



ITM 6.x Problem Determination and Troubleshooting

Summarization and Pruning Agent - Overview

Tivoli software

ON DEMAND BUSINESS™

Overview of Summarization and Pruning Agent

- Summarization and Pruning Agent (S&P) is responsible for summarizing and pruning historical data that is contained within the Tivoli Data Warehouse (TDW)
- S&P runs on:
 - ▶ All platforms
- TDW supports the following Relational Databases:
 - ▶ DB2 V8.2 FP10 and higher
 - ▶ DB2 V9.1
 - ▶ Oracle 9.2
 - ▶ Oracle 10g Release 1
 - ▶ Oracle 10g Release 2
 - ▶ MS SQL Server 2000
 - ▶ MS SQL Server 2003

Note: For the exact list of platforms see product documentation and/or the certification matrix

http://www-306.ibm.com/software/sysmgmt/products/support/Tivoli_Supported_Platforms.html



Overview of Summarization and Pruning Agent - continued

- TDW database can be local or remote – running S&P agent on same system where TDW database resides is strongly recommended
- There can only be one S&P per TDW
- S&P does NOT communicate with other agents (i.e. OS agents, database agents, etc) or the Warehouse Proxy agent
- S&P does communicate with TEPS and the TDW
- S&P needs to be scheduled to run, it can be configured to run at a fixed or flexible time.



Overview of Summarization and Pruning Agent - continued

- All summarization and pruning logic is implemented in Java
- All summarizations are performed before pruning is evaluated and performed.
- Increasing the KSY_MAX_WORKER_THREADS configuration variable found in KSYENV (Windows) or sy.ini (Non-Windows) can drastically improve performance if the extra horsepower is available on the system
- Tables created and maintained in TDW by the S&P agent:
 - ▶ WAREHOUSEAGGREGLOG
Log information regarding work performed by S&P agent
 - ▶ WAREHOUSEID
Contains mappings between long names and short names which have been substituted by WPA or S&P agent for table names or column names. Some RDBMS have limits on the table names and column names. When a limit is reached the WPA and S&P agent will use the short version of the name which is specified in the ODI file – TABLE and COLUMN respectively
 - ▶ WAREHOUSEMARKER
Contains information that the S&P agent utilizes to know where it has left off in the summarization calculations for each summarization for each agent/node

Overview of Summarization and Pruning Agent - continued

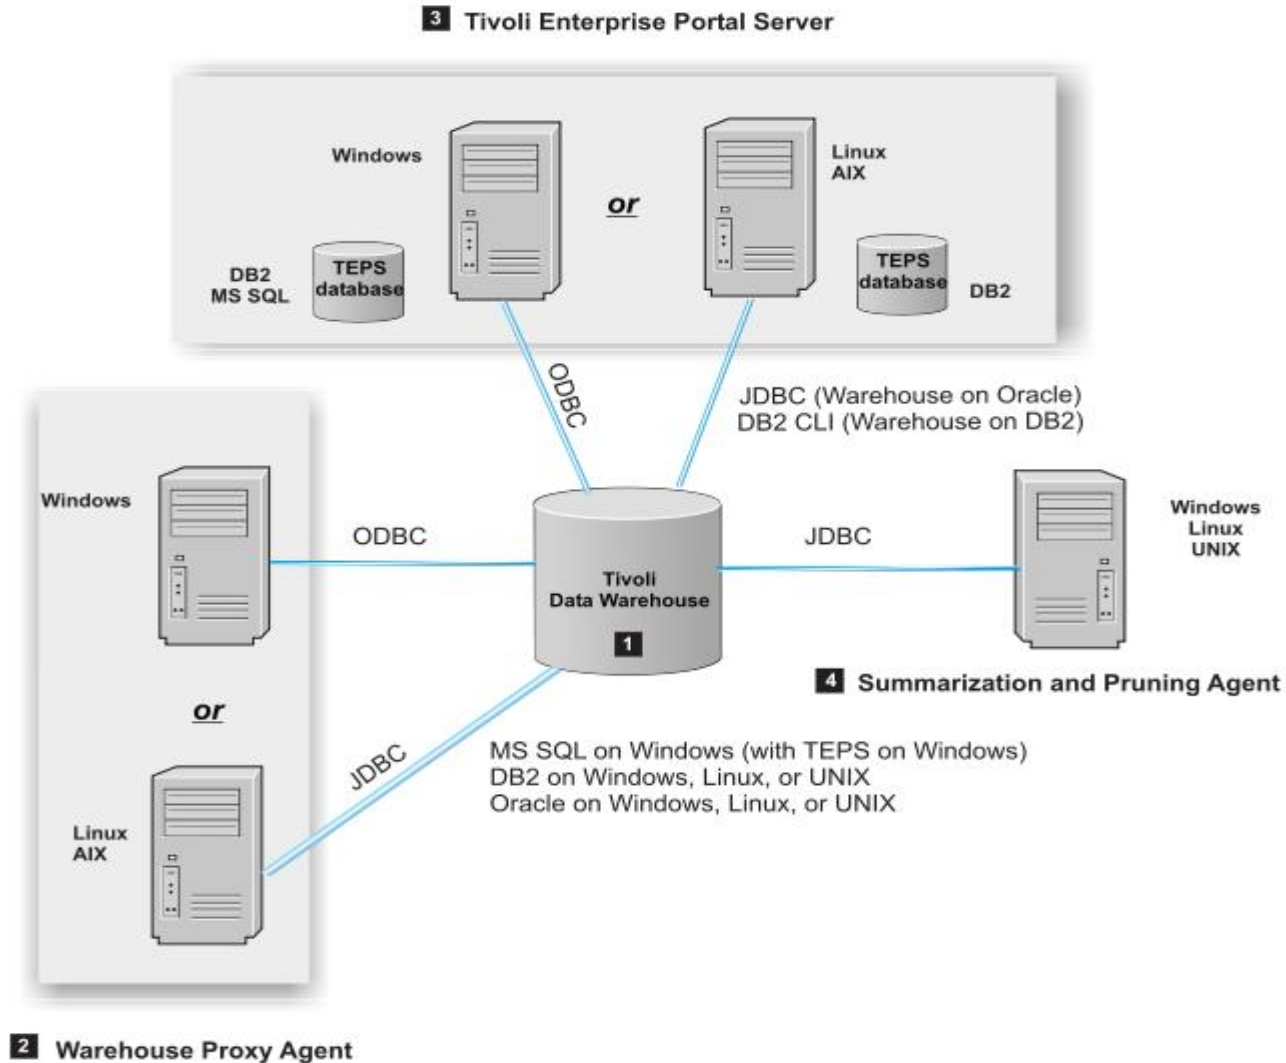
Beginning with V6.2.2 Fix Pack 2, historical data collection and summarization and pruning settings are stored in a table WAREHOUSESUMPRUNE in the Tivoli Data Warehouse database, instead of in the Tivoli Enterprise Portal database.

Tivoli Enterprise Portal Server must be able to connect to that table. Otherwise, users (including possibly your database administrator) will be unable to even view the Summarization and Pruning parameters

- From ITM 6.2.3, Statements from the Summarization and Pruning Agent into the WAREHOUSEAGGREGLOG table are now disabled by default.
- Tivoli Data Warehouse log tables can grow very large and require regular pruning.
- The self-monitoring workspaces provided by the Summarization and Pruning Agent provide sufficient information to determine if agent is operating correctly
- To restore the old behavior, you must edit the Summarization and Pruning Agent configuration file and change the KSY_WHLOG_ENABLE variable.



Summarization and Pruning Agent Architecture



Summarization and Pruning Agent Architecture (cont.)

- Summarization and Pruning agent Java component utilizes TEP client interface to get metadata about attribute groups, attributes and the current summarization and prune configuration settings. The metadata is contained in the ODI file. The configuration settings are stored in the TEPS database.
- Summarization and Pruning agent communicates with the TDW via JDBC to create summarization tables, get raw data, summarizes it and then put it back into the TDW. Pruning is performed after summarization and is done with SQL DELETE statements.
- User chooses to view summarized data with their TEP client. Once the TEPS gathers all the data, the TEPS gives the data to the TEP to display.
- The TEPS communicates with the TDW via ODBC/JDBC to provide the requested data. If the user uses the time span dialog to indicate they want to see summarized data, then TEPS dynamically builds the SQL that it executes against the TDW. It constructs this SQL based on the ODI file (metadata). If the user uses a Custom SQL query to view data then TEPS will substitute the variables and execute the SQL against the TDW. The author of the Custom SQL query is responsible for making sure the SQL is syntactically and technically correct.



Summarization and Pruning Agent Architecture (6.2.2FP2)

- If SPA configured to run in autonomous mode, then Summarization and Pruning agent Java component will look for required application support files in a specified location instead of obtaining the information from the portal server. This means that the SPA can operate without connecting to the Tivoli Enterprise Portal Server.
- The Summarization and Pruning agent can operate without a connection to a Tivoli Enterprise Portal Server, but a portal server must be installed and application support for all agents that will be collecting historical data must be installed on the portal server for required application support files to be available.
- To enable the Summarization and Pruning agent to run without connecting to the Tivoli Enterprise Portal Server, set KSY_AUTONOMOUS=YES



ITM 6.x Problem Determination and Troubleshooting

Summarization and Pruning Agent - Troubleshooting techniques

Tivoli software

ON DEMAND BUSINESS™

Log and Trace files

- S&P RAS1 log

- ▶ **Name:**

- Windows and Non-Windows

- `<hostname>_sy_<nnnnnnnnnn>-<xx>.log` – contains trace of C code

- `<hostname>_sy_java_<nnnnnnnnnn>-<xx>.log` – contains trace of Java code

nnnnnnnnnn is the timestamp in hex form

xx is the trace file identifier, utilized for wrapping feature

- ▶ **Locations:**

- Windows

- `%CANDLE_HOME%\logs`

- Non-Windows

- `$CANDLEHOME/logs`

- ▶ Contains trace statements from execution of S&P agent



Operation Logs

- Name:
 - ▶ Windows
 - `<hostname>_SY.LG0` - most current log
 - `<hostname>_SY.LG1` - previous log
 - ▶ Non-Windows
 - `<hostname>:SY.LG0` - most current log
 - `<hostname>:SY.LG1` - previous log
- Locations :
 - ▶ Windows
 - `%CANDLE_HOME%\TMAITM6\logs`
 - ▶ Non-Windows
 - `$CANDLEHOME/logs`
- Contains an audit trail for each export written to the warehouse database

Configuration Files

- Windows
 - ▶ KSYENV which resides in %CANDLE_HOME%/tmait6
- UNIX/Linux
 - ▶ sy.ini which resides in %CANDLE_HOME/config
 - ▶ sy.config which resides in %CANDLE_HOME/config



Warehouse Tables

- WAREHOUSEAGGREGLOG

- ▶ This table allows you to keep track of the data it has summarized.
- ▶ Running this query :

```
SELECT * FROM WAREHOUSEAGRREGLOG where ROWINSERTED = 0;
```

Will tell you which agent (originnode) , which attribute group (Object) at what time (EXPORTTIME) got an error (ERROR MSG)



Trace options

- Possible trace levels:
 - ▶ Trace general errors
 - KBB_RAS1=ERROR
 - ▶ Trace agent startup - C Code only
 - KBB_RAS1=ERROR (UNIT:ksz ALL)
 - ▶ Minimum level trace for summarization
 - KBB_RAS1=ERROR (COMP:com.tivoli.twh.ksy ALL)
 - ▶ Medium level trace for summarization
 - KBB_RAS1=ERROR (UNIT:ksy1 ALL)
 - ▶ Connection level trace for summarization
 - KBB_RAS1=ERROR (UNIT:ksy2 ALL)
 - ▶ Statement level trace for summarization
 - KBB_RAS1=ERROR (UNIT:ksy3 ALL)
 - ▶ ResultSet level trace for summarization
 - KBB_RAS1=ERROR (UNIT:ksy4 ALL)
 - ▶ Column value level trace for summarization
 - KBB_RAS1=ERROR (UNIT:ksy5 ALL)
 - ▶ Trace every SQL statements being executed.
 - KBB_RAS1=ERROR (UNIT:ksysql ALL)
 - ▶ Trace every SQL statements and includes all the parameters used.
 - KBB_RAS1=ERROR (UNIT:ksysql1 ALL)



Trace options (cont.)

- ▶ To control the size and number of trace files to use, you need to add the following to the KSZ_JAVA_ARGS option in the KSYENV (Windows) or sy.ini (Non-Windows) file
 - `-Dibm.tdw.maxNumberDetailTraceFiles=n` (default=5)
 - `-Dibm.tdw.maxLinesForDetailTraceFile=n` (default=300000)
- ▶ **Note:** Trace file will wrap quickly, particularly when KBB_RAS1=ERROR (UNIT:ksy[1-5] ALL) is specified
- ▶ Search for keyword 'Exception' in the `*sy_java*.log` file. There are some errors that contain this keyword in the log file which are acceptable and are therefore ignored. Typically there is some verbiage in the log which states that the error is being ignored and thus okay.
- ▶ At the end of execution, a line like the following will be displayed. It should report that zero tables failed:

```
== 78241 t=main Tables attempted : (2) Tables succeeded : (2) Tables failed : (0)
== 78242 t=main Summarization and pruning agent successfully ended
== 2006-11-27 13.55.13.775 -0500 : Trace paused
```



SPA Troubleshooting

- Ensure that the process ksy610.exe is running
- There are 2 distinct parts to the SPA – Summarization and Pruning.
 - ▶ Did the SPA connect to the HUB TEMS.
 - ▶ Did the SPA connect to the correct Warehouse Database
 - ▶ Did the SPA connect to the TEPS
 - ▶ Has the SPA been running whilst a scheduled time has elapsed.
 - ▶ Is the table that is not getting summarised records created or records pruned correctly configured.
 - ▶ Do raw records exist for the table
- Setting appropriate traces
 - ▶ `KBB_RAS1=(UNIT:KSYn ALL)`
 - ▶ `KBB_RAS1=(UNIT:KSZ ALL)`



SPA Troubleshooting

Use TEP to see if either agent had any problems

Check the Summarization and Pruning workspace which contains the Error table – situation will fire by default on any summarization or pruning failure

Most errors on S&P are due to database problem or bad data from agents

Database errors must be fixed – for instance, insufficient log space, permissions, etc...

Invalid data problems should be avoided by this new version of the code

Bad data will be replaced with <ERROR> string or an old timestamp

Warehouse Proxy functionality can be verified by enabling the trace setting (UNIT:khdxdbex OUTPUT)

Check logs if not enough information in TEP

For Summarization and Pruning Agent:

<hostname>_sy_*.log and <hostname>_sy_java_*.log (start with the sy_java log first)



SPA Facts

- When the SPA schedule becomes active the SPA process can run for quiet a period of time.
- You should attempt to locate the SPA on the same box as the warehouse database.



ITM 6.x Problem Determination and Troubleshooting

SPA – Special/Unique Situation

Tivoli software

The IBM On Demand Business logo, featuring the word "ON" in a red square followed by the words "DEMAND BUSINESS" in white text on a blue background.

Not able to view summarized data in TEP and summarized tables do not exist in the TDW problems

- **Solution 1 (ITM6.1 only):** Make sure the S&P agent situation is seeded into Hub TEMS*.
 - ▶ Run: `tacmd viewSit -s UADVISOR_KSY_ENABLE`
 - ▶ If the situation does not exist, re-seed the Hub TEMS, recycle the Hub TEMS and then verify that the S&P agent situation now exists.
- **Solution 2 (ITM6.1 only):** Make sure the S&P agent operations log says that it has connected to the TEMS and that the **UADVISOR_KSY_ENABLE** situation has started. For example you should see the following in the
`<hostname>_SY.LG0 (Windows) <hostname>_SY.LG0 (Non-Windows)`
`1080930204800032KRAREG000 Connecting to CMS HUB_IBM-F8DEA58B00A`
`1061120225111000KRAIRA000 Starting UADVISOR_KSY_ENABLE <7340213,1048802> for`
`KSJ.ENABLE.IRA Constructor`



Not able to view summarized data in TEP and summarized tables do not exist in the TDW problems (*cont.*)

- **Solution 3:** Make sure the S&P agent has been configured to run and the scheduled time has occurred. The S&P agent will only run when it's scheduled in the configuration. If that scheduled time happens when the agent was not running, then it won't perform any work. You can check the `*sy_java*.log` file to see when the next scheduled time is and at what times the agent has evaluated whether it was time to run. For example you might find this in the log indicating the next scheduled run:

```
== 122 t=Thread-0 Scheduled time : Yr=2008 Mon=10 Day=1 Hr=18 Min=15 Sec=0
MSc=341
```

- **Solution 4:** Make sure the S&P agent can communicