS	This table provides a list of valid environments. In essence, the ENV_TYPE is like the short nickname for a host. For example, DEV1 is the short-name the development server. This is used by various email notification procedures.
METRICS_JOB_FILTERS	This table is used to hold the job filters used by “P_METRICS_ALL_TABLES” whereby the metrics_resource_usage collection table filters out rows based on user, domain and resourcekind. This capability allows rows to be filtered out before they get to the history table thus reducing the overall burden. Without this feature, the database would be overwhelmed by millions of unnecessary rows as absolutely everything in DV is reported.
METRICS_LDAP_PERSON	This LDAP_PERSON table is used to pre-cache LDAP user information. It is processed using the system interface lineage: <u>TRIGGER:</u> /KPImetrics/Physical/Metadata/System/ClusterSafeTriggers/ kpimetricsTrig_03_Cache_LDAP_PERSON→ <u>PROCEDURE:</u> /KPImetrics/Physical/Metadata/System/ClusterSafeCache/Cache_LDAP_PERSON→ <u>READ:</u> /KPImetrics/Physical/Metadata/System/LDAP_PERSON <u>INSERT:</u> synchronous cache refresh executed on one of applicable tables: /KPImetrics/Physical/Physical/KPI_oracle/METRICS_LDAP_PERSON /KPImetrics/Physical/Physical/KPI_sqlserver/METRICS_LDAP_PERSON
METRICS_LOG_DISK	This table stores logs of available disk space incrementally cached from the DV system table LOG_DISK. It is processed using the system interface lineage: <u>TRIGGER:</u> /KPImetrics/Physical/Metadata/System/ClusterSafeTriggers/ kpimetricsTrig_06_Cache_LOG_DISK→ <u>PROCEDURE:</u> /KPImetrics/Physical/Metadata/System/ClusterSafeCache/Cache_LOG_DISK→ <u>READ:</u> /KPImetrics/Physical/Metadata/System/LOG_DISK <u>INSERT:</u> /KPImetrics/Physical/Physical/Abstraction/METRICS_LOG_DISK
METRICS_LOG_IO	This table stores IO logs incrementally cached from the DV system table LOG_IO. It is processed using the system interface lineage: <u>TRIGGER:</u> /KPImetrics/Physical/Metadata/System/ClusterSafeTriggers/ kpimetricsTrig_07_Cache_LOG_IO→ <u>PROCEDURE:</u> /KPImetrics/Physical/Metadata/System/ClusterSafeCache/Cache_LOG_IO→

	<p><u>READ:</u> /KPImetrics/Physical/Metadata/System/LOG_IO</p> <p><u>INSERT:</u> /KPImetrics/Physical/Physical/Abstraction/METRICS_LOG_IO</p>
METRICS_LOG_MEMORY	<p>This table stores jvm memory logs incrementally cached from the DV system table LOG_MEMORY. It is processed using the system interface lineage:</p> <p><u>TRIGGER:</u> /KPImetrics/Physical/Metadata/System/ClusterSafeTriggers/ kpimetricsTrig_08_Cache_LOG_MEMORY →</p> <p><u>PROCEDURE:</u> /KPImetrics/Physical/Metadata/System/ClusterSafeCache/Cache_LOG_MEMORY →</p> <p><u>READ:</u> /KPImetrics/Physical/Metadata/System/LOG_MEMORY</p> <p><u>INSERT:</u> /KPImetrics/Physical/Physical/Abstraction/METRICS_LOG_MEMORY</p>
metrics_requests	<p>This is the DV out-of-the-box requests table. It cannot contain any indexes. Its purpose is simply to be a collector of DV metrics requests.</p>
metrics_requests_hist	<p>This is the KPImetrics historical requests table. It may be partitioned and contain indexes for better query performance. It is updated periodically from the metrics_requests table. If partitioned by month, it allows a more efficient way to purge data by dropping a partition rather than deleting data.</p>
metrics_requests_stg	<p>This is a KPImetrics stage tables used to perform updates and deletes on metrics_requests collection data before inserting into metrics_requests_hist.</p>
metrics_requests_stg_upd	<p>This is a KPImetrics historical stage table that is used during the data transfer procedure for doing a mass update on the metrics_requests_stg table with Oracle.</p>
metrics_resources_usage	<p>This is the DV out-of-the-box resources usage table. It cannot contain any indexes. Its purpose is simply to be a collector of DV metrics resources usage.</p>
metrics_resources_usage_stg	<p>This is a KPImetrics stage tables used to perform updates on metrics_resources_usage collection data before inserting into metrics_resources_usage_hist.</p>
metrics_resources_usage_hist	<p>This is the KPImetrics historical resources usage table. It may be partitioned and contain indexes for better query performance. It is updated periodically from the metrics_resources_usage table. If partitioned by month, it allows a more efficient way to purge data by dropping a partition rather than deleting data.</p>
metrics_sessions	<p>This is the DV out-of-the-box sessions table. It cannot contain any indexes. Its purpose is simply to be a collector of DV metrics sessions.</p>
metrics_sessions_stg	<p>This is a KPImetrics stage tables used to perform updates on metrics_sessions collection data before inserting into metrics_sessions_hist.</p>
metrics_sessions_hist	<p>This is the KPImetrics historical sessions table. It may be partitioned and contain indexes for better query performance. It is updated periodically from the metrics_sessions table. If partitioned by month, it allows a more efficient way to purge data by dropping a partition rather than deleting data.</p>
METRICS_SQL_COLUMNS	<p>This table holds the parsed SQL columns for the query statement. It is processed using the system interface lineage:</p> <p><u>TRIGGER:</u> /KPImetrics/Physical/Metadata/System/ClusterSafeTriggers/ kpimetricsTrig_10_Cache_METRICS_SQL_REQUEST and kpimetricsTrig_11_Cache_METRICS_SQL_REQUEST_REPROCESS→</p>

	<p><u>PROCEDURE:</u> /KPImetrics/Physical/Metadata/System/ClusterSafeCache/ Cache_METRICS_SQL_REQUEST_EXEC → Cache_METRICS_SQL_REQUEST_GENSQL and Cache_METRICS_SQL_REQUEST_EXEC_REPROCESS→ Cache_METRICS_SQL_REQUEST_GENSQL_REPROCESS</p> <p><u>READ:</u> /KPImetrics/Physical/Formatting/metrics_requests_hist + /KPImetrics/Physical/Formatting/METRICS_SQL_REQUEST</p> <p><u>INSERT:</u> /KPImetrics/Physical/Physical/Abstraction/METRICS_SQL_COLUMNS</p>
METRICS_SQL_REQUEST	<p>This table holds the parsed SQL status for the query statement. It is processed using the system interface lineage:</p> <p><u>TRIGGER:</u> /KPImetrics/Physical/Metadata/System/ClusterSafeTriggers/ kpimetricsTrig_10_Cache_METRICS_SQL_REQUEST and kpimetricsTrig_11_Cache_METRICS_SQL_REQUEST_REPROCESS→</p> <p><u>PROCEDURE:</u> /KPImetrics/Physical/Metadata/System/ClusterSafeCache/ Cache_METRICS_SQL_REQUEST_EXEC → Cache_METRICS_SQL_REQUEST_GENSQL and Cache_METRICS_SQL_REQUEST_EXEC_REPROCESS→ Cache_METRICS_SQL_REQUEST_GENSQL_REPROCESS</p> <p><u>READ:</u> /KPImetrics/Physical/Formatting/metrics_requests_hist + /KPImetrics/Physical/Formatting/METRICS_SQL_REQUEST</p> <p><u>INSERT:</u> /KPImetrics/Physical/Physical/Abstraction/METRICS_SQL_REQUEST</p>
METRICS_SQL_CONTROL	<p>GATEKEEPER - SEMAPHORE: Used by Cache_METRICS_SQL_REQUEST_EXEC.</p> <p>This section of code acts a semaphore to determine whether to process records or not. This procedure is allowed to process the records that were generated from a different node in order to more efficiently share in the workload. This is important as some nodes in a cluster may generate more rows than others due to an imbalance in the load balancer due to “sticky” setting instead of a true “round-robin” setting.</p> <p>Upon entering this procedure, if there are rows in the METRICS_SQL_REQUEST table where KPI_PROCESS_TIME is null and the PROCESSED_NODE_HOST and PROCESSED_NODE_PORT matches this nodes nodehost and nodeport, then it must complete its current work before starting any new work.</p> <p>This procedure can only execute the pre-processing by one node at a time within a cluster because this procedure will work on data from other nodes if the current node has no more work to do. It is imperative that each node be allowed time to pre-insert the number of rows identified by the variable "numRowsToProcessBeforeExiting" into the METRICS_SQL_REQUEST table. For example, 500 rows would be pre-inserted with the KPI_PROCESS_TIME being set to null. The actual processing will take place based on those rows.</p> <p>As soon as the pre-processing has completed, the control record will be removed from the METRICS_SQL_CONTROL table thus allowing another node to perform its pre-processing.</p>
METRICS_SQL_CONTROL_LOG	<p>Contains a log of the Gatekeeper code block within Cache_METRICS_SQL_REQUEST_EXEC when debugGatekeeper = '1'.</p>

	<p>During normal operation, this will be turned off so that no rows are produced. It is only useful for debugging purposes to ensure that in a clustered environment, the nodes are taking their turn initializing their own set of rows when doing parallel processing on one of the node's data.</p>
METRICS_SQL_RESOURCE	<p>This table holds the parsed SQL resource for the query statement. It is processed using the system interface lineage:</p> <p>TRIGGER: /KPImetrics/Physical/Metadata/System/ClusterSafeTriggers/ kpiMetricsTrig_11_Cache_METRICS_SQL_REQUEST and kpiMetricsTrig_12_Cache_METRICS_SQL_REQUEST_REPROCESS →</p> <p>PROCEDURE: /KPImetrics/Physical/Metadata/System/ClusterSafeCache/ Cache_METRICS_SQL_REQUEST_EXEC → Cache_METRICS_SQL_REQUEST_GENSQL and Cache_METRICS_SQL_REQUEST_EXEC_REPROCESS → Cache_METRICS_SQL_REQUEST_GENSQL_REPROCESS</p> <p>READ: /KPImetrics/Physical/Formatting/metrics_requests_hist + /KPImetrics/Physical/Formatting/METRICS_SQL_REQUEST</p> <p>INSERT: /KPImetrics/Physical/Physical/Abstraction/METRICS_SQL_RESOURCE</p>
METRICS_SQL_RESOURCE_LINEAGE	<p>This table holds the parsed SQL data source resource for each resource in the query statement. It is processed using the system interface lineage:</p> <p>TRIGGER: /KPImetrics/Physical/Metadata/System/ClusterSafeTriggers/ kpiMetricsTrig_23_Cache_METRICS_SQL_RESOURCE_LINEAGE</p> <p>PROCEDURE: /KPImetrics/Physical/Metadata/System/ClusterSafeCache/ Cache_METRICS_SQL_RESOURCE_LINEAGE</p> <p>READ: /KPImetrics/Physical/Physical/Abstraction/METRICS_SQL_RESOURCE + /KPImetrics/Physical/Physical/Abstraction/METRICS_SQL_RESOURCE_LINEAGE</p> <p>INSERT: /KPImetrics/Physical/Physical/Abstraction/METRICS_SQL_RESOURCE_LINEAGE</p>
METRICS_SYS_CACHES	<p>This table stores a list of all cached resources and their current statuses incrementally cached from the DV system table SYS_CACHES. It is processed using the system interface lineage:</p> <p>TRIGGER: /KPImetrics/Physical/Metadata/System/ClusterSafeTriggers/ kpiMetricsTrig_12_Cache_SYS_CACHES →</p> <p>PROCEDURE: /KPImetrics/Physical/Metadata/System/ClusterSafeCache/Cache_SYS_CACHES →</p> <p>READ: /KPImetrics/Physical/Metadata/System/SYS_CACHES</p> <p>INSERT: /KPImetrics/Physical/Physical/Abstraction/METRICS_SYS_CACHES</p>

METRICS_SYS_DATASOURCES	<p>This table stores a list of all data sources and their current statuses incrementally cached from the DV system table SYS_DATASOURCES. It is processed using the system interface lineage:</p> <p><u>TRIGGER:</u> /KPImetrics/Physical/Metadata/System/ClusterSafeTriggers/ kpimetricsTrig_13_Cache_SYS_DATASOURCES →</p> <p><u>PROCEDURE:</u> /KPImetrics/Physical/Metadata/System/ClusterSafeCache/Cache_SYS_DATASOURCES →</p> <p><u>READ:</u> /KPImetrics/Physical/Metadata/System/SYS_DATASOURCES</p> <p><u>INSERT:</u> /KPImetrics/Physical/Physical/Abstraction/METRICS_SYS_DATASOURCES</p>
-------------------------	--

Metadata System Triggers and Load Scripts

Location: /shared/ASAssets/KPImetrics/Physical/Metadata/System

/ClusterSafeCache

/ClusterSafeTriggers

/Helpers

This section provides a quick summary of all triggers, their schedules and how they execute in a cluster.

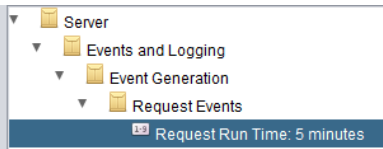
Trigger Name	Trigger Schedule	Trigger Period	Cluster execution
kpimetricsTrig_00_CheckMetricsActivityDebug	1:00 AM	1 hour	all nodes
kpimetricsTrig_01_Cache_ALL_RESOURCES	1:00 AM	1 hour	all nodes
kpimetricsTrig_02_Cache_ALL_USERS	1:00 AM	1 hour	all nodes
kpimetricsTrig_03_Cache_LDAP_PERSON	7:00 AM	1 day	only once per cluster
kpimetricsTrig_04_Cache_CIS_SYSTEM_RESOURCES	1:15 AM	1 hour	all nodes
kpimetricsTrig_05_Cache_CPU_MEMORY_CHECKER	1:00 AM	1 hour	all nodes
kpimetricsTrig_06_Cache_LOG_DISK	1:00 AM	1 hour	all nodes
kpimetricsTrig_07_Cache_LOG_IO	1:00 AM	1 hour	all nodes
kpimetricsTrig_08_Cache_LOG_MEMORY	1:00 AM	1 hour	all nodes
kpimetricsTrig_10_Cache_METRICS_SQL_REQUEST	1:15 AM	30 min	all nodes
kpimetricsTrig_11_Cache_METRICS_SQL_REQUEST_REPROCESS	1:00 AM	1 day	all nodes
kpimetricsTrig_12_Cache_SYS_CACHES	1:00 AM	6 hours	all nodes
kpimetricsTrig_13_Cache_SYS_DATASOURCES	6:00 AM	12 hours	all nodes
kpimetricsTrig_14_CheckCISWorkflowStatusFail	1:30 AM	1 hour	all nodes
kpimetricsTrig_15_CheckMetricsActivity	12:59:50 AM	1 day	all nodes
kpimetricsTrig_16_PurgeHistoryData	4:30 AM	1 day	all nodes
kpimetricsTrig_17_CheckExceedMemoryPercentRequests	1:00 AM	1 hour	all nodes
kpimetricsTrig_18_CheckLongRunningRequests	system event	RequestRunForTooLong	all nodes
kpimetricsTrig_19_AllCustom_AccessByUserOverTime	12:15 AM	1 day	all nodes
kpimetricsTrig_20_AllCustom_ActiveResourcesOverPeriodOfTime	12:15 AM	1 day	all nodes
kpimetricsTrig_21_AllCustom_ResourceCount_Details	12:15 AM	1 day	all nodes
kpimetricsTrig_22_AllCustom_ResourceCount_Total	12:15 AM	1 day	all nodes
kpimetricsTrig_23_Cache_METRICS_SQL_RESOURCE_LINEAGE	1:05 AM	1 hour	all nodes
kpimetricsTrig_30_DBMSScheduler	12:30 AM	2 hours	all nodes
kpimetricsTrig_31_DBMSSchedulerError	1:30 AM	2 hours	only once per cluster
kpimetricsTrig_32_DBMSPartitionManager	12:00 AM	1 day	only once per cluster
kpimetricsTrig_33_DeleteCollection	1:05 AM	30 min	all nodes

This section lists all triggers and load scripts that have been defined to execute various KPI metrics procedures at regular intervals. The default execution frequencies are listed for each trigger. The load scripts have been created to load and aggregate raw data into processed KPI metrics metrics.

Trigger [schedule] → Script Name → View name	Description
<p>Schedule: [1 hour, 1:00 am] kpimetricsTrig_00_CheckMetricsActivityDebug 'Y',60,06:00:00,20:00:00 → pCheckMetricsActivity → pGetEmailSubscriptions</p>	<p>This script is used for debugging the DV native metrics and checks for activity on an hourly basis and alerts the configured user whether there was any inactivity. This can be a useful gauge of the health of the DV system. It is configured hourly, but it may make sense to alter the timing to every 2 or more hours depending activity-levels. It is also configured to check between the hours of 6 am to 8 pm daily. It will look at the current timestamp and compare with MAX(starttime) for the collection tables [metrics_requests, metrics_resources_usage] and MAX(logintime) for [metrics_sessions]. It also looks to see if the data transfer jobs have been running by looking at METRICS_JOB_DETAILS. Email integration must be configured on the DV server for this to work.</p>
<p>Schedule: [1 hour, 1:00 am] kpimetricsTrig_01_Cache_ALL_RESOURCES → Cache_ALL_RESOURCES → /System/ALL_RESOURCES → [/services/databases/system/ALL_RESOURCES, l_hostname, l_port] → METRICS_ALL_RESOURCES_STG → EXCEPT_ALL_RESOURCES → METRICS_ALL_RESOURCES</p>	<p>Cache ALL_RESOURCES to make joining with other KPI metrics tables more efficient. Additionally, the METRICS_ALL_RESOURCES table contains a historical record of all “new” resources. It does not delete old ones. Insert into METRICS_ALL_RESOURCES select from EXCEPT_ALL_RESOURCES [select from METRICS_ALL_RESOURCES_STG where not exist in METRICS_ALL_RESOURCES]</p>
<p>Schedule: [1 hour, 1:00 am] kpimetricsTrig_02_Cache_ALL_USERS → Cache_ALL_USERS → /System/ALL_USERS → [/services/databases/system/ALL_USERS, l_hostname, l_port] → METRICS_ALL_USERS</p>	<p>Cache ALL_USERS to make joining with other KPI metrics tables more efficient. Additionally, the METRICS_ALL_USERS table contains a historical record of all “new” resources. It does not delete old ones. Insert into METRICS_ALL_USERS select from EXCEPT_ALL_USERS [Select from the /System/All_USERS where not exist in METRICS_ALL_USERS]</p>
<p>Schedule: [1 day, 7:00 am] kpimetricsTrig_03_Cache_LDAP_PERSON → Cache_LDAP_PERSON → ./Metadata/System/LDAP_PERSON → /KPI metrics/Physical/Metadata/LDAP/organizationalPerson</p>	<p>Cache LDAP person data once a day. This is not a historical view of users. It gets a new copy each day. Insert into METRICS_LDAP_PERSON select from ./Metadata/System/LDAP_PERSON</p>
<p>Schedule: [1 hour, 1:15 am] kpimetricsTrig_04_Cache_CIS_SYSTEM_RESOURCES → Cache_CIS_SYSTEM_RESOURCES → p15MinutesIncrements</p>	<p>Insert 15 min increment rows into METRICS_CIS_SYSTEM_RESOURCES select from METRICS_LOG_MEMORY, METRICS_LOG_IO, METRICS_LOG_DISK, METRICS_CPU_MEMORY_CHECKER.</p>
<p>Schedule: [1 hour, 1:00 am] kpimetricsTrig_05_Cache_CPU_MEMORY_CHECKER → Cache_CPU_MEMORY_CHECKER → ./Metadata/System/CPU_MEMORY_CHECKER →</p>	<p>Cache the system CPU and memory values. Insert into METRICS_CPU_MEMORY_CHECKER select from ./Metadata/System/CPU_MEMORY_CHECKER.</p>

<p>/KPImetrics/Physical/Metadata/ CPUAndMemChecker/CpuAndMemCheckerCjp</p>	
<p>Schedule: [1 hour, 1:00 am] kpimetricsTrig_06_Cache_LOG_DISK → Cache_LOG_DISK → ./Metadata/System/LOG_DISK → ./services/databases/system/LOG_DISK</p>	<p>Cache DV system log disk information. Insert into METRICS_LOG_DISK select from ./Metadata/System/LOG_DISK</p>
<p>Schedule: [1 hour, 1:00 am] kpimetricsTrig_07_Cache_LOG_IO [→ Cache_LOG_IO → ./Metadata/System/LOG_IO → ./services/databases/system/LOG_IO</p>	<p>Cache DV system IO information. Insert into METRICS_LOG_IO select from ./Metadata/System/LOG_IO</p>
<p>Schedule: [1 hour, 1:00 am] kpimetricsTrig_08_Cache_LOG_MEMORY → Cache_LOG_MEMORY → ./Metadata/System/LOG_MEMORY → ./services/databases/system/LOG_MEMORY</p>	<p>Cache DV system log memory information. Insert into METRICS_LOG_MEMORY select from ./Metadata/System/LOG_MEMORY</p>
<p>Schedule: [30 min, 1:15 am] – Primary configuration. Schedule: [15 min, 1:15 am] – Alternative if more frequent processing is needed. kpimetricsTrig_10_Cache_METRICS_SQL_REQUEST → Cache_METRICS_SQL_REQUEST_EXEC → Cache_METRICS_SQL_REQUEST_GENSQL</p>	<p>Execute the SQL Parser to parse the SQL request (description) originating from metrics_requests_hist. Retrieve the list of tables and columns and update the METRICS_SQL_REQUEST, METRICS_SQL_RESOURCE and METRICS_SQL_COLUMNS tables.</p> <p>The schedule is designed to run 15 min after kpimetricsTrig_30_DBMSScheduler allowing that trigger to execute pMETRICS_ALL_TABLES_exec which updates the metrics_requests_hist table. The data from metrics_requests_hist is what feeds this trigger and procedure. This procedure runs every 30 min because it often runs behind in its work due to the intensity level at which runs.</p> <p>This trigger blocks (does not run) if Cache_METRICS_SQL_RESOURCE_LINEAGE is currently executing. The two procedures running together causes DV and the database to become overwhelmed.</p> <p>This procedure is architected for multi-host processing. If there are no rows to process for the current node, it will look to see which node (if cluster) has the most rows to process and help process those rows. This code uses a semaphore to reserve rows for processing. The row reservation can only be executed by one node at a time. The actual processing is done in parallel.</p> <p>NOTE: It may be necessary to set the trigger to 15 min if more frequent processing is required and it can still keep up with the number rows to process in the time period. The number of rows can be found in commonValues.numRowsToProcessBeforeExiting. Adjust the default value of 500 to whatever makes sense for each environment.</p> <p>When running in a cluster, this trigger/procedure will attempt to process another nodes SQL once it has processed all of its own. This parallel processing is important as some nodes may be slower than others and fall behind or the load balancer may be set to “sticky” instead of a true “round-robin” and the majority of requests end up on a few of the nodes in the cluster.</p>

<p>[No trigger – manual] Cache_METRICS_SQL_REQUEST_EXEC_ADHOC → Cache_METRICS_SQL_REQUEST_GENSQL_ADHOC</p>	<p>Provides a way to manually re-process successfully parsed requests. This is only necessary if a code-patch has been provided that changes the values that are inserted. Normally, this will never be executed.</p>
<p>Schedule: [1 day, 1:00 am] kpimetricsTrig_11_Cache_METRICS_SQL_REQUEST_REPROCESS → Cache_METRICS_SQL_REQUEST_EXEC_REPROCESS → Cache_METRICS_SQL_REQUEST_GENSQL_REPROCESS (0)</p>	<p>Re-process errors and execute the SQL Parser to parse the SQL request (description) originating from metrics_requests_hist. Retrieve the list of tables and columns and update the METRICS_SQL_REQUEST, METRICS_SQL_RESOURCE and METRICS_SQL_COLUMNS tables.</p> <p>Typically, this trigger is not turned off unless deemed necessary.</p> <p>This trigger blocks (does not run) if Cache_METRICS_SQL_RESOURCE_LINEAGE is currently executing. The two procedures running together causes DV and the database to become overwhelmed.</p>
<p>Schedule: [6 hours, 1:00 am] kpimetricsTrig_12_Cache_SYS_CACHES → Cache_SYS_CACHES → ./Metadata/System/SYS_CACHES → /services/databases/system/SYS_CACHES</p>	<p>Cache the system cache status information. Insert into METRICS_SYS_CACHES select from / ./Metadata/System/SYS_CACHES.</p>
<p>Schedule: [12 hours, 6:00 am] kpimetricsTrig_13_Cache_SYS_DATASOURCES [→ Cache_SYS_DATASOURCES → ./Metadata/System/SYS_DATASOURCES → /services/databases/system/SYS_DATASOURCES</p>	<p>Cache the system data source information. Insert into METRICS_SYS_DATASOURCES select from ./Metadata/System/SYS_DATASOURCES.</p>
<p>Schedule: [1 hour, 1:30 am] kpimetricsTrig_14_CheckCISWorkflowStatusFail [→ pCheckCISWorkflowStatusFail → [METRICS_CIS_WORKFLOW, pGetEmailSubscriptions]</p>	<p>Check for WORKFLOW_STATUS=F in the METRICS_CIS_WORKFLOW table since the last check. Each time this procedure is called it puts a marker row in the table with WORKFLOW_NAME=CHECK_WORKFLOW_STATUS</p> <p>Email integration must be configured on the DV server for this to work.</p>
<p>Schedule: [1 day, 12:59:50 am] kpimetricsTrig_15_CheckMetricsActivity 'N',60,00:00:00.000,23:59:59.999 → pCheckMetricsActivity → pGetEmailSubscriptions</p>	<p>This script checks for activity on a daily basis just before midnight and alerts the configured user whether there was any inactivity during the day. It is configured to look for inactivity between the hours of 12 am – 12 pm in the current day. This can be a useful gauge of the health of the DV system. It will look at the current timestamp and compare with MAX(starttime) for the collection tables [metrics_requests, metrics_resources_usage] and MAX(logintime) for [metrics_sessions]. It also looks to see if the data transfer jobs have been running by looking at METRICS_JOB_DETAILS. Email integration must be configured on the DV server for this to work.</p>
<p>Schedule: [1 day, 4:30 am] kpimetricsTrig_16_PurgeHistoryData → pPurgeData</p>	<p>This script purges old data from METRICS tables by executing series of DELETE statements. The purge period for each delete is defined within /Configuration/commonValues script.</p> <p>purgeWorkflowData 120 = 4 months - Purge tables: METRICS_CIS_WORKFLOW</p> <p>purgeSQLRequests 120 = 4 months - Purge tables: METRICS_SQL_COLUMNS,</p>

	<p>METRICS_SQL_RESOURCE, METRICS_SQL_REQUEST</p> <p>purgeResourceUsage 120 = 4 months - Purge tables: METRICS_CIS_SYSTEM_RESOURCES, METRICS_CPU_MEMORY_CHECKER, METRICS_LOG_DISK, METRICS_LOG_IO, METRICS_LOG_MEMORY, METRICS_SYS_DATASOURCES</p> <p>Email integration must be configured on the DV server for this to work.</p>
<p>Schedule: [1 hour, 1:00 am] kpimetricsTrig_17_CheckExceedMemoryPercentRequests → pCheckExceedMemoryPercentRequests → pGetEmailSubscriptions</p>	<p>This procedure queries this procedure in real-time /shared/ASAssets/KPImetrics/Business/Business/requests/pExceededMemoryPercentRequests to generate a list of queries exceeding memory percent per request. It generates an html table containing each of the requests and emails to the subscriber of the event [EXCEEDED_MEMORY]. Email integration must be configured on the DV server for this to work.</p>
<p>Schedule: [system event=requestRunForTooLong] kpimetricsTrig_18_CheckLongRunningRequests → pCheckLongRunningRequests → pGetEmailSubscriptions</p>	<p>This procedure queries this procedure in real-time /shared/ASAssets/KPImetrics/Business/requests/pLongRunningRequests to generate a list of long running requests. It generates an html table containing each of the requests and emails to the subscriber of the event [LONG_RUNNING]. Email integration must be configured on the DV server for this to work. The trigger is activated by the system request event: “Request Run Time”. This is set in the Administration Configuration.</p> 
<p>Schedule: [1 day, 12:15 am] kpimetricsTrig_19_AllCustom_AccessByUserOvertime → Cache_AllCustom_AccessByUserOvertime → /Metadata/Physical/Abstraction/EXCEPT_ACR_AccessByUserOvertime → [metrics_resources_usage_hist NOT EXISTS METRICS_ACR_ABUOT]</p>	<p>This procedure invokes “ACR_AccessByUserOvertime” once a day to improve overall query performance for this report. It finds the delta/difference in rows using the original underlying view metrics_resources_usage_hist and NOT EXISTS in METRICS_ACR_ABUOT. The acronym for the cache table must be short and thus it stands for “All Custom Reports Access By User Over Time”.</p>
<p>Schedule: [1 day, 12:15 am] kpimetricsTrig_20_AllCustom_ActiveResourcesOverPeriodOfTime → Cache_AllCustom_ActiveResourcesOverPeriodOfTime → /Metadata/Physical/Abstraction/EXCEPT_ACR_ActiveResourcesOverPeriodOfTime → [metrics_resources_usage_hist NOT EXISTS METRICS_ACR_ABUOT]</p>	<p>This procedure invokes “EXCEPT_ACR_ActiveResourcesOverPeriodOfTime” once a day to improve overall query performance for this report. It finds the delta/difference in rows using the original underlying view metrics_resources_usage_hist and NOT EXISTS in METRICS_ACR_AROPOT. The acronym for the cache table must be short and thus it stands for “All Custom Reports Active Resources Over Period Of Time”.</p>
<p>Schedule: [1 day, 12:15 am] kpimetricsTrig_21_AllCustom_ResourceCount_Details → Cache_AllCustom_ResourceCount_Details → /AllCustomReports/ResourceCount_DetailsRT → vResourceUsageUD → metrics_resources_usage_hist</p>	<p>This procedure caches the real-time query “ACR_ResourceCount_DetailsRT” once a day to improve overall query performance for this report. It caches a completely new set of rows to METRICS_ACR_RCD. The acronym for the cache</p>

	<p>table must be short and thus it stands for “All Custom Reports Resource Count Details”.</p>
<p>Schedule: [1 day, 12:15 am] kpimetricsTrig_22_AllCustom_ResourceCount_Total → Cache_AllCustom_ResourceCount_Total → /AllCustomReports/ResourceCount_TotalRT → vResourceUsageUD → metrics_resources_usage_hist</p>	<p>This procedure caches the real-time query “ACR_ResourceCount_TotalRT” once a day to improve overall query performance for this report. It caches a completely new set of rows to METRICS_ACR_RCT. The acronym for the cache table must be short and thus it stands for “All Custom Reports Resource Count Total”.</p>
<p>Schedule: [1 hour, 1:05 am] kpimetricsTrig_23_Cache_METRICS_SQL_RESOURCE_LINEAGE → Cache_METRICS_SQL_RESOURCE_LINEAGE → METRICS_SQL_RESOURCE and METRICS_SQL_RESOURCE_LINEAGE</p>	<p>This procedure will insert the new data into METRICS_SQL_RESOURCE_LINEAGE based on the START_TIME value from the table METRICS_SQL_RESOURCE. The lineage is derived by either copying a similar resource from the database table or by generating the lineage using getResourceLineageDirectRecursive().</p> <p>This procedure will block and not execute if Cache_METRICS_SQL_REQUEST [EXEC or_ADHOC or_REPROCESS] is running. The reason is that the two procedures combined execution can overwhelm both the database and the DV node.</p> <p>This procedure is architected for multi-host processing. If there are no rows to process for the current node, it will look to see which node (if cluster) has the most rows to process and help process those rows. This code uses a semaphore to reserve rows for processing. The row reservation can only be executed by one node at a time. The actual processing is done in parallel.</p> <p>The lineage is generated for each non-null RESOURCE_PATH and RESOURCE_TYPE found in the METRICS_SQL_RESOURCE table.</p> <p>METRICS_SQL_REQUEST [parent] --> [0 to many] METRICS_SQL_RESOURCE [child] --> [0 to many] METRICS_SQL_RESOURCE_LINEAGE [child]</p> <ol style="list-style-type: none"> 1. A request may have 0 or more resources. 2. A resource may have 0 or more data source resources. <ol style="list-style-type: none"> a. A single resource lineage row contains the Ancestor who uses the data source resource, the data source table/procedure/tree resource and the data source itself.
<p>Schedule: [2 hours, 12:30 am] kpimetricsTrig_30_DBMScheduler → pMETRICS_ALL_TABLES_exec → P_METRICS_ALL_TABLES()</p>	<p>This trigger executes the pMETRICS_ALL_TABLES_exec PLSQL procedure to transfer data from the metrics collection tables to the metrics history tables. It inserts a record into the METRICS_JOB_DETAILS table when it starts with a JOB_TABLE_NAME='DBMS_SCHEDULER'. It updates the same row with a STATUS='SUCCESS' or 'FAILURE'. If 'FAILURE' then update the ADDITIONAL_INFO field with the database error.</p> <p>Operational insights:</p> <ol style="list-style-type: none"> 1. This procedure will <u>not</u> execute the delete collection part of the script if P_METRICS_DELETE_COLLECTION is executing. This is determined by getting the STATUS for a given MAX(REQ_START_DATE) and nodehost and nodeport and

	<p>JOB_TABLE_NAME='DELETE_COLLECTION' and STATUS='BEGIN JOB' from METRICS_JOB_DETAILS.</p>
<p>Schedule: [2 hours, 1:30 am] kpimetricsTrig_31_DBMSSchedulerError → pCheckDBMSSchedulerError → pGetEmailSubscriptions</p>	<p>Send an email if there is a database PLSQL data transfer error that gets generated. Select details from /Abstraction/METRICS_JOB_DETAILS. The timing of 2 hours on the odd hour is based on the fact that the DBMS Scheduler trigger runs every 2 hours on the even hour. Therefore, this trigger runs an hour later to allow the PLSQL data transfer script to complete and post any issues or not.</p> <p>Note: This trigger only runs once per cluster because it finds all errors for all nodes if there is a cluster.</p>
<p>Schedule: [1 day, 12:00 am] kpimetricsTrig_32_DBMSPartitionManager → pPARTITION_MANAGER_exec →</p> <p>Oracle /KPImetrics/Physical/Metadata/DDL/Oracle/ [03_pqCreateDrop_KPI_Tables_oracle_metrics_history_tables_ADD, 03_pqCreateDrop_KPI_Tables_oracle_metrics_history_tables_DROP]</p> <p>SqlServer /KPImetrics/Physical/Metadata/DDL/SqlServer/ [03_pqCreateDrop_KPI_Tables_sqlserver_metrics_history_tables_ADD, 03_pqCreateDrop_KPI_Tables_sqlserver_metrics_history_tables_DROP]</p>	<p>Partition management is required for the metrics history tables when commonValues.partitionNumber and partitionStartDate are configured. The partition manager trigger wakes up once a day at 12 am and determines if a partition needs to be added or dropped. Technically, the only time any actual action will take place is the 1st day of the month unless DV is down at 12 am on the 1st. This is why it is scheduled to run every day to address any downtime. For every day except the 1st, it will simply find no partitions to add or drop based on what is currently in place and how the partitionNumber is configured. For adding a partition, it always looks at the current Year/Month it executes in and calculates the partition for the next month and determines if it exists or not. For dropping a partition, it counts the current number+1 for next month and compares with the partitionNumber to determine if it should drop the oldest partition.</p>
<p>Schedule: [30 min, 1:05 am] – Primary configuration. Schedule: [15 min, 1:05 am] – Alternative if more frequent processing is needed. kpimetricsTrig_33_DeleteCollection → pMETRICS_DELETE_COLLECTION_exec → P_METRICS_DELETE_COLLECTION()</p>	<p>This trigger executes the PLSQL procedure pMETRICS_DELETE_COLLECTION_exec to delete non-essential and unwanted rows from the two collection tables: metrics_resources_usage and metrics_requests. It inserts a record into the METRICS_JOB_DETAILS table when it starts with a JOB_TABLE_NAME='DELETE_COLLECTION'. It updates the same row with a STATUS='SUCCESS' or 'FAILURE'. If 'FAILURE' then update the ADDITIONAL_INFO field with the database error.</p> <p>It is recommended to have this turned on if the system is very active and producing 100K or more transactions in 1 hour or if the commonValues.queryRequestTemplate='Y' which will produce 100K or more transactions itself within 1 hour. This setting is used to keep the collection tables as small as possible of unwanted data based on METRICS_JOB_FILTERS rows.</p> <ol style="list-style-type: none"> 1. metrics_resources_usage – delete from this collection table where a row matches METRICS_JOB_FILTERS based on user, domain and resourcekind. This delete gets executed first. 2. metrics_requests – delete from this collection table when a row does not exist in metrics_resources_usage collection table. This delete gets executed second. <p>Operational insights:</p>

	<p>1. This procedure will <u>not</u> execute if P_METRICS_ALL_TABLES is executing. This is determined by getting the STATUS for a given MAX(REQ_START_DATE) and nodehost and nodeport and JOB_TABLE_NAME='DBMS_SCHEDULER' and STATUS='BEGIN JOB' from METRICS_JOB_DETAILS.</p> <p>2. If this procedure finds another record for the same nodehost and nodeport and JOB_TABLE_NAME='DELETE_COLLECTION' and STATUS='BEGIN JOB' then it determines that the record is stranded due to some unforeseen event and marks that record with a STATUS='FAILURE'. It then continues on as normal.</p> <p>3. The timing of the kpimetricsTrig_33_DeleteCollection [DC] trigger is important. It is set to execute after the kpimetricsTrig_10_Cache_METRICS_SQL_REQUEST [MSR] trigger and before the kpimetricsTrig_30_DBMSScheduler [DBS] trigger. It is recommended that the period be the same between kpimetricsTrig_33_DeleteCollection and kpimetricsTrig_10_Cache_METRICS_SQL_REQUEST. If one gets changed from 30 to 15, then the other one should be changed also. The objective is to delete unwanted rows produced by the first trigger and delete as much as possible for DBMSScheduler runs so that it does not have to spend time querying through 100K of unwanted rows. Ultimately, DBMSScheduler to run as quickly as possible. The only way to do that is to keep the data set as small as possible. The chart below shows the time interaction between the three triggers that are responsible for the following: [MSR] – SQL parsing with a byproduct of creating metrics request volume, [DBS] – transferring data from collection to history and deleting collection. [DC] – deleting collection data with the objective of keeping collection tables small and manageable.</p> <p>30 min period on MSR and DC triggers: [MSR] [DBS] [DC] 12:00 12:05 ←del collection just before processing 12:30 12:30 12:35</p> <p>15 min period on MSR and DC triggers: [MSR] [DBS] [DC] 12:00 12:05 12:15 12:20 ←del collection just before processing 12:30 12:30 12:35 12:45 12:50</p>
--	---

Metadata System Helpers Scripts

Location: /shared/ASAssets/KPImetrics/Physical/Metadata/System/Helpers

This section lists all triggers and load scripts that have been defined to execute various KPImetrics procedures at regular intervals. The default execution frequencies are listed for each trigger. The load scripts have been created to load and aggregate raw data into processed KPImetrics metrics.

Script Name → Resource name	Description
getClusterHostnamesDiff	The purpose of this procedure is to compare the current hostname for the current server with names in the cluster to determine the difference of those names. This difference will be used to create and use stage tables for the P_METRICS_ALL_TABLES procedure to use. It requires a unique stage table for each node in the cluster because it performs a table truncate during its processing. hostname1 hostname2 The difference for the current hostname=1 hostname_a hostname_b The difference for the current hostname=a
getCurrentTimestamp	This procedure is used to get the current timestamp which was a workaround for a bug in a previous version of DV.
getDatasourceConfiguration → /shared/ASAssets/Utilities/.../ getBasicResourceCursor → getBasicResourceXSLT	This procedure is used to get the current data source configuration for various commonValue settings, paths and derived paths. It invokes a couple of other procedures to get the data source type.
p15MinutesIncrements	This procedure returns a cursor of 15 minute increments given a starting timestamp that is passed in. It generates one full day of 15 increments for a total of 53 records.
pGetEmailSubscriptions → pGetDomainUsers → [pGetDomainsXSLT,pGetDomainUsersXSLT]	This procedure constructs a list of emails which is used to send error or informational emails based on subscription to DV groups or LDAP groups.
pGetSystemInformation	This script is used to get cluster name and server name from /lib/util/getProperties() built in function. This function is used in load scripts and by custom logger to get server name.
pStartWorkflow	This script is invoked from each of the load scripts and marks the start of a workflow in the METRICS_CIS_WORKFLOW table. For a given workflow, the script takes workflow name as an input and returns the next workflow start and end time.
pEndWorkflow	Like pStartWorkflow script, this script is also called from other load scripts and marks the end of a workflow by updating the METRICS_CIS_WORKFLOW table when a workflow finishes. This script takes workflow name, workflow start and end times, workflow status and number of rows affected as input and updates the METRICS_CIS_WORKFLOW table.
pUpdateEventRegLog	This procedure is used to insert/update rows in the METRICS_EVENT_REG_LOG and METRICS_EVENT_REG_LOG_LINEAGE tables.

Physical Oracle Data Transfer Script

Location: /shared/ASAssets/KPImetrics/Physical/Metadata/DDL/Oracle/
06_pqCreateDrop_KPI_Plsq_oracle_data_xfer_script

This section describes the actions and SQL taken in the P_METRICS_ALL_TABLES data transfer script. The purpose of this script is to transfer data from the collection tables to the history tables using native Oracle SQL script. Stage tables are used to prepare and augment the collection data prior to inserting into the history tables. The result of this is that only “inserts” are performed on the history tables. All of the insert/update/deletes that are required are performed on the stage tables.

metrics_sessions

1. Get the min/max sessionid and logintime

```
select min("sessionid") minsessionid, min("logintime") minlogintime, max("sessionid") maxsessionid,
max("logintime") maxlogintime from "||dataSourceSchemaName||"."||metrics_sessions_collection||" where
"nodehost" = c_node_host and "nodeport" = c_node_port;
```

2. Insert into the metrics_sessions_stg stage table from the metrics_sessions collection table

```
insert into "||dataSourceSchemaName||"."||metrics_sessions_stg||"
select ms.*,
LU."USER_ID", LU."user", LU."domain",
-- Perform LDAP userkey join transformation. The transformation value in userKeyTransformation is
provided at the time this procedure is created which is during installation.
'||REPLACE(userKeyTransform, "user", 'LU."user"')||' as "userkey",
'||REPLACE(domainKeyTransform, "domain", 'LU."domain"')||' as "domainkey",
null "fullName", null "lastName", null "firstName", null "mail", null "employeeNumber", null
"employeeID", null "telephoneNumber", null "friendlyCountryName"
from "||dataSourceSchemaName||"."||metrics_sessions_collection||" ms
-- Join with metrics_resources_usage to get the user
LEFT OUTER JOIN (
select distinct mruc."nodehost", mruc."nodeport", mruc."sessionid", mruc."user", mruc."domain",
mau."USER_ID"
FROM "||dataSourceSchemaName||"."||metrics_resources_usage_collection||" mruc
-- Join with METRICS_ALL_USERS to get the latest (max LOAD_TIME) userid
LEFT OUTER JOIN
- "||dataSourceSchemaName||"."METRICS_ALL_USERS" mau
(SELECT * FROM "||dataSourceSchemaName||"."METRICS_ALL_USERS" au1
WHERE LOAD_TIME = ( SELECT MAX(LOAD_TIME) FROM
"||dataSourceSchemaName||"."METRICS_ALL_USERS" au2
WHERE au1.USERNAME = au2.USERNAME AND au1.DOMAIN_NAME = au2.DOMAIN_NAME
AND au1.NODE_HOST = au2.NODE_HOST AND au1.NODE_PORT = au2.NODE_PORT )
) mau
ON mau."USERNAME" = mruc."user"
AND mau."DOMAIN_NAME" = mruc."domain"
AND mau."NODE_HOST" = mruc."nodehost"
AND mau."NODE_PORT" = mruc."nodeport"
) LU
ON LU."sessionid" = ms."sessionid"
AND LU."nodehost" = ms."nodehost"
AND LU."nodeport" = ms."nodeport"
where ms."logintime" <= maxlogintime
and ms."sessionid" <= maxsessionid
and ms."nodehost" = in_node_host
and ms."nodeport" = in_node_port
and NOT EXISTS (
```

```

select *
from "||dataSourceSchemaName||"."||metrics_sessions_hist||" ms1
where ms1."nodehost" = ms."nodehost"
  and ms1."nodeport" = ms."nodeport"
  and ms1."sessionid" = ms."sessionid"
  and ms1."logintime" = ms."logintime"
  and ms1."type" = ms."type"
);

```

3. Update the metrics_sessions_stg stage table with user information from METRICS_LDAP_PERSON using the "userkey" to join.

```

UPDATE "||dataSourceSchemaName||"."||metrics_sessions_stg||" msh1
  SET ("fullname", "lastname", "firstname", "mail", "employeenumber", "employeeid", "telephonenumber",
"friendlycountryname") =
  (SELECT DISTINCT mlp."fullName", mlp."lastName", mlp."firstName", mlp."mail",
mlp."employeeNumber", mlp."employeeID", mlp."telephoneNumber", mlp."friendlyCountryName"
  FROM "||dataSourceSchemaName||"."||metrics_sessions_stg||" msh2
  LEFT OUTER JOIN (
    SELECT mlp1."userkey", mlp1."domainkey", mlp1."fullName", mlp1."lastName", mlp1."firstName",
mlp1."mail", mlp1."employeeNumber", mlp1."employeeID", mlp1."telephoneNumber",
mlp1."friendlyCountryName"
  FROM "||dataSourceSchemaName||"."METRICS_LDAP_PERSON" mlp1
  WHERE mlp1."loadtime" =
    ( SELECT MAX(mlp2."loadtime")
      FROM "||dataSourceSchemaName||"."METRICS_LDAP_PERSON" mlp2
      WHERE mlp1."userkey" = mlp2."userkey" AND mlp1."domainkey" = mlp2."domainkey")
    ) mlp
  ON mlp."userkey" = msh2."userkey"
  AND mlp."domainkey" = msh2."domainkey"
  WHERE msh1."nodehost" = msh2."nodehost"
  AND msh1."nodeport" = msh2."nodeport"
  AND msh1."sessionid" = msh2."sessionid"
  AND msh1."logintime" = msh2."logintime"
  AND msh1."type" = msh2."type"
)
WHERE msh1."logintime" >= minlogintime
  AND msh1."sessionid" >= minsessionid
  AND msh1."nodehost" = in_node_host
  AND msh1."nodeport" = in_node_port
  AND msh1."user" is not null
  AND msh1."userkey" is not null
  AND msh1."mail" is null
  AND msh1."lastname" is null;

```

4. Insert into metrics_sessions_hist history from metrics_sessions_stg stage

```

INSERT INTO "||dataSourceSchemaName||"."||metrics_sessions_hist||"
  SELECT * FROM "||dataSourceSchemaName||"."||metrics_sessions_stg||"
  where "nodehost" = in_node_host
  and "nodeport" = in_node_port;

```

5. Delete this batch of data from the metrics_sessions collection table that was just processed being careful not to delete any new rows.

```

delete from "||dataSourceSchemaName||"."||metrics_sessions_collection||"
  where "logintime" <= maxlogintime
  and "sessionid" <= maxsessionid
  and "nodehost" = in_node_host

```

```
and "nodeport" = in_node_port;
```

metrics_resources_usage

1. Get the min/max requestid and starttime

```
select min("requestid") minrequestid, min("starttime") minstarttime, max("requestid") maxrequestid,
       max("starttime") maxstarttime
       from "||dataSourceSchemaName||"."||metrics_resources_usage_collection||"
       where "nodehost" = c_node_host and "nodeport" = c_node_port;
```

2. Insert the new rows into the metrics_resources_usage_stg stage from the metrics_resources_usage collection where not exists in metrics_resources_usage_hist history and not in METRICS_JOB_FILTERS.

```
INSERT INTO "||dataSourceSchemaName||"."||metrics_resources_usage_stg||"
SELECT
  -- Insert collection columns
  mruc."cluster"
  , mruc."nodehost"
  , mruc."nodeport"
  , mruc."sessionid"
  , mruc."user"
  , mruc."domain"
  , mruc."group"
  , mruc."requestid"
  , mruc."parentid"
  , mruc."datasourcepath"
  , mruc."datasourcetype"
  , mruc."resourcepath"
  , mruc."resourcetype"
  , mruc."resourceguid"
  , mruc."resourcekind"
  , mruc."starttime"
  , mruc."endtime"
  -- Insert expanded information
  , mar."RESOURCE_ORIGIN" "resourceorigin"
  , mar."RESOURCE_ID" "resourceid"
  , mar."DATASERVICE_NAME" "dataservicename"
  , mar."RESOURCE_NAME" "resourcename"
  , mar."SCHEMA_NAME" "parentname"
  , mar."CATALOG_NAME" "grandparentname"
  , mar."CATEGORY_NAME" "categoryname"
  , mar."PARENT_PATH" "parentpath"
  -- Insert expanded user information
  , mau."USER_ID" "userid"
  , '||REPLACE(userKeyTransform, "user", mruc."user")||' as "userkey"
  , '||REPLACE(domainKeyTransform, "domain", mruc."domain")||' as "domainkey"
  , null "fullName", null "lastName", null "firstName", null "mail", null "employeeNumber", null "employeeID", null
  "telephoneNumber", null "friendlyCountryName"
FROM "||dataSourceSchemaName||"."||metrics_resources_usage_collection||" mruc
-- METRICS_ALL_USERS
LEFT OUTER JOIN
(SELECT *
 FROM "||dataSourceSchemaName||"."METRICS_ALL_USERS" au1
 WHERE "LOAD_TIME" = (SELECT MAX("LOAD_TIME") FROM
 ||dataSourceSchemaName||"."METRICS_ALL_USERS" au2
 WHERE au1."USERNAME" = au2."USERNAME"
 AND au1."DOMAIN_NAME" = au2."DOMAIN_NAME"
 AND au1."NODE_HOST" = au2."NODE_HOST"
 AND au1."NODE_PORT" = au2."NODE_PORT" )
) mau
ON mau."USERNAME" = mruc."user"
AND mau."DOMAIN_NAME" = mruc."domain"
AND mau."NODE_HOST" = mruc."nodehost"
```

```

AND mau."NODE_PORT" = mruc."nodeport"
-- METRICS_ALL_RESOURCES
LEFT OUTER JOIN
(SELECT mar1."RESOURCE_ORIGIN", mar1."RESOURCE_ID", mar1."RESOURCE_NAME",
mar1."RESOURCEPATH", mar1."RESOURCE_TYPE", mar1."PARENT_PATH",
mar1."DATASERVICE_NAME", mar1."CATALOG_NAME", mar1."SCHEMA_NAME",
mar1."CATEGORY_NAME", mar1."NODE_HOST", mar1."NODE_PORT"
FROM "||dataSourceSchemaName||"."METRICS_ALL_RESOURCES" mar1
WHERE mar1."LOAD_TIME" = (SELECT MAX(mar2."LOAD_TIME")
FROM "||dataSourceSchemaName||"."METRICS_ALL_RESOURCES" mar2
WHERE mar1."RESOURCEPATH" = mar2."RESOURCEPATH"
AND mar1."RESOURCE_TYPE" = mar2."RESOURCE_TYPE"
AND mar1."NODE_HOST" = mar2."NODE_HOST"
AND mar1."NODE_PORT" = mar2."NODE_PORT" )
) mar
ON mar."RESOURCEPATH" = mruc."resourcepath"
AND mar."RESOURCE_TYPE" = mruc."resourcetype"
AND mar."NODE_HOST" = mruc."nodehost"
AND mar."NODE_PORT" = mruc."nodeport"
WHERE mruc."nodehost" = in_node_host
AND mruc."nodeport" = in_node_port
AND mruc."starttime" <= maxstarttime_mru
AND mruc."requestid" <= maxrequestid_mr
-- Do not insert rows matching the filter rows in METRICS_JOB_FILTERS
AND (mruc."user", mruc."domain", mruc."resourcekind") NOT IN (select "USER", "DOMAIN", "RESOURCE_KIND"
FROM "||dataSourceSchemaName||"."METRICS_JOB_FILTERS" WHERE "ENV_TYPE" = in_env_type)
-- Do not insert rows if they exist in metrics_resources_usage_hist
AND NOT EXISTS (
SELECT 1
FROM "||dataSourceSchemaName||"."metrics_resources_usage_hist" mruh
WHERE mruc."nodehost" = mruh."nodehost"
AND mruc."nodeport" = mruh."nodeport"
AND mruc."starttime" = mruh."starttime"
AND mruc."requestid" = mruh."requestid"
);

```

6. Update the metrics_resources_usage_stg stage table with user information from METRICS_LDAP_PERSON using the “userkey” to join.

```

UPDATE "||dataSourceSchemaName||"."metrics_resources_usage_stg" mruh1
SET ("fullname", "lastname", "firstname", "mail", "employeenumber", "employeeid", "telephonenumber",
"friendlycountryname") =
(SELECT DISTINCT mlp."fullName", mlp."lastName", mlp."firstName", mlp."mail",
mlp."employeeNumber", mlp."employeeID", mlp."telephoneNumber", mlp."friendlyCountryName"
FROM "||dataSourceSchemaName||"."metrics_resources_usage_stg" mruh2
LEFT OUTER JOIN (
SELECT mlp1."userkey", mlp1."domainkey", mlp1."fullName", mlp1."lastName", mlp1."firstName",
mlp1."mail", mlp1."employeeNumber", mlp1."employeeID", mlp1."telephoneNumber",
mlp1."friendlyCountryName"
FROM "||dataSourceSchemaName||"."METRICS_LDAP_PERSON" mlp1
WHERE mlp1."loadtime" =
( SELECT MAX(mlp2."loadtime")
FROM "||dataSourceSchemaName||"."METRICS_LDAP_PERSON" mlp2
WHERE mlp1."userkey" = mlp2."userkey" AND mlp1."domainkey" = mlp2."domainkey")
) mlp
ON mlp."userkey" = mruh2 ."userkey"
AND mlp."domainkey" = mruh2 ."domainkey"
WHERE mruh1 ."nodehost" = mruh2 ."nodehost"
AND mruh1 ."nodeport" = mruh2 ."nodeport"
AND mruh1 ."requestid" = mruh2 ."requestid"
AND mruh1 ."starttime" = mruh2 ."starttime"
)
WHERE mruh1 ."starttime" >= minstarttime_mru
AND mruh1 ."requestid" >= minrequestid_mru

```

```

AND mruh1 ."nodehost" = in_node_host
AND mruh1 ."nodeport" = in_node_port
AND mruh1 ."user" is not null
AND mruh1 ."userkey" is not null
AND mruh1 ."mail" is null
AND mruh1 ."lastname" is null;

```

7. Insert into metrics_resources_usage_hist history from metrics_resources_usage_stg stage

```

INSERT INTO "||dataSourceSchemaName||"."||metrics_resources_usage_hist||"
  SELECT * FROM "||dataSourceSchemaName||"."||metrics_resources_usage_stg||"
  where "nodehost" = in_node_host
  and "nodeport" = in_node_port;

```

8. Delete the metrics_resources_usage collection remaining data

- a. **Note: if the kpimetricsTrig_33_DeleteCollection trigger is still running then this section is skipped.**

```

delete from "||dataSourceSchemaName||"."||metrics_resources_usage_collection||"
  where "nodehost" = in_node_host
  and "nodeport" = in_node_port
  and "starttime" <= maxstarttime_mru
  and "requestid" <= maxrequestid_mru;

```

metrics_requests

1. Get the min/max requestid and starttime

```

select min("requestid") minrequestid, min("starttime") minstarttime, max("requestid") maxrequestid,
  max("starttime") maxstarttime
  from "||dataSourceSchemaName||"."||metrics_requests_collection||"
  where "nodehost" = c_node_host and "nodeport" = c_node_port;

```

2. Insert the new rows into the metrics_requests_stg stage from the metrics_requests collection where not exists in metrics_requests_hist history and exists in metrics_resources_usage_stg.

-- A cursor is created in order to be able to loop through the records and commit a batch at a time. This was done so as to minimize the UNDO tablespace and archiver rollback logs. When Oracle archiver reaches a configured length of time and the query has not finished it tries to rollback. If there is not enough UNDO tablespace, the query fails. For this particular query it may be necessary to greatly increase the UNDO tablespace and increase the time for the archiver to 2 hours.

```

cursor c_metrics_requests(c_node_host varchar2, c_node_port number, c_maxrequestid number,
  c_maxstarttime timestamp) is
  SELECT mrc.*,
    null "dataservicename", null "resourcekind",
    LU."USER_ID", LU."user", LU."domain",
    '||REPLACE(userKeyTransform, "user", 'LU."user")||' as "userkey",
    '||REPLACE(domainKeyTransform, "domain", 'LU."domain")||' as "domainkey",
    null "fullName", null "lastName", null "firstName", null "mail", null "employeeNumber", null
    "employeeID", null "telephoneNumber", null "friendlyCountryName"
  FROM "||dataSourceSchemaName||"."||metrics_requests_collection||" mrc
  -- METRICS_ALL_USERS
  LEFT OUTER JOIN
    (SELECT DISTINCT mruc."nodehost", mruc."nodeport", mruc."requestid", mruc."user",
    mruc."domain", mau."USER_ID"
    FROM "||dataSourceSchemaName||"."||metrics_resources_usage_collection||" mruc

```

```

LEFT OUTER JOIN
  (SELECT * FROM "||dataSourceSchemaName||"."METRICS_ALL_USERS" au1
   WHERE "LOAD_TIME" = (
     SELECT MAX("LOAD_TIME")
     FROM "||dataSourceSchemaName||"."METRICS_ALL_USERS" au2
     WHERE au1."USERNAME" = au2."USERNAME"
     AND au1."DOMAIN_NAME" = au2."DOMAIN_NAME"
     AND au1."NODE_HOST" = au2."NODE_HOST"
     AND au1."NODE_PORT" = au2."NODE_PORT" )
   ) mau
  ON mau."USERNAME" = mruc."user"
  AND mau."DOMAIN_NAME" = mruc."domain"
  AND mau."NODE_HOST" = mruc."nodehost"
  AND mau."NODE_PORT" = mruc."nodeport"
) LU
ON LU."requestid" = mrc."requestid"
AND LU."nodehost" = mrc."nodehost"
AND LU."nodeport" = mrc."nodeport"
WHERE mrc."nodehost" = c_node_host
AND mrc."nodeport" = c_node_port
AND mrc."requestid" <= c_maxrequestid
AND mrc."starttime" <= c_maxstarttime
-- Only insert metrics_requests collection rows with a corresponding row in the
metrics_resources_usage_stg
-- No point in capturing these rows without relevant resource usage information.
AND EXISTS (
  SELECT 1 FROM
  "||dataSourceSchemaName||"."||metrics_resources_usage_stg||" mru
  WHERE mru."nodehost" = mrc."nodehost"
  AND mru."nodeport" = mrc."nodeport"
  AND mru."requestid" = mrc."requestid"
  AND mru."starttime" = mrc."starttime"
);

```

```

TYPE c_metrics_requests_T is table of c_metrics_requests%rowtype;
c_metrics_requests_array c_metrics_requests_T;

```

```

-- Perform the insert in batches that get committed every "batchInsertMax" to free up UNDO space.
-- The fetch size and commit size are the same to keep the logic clean.

```

```

batchcount := 0;
begintimestamp3 := SYSTIMESTAMP;
open c_metrics_requests(in_node_host, in_node_port, maxrequestid_mr, maxstarttime_mr);

```

```

loop fetch c_metrics_requests bulk collect into c_metrics_requests_array limit batchInsertMax;
  -- Calculate fetch collection duration
  endtimestamp := SYSTIMESTAMP;
  fetchDuration_mr := endtimestamp - begintimestamp3;
  durationStr := replace(replace(cast(fetchDuration_mr as varchar2), "000000000 ", "0 "), "+", "");
  fetchDurationStr_mr := substr(durationStr, 1, instr(durationStr, ".") + 3);

```

```

  -- Insert the rows a batch at a time
  begintimestamp3 := SYSTIMESTAMP;
  forall i in 1 .. c_metrics_requests_array.count
    INSERT INTO "||dataSourceSchemaName||"."||metrics_requests_stg||" VALUES
  c_metrics_requests_array(i);

```

```

  uncommitted := uncommitted + sql%rowcount;
  insertrows_mr := insertrows_mr + sql%rowcount;
  batchcount := batchcount + 1;
  -- Calculate insert stage duration

```

```

endtimestamp := SYSTIMESTAMP;
insertStgDuration_mr := endtimestamp - begintimestamp3;
durationStr := replace(replace(cast(insertStgDuration_mr as varchar2), "000000000 ", "0 "), "+", "");
insertStgDurationStr_mr := substr(durationStr, 1, instr(durationStr, ".") + 3);
-- Update a status row into the METRICS_JOB_DETAILS table for number of rows affected
statusmsg := "INSERT STAGE: '||metrics_requests_stg||' B#='||batchcount||' T#='||insertrows_mr||'"
FETCH='||fetchDurationStr_mr||' INSERT='||insertStgDurationStr_mr;

```

```

update '||dataSourceSchemaName||'. "METRICS_JOB_DETAILS"
    set CURRENT_OPERATION = statusmsg
    where REQ_START_DATE = begintimestamp
    and JOB_TABLE_NAME = '||metrics_requests_collection||'
    and NODE_HOST = in_node_host
    and NODE_PORT = in_node_port;
exit when c_metrics_requests_array.count < batchInsertMax;
if (uncommitted >= batchInsertMax) then
    commit;
    uncommitted := 0;
end if;
begintimestamp3 := SYSTIMESTAMP;
end loop;
if (uncommitted > 0) then
    commit;
end if;
close c_metrics_requests;

```

-- Delete any rows from stage that exist in history. This logic was broken out as a separate piece instead of embedding in the cursor query above because in most cases there should be no rows. This is here as more of an insurance policy.

```

DELETE FROM '||dataSourceSchemaName||'. "||metrics_requests_stg||" mrc
-- Remove any rows that already exist in metrics_requests_hist
WHERE mrc."nodehost" = in_node_host
AND mrc."nodeport" = in_node_port
AND EXISTS (
    SELECT 1
    FROM '||dataSourceSchemaName||'. "||metrics_requests_hist||" mrh
    WHERE mrc."nodehost" = mrh."nodehost"
    AND mrc."nodeport" = mrh."nodeport"
    AND mrc."starttime" = mrh."starttime"
    AND mrc."requestid" = mrh."requestid"
);
deleterows_mr := sql%rowcount;
commit;

```

3. Update the new rows from the metrics_requests_stg stage with METRICS_LDAP_PERSON using the “userkey” to join.

```

UPDATE '||dataSourceSchemaName||'. "||metrics_requests_stg||" mrh1
    SET ("fullname", "lastname", "firstname", "mail", "employeenumber", "employeeid", "telephonenumber",
        "friendlycountryname") =
        (SELECT DISTINCT mlp."fullName", mlp."lastName", mlp."firstName", mlp."mail",
            mlp."employeeNumber", mlp."employeeID", mlp."telephoneNumber", mlp."friendlyCountryName"
        FROM '||dataSourceSchemaName||'. "||metrics_requests_stg||" mrh2
        LEFT OUTER JOIN (
            SELECT mlp1."userkey", mlp1."domainkey", mlp1."fullName", mlp1."lastName", mlp1."firstName",
                mlp1."mail", mlp1."employeeNumber", mlp1."employeeID", mlp1."telephoneNumber",
                mlp1."friendlyCountryName"
            FROM '||dataSourceSchemaName||'. "METRICS_LDAP_PERSON" mlp1
            WHERE mlp1."loadtime" =

```



```

( SELECT MAX(mlp2."loadtime")
  FROM ""||dataSourceSchemaName||"."METRICS_LDAP_PERSON" mlp2
  WHERE mlp1."userkey" = mlp2."userkey" AND mlp1."domainkey" = mlp2."domainkey" )
) mlp
ON mlp."userkey" = mrh2."userkey"
AND mlp."domainkey" = mrh2."domainkey"
WHERE mrh1."nodehost" = mrh2."nodehost"
  AND mrh1."nodeport" = mrh2."nodeport"
  AND mrh1."requestid" = mrh2."requestid"
  AND mrh1."starttime" = mrh2."starttime"
)
WHERE mrh1."starttime" >= minstarttime_mr
  AND mrh1."requestid" >= minrequestid_mr
  AND mrh1."nodehost" = in_node_host
  AND mrh1."nodeport" = in_node_port
  AND mrh1."user" is not null
  AND mrh1."userkey" is not null
  AND mrh1."mail" is null
  AND mrh1."lastname" is null;

```

4. Update metrics_requests_stg stage [resourcekind and dataservicename] joined with metrics_resources_usage_stg.

- Use the resourcekind and dataservicename from metrics_resources_usage_stg.
- This process is to insure that user defined requests are updated first in the event that metrics_resoruces_usage contains multiple rows for the same requestid where the resourcekind spans both [user defined] and [system].
- If multiple distinct "requestid, resourcekind and dataservicename" rows are found the algorithm orders by mruh.requestid, mruh.resourcekind desc, mruh.dataservicename so that "user defined" is sorted before "system" as "user defined" has higher priority.
- Only one record can be chosen.

Loop through the records using the following SQL

```

select distinct mruh."requestid", mrh."starttime", mrh."endtime", mrh."status", mruh."resourcekind",
mruh."dataservicename"
from ""||dataSourceSchemaName||"."||metrics_requests_stg||" mrh
  inner join ""||dataSourceSchemaName||"."||metrics_resources_usage_stg||" mruh
    on mrh."requestid" = mruh."requestid"
    and mrh."starttime" = mruh."starttime"
    and mrh."nodehost" = mruh."nodehost"
    and mrh."nodeport" = mruh."nodeport"
    where mruh."resourcekind" is not null
    and mruh."dataservicename" is not null
    and mrh."nodehost" = c_node_host
    and mrh."nodeport" = c_node_port
    and (mrh."resourcekind" is null or mrh."dataservicename" is null)
    and mrh."requestid" >= c_minrequestid and mrh."requestid" <= c_maxrequestid
    and mrh."starttime" >= c_minstarttime and mrh."starttime" <= c_maxstarttime
  order by mruh."requestid", mruh."resourcekind" desc, mruh."dataservicename";

```

Insert records into metrics_requests_stg table in pre-defined batches

```

insert into ""||dataSourceSchemaName||"."||metrics_requests_stg_upd||" values(t."requestid",
t."starttime", t."endtime", t."status", t."resourcekind", t."dataservicename", in_node_host, in_node_port);

```

Update the metrics_requests_stg with the single row for each requestid using the batch in the metrics_request_stg_upd stage table. The min and max requestid and starttime are important

to establish boundaries for the update. Without them the resourcekind and dataservicename for rows not in the stage update table are set to null.

```
update "||dataSourceSchemaName||"."||metrics_requests_stg||" mrh
  set (mrh."resourcekind", mrh."dataservicename") =
    (select mrht."resourcekind", mrht."dataservicename"
      from "||dataSourceSchemaName||"."||metrics_requests_stg_upd||" mrht
      where mrh."requestid" = mrht."requestid"
        and mrh."starttime" = mrht."starttime"
        and mrh."endtime" = mrht."endtime"
        and mrh."status" = mrht."status"
        and mrh."nodehost" = mrht."nodehost"
        and mrh."nodeport" = mrht."nodeport"
        and mrht."nodehost" = in_node_host
        and mrht."nodeport" = in_node_port
    )
  where mrh."resourcekind" is null
    and mrh."dataservicename" is null
    and mrh."nodehost" = in_node_host
    and mrh."nodeport" = in_node_port
    and mrh."requestid" >= minrequestid_mrb
    and mrh."requestid" <= maxrequestid_mrb
    and mrh."starttime" >= minstarttime_mrb
    and mrh."starttime" <= maxstarttime_mrb;
```

Delete only rows associated with the nodehost and nodeport that is executing this script

```
delete from "||dataSourceSchemaName||"."||metrics_requests_stg_upd||" where "nodehost" =
in_node_host and "nodeport" = in_node_port;
```

End of Loop

5. Insert into the metrics_requests_hist history from the metrics_requests_stg stage

```
INSERT INTO "||dataSourceSchemaName||"."||metrics_requests_hist||"
  SELECT * FROM "||dataSourceSchemaName||"."||metrics_requests_stg||"
  where "nodehost" = in_node_host
    and "nodeport" = in_node_port;
```

6. Delete the remainder of metrics_requests collection rows

- a. **Note: if the kpimetricsTrig_33_DeleteCollection trigger is still running then this section is skipped.**

```
delete from "||dataSourceSchemaName||"."||metrics_requests_collection||"
  where "nodehost" = in_node_host
    and "nodeport" = in_node_port
    and "starttime" <= maxstarttime_mr
    and "requestid" <= maxrequestid_mr;
```

Physical SQL Server Data Transfer Script

Location: /shared/ASAssets/KPImetrics/Physical/Metadata/DDL/SqlServer/
06_pqCreateDrop_KPI_Plsq_sqlserver_data_xfer_script

This section describes the actions and SQL taken in the P_METRICS_ALL_TABLES data transfer script. The purpose of this script is to transfer data from the collection tables to the history tables using native SQL Server SQL script. Stage tables are used to prepare and augment the collection data prior to inserting into the history tables. The result of this is that only "inserts" are performed on the history tables. All of the insert/update/deletes that are required are performed on the stage tables.

metrics_sessions

1. Get the min/max sessionid and logintime

```
select min("sessionid") minsessionid, min("logintime") minlogintime, max("sessionid") maxsessionid,
max("logintime") maxlogintime
  from "||dataSourceCatalogName||"."||dataSourceSchemaName||"."||metrics_sessions_collection||"
  where "nodehost" = @in_node_host and "nodeport" = @in_node_port;
```

2. Insert from metrics_sessions collection table into metrics_sessions_stg stage table

```
insert into "||dataSourceCatalogName||"."||dataSourceSchemaName||"."||metrics_sessions_stg||"
  select ms.*,
         LU."USER_ID", LU."user", LU."domain",
  -- Perform LDAP userkey join transformation. The transformation value in userKeyTransformation is provided
  -- at the time this procedure is created which is during installation.
         '||REPLACE(userKeyTransform, "user", 'LU."user"||' as "userkey",
         '||REPLACE(domainKeyTransform, "domain", 'LU."domain"||' as "domainkey",
         null "fullName", null "lastName", null "firstName", null "mail", null "employeeNumber", null "employeeID",
         null "telephoneNumber", null "friendlyCountryName"
  from "||dataSourceCatalogName||"."||dataSourceSchemaName||"."||metrics_sessions_collection||" ms
  -- Join with metrics_resources_usage to get the user
  LEFT OUTER JOIN (
    select distinct mruc."nodehost", mruc."nodeport", mruc."sessionid", mruc."user", mruc."domain",
    mau."USER_ID"
  FROM
    "||dataSourceCatalogName||"."||dataSourceSchemaName||"."||metrics_resources_usage_collection||" mruc
  -- Join with METRICS_ALL_USERS to get the latest (max LOAD_TIME) userid
  LEFT OUTER JOIN
    (SELECT * FROM
    "||dataSourceCatalogName||"."||dataSourceSchemaName||"."METRICS_ALL_USERS" au1
    WHERE LOAD_TIME = ( SELECT MAX(LOAD_TIME) FROM
    "||dataSourceCatalogName||"."||dataSourceSchemaName||"."METRICS_ALL_USERS" au2
    WHERE au1.USERNAME = au2.USERNAME AND au1.DOMAIN_NAME = au2.
    DOMAIN_NAME
    AND au1.NODE_HOST = au2.NODE_HOST AND au1.NODE_PORT = au2.NODE_PORT
    )
  ) mau
  ON mau."USERNAME" = mruc."user"
  AND mau."DOMAIN_NAME" = mruc."domain"
  AND mau."NODE_HOST" = mruc."nodehost"
  AND mau."NODE_PORT" = mruc."nodeport"
  ) LU
  ON LU."sessionid" = ms."sessionid"
  AND LU."nodehost" = ms."nodehost"
  AND LU."nodeport" = ms."nodeport"
  where ms."logintime" <= @maxlogintime
```

```

and ms."sessionid" <= @maxsessionid
and ms."nodehost" = @in_node_host
and ms."nodeport" = @in_node_port
and NOT EXISTS (
  select *
  from "||dataSourceCatalogName||"."||dataSourceSchemaName||"."||metrics_sessions_hist||" ms1
  where ms1."nodehost" = ms."nodehost"
  and ms1."nodeport" = ms."nodeport"
  and ms1."sessionid" = ms."sessionid"
  and ms1."logintime" = ms."logintime"
  and ms1."type" = ms."type"
);

```

3. Update the `metrics_sessions_stg` table with user information from `METRICS_LDAP_PERSON` using the “userkey” to join.

```

UPDATE "||dataSourceCatalogName||"."||dataSourceSchemaName||"."||metrics_sessions_stg||"
  SET "fullname" = mlp."fullName",
      "lastname" = mlp."lastName",
      "firstname" = mlp."firstName",
      "mail" = mlp."mail",
      "employeenumber" = mlp."employeeNumber",
      "employeeid" = mlp."employeeID",
      "telephonenumber" = mlp."telephoneNumber",
      "friendlycountryname" = mlp."friendlyCountryName"
  FROM "||dataSourceCatalogName||"."||dataSourceSchemaName||"."||metrics_sessions_stg||" msh
  LEFT OUTER JOIN (
    SELECT mlp1."userkey", mlp1."domainkey", mlp1."fullName", mlp1."lastName", mlp1."firstName",
    mlp1."mail", mlp1."employeeNumber", mlp1."employeeID", mlp1."telephoneNumber",
    mlp1."friendlyCountryName"
    FROM "||dataSourceCatalogName||"."||dataSourceSchemaName||"."METRICS_LDAP_PERSON"
    mlp1
    WHERE mlp1."loadtime" =
      ( SELECT MAX(mlp2."loadtime")
        FROM
        "||dataSourceCatalogName||"."||dataSourceSchemaName||"."METRICS_LDAP_PERSON" mlp2
        WHERE mlp1."userkey" = mlp2."userkey" AND mlp1."domainkey" = mlp2."domainkey")
    ) mlp
  ON mlp."userkey" = msh."userkey"
  AND mlp."domainkey" = msh."domainkey"
  WHERE msh."logintime" >= @minlogintime
  AND msh."sessionid" >= @minsessionid
  AND msh."nodehost" = @in_node_host
  AND msh."nodeport" = @in_node_port
  AND msh."user" is not null
  AND msh."userkey" is not null
  AND msh."mail" is null
  AND msh."lastname" is null;

```

4. Insert into `metrics_sessions_hist` history from `metrics_sessions_stg` stage

```

INSERT INTO "||dataSourceCatalogName||"."||dataSourceSchemaName||"."||metrics_sessions_hist||"
  SELECT * FROM "||dataSourceCatalogName||"."||dataSourceSchemaName||"."||metrics_sessions_stg||"
  where "nodehost" = in_node_host
  and "nodeport" = in_node_port;

```

5. Delete this batch of data from the `metrics_sessions` collection table that was just processed being careful not to delete any new rows.

```
delete from "||dataSourceCatalogName||"."||dataSourceSchemaName||"."||metrics_sessions_collection||"
where "logintime" <= @maxlogintime
and "sessionid" <= @maxsessionid
and "nodehost" = @in_node_host
and "nodeport" = @in_node_port;
```

metrics_resources_usage

1. Get the min/max requestid and starttime

```
select min("requestid") minrequestid, min("starttime") minstarttime, max("requestid") maxrequestid,
max("starttime") maxstarttime
from
"||dataSourceCatalogName||"."||dataSourceSchemaName||"."||metrics_resources_usage_collection||"
where "nodehost" = @in_node_host and "nodeport" = @in_node_port;
```

2. Insert the new rows into the metrics_resources_usage_stg stage from the metrics_resources_usage collection where not exists in metrics_resources_usage_hist history and not in METRICS_JOB_FILTERS.

```
INSERT INTO "||dataSourceCatalogName||"."||dataSourceSchemaName||"."||metrics_resources_usage_stg||"
SELECT
-- Insert collection columns
  mruc."cluster"
  , mruc."nodehost"
  , mruc."nodeport"
  , mruc."sessionid"
  , mruc."user"
  , mruc."domain"
  , mruc."group"
  , mruc."requestid"
  , mruc."parentid"
  , mruc."datasourcepath"
  , mruc."datasourcetype"
  , mruc."resourcepath"
  , mruc."resourcetype"
  , mruc."resourceguid"
  , mruc."resourcekind"
  , mruc."starttime"
  , mruc."endtime"
-- Insert expanded information
  , mar."RESOURCE_ORIGIN" "resourceorigin"
  , mar."RESOURCE_ID" "resourceid"
  , mar."DATASERVICE_NAME" "dataservicename"
  , mar."RESOURCE_NAME" "resourcenname"
  , mar."SCHEMA_NAME" "parentname"
  , mar."CATALOG_NAME" "grandparentname"
  , mar."CATEGORY_NAME" "categoryname"
  , mar."PARENT_PATH" "parentpath"
-- Insert expanded user information
  , mau."USER_ID" "userid"
  , '||REPLACE(userKeyTransform, "user", 'mruc."user"')||' as "userkey"
  , '||REPLACE(domainKeyTransform, "domain", 'mruc."domain"')||' as "domainkey"
  , null "fullName", null "lastName", null "firstName", null "mail", null "employeeNumber", null "employeeID", null
"telephoneNumber", null "friendlyCountryName"
FROM "||dataSourceCatalogName||"."||dataSourceSchemaName||"."||metrics_resources_usage_collection||" mruc
-- METRICS_ALL_USERS
LEFT OUTER JOIN
(SELECT * FROM "||dataSourceCatalogName||"."||dataSourceSchemaName||"."METRICS_ALL_USERS" au1
WHERE "LOAD_TIME" = (SELECT MAX("LOAD_TIME") FROM
"||dataSourceCatalogName||"."||dataSourceSchemaName||"."METRICS_ALL_USERS" au2
WHERE au1."USERNAME" = au2."USERNAME"
AND au1."DOMAIN_NAME" = au2."DOMAIN_NAME"
AND au1."NODE_HOST" = au2."NODE_HOST"
```

```

        AND au1."NODE_PORT" = au2."NODE_PORT" )
    ) mau
    ON mau."USERNAME" = mruc."user"
    AND mau."DOMAIN_NAME" = mruc."domain"
    AND mau."NODE_HOST" = mruc."nodehost"
    AND mau."NODE_PORT" = mruc."nodeport"
-- METRICS_ALL_RESOURCES
    LEFT OUTER JOIN
    (SELECT mar1."RESOURCE_ORIGIN", mar1."RESOURCE_ID", mar1."RESOURCE_NAME",
    mar1."RESOURCEPATH", mar1."RESOURCE_TYPE", mar1."PARENT_PATH",
    mar1."DATASERVICE_NAME", mar1."CATALOG_NAME", mar1."SCHEMA_NAME",
    mar1."CATEGORY_NAME", mar1."NODE_HOST", mar1."NODE_PORT"
    FROM "||dataSourceCatalogName||"."||dataSourceSchemaName||"."METRICS_ALL_RESOURCES" mar1
    WHERE mar1."LOAD_TIME" =
    ( SELECT MAX(mar2."LOAD_TIME")
    FROM "||dataSourceCatalogName||"."||dataSourceSchemaName||"."METRICS_ALL_RESOURCES" mar2
    WHERE mar1."RESOURCEPATH" = mar2."RESOURCEPATH"
    AND mar1."RESOURCE_TYPE" = mar2."RESOURCE_TYPE"
    AND mar1."NODE_HOST" = mar2."NODE_HOST"
    AND mar1."NODE_PORT" = mar2."NODE_PORT" )
    ) mar
    ON mar."RESOURCEPATH" = mruc."resourcepath"
    AND mar."RESOURCE_TYPE" = mruc."resourcetype"
    AND mar."NODE_HOST" = mruc."nodehost"
    AND mar."NODE_PORT" = mruc."nodeport"
WHERE mruc."nodehost" = @in_node_host
AND mruc."nodeport" = @in_node_port
AND mruc."starttime" <= @maxstarttime_mru
AND mruc."requestid" <= @maxrequestid_mr
-- Do not insert rows matching the filter rows in METRICS_JOB_FILTERS
AND mruc."requestid" NOT IN
(SELECT mru."requestid"
FROM "||dataSourceCatalogName||"."||dataSourceSchemaName||"."metrics_resources_usage_collection||"
mru
JOIN "||dataSourceCatalogName||"."||dataSourceSchemaName||"."METRICS_JOB_FILTERS" mjf
ON mru."user" = mjf."USER"
AND mru."domain" = mjf."DOMAIN"
AND mru."resourcekind" = mjf."RESOURCE_KIND"
AND mru."nodeport" = @in_node_port AND mru."nodehost" = @in_node_host
WHERE mjf."ENV_TYPE" = @in_env_type )
-- Do not insert rows if they exist in metrics_resources_usage_hist
AND NOT EXISTS (
SELECT 1
FROM "||dataSourceCatalogName||"."||dataSourceSchemaName||"."metrics_resources_usage_hist||" mruh
WHERE mruc."nodehost" = mruh."nodehost"
AND mruc."nodeport" = mruh."nodeport"
AND mruc."starttime" = mruh."starttime"
AND mruc."requestid" = mruh."requestid"
);

```

3. Update the metrics_resources_usage_stg table with user information from METRICS_LDAP_PERSON using the “userkey” to join.

```

UPDATE "||dataSourceCatalogName||"."||dataSourceSchemaName||"."metrics_resources_usage_stg||"
SET "fullname" = mlp."fullName",
    "lastname" = mlp."lastName",
    "firstname" = mlp."firstName",
    "mail" = mlp."mail",
    "employeenumber" = mlp."employeeNumber",
    "employeeid" = mlp."employeeID",
    "telephonenumber" = mlp."telephoneNumber",
    "friendlycountryname" = mlp."friendlyCountryName"
FROM "||dataSourceCatalogName||"."||dataSourceSchemaName||"."metrics_resources_usage_stg||"
mruh

```

```

LEFT OUTER JOIN (
    SELECT mlp1."userkey", mlp1."domainkey", mlp1."fullName", mlp1."lastName", mlp1."firstName",
    mlp1."mail", mlp1."employeeNumber", mlp1."employeeID", mlp1."telephoneNumber",
    mlp1."friendlyCountryName"
    FROM "||dataSourceCatalogName||"."||dataSourceSchemaName||"."METRICS_LDAP_PERSON"
    mlp1
    WHERE mlp1."loadtime" =
    ( SELECT MAX(mlp2."loadtime")
    FROM
    "||dataSourceCatalogName||"."||dataSourceSchemaName||"."METRICS_LDAP_PERSON" mlp2
    WHERE mlp1."userkey" = mlp2."userkey" AND mlp1."domainkey" = mlp2."domainkey")
) mlp
ON mlp."userkey" = mruh."userkey"
AND mlp."domainkey" = mruh."domainkey"
WHERE mruh."starttime" >= @minstarttime_mru
AND mruh."requestid" >= @minrequestid_mru
AND mruh."nodehost" = @in_node_host
AND mruh."nodeport" = @in_node_port
AND mruh."user" is not null
AND mruh."userkey" is not null
AND mruh."mail" is null
AND mruh."lastname" is null;

```

6. Insert into metrics_resources_usage_hist history from metrics_resources_usage_stg stage

```

INSERT INTO
"||dataSourceCatalogName||"."||dataSourceSchemaName||"."metrics_resources_usage_hist||"
SELECT * FROM
"||dataSourceCatalogName||"."||dataSourceSchemaName||"."metrics_resources_usage_stg||"
where "nodehost" = in_node_host
and "nodeport" = in_node_port;

```

4. Delete the metrics_resources_usage collection remaining data

- a. **Note: if the kpimetricsTrig_33_DeleteCollection trigger is still running then this section is skipped.**

```

delete from
"||dataSourceCatalogName||"."||dataSourceSchemaName||"."metrics_resources_usage_collection||"
where "nodehost" = @in_node_host
and "nodeport" = @in_node_port
and "starttime" <= @maxstarttime_mru
and "requestid" <= @maxrequestid_mru;

```

metrics_requests

1. Get the min/max requestid and starttime

```

select min("requestid") minrequestid, min("starttime") minstarttime, max("requestid") maxrequestid,
max("starttime") maxstarttime
from "||dataSourceCatalogName||"."||dataSourceSchemaName||"."metrics_requests_collection||"
where "nodehost" = @in_node_host and "nodeport" = @in_node_port;

```

2. Insert the new rows into the metrics_requests_stg stage from the metrics_requests collection where not exists in metrics_requests_hist history and exists in metrics_resources_usage_stg.

```

INSERT INTO "||dataSourceCatalogName||"."||dataSourceSchemaName||"."metrics_requests_stg||"
SELECT mrc.*,
null "dataservicename", null "resourcekind",

```

```

    LU."USER_ID", LU."user", LU."domain",
    '||REPLACE(userKeyTransform, "user", 'LU."user"||' as "userkey",
    null "fullName", null "lastName", null "firstName", null "mail", null "employeeNumber", null
"employeeID", null "telephoneNumber", null "friendlyCountryName"
FROM "||dataSourceCatalogName||"."||dataSourceSchemaName||"."||metrics_requests_collection||" mrc
LEFT OUTER JOIN (
    SELECT DISTINCT mruc."nodehost", mruc."nodeport", mruc."requestid", mruc."user", mruc."domain",
    mau."USER_ID"
    FROM
    "||dataSourceCatalogName||"."||dataSourceSchemaName||"."||metrics_resources_usage_collection||"
    mruc
    LEFT OUTER JOIN
    (SELECT * FROM
    "||dataSourceCatalogName||"."||dataSourceSchemaName||"."METRICS_ALL_USERS" au1
    WHERE "LOAD_TIME" = (SELECT MAX("LOAD_TIME") FROM
    "||dataSourceCatalogName||"."||dataSourceSchemaName||"."METRICS_ALL_USERS" au2
    WHERE au1."USERNAME" = au2."USERNAME"
    AND au1."DOMAIN_NAME" = au2."DOMAIN_NAME"
    AND au1."NODE_HOST" = au2."NODE_HOST"
    AND au1."NODE_PORT" = au2."NODE_PORT" )
    ) mau
    ON mau."USERNAME" = mruc."user"
    AND mau."DOMAIN_NAME" = mruc."domain"
    AND mau."NODE_HOST" = mruc."nodehost"
    AND mau."NODE_PORT" = mruc."nodeport"
    ) LU
    ON LU."requestid" = mrc."requestid"
    AND LU."nodehost" = mrc."nodehost"
    AND LU."nodeport" = mrc."nodeport"
WHERE mrc."nodehost" = @in_node_host
AND mrc."nodeport" = @in_node_port
AND mrc."starttime" <= @maxstarttime_mr
AND mrc."requestid" <= @maxrequestid_mr
AND NOT EXISTS (
    SELECT 1
    FROM "||dataSourceCatalogName||"."||dataSourceSchemaName||"."||metrics_requests_hist||" mrh
    WHERE mrc."nodehost" = mrh."nodehost"
    AND mrc."nodeport" = mrh."nodeport"
    AND mrc."starttime" = mrh."starttime"
    AND mrc."requestid" = mrh."requestid"
    )
-- Only insert metrics_requests collection rows when a corresponding row exists in the
metrics_resources_usage_stg
-- No point in capturing these rows without relevant resource usage information.
AND EXISTS (
    SELECT 1 FROM
    "||dataSourceCatalogName||"."||dataSourceSchemaName||"."||metrics_resources_usage_stg||" mru
    WHERE mru."nodehost" = mrc."nodehost"
    AND mru."nodeport" = mrc."nodeport"
    AND mru."requestid" = mrc."requestid"
    AND mru."starttime" = mrc."starttime"
    );

```

3. Update the new rows from the metrics_requests_stg stage with METRICS_LDAP_PERSON data

```

UPDATE "||dataSourceCatalogName||"."||dataSourceSchemaName||"."||metrics_requests_stg||"
SET "fullname" = mlp."fullName",
    "lastname" = mlp."lastName",
    "firstname" = mlp."firstName",

```



```

"mail" = mlp."mail",
"employeenumber" = mlp."employeeNumber",
"employeeid" = mlp."employeeID",
"telephonenumber" = mlp."telephoneNumber",
"friendlycountryname" = mlp."friendlyCountryName"
FROM ""||dataSourceCatalogName||"."||dataSourceSchemaName||"."||metrics_requests_stg||" mrh
LEFT OUTER JOIN (
    SELECT mlp1."userkey", mlp1."domainkey", mlp1."fullName", mlp1."lastName", mlp1."firstName",
    mlp1."mail", mlp1."employeeNumber", mlp1."employeeID", mlp1."telephoneNumber",
    mlp1."friendlyCountryName"
    FROM ""||dataSourceCatalogName||"."||dataSourceSchemaName||"."METRICS_LDAP_PERSON"
    mlp1
    WHERE mlp1."loadtime" =
        ( SELECT MAX(mlp2."loadtime")
        FROM
""||dataSourceCatalogName||"."||dataSourceSchemaName||"."METRICS_LDAP_PERSON" mlp2
        WHERE mlp1."userkey" = mlp2."userkey" AND mlp1."domainkey" = mlp2."domainkey")
    ) mlp
    ON mlp."userkey" = mrh."userkey"
    AND mlp."domainkey" = mrh."domainkey"
WHERE mrh."starttime" >= @minstarttime_mr
AND mrh."requestid" >= @minrequestid_mr
AND mrh."nodehost" = @in_node_host
AND mrh."nodeport" = @in_node_port
AND mrh."user" is not null
AND mrh."userkey" is not null
AND mrh."mail" is null
AND mrh."lastname" is null;

```

4. Update metrics_requests_stg stage [resourcekind and dataservicename] joined with metrics_resources_usage_stg.

- Use the resourcekind and dataservicename from metrics_resources_usage_stg.
- This process is to ensure that user defined requests are updated first in the event that metrics_resources_usage contains multiple rows for the same requestid where the resourcekind spans both [user defined] and [system].
- If multiple distinct "requestid, resourcekind and dataservicename" rows are found the algorithm orders by mruh.requestid, mruh.resourcekind desc, mruh.dataservicename so that "user defined" is sorted before "system" as "user defined" has higher priority.
- Only one record can be chosen.

Update where resourcekind='user defined'

```

update ""||dataSourceCatalogName||"."||dataSourceSchemaName||"."||metrics_requests_stg||"
    set "resourcekind" = mruh."resourcekind",
        "dataservicename" = mruh."dataservicename"
    from ""||dataSourceCatalogName||"."||dataSourceSchemaName||"."||metrics_requests_stg||" mrh
    inner join
""||dataSourceCatalogName||"."||dataSourceSchemaName||"."||metrics_resources_usage_stg||" mruh
    on mruh."requestid" = mrh."requestid"
    and mruh."starttime" = mrh."starttime"
    and mruh."nodehost" = mrh."nodehost"
    and mruh."nodeport" = mrh."nodeport"
    where mrh."nodehost" = @in_node_host
    and mrh."nodeport" = @in_node_port
    and mrh."starttime" <= @maxstarttime_mr
    and mrh."requestid" <= @maxrequestid_mr
    and mrh."resourcekind" is null
    and mrh."dataservicename" is null
    and mruh."dataservicename" is not null

```

```
and mruh."resourcekind" = "user defined"; -- [user defined, system]
```

Update where resourcekind='system'

```
update ""||dataSourceCatalogName||"."||dataSourceSchemaName||"."||metrics_requests_stg||"
    set "resourcekind" = mruh."resourcekind",
        "dataservicename" = mruh."dataservicename"
    from ""||dataSourceCatalogName||"."||dataSourceSchemaName||"."||metrics_requests_stg||" mruh
    inner join
""||dataSourceCatalogName||"."||dataSourceSchemaName||"."||metrics_resources_usage_stg||" mruh
    on mruh."requestid" = mrh."requestid"
    and mruh."starttime" = mrh."starttime"
    and mruh."nodehost" = mrh."nodehost"
    and mruh."nodeport" = mrh."nodeport"
    where mrh."nodehost" = @in_node_host
    and mrh."nodeport" = @in_node_port
    and mrh."starttime" <= @maxstarttime_mr
    and mrh."requestid" <= @maxrequestid_mr
    and mrh."resourcekind" is null
    and mrh."dataservicename" is null
    and mruh."dataservicename" is not null
    and mruh."resourcekind" = "system"; -- [user defined, system]
```

7. Insert into metrics_requests_hist history from metrics_requests_stg stage

```
INSERT INTO ""||dataSourceCatalogName||"."||dataSourceSchemaName||"."||metrics_requests_hist||"
    SELECT * FROM ""||dataSourceCatalogName||"."||dataSourceSchemaName||"."||metrics_requests_stg||"
    where "nodehost" = in_node_host
    and "nodeport" = in_node_port;
```

5. Delete the remainder of metrics_requests collection rows

- a. **Note: if the kpimetricsTrig_33_DeleteCollection trigger is still running then this section is skipped.**

```
delete from ""||dataSourceCatalogName||"."||dataSourceSchemaName||"."||metrics_requests_collection||"
    where "nodehost" = @in_node_host
    and "nodeport" = @in_node_port
    and "starttime" <= @maxstarttime_mr
    and "requestid" <= @maxrequestid_mr;
```

7 Release Notes

This section provides information on what has changed in this release.

Added or Modified in this Release

This section provides bullet points on what has been added or changed in this release.

Release 2019Q102 [Mar 12 2019]

- Release Objective: Added upgrade scripts. Added "group" to metrics_requests and metrics_sessions for 7.0 to bring up to par with 8.0.
 - Added upgrade scripts starting with 2018Q103 for oracle and sqlserver.
 - Allows upgrading from 2018Q103 or higher on 7.x to 7.x.
 - Allows upgrading from 2018Q103 or higher on 7.x to 8.x.
 - Added indexes for METRICS_CIS_WORKFLOW:
 - METRICS_CIS_WORKFLOW_IDX1 and METRICS_CIS_WORKFLOW_IDX2
 - Modified schemas [collection, staging] for 7.0.
 - "group" was added to metrics_sessions and metrics_requests for 7.0.
 - Modified schemas [history] to work with 8.0 modified metrics schemas.
 - "group" was added to metrics_sessions_hist and metrics_requests_hist for 7.0 as "user" and "domain" already exists.
 - Modified Oracle DDL Scripts:
 - Modified 02_pqCreateDrop_KPI_Tables_oracle_metrics_collection_stage_tables_700 to add "group"
 - Modified 03_pqCreateDrop_KPI_Tables_oracle_metrics_history_tables_700 to add "group"
 - Modified 06_pqCreateDrop_KPI_Plsql_oracle_data_xfer_script_700 to add "group"
 - Modified SQL Server DDL Scripts:
 - Modified 02_pqCreateDrop_KPI_Tables_sqlserver_metrics_collection_stage_tables_700 to add "group"
 - Modified 03_pqCreateDrop_KPI_Tables_sqlserver_metrics_history_tables_700 to add "group"
 - Modified 06_pqCreateDrop_KPI_Plsql_sqlserver_data_xfer_script_700 to add "group"
 - Renamed updateImpactedResources to updateImpactedResources (missing d in Impacted).
 - Modified /shared/ASAssets/KPImetrics_installation to account for updateImpactedResources.
 - Modified all the scripts in /shared/ASAssets/KPImetrics/Physical/Metadata/System/ClusterSafeCache that contained "INSERT INTO" statements to specify INSERT and SELECT columns and not use SELECT *.

Release 2019Q101 [Jan 30 2019]

- Release Objective: Added support for DV 8.0 as the out-of-the-box metrics tables [metrics_requests and metrics_sessions] added user, domain and group. KPI installation automatically checks for DV version 7.0 or 8.0 and executes the appropriate scripts to install KPImetrics.
 - Bug Fixes:
 - sqlParserV1_2 - Various scripts. Fixed a bug when server IGNORE_TRAILING_SPACES=true. Changed INSTR functions with a trailing space to use SELECT {OPTION IGNORE_TRAILING_SPACES="FALSE"} INSTR(sqlScript, ' ')
 - Modified KPImetrics_installation scripts to install for 7.0 or 8.0 automatically.
 - Modified schemas [collection, staging, history] to work with 8.0 modified metrics schemas. The columns "user", "domain", "group" were added to metrics_sessions and metrics_requests for 8.0.
 - Modified Oracle DDL Scripts:
 - Renamed 02_pqCreateDrop_KPI_Tables_oracle_metrics_collection_stage_tables to 02_pqCreateDrop_KPI_Tables_oracle_metrics_collection_stage_tables_700
 - Added 02_pqCreateDrop_KPI_Tables_oracle_metrics_collection_stage_tables_800
 - Renamed 02_pqCreateDrop_KPI_Tables_oracle_metrics_collection_tables to 02_pqCreateDrop_KPI_Tables_oracle_metrics_collection_tables_700
 - Added 02_pqCreateDrop_KPI_Tables_oracle_metrics_collection_tables_800
 - Renamed 03_pqCreateDrop_KPI_Tables_oracle_metrics_history_tables to 03_pqCreateDrop_KPI_Tables_oracle_metrics_history_tables_700
 - Added 03_pqCreateDrop_KPI_Tables_oracle_metrics_history_tables_800
 - Renamed 06_pqCreateDrop_KPI_Plsql_oracle_data_xfer_script to 06_pqCreateDrop_KPI_Plsql_oracle_data_xfer_script_700
 - Added 06_pqCreateDrop_KPI_Plsql_oracle_data_xfer_script_800
 - Modified SQL Server DDL Scripts:
 - Renamed 02_pqCreateDrop_KPI_Tables_oracle_metrics_collection_stage_tables to 02_pqCreateDrop_KPI_Tables_oracle_metrics_collection_stage_tables_700
 - Renamed 02_pqCreateDrop_KPI_Tables_sqlserver_metrics_collection_stage_tables to 02_pqCreateDrop_KPI_Tables_sqlserver_metrics_collection_stage_tables_700
 - Added 02_pqCreateDrop_KPI_Tables_sqlserver_metrics_collection_stage_tables_800
 - Renamed 02_pqCreateDrop_KPI_Tables_sqlserver_metrics_collection_tables to 02_pqCreateDrop_KPI_Tables_sqlserver_metrics_collection_tables_700
 - Added 02_pqCreateDrop_KPI_Tables_sqlserver_metrics_collection_tables_800
 - Renamed 03_pqCreateDrop_KPI_Tables_sqlserver_metrics_history_tables to 03_pqCreateDrop_KPI_Tables_sqlserver_metrics_history_tables_700
 - Added 03_pqCreateDrop_KPI_Tables_sqlserver_metrics_history_tables_800
 - Renamed 06_pqCreateDrop_KPI_Plsql_sqlserver_data_xfer_script to 06_pqCreateDrop_KPI_Plsql_sqlserver_data_xfer_script_700

- Added 06_pqCreateDrop_KPI_Plsql_sqlserver_data_xfer_script_800

Release 2019Q100 [Jan 7 2019]

- Release Objective: Fixed installation scripts. Added "domainkey" to METRICS_LDAP_PERSON and history tables.
 - Fixed KPImetrics_installation scripts.
 - Modified rebindPhysicalDatabaseType to work with installation scripts.
 - Modified Cache_CPU_MEMORY_CHECKER to put double quotes around windows path.
 - Added domainkey column to qualify a user. The same user may exist in different domains.
 - **** NOTE ****
 - [Perform these actions in order for an existing installation of KPImetrics which contains history data.]
 - [This is not required for new installations.]
 - 1. If the history tables already exist then use alter statements found in /Physical/Physical/KPI_[oracle|sqlserver]/00_fix
 - The alter statements will also modify METRICS_LDAP_PERSON
 - 2. It is required to drop and recreate the staging tables using 02_pqCreateDrop_KPI_Tables_....
 - 3. It is required to drop and recreate the stored procedure using 06_pqCreateDrop_KPI_Plsql_...
 - 4. Update the history tables "domainkey" column with the domain name for all rows.
 - Update the METRICS_LDAP_PERSON table "domainkey" column with the domain name for all rows.
 - Provided the ability to map to multiple LDAP domains
 - Renamed data source /shared/ASAssets/KPImetrics/Physical/Metadata/LDAP to LDAP1
 - Added data source /shared/ASAssets/KPImetrics/Physical/Metadata/LDAP2
 - Modified /Metadata/System/LDAP_PERSON - added "domainkey".
 - Modified /Metadata/System/ClusterSafeCache/Cache_LDAP_PERSON
 - Modified /Physical/Formatting/metrics_resources_usage_hist - added "domainkey".
 - Modified /Business/Logical/userUsage/vLdapPerson - added "domainkey" to the where clause to qualify "user".
 - Modified /shared/ASAssets/KPImetrics/Business/Business/requests/pExceededMemoryPercentRequests to return domain name.
 - Modified /shared/ASAssets/KPImetrics/Physical/Metadata/System/ClusterSafeCache/pCheckExceedMemoryPercentRequests to qualify userkey with domainkey.
 - Modified /shared/ASAssets/KPImetrics/Business/Business/requests/pLongRunningRequests to return domain name.
 - Modified /shared/ASAssets/KPImetrics/Physical/Metadata/System/ClusterSafeCache/pCheckLongRunningRequests to qualify userkey with domainkey.

- Modified /Metadata/System/ClusterSafeCache/ Cache_AllCustom_AccessByUserOverTime
 - added "domain" to the where clause to qualify "user".
- Modified /Physical/Physical/KPI_oracle/00_fix - alter statements for tables.
- Modified /Physical/Physical/KPI_sqlserver/00_fix - alter statements for tables.
- Modified /Metadata/DDL/Oracle/02_pqCreateDrop_KPI_Tables_oracle_metrics_collection_stage_tables
 - added "domainkey" to metrics_requests_stg, metrics_resources_usage_stg, metrics_sessions_stg
- Modified /Metadata/DDL/Oracle/03_pqCreateDrop_KPI_Tables_oracle_metrics_history_tables
 - added "domainkey" to metrics_requests_hist, metrics_resources_usage_hist, metrics_sessions_hist
- Modified /Metadata/DDL/Oracle/04_pqCreateDrop_KPI_Tables_oracle_kpi_tables
 - added "domainkey" to METRICS_LDAP_PERSON
- Modified /Metadata/DDL/Oracle/06_pqCreateDrop_KPI_Plsql_oracle_data_xfer_script to use "domainkey"
- Modified /Metadata/DDL/SqlServer/02_pqCreateDrop_KPI_Tables_sqlserver_metrics_collection_stage_tables
 - added "domainkey" to metrics_requests_stg, metrics_resources_usage_stg, metrics_sessions_stg
- Modified /Metadata/DDL/SqlServer/03_pqCreateDrop_KPI_Tables_sqlserver_metrics_history_tables
 - added "domainkey" to metrics_requests_hist, metrics_resources_usage_hist, metrics_sessions_hist
- Modified /Metadata/DDL/SqlServer/04_pqCreateDrop_KPI_Tables_sqlserver_kpi_tables
 - added "domainkey" to METRICS_LDAP_PERSON
- Modified /Metadata/DDL/SqlServer/06_pqCreateDrop_KPI_Plsql_sqlserver_data_xfer_script to use "domainkey"

Release 2018Q401 [Dec 4 2018]

- Release Objective: Deprecated/Removed METRICS_RESOURCES_USAGE_UD.
 - This was a copy of the data from metrics_resources_usage_hist where resourcekind='user defined'. Removed all processing associated with this table and made it a native access.
 - Modified /logical/resourceUsage/vResourceUsageUD to point to metrics_resources_usage_hist.
 - Modified /logical/resourceUsage/vResourceUsageAll to select categoryname with no coalesce (bug fix).
 - Modified /Physical/KPI_[oracle|sqlserver]/EXCEPT_ACR_AccessByUserOverTime to use metrics_resources_usage_hist.
 - Modified /Physical/KPI_[oracle|sqlserver]/EXCEPT_ACR_ActiveResourcesOverPeriodOfTime to use metrics_resources_usage_hist.
 - Modified 03_pqCreateDrop_KPI_Tables_oracle_metrics_history_tables to add index "mru_hist_reskind_dsname".

- Modified 04_pqCreateDrop_KPI_Tables_oracle_kpi_tables to remove references to METRICS_RESOURCES_USAGE_UD.
- Modified /Oracle_Larger_DataTypes/04_pqCreateDrop_KPI_Tables_oracle_kpi_tables to remove references to METRICS_RESOURCES_USAGE_UD.
- Modified 06_pqCreateDrop_KPI_Plsql_oracle_data_xfer_script to remove references to METRICS_RESOURCES_USAGE_UD.
- Modified 03_pqCreateDrop_KPI_Tables_sqlserver_metrics_history_tables to add index "mru_hist_reskind_dsname".
- Modified 04_pqCreateDrop_KPI_Tables_sqlserver_kpi_tables to remove references to METRICS_RESOURCES_USAGE_UD.
- Modified 06_pqCreateDrop_KPI_Plsql_sqlserver_data_xfer_script to remove references to METRICS_RESOURCES_USAGE_UD.
- Modified pPurgeData to remove references to METRICS_RESOURCES_USAGE_UD.
- Removed /Formatting/METRICS_RESOURCES_USAGE_UD
- Removed /Physical/Abstraction/METRICS_RESOURCES_USAGE_UD
- Removed /Physical/KPI_oracle/METRICS_RESOURCES_USAGE_UD
- Removed /Metadata/KPI_oracle_11g/CIS_KPI/METRICS_RESOURCES_USAGE_UD
- Removed /Metadata/KPI_oracle_12c/CIS_KPI/METRICS_RESOURCES_USAGE_UD
- Removed /Physical/KPI_sqlserver/METRICS_RESOURCES_USAGE_UD
- Removed /Metadata/KPI_sqlserver_2012/CIS_KPI/dbo/METRICS_RESOURCES_USAGE_UD
- Removed /Metadata/KPI_sqlserver_2014/CIS_KPI/dbo/METRICS_RESOURCES_USAGE_UD

Release 2018Q400 [Dec 1 2018]

- Release Objective: Correlation of published view with datasource lineage.
 - Added processing for SQL request datasource lineage thus allowing reporting on what physical data source resources are related to which published resources.
 - Procedure tuning for Cache_METRICS_SQL_REQUEST_EXEC to block execution when Cache_METRICS_SQL_RESOURCE_LINEAGE is executing so as to not overwhelm the DV node. The opposite is true for Cache_METRICS_SQL_RESOURCE_LINEAGE which checks to see if Cache_METRICS_SQL_REQUEST_EXEC is running before executing.
- Lineage Changes:
- Altered tables METRICS_SQL_RESOURCE and METRICS_SQL_COLUMNS to add START_TIME TIMESTAMP column.
- Altered table METRICS_SQL_RESOURCE to add MESSAGE_TYPE, MESSAGE, PROCESS_NODE_HOST and PROCESSED_NODE_PORT.
- Altered table METRICS_SQL_CONTROL and METRICS_SQL_CONTROL_LOG to add column CONTROL_NAME.
- Added new table METRICS_SQL_RESOURCE_LINEAGE - datasource lineage for each resource found in METRICS_SQL_RESOURCE.
- Added table DDL to 04_pqCreateDrop_KPI_Tables_oracle_kpi_tables and 04_pqCreateDrop_KPI_Tables_sqlserver_kpi_tables.
- Added views and published KPImetrics.requests.vMetricsSqlResourceLineage and vMetricsSqlResourceLineageCountReport.

- Added trigger `kpimetricsTrig_23_Cache_METRICS_SQL_RESOURCE_LINEAGE`.
- Added cluster safe procedure `Cache_METRICS_SQL_RESOURCE_LINEAGE` which allows for multi-host processing. This means that node a can process node b's lineage as long as node a has no more rows to process.
- Added variables to `commonValues`: `queryTableLookupDefault`, `queryTableLookupInterval`, `getResourceLineagePathExclusionList`, `numRowsToProcessBeforeExitingLineage`
- Modified `defaultTriggersToEnable` to include new trigger.
- Added index `METRICS_SQL_RES_LINEAGE_IDX1` and `_IDX2`.
- Added index `METRICS_SQL_COLUMNS_IDX3`.
- Modified `pPurgeData` to delete from `METRICS_SQL_RESOURCE_LINEAGE`.
- General Fixes:
- Fixed a bug in the following view to eliminate `ALL_COLUMNS` and `ALL_PARAMETERS` from the result set: `/Business/Business/resourceList/vAllPublishedResources`.
- Modified `pEndWorkflow` to better hand error messages with single quotes.
- Fixed `updateTriggers` which execute the cache when turning off triggers.

Release 2018Q302 [Oct 12 2018]

- Fixed a collection table row volume problem by adding an independent delete trigger [`kpimetricsTrig_33_DeleteCollection`] for deleting unwanted rows from the collection tables: `metrics_requests` and `metrics_resources_usage`. By deleting unwanted rows more often it allows the data transfer script to work more efficiently and quickly.
- Added indexes `metrics_resources_usage_idx4`, `metrics_resources_usage_idx5`, `metrics_requests_idx1`, `metrics_requests_idx2` to improve to different styles of delete.
- Modified index `METRICS_SQL_REQUEST_IDX4` to try and improve lookup speed.
- Added indexes `METRICS_SQL_RESOURCE_IDX2`, `METRICS_SQL_COLUMNS_IDX2` to try and improve joins.

Release 2018Q301 [Oct 1 2018]

- Modified `parseSqlScriptComplex` and `parseSqlScriptTemplate` to change template select from `SELECT DISTINCT` to `SELECT TOP 1` to make it more efficient.

Release 2018Q3 [Sep 2018]

- Modified to allow datasource names in the format of `KPI_oracle_11g`, `KPI_oracle_12c`, `KPI_sqlserver_2012` or `KPI_sqlserver_2014` in order to allow for adding a datasource.
- Modified `sqlParserV1_2` to allow for parallel processing of SQL when a cluster is present. A node can process rows other than its own rows when it has no work to do.
- Modified tables:
 - `METRICS_SQL_REQUEST` - added `PROCESSED_NODE_HOST` and `PROCESSED_NODE_PORT`
 - `metrics_resources_usage_hist` - added `resourceorigin`
 - `METRICS_SQL_RESOURCE` - added `RESOURCE_ORIGIN`
 - `METRICS_ALL_RESOURCES` - added `RESOURCE_ORIGIN`
 - `METRICS_RESOURCES_USAGE_UD` - added `RESOURCE_ORIGIN`
- Added tables:

- metrics_sessions_stg - staging table for metrics_sessions_hist
- metrics_requests_stg - staging table for metrics_requests_hist
- metrics_resources_usage_stg - staging table for metrics_resources_usage_hist
- METRICS_SQL_CONTROL - used to control parallel processing for sqlParserV1_2
- METRICS_SQL_CONTROL_LOG - logs the activity if turned if debugging is turned on.
- Modified sqlParserV1_2:
 - Allow for parallel processing of SQL when a cluster is present. A node can process rows other than its own rows when it has no work to do.
 - Added ability to process web service calls for operations that are TABLES or PROCEDURES.
- Modified /shared/ASAssets/KPImetrics/Business/Logical/metrics *_hist views to use where starttime > TO_TIMESTAMP('1900-01-01 00:00:00.000') to force the use of parallel queries.
- Consolidated P_METRICS_ALL_TABLES_exec from KPI_oracle/KPI_sqlserver to /ClusterSafeCache/pMETRICS_ALL_TABLES_exec.
- Consolidated P_PARTITION_MANAGER_exec from KPI_oracle/KPI_sqlserver to /ClusterSafeCache/pPARTITION_MANAGER_exec.
- Consolidated the 3 DBMS triggers each from KPI_oracle/KPI_sqlserver to a single set under ClusterSafeTriggers.
- Removed GetSequenceValueProc from KPI_oracle/KPI_sqlserver and consolidated the functionality into /Abstraction/GetSequenceValueProc.
- Moved the "update METRICS_RESOURCES_USAGE_UD where RESOURCE_ID IS NULL" logic from Cache_ALL_RESOURCES to the data transfer script.
- Removed the "update metrics_resources_usage_hist where RESOURCE_ID IS NULL" logic from Cache_ALL_RESOURCES altogether to avoid updates on the history table.
- Consolidated 05_pqCreateDrop_KPI_Tables_[oracle|sqlserver]_kpi_application_tables into 04_pqCreateDrop_KPI_Tables_[oracle|sqlserver]_kpi_tables
- Renamed 06_pqCreateDrop_KPI_Tables_[oracle|sqlserver]_kpi_sequence to 05_pqCreateDrop_KPI_Tables_[oracle|sqlserver]_kpi_sequence
- Renamed 07_pqCreateDrop_KPI_Plsq_[oracle|sqlserver]_data_xfer_script to 06_pqCreateDrop_KPI_Plsq_[oracle|sqlserver]_data_xfer_script
- Modified 06_pqCreateDrop_KPI_Plsq_[oracle|sqlserver]_data_xfer_script: [inserts are faster than deletes]
 - Modified the Oracle and SQL Server data transfer scripts to use staging tables for insert/update/deletes in order to only perform inserts into history tables.
 - Eliminated the delete job filter SQL on metrics_resources_usage by including it in the insert stage where clause.
 - Eliminated the delete not exist SQL on metrics_requests by including a WHERE EXISTS in the insert stage where clause.
- Renamed 08_pqInsert_KPI_Tables_METRICS_JOB_tables to 07_pqInsert_KPI_Tables_METRICS_JOB_tables
- Renamed 09_pqInsert_KPI_Tables_METRICS_EVENT_REGISTRATION to 08_pqInsert_KPI_Tables_METRICS_EVENT_REGISTRATION
- Modified commonValues:
 - Removed dataSourceSchemaPath_oracle, dataSourceSchemaPath_sqlserver.

- Removed cisCaseSensitivity as it is no longer needed.
- Added dataSourceCatalog, dataSourceSchema, historyTableCompression.
- Added dataSourceCollation_sqlserver which gets applied at the time of creating the SQL Server tables to allow for case sensitive searches and joins on columns.

8 Appendix A – Partitioning Schemes

This section provides information on the various database partition schemes.

Oracle Partition Scheme

This section describes how Oracle partitioning is utilized. Oracle by far provides the most elegant and easy to implement solution for partitioning. There are very few moving parts and the commands are all inclusive. The following description demonstrates the “Oracle Partition Management Sliding Window Scenario”. It shows the SQL statements that get executed for creation, adding and dropping of partitions.

Step 1. Configure commonValues

Setup of the \Configuration\commonValues is a key aspect for partitioning. The partitionNumber and partitionStartDate define how many partitions will be initially created and managed and when the first partition start date is configured for. In the example below there will be 3 partitions created starting with 20170501.

```
commonValues.partitionNumber=3
commonValues.partitionStartDate='2017-05-01'
```

Therefore, the history tables will be created as follows:

<u>Metrics History Table</u>	<u>Partition Name</u>	<u>Partition Rule</u>
metrics_requests_hist	MR201705	< 2017-06-01 00:00:00
	MR201706	< 2017-07-01 00:00:00
	MR201707	< 2017-08-01 00:00:00
metrics_resources_usage_hist	MRU201705	< 2017-06-01 00:00:00
	MRU201706	< 2017-07-01 00:00:00
	MRU201707	< 2017-08-01 00:00:00
metrics_sessions_hist	MS201705	< 2017-06-01 00:00:00
	MS201706	< 2017-07-01 00:00:00
	MS201707	< 2017-08-01 00:00:00

Step 2. Create initial history tables, partition strategy and indexes:

```
[03_pqCreateDrop_KPI_Tables_oracle_metrics_history_tables]
```

Create the history table for the respective tables. For Oracle, the create table statement contains the syntax for creating the partitions. Notice that the PARTITION BY RANGE is used on the “starttime” column. Since the partitionNumber=3 then 3 partitions are created for the initial partition starting with partitionStartDate=2017-05-01. Each partition has a unique name which describes what bucket of data it contains. The partition also contains a rule to compare the data to determine which bucket it goes in.

```
CREATE TABLE "CIS_KPI"."metrics_requests_hist" (<column_list>
```

LOB ("description") STORE AS (TABLESPACE "METRICS_DATA_HIST" ENABLE STORAGE IN ROW
CHUNK 8192 PCTVERSION 10 NOCACHE NOLOGGING)

LOB ("message") STORE AS (TABLESPACE "METRICS_DATA_HIST" ENABLE STORAGE IN ROW
CHUNK 8192 PCTVERSION 10 NOCACHE NOLOGGING)

NOCOMPRESS TABLESPACE "METRICS_DATA_HIST" RESULT_CACHE (MODE DEFAULT) PCTUSED
0 PCTFREE 10 INITRANS 1 MAXTRANS 255

STORAGE (BUFFER_POOL DEFAULT FLASH_CACHE DEFAULT CELL_FLASH_CACHE DEFAULT)

PARTITION BY RANGE ("starttime")

(

PARTITION MR201705 VALUES LESS THAN (TIMESTAMP' 2017-06-01 00:00:00')

LOGGING NOCOMPRESS TABLESPACE "METRICS_DATA_HIST"

LOB ("description") STORE AS (TABLESPACE "METRICS_DATA_HIST" ENABLE STORAGE IN ROW
CHUNK 8192 RETENTION NOCACHE LOGGING

STORAGE (INITIAL 8M NEXT 1M MINEXTENTS 1 MAXEXTENTS UNLIMITED PCTINCREASE 0
BUFFER_POOL DEFAULT FLASH_CACHE DEFAULT CELL_FLASH_CACHE DEFAULT))

LOB ("message") STORE AS (TABLESPACE "METRICS_DATA_HIST" ENABLE STORAGE IN ROW
CHUNK 8192 RETENTION NOCACHE LOGGING

STORAGE (INITIAL 8M NEXT 1M MINEXTENTS 1 MAXEXTENTS UNLIMITED PCTINCREASE 0
BUFFER_POOL DEFAULT FLASH_CACHE DEFAULT CELL_FLASH_CACHE DEFAULT))

PCTFREE 10 INITRANS 1 MAXTRANS 255

STORAGE (INITIAL 8M NEXT 1M MAXSIZE UNLIMITED MINEXTENTS 1 MAXEXTENTS
UNLIMITED BUFFER_POOL DEFAULT FLASH_CACHE DEFAULT CELL_FLASH_CACHE DEFAULT),

PARTITION MR201706 VALUES LESS THAN (TIMESTAMP' 2017-07-01 00:00:00')

LOGGING NOCOMPRESS TABLESPACE "METRICS_DATA_HIST"

LOB ("description") STORE AS (TABLESPACE "METRICS_DATA_HIST" ENABLE STORAGE IN ROW
CHUNK 8192 RETENTION NOCACHE LOGGING

STORAGE (INITIAL 8M NEXT 1M MINEXTENTS 1 MAXEXTENTS UNLIMITED PCTINCREASE 0
BUFFER_POOL DEFAULT FLASH_CACHE DEFAULT CELL_FLASH_CACHE DEFAULT))

LOB ("message") STORE AS (TABLESPACE "METRICS_DATA_HIST" ENABLE STORAGE IN ROW
CHUNK 8192 RETENTION NOCACHE LOGGING

STORAGE (INITIAL 8M NEXT 1M MINEXTENTS 1 MAXEXTENTS UNLIMITED PCTINCREASE 0
BUFFER_POOL DEFAULT FLASH_CACHE DEFAULT CELL_FLASH_CACHE DEFAULT))

PCTFREE 10 INITRANS 1 MAXTRANS 255

STORAGE (INITIAL 8M NEXT 1M MAXSIZE UNLIMITED MINEXTENTS 1 MAXEXTENTS
UNLIMITED BUFFER_POOL DEFAULT FLASH_CACHE DEFAULT CELL_FLASH_CACHE DEFAULT),

PARTITION MR201707 VALUES LESS THAN (TIMESTAMP' 2017-08-01 00:00:00')

LOGGING NOCOMPRESS TABLESPACE "METRICS_DATA_HIST"

```

LOB ("description") STORE AS (TABLESPACE "METRICS_DATA_HIST" ENABLE STORAGE IN ROW
CHUNK 8192 RETENTION NOCACHE LOGGING

STORAGE (INITIAL 8M NEXT 1M MINEXTENTS 1 MAXEXTENTS UNLIMITED PCTINCREASE 0
BUFFER_POOL DEFAULT FLASH_CACHE DEFAULT CELL_FLASH_CACHE DEFAULT))

LOB ("message") STORE AS (TABLESPACE "METRICS_DATA_HIST" ENABLE STORAGE IN ROW
CHUNK 8192 RETENTION NOCACHE LOGGING

STORAGE (INITIAL 8M NEXT 1M MINEXTENTS 1 MAXEXTENTS UNLIMITED PCTINCREASE 0
BUFFER_POOL DEFAULT FLASH_CACHE DEFAULT CELL_FLASH_CACHE DEFAULT))

PCTFREE 10 INITRANS 1 MAXTRANS 255

STORAGE (INITIAL 8M NEXT 1M MAXSIZE UNLIMITED MINEXTENTS 1 MAXEXTENTS
UNLIMITED BUFFER_POOL DEFAULT FLASH_CACHE DEFAULT CELL_FLASH_CACHE DEFAULT)
)
NOCACHE NOPARALLEL MONITORING;

```

Create the indexes for the history tables. Indexes are created/managed on the main table. Notice that the tablespace can be different for indexes than for the history tables.

```

CREATE INDEX "mr_hist_rid" ON "CIS_KPI"."metrics_requests_hist" ("requestid", "nodehost", "nodeport")
TABLESPACE "METRICS_DATA_IDX";

CREATE INDEX "mr_hist_rid_time" ON "CIS_KPI"."metrics_requests_hist" ("requestid", "starttime",
"nodehost", "nodeport") TABLESPACE "METRICS_DATA_IDX";

```

Repeat all creation operations shown above for “metrics_requests_usage_hist” and “metrics_sessions_hist”.

This shows the partition distribution for the history tables.

table_name	partition_number	num_rows
"MR201705"	1	0
"MR201706"	2	0
"MR201707"	3	0
"MRU201705"	1	0
"MRU201706"	2	0
"MRU201707"	3	0
"MS201705"	1	0
"MS201706"	2	0
"MS201707"	3	0

Step 3. Add next month partition:

```
[03_pqCreateDrop_KPI_Tables_oracle_metrics_history_tables_ADD]
```

Notice how a new partition table is created simply with an ALTER TABLE statement. No other operation is required. Indexes are automatically created.

```

ALTER TABLE "CIS_KPI"."metrics_requests_hist" ADD

PARTITION "MR201708" VALUES LESS THAN (TIMESTAMP '2017-09-01 00:00:00')

```

```
LOGGING NOCOMPRESS TABLESPACE "METRICS_DATA_COLL"
```

```
LOB ("description") STORE AS (TABLESPACE "METRICS_DATA_COLL" ENABLE STORAGE IN ROW
CHUNK 8192 RETENTION NOCACHE LOGGING
```

```
STORAGE (INITIAL 8M NEXT 1M MINEXTENTS 1 MAXEXTENTS UNLIMITED PCTINCREASE 0
BUFFER_POOL DEFAULT FLASH_CACHE DEFAULT CELL_FLASH_CACHE DEFAULT))
```

```
LOB ("message") STORE AS (TABLESPACE "METRICS_DATA_COLL" ENABLE STORAGE IN ROW
CHUNK 8192 RETENTION NOCACHE LOGGING
```

```
STORAGE (INITIAL 8M NEXT 1M MINEXTENTS 1 MAXEXTENTS UNLIMITED PCTINCREASE 0
BUFFER_POOL DEFAULT FLASH_CACHE DEFAULT CELL_FLASH_CACHE DEFAULT))
```

```
PCTFREE 10 INITRANS 1 MAXTRANS 255
```

```
STORAGE (INITIAL 8M NEXT 1M MAXSIZE UNLIMITED MINEXTENTS 1 MAXEXTENTS
UNLIMITED BUFFER_POOL DEFAULT FLASH_CACHE DEFAULT CELL_FLASH_CACHE DEFAULT);
```

Repeat all creation operations shown above for "metrics_requests_usage_hist" and "metrics_sessions_hist".

This shows the partition distribution for the history tables. Notice how 201708 partitions were added.

table_name	partition_number	num_rows
"MR201705"	1	0
"MR201706"	2	0
"MR201707"	3	0
"MR201708"	4	0
"MRU201705"	1	0
"MRU201706"	2	0
"MRU201707"	3	0
"MRU201708"	4	0
"MS201705"	1	0
"MS201706"	2	0
"MS201707"	3	0
"MS201708"	4	0

Step 4. Insert test rows:

```
[/shared/ASAssets/KPImetrics/Physical/Metadata/DML/test_insert_metrics_requests_hist]
```

Notice how there is 1 row in each metrics_requests_hist partitions.

table_name	partition_number	num_rows
"MR201705"	1	1
"MR201706"	2	1
"MR201707"	3	1
"MR201708"	4	1
"MRU201705"	1	0
"MRU201706"	2	0
"MRU201707"	3	0
"MRU201708"	4	0
"MS201705"	1	0
"MS201706"	2	0
"MS201707"	3	0
"MS201708"	4	0

Step 5. Drop the oldest partition.

Test the Drop Scenario by forcing the number of partitions to be 1 less than before with the date incremented by 1 month

```
commonValues.partitionNumber=2
```

```
commonValues.partitionStartDate='2017-06-01'
```

Drop the oldest partition: [03_pqCreateDrop_KPI_Tables_oracle_metrics_history_tables_DROP]

The strategy for dropping a partition is called the sliding window. It is much more efficient than deleting rows. KPI metrics simply drops the partition. It is very simple.

```
ALTER TABLE "CIS_KPI"."metrics_requests_hist" DROP PARTITION "MR201705";
```

Repeat the same steps for "metrics_resources_usage_hist" and "metrics_sessions_hist":

As shown below, the number of partitions has been reduced from 4 to 3. The Jun 1 2017 boundary representing May 2017 data has been dropped. The oldest partition was dropped. The remaining data shown below is still properly partitioned as expected.

cluster	nodehost	nodeport	requestid	parentid	sessionid	requesttype	description	starttime	endtime	totaldurati...
[NULL]	Isdnevcbe...	9500	2	1	1	JDBC	[NULL]	2017-06-01 07:00...	2017-06-...	0
[NULL]	Isdnevcbe...	9500	3	1	1	JDBC	[NULL]	2017-07-01 07:00...	2017-07-...	0
[NULL]	Isdnevcbe...	9500	4	1	1	JDBC	[NULL]	2017-08-01 07:00...	2017-08-...	0

table_name	partition_number	num_rows
"MR201706"	1	1
"MR201707"	2	1
"MR201708"	3	1
"MRU201706"	1	0
"MRU201707"	2	0
"MRU201708"	3	0
"MS201706"	1	0
"MS201707"	2	0
"MS201708"	3	0

Step 6. Repeat Drop oldest partition.

Now, repeat the drop exercise one more time

```
commonValues.partitionNumber=1  
commonValues.partitionStartDate='2017-07-01'
```

Drop the oldest partition: [03_pqCreateDrop_KPI_Tables_oracle_metrics_history_tables_DROP]

```
ALTER TABLE "CIS_KPI"."metrics_requests_hist" DROP PARTITION "MR201706";
```

As shown below, the number of partitions has been reduced from 3 to 2. The Jul 1 2017 boundary representing June 2017 data has been dropped. The table representing the oldest partition was dropped. The remaining data shown below is still properly partitioned as expected.

cluster	nodehost	nodeport	requestid	parentid	sessionid	requesttype	description	starttime	endtime	totaldurati...
[NULL]	Isdnecvbe...	9500	3	1	1	JDBC	[NULL]	2017-07-01 07:00:...	2017-07-...	0
[NULL]	Isdnecvbe...	9500	4	1	1	JDBC	[NULL]	2017-08-01 07:00:...	2017-08-...	0

table_name	partition_number	num_rows
"MR201707"	1	1
"MR201708"	2	1
"MRU201707"	1	0
"MRU201708"	2	0
"MS201707"	1	0
"MS201708"	2	0

SQL Server Partition Scheme

This section describes how SQL Server partitioning is utilized. The scenario that is explained here is referred to as the “SQL Server Partition Management Sliding Window Scenario”. It shows the SQL statements that get executed for creation, adding and dropping of partitions.

Step 1. Configure commonValues

Setup of the \Configuration\commonValues is a key aspect for partitioning. The partitionNumber and partitionStartDate define how many partitions will be initially created and managed and when the first partition start date is configured for. In the example below there will be 3 partitions created starting with 20170501.

```
commonValues.partitionNumber=3
commonValues.partitionStartDate='2017-05-01'
```

Therefore the history tables will be created as follows:

<u>Metrics History Table</u>	<u>Derived Partition Name</u>	<u>Partition Function Rule</u>
metrics_requests_hist	MR201705	< 2017-06-01 00:00:00
	MR201706	< 2017-07-01 00:00:00
	MR201707	< 2017-08-01 00:00:00
metrics_resources_usage_hist	MRU201705	< 2017-06-01 00:00:00
	MRU201706	< 2017-07-01 00:00:00
	MRU201707	< 2017-08-01 00:00:00
metrics_sessions_hist	MS201705	< 2017-06-01 00:00:00
	MS201706	< 2017-07-01 00:00:00
	MS201707	< 2017-08-01 00:00:00

Note that in SQL Server there is no such thing as a partition name. It is simply shown here as the “Derived Partition Name” to describe how the partitions of data are distributed. It is shown for reporting purposes in the procedure

```
03_pqCreateDrop_KPI_Tables_sqlserver_metrics_history_tables_ROW_DISTRIBUTION.
```

Step 2. Create initial history tables, partition strategy and indexes:

```
[03_pqCreateDrop_KPI_Tables_sqlserver_metrics_history_tables]
```

Create the partition function and partition scheme. This partition strategy uses “RANGE RIGHT” as it makes it easier to define a monthly bucket based on the 1st of the month at midnight. All data less than that value is placed in the bucket for the current month. Also notice that the partitioning scheme applies the filegroup. This provides the flexibility of placing each month partition in its own filegroup and potentially disk spindle. However, for the sake of ease of implementation, KPImetrics only implements a single file group for all partitions. These days, the underlying disk architecture is usually hidden anyway with NFS mounted drives.

```
CREATE PARTITION FUNCTION "mr_hist_partition_function" (DATETIME2(3)) AS RANGE RIGHT FOR
VALUES ('2017-06-01 00:00:00', '2017-07-01 00:00:00', '2017-08-01 00:00:00');
```

```
CREATE PARTITION SCHEME "mr_hist_partition_scheme" AS PARTITION "mr_hist_partition_function" ALL
TO ((METRICS_DATA_HIST));
```

Create the partitioned history table, partitioned archive table and partitioned indexes for the respective tables. The archive table must be created exactly like the history table in order for the SWITCH to take place. This includes the partition scheme and indexes. Notice how the tables are created based on the partition scheme and not the filegroup. The filegroup is actually assigned to the partition scheme as shown previously.

```
CREATE TABLE "dbo"."metrics_requests_hist_arch" (<column_list>) ON
"mr_hist_partition_scheme"("starttime");
```

```
CREATE TABLE "dbo"."metrics_requests_hist" (<column_list>) ON "mr_hist_partition_scheme"("starttime");
```

Create partitioned indexes on history table. The indexes are created on the same partitioning scheme as the tables.

```
IF NOT EXISTS(SELECT * FROM sys.indexes WHERE object_id = object_id(N'metrics_requests_hist') AND
NAME ='mr_hist_rid') CREATE NONCLUSTERED INDEX "mr_hist_rid" ON "dbo"."metrics_requests_hist"
("requestid", "nodehost", "nodeport") ON "mr_hist_partition_scheme"("starttime");
```

```
IF NOT EXISTS(SELECT * FROM sys.indexes WHERE object_id = object_id(N'metrics_requests_hist') AND
NAME ='mr_hist_rid_time') CREATE NONCLUSTERED INDEX "mr_hist_rid_time" ON
"dbo"."metrics_requests_hist" ("requestid", "starttime", "nodehost", "nodeport") ON
"mr_hist_partition_scheme"("starttime");
```

Create partitioned indexes on archive table. The archive table must look exactly like the history table to perform SWITCH.

```
IF NOT EXISTS(SELECT * FROM sys.indexes WHERE object_id = object_id(N'metrics_requests_hist_arch')
AND NAME ='mr_hist_rid_arch') CREATE NONCLUSTERED INDEX "mr_hist_rid_arch" ON
"dbo"."metrics_requests_hist_arch" ("requestid", "nodehost", "nodeport") ON
"mr_hist_partition_scheme"("starttime");
```

```
IF NOT EXISTS(SELECT * FROM sys.indexes WHERE object_id = object_id(N'metrics_requests_hist_arch')
AND NAME ='mr_hist_rid_time_arch') CREATE NONCLUSTERED INDEX "mr_hist_rid_time_arch" ON
"dbo"."metrics_requests_hist_arch" ("requestid", "starttime", "nodehost", "nodeport") ON
"mr_hist_partition_scheme"("starttime");
```

Repeat table and index creation for “metrics_requests_usage_hist” and “metrics_sessions_hist”.

This shows the partition distribution for the history tables

table_name	boundary_definition	partition_number	num_rows
metrics_requests_hist	Jun 1 2017 12:00AM	1	0
metrics_requests_hist	Jul 1 2017 12:00AM	2	0
metrics_requests_hist	Aug 1 2017 12:00AM	3	0
metrics_resources_usage_hist	Jun 1 2017 12:00AM	1	0
metrics_resources_usage_hist	Jul 1 2017 12:00AM	2	0
metrics_resources_usage_hist	Aug 1 2017 12:00AM	3	0
metrics_sessions_hist	Jun 1 2017 12:00AM	1	0
metrics_sessions_hist	Jul 1 2017 12:00AM	2	0
metrics_sessions_hist	Aug 1 2017 12:00AM	3	0

Step 3. Add next month partition:

```
[03_pqCreateDrop_KPI_Tables_sqlserver_metrics_history_tables_ADD]
```

Notice how the SPLIT RANGE is used to add the next month range. If the current month is 2017-07 [Jul] then the boundary is 2017-08. The next month is 2017-08 [Aug] and the boundary is 2017-09 [Sep]. The following statement alters the scheme with NEXT USED [filegroup] to force the new partition range to take affect.

```
ALTER PARTITION FUNCTION "mr_hist_partition_function"() SPLIT RANGE ('2017-09-01 00:00:00');
```

```
ALTER PARTITION SCHEME "mr_hist_partition_scheme" NEXT USED [METRICS_DATA_HIST];
```

```
ALTER PARTITION FUNCTION "mru_hist_partition_function"() SPLIT RANGE ('2017-09-01 00:00:00');
```

```
ALTER PARTITION SCHEME "mru_hist_partition_scheme" NEXT USED [METRICS_DATA_HIST];
```

```
ALTER PARTITION FUNCTION "ms_hist_partition_function"() SPLIT RANGE ('2017-09-01 00:00:00');
```

```
ALTER PARTITION SCHEME "ms_hist_partition_scheme" NEXT USED [METRICS_DATA_HIST];
```

Step 4. Insert test rows:

```
[/shared/ASAssets/KPImetrics/Physical/Metadata/DML/test_insert_metrics_requests_hist]
```

table_name	boundary_definition	partition_number	num_rows
metrics_requests_hist	Jun 1 2017 12:00AM	1	1
metrics_requests_hist	Jul 1 2017 12:00AM	2	1
metrics_requests_hist	Aug 1 2017 12:00AM	3	1
metrics_requests_hist	Sep 1 2017 12:00AM	4	1
metrics_resources_usage_hist	Jun 1 2017 12:00AM	1	0
metrics_resources_usage_hist	Jul 1 2017 12:00AM	2	0
metrics_resources_usage_hist	Aug 1 2017 12:00AM	3	0
metrics_resources_usage_hist	Sep 1 2017 12:00AM	4	0
metrics_sessions_hist	Jun 1 2017 12:00AM	1	0
metrics_sessions_hist	Jul 1 2017 12:00AM	2	0
metrics_sessions_hist	Aug 1 2017 12:00AM	3	0
metrics_sessions_hist	Sep 1 2017 12:00AM	4	0

Notice that there is 1 row in each partition

cluster	nodehost	nodeport	requestid	parentid	sessionid	requesttype	description	starttime	endtime	totaldurati...
[NULL]	Isdnecvbe...	9500	1	1	1	JDBC	[NULL]	2017-05-22 07:57:21.407	2017-05-...	0
[NULL]	Isdnecvbe...	9500	1	1	1	JDBC	[NULL]	2017-06-22 07:57:21.407	2017-06-...	0
[NULL]	Isdnecvbe...	9500	2	1	1	JDBC	[NULL]	2017-07-22 07:57:21.407	2017-07-...	0
[NULL]	Isdnecvbe...	9500	3	1	1	JDBC	[NULL]	2017-08-22 07:57:21.407	2017-08-...	0

Step 5. Drop the oldest partition.

Test the Drop Scenario by forcing the number of partitions to be 1 less than before with the date incremented by 1 month

```
commonValues.partitionNumber=2
```

```
commonValues.partitionStartDate='2017-06-01'
```

Drop the oldest partition:

```
[03_pqCreateDrop_KPI_Tables_sqlserver_metrics_history_tables_DROP]
```

The strategy for dropping a partition is called the sliding window. It is much more efficient than deleting rows. The idea is to SWITCH the oldest partition with an empty archive table which is configured with the same structure and indexes. This is accomplished with the ALTER TABLE SWITCH PARTITION shown below.

```
ALTER TABLE "dbo"."metrics_requests_hist" SWITCH PARTITION 1 TO "dbo"."metrics_requests_hist_arch" PARTITION 1;
```

The next part of the strategy is to perform the ALTER FUNCTION MERGE RANGE to merge the empty partition with the next oldest partition. This will be very fast since the oldest partition designated by the boundary range "2017-06-01 00:00:00" is empty.

```
ALTER PARTITION FUNCTION mr_hist_partition_function() MERGE RANGE ('2017-06-01 00:00:00');
```

Finally, the archive table which now contains the rows from the oldest partition of the history table is truncated to make it empty. Truncating is very fast compared to deleting rows. Theoretically, if it is required, the rows could be archived off to off-line storage or a Big Data solution like Hadoop.

```
TRUNCATE TABLE "dbo"."metrics_requests_hist_arch";
```

Repeat the same steps for "metrics_resources_usage_hist" and "metrics_sessions_hist":

```
ALTER TABLE "dbo"."metrics_resources_usage_hist" SWITCH PARTITION 1 TO "dbo"."metrics_resources_usage_hist_arch" PARTITION 1;
```

```
ALTER PARTITION FUNCTION mru_hist_partition_function() MERGE RANGE ('2017-06-01 00:00:00');
```

```
TRUNCATE TABLE "dbo"."metrics_resources_usage_hist_arch";
```

```
ALTER TABLE "dbo"."metrics_sessions_hist" SWITCH PARTITION 1 TO "dbo"."metrics_sessions_hist_arch"
PARTITION 1;
```

```
ALTER PARTITION FUNCTION ms_hist_partition_function() MERGE RANGE ('2017-06-01 00:00:00');
```

```
TRUNCATE TABLE "dbo"."metrics_sessions_hist_arch";
```

As shown below, the number of partitions has been reduced from 4 to 3. The Jun 1 2017 boundary representing May 2017 data has been dropped. The data in the partition was switched with an empty partitioned archive table. The data in the archive table was truncated.

The remaining data shown below is still properly partitioned as expected.

cluster	nodehost	nodeport	requestid	parentid	sessionid	requesttype	description	starttime	endtime	totaldurati...
[NULL]	Isdnecybe...	9500	1	1	1	JDBC	[NULL]	2017-06-22 07:57:21.407	2017-06-...	0
[NULL]	Isdnecybe...	9500	2	1	1	JDBC	[NULL]	2017-07-22 07:57:21.407	2017-07-...	0
[NULL]	Isdnecybe...	9500	3	1	1	JDBC	[NULL]	2017-08-22 07:57:21.407	2017-08-...	0

table_name	boundary_definition	partition_number	num_rows
metrics_requests_hist	Jul 1 2017 12:00AM	1	1
metrics_requests_hist	Aug 1 2017 12:00AM	2	1
metrics_requests_hist	Sep 1 2017 12:00AM	3	1
metrics_resources_usage_hist	Jul 1 2017 12:00AM	1	0
metrics_resources_usage_hist	Aug 1 2017 12:00AM	2	0
metrics_resources_usage_hist	Sep 1 2017 12:00AM	3	0
metrics_sessions_hist	Jul 1 2017 12:00AM	1	0
metrics_sessions_hist	Aug 1 2017 12:00AM	2	0
metrics_sessions_hist	Sep 1 2017 12:00AM	3	0

Step 6. Repeat Drop oldest partition.

Now, repeat the drop exercise one more time

```
commonValues.partitionNumber=1
```

```
commonValues.partitionStartDate='2017-07-01'
```

Drop the oldest partition:

```
[03_pqCreateDrop_KPI_Tables_sqlserver_metrics_history_tables_DROP]
```

```
ALTER TABLE "dbo"."metrics_requests_hist" SWITCH PARTITION 1 TO "dbo"."metrics_requests_hist_arch"
PARTITION 1;
```

```
ALTER PARTITION FUNCTION mr_hist_partition_function() MERGE RANGE ('2017-07-01 00:00:00');
```

```
TRUNCATE TABLE "dbo"."metrics_requests_hist_arch";
```

```
ALTER TABLE "dbo"."metrics_resources_usage_hist" SWITCH PARTITION 1 TO
"dbo"."metrics_resources_usage_hist_arch" PARTITION 1;
```

```
ALTER PARTITION FUNCTION mru_hist_partition_function() MERGE RANGE ('2017-07-01 00:00:00');
```

```
TRUNCATE TABLE "dbo"."metrics_resources_usage_hist_arch";
```

```
ALTER TABLE "dbo"."metrics_sessions_hist" SWITCH PARTITION 1 TO "dbo"."metrics_sessions_hist_arch"  
PARTITION 1;
```

```
ALTER PARTITION FUNCTION ms_hist_partition_function() MERGE RANGE ('2017-07-01 00:00:00');
```

```
TRUNCATE TABLE "dbo"."metrics_sessions_hist_arch";
```

As shown below, the number of partitions has been reduced from 3 to 2. The Jul 1 2017 boundary representing June 2017 data has been dropped. The data in the partition was switched with an empty partitioned archive table. The data in the archive table was truncated.

The remaining data shown below is still properly partitioned as expected.

cluster	nodehost	nodeport	requestid	parentid	sessionid	requesttype	description	starttime	endtime	totaldurati...
[NULL]	Isdncvbe...	9500	2	1	1	JDBC	[NULL]	2017-07-22 07:57:21.407	2017-07-...	0
[NULL]	Isdncvbe...	9500	3	1	1	JDBC	[NULL]	2017-08-22 07:57:21.407	2017-08-...	0

table_name	boundary_definition	partition_number	num_rows
metrics_requests_hist	Aug 1 2017 12:00AM	1	1
metrics_requests_hist	Sep 1 2017 12:00AM	2	1
metrics_resources_usage_hist	Aug 1 2017 12:00AM	1	0
metrics_resources_usage_hist	Sep 1 2017 12:00AM	2	0
metrics_sessions_hist	Aug 1 2017 12:00AM	1	0
metrics_sessions_hist	Sep 1 2017 12:00AM	2	0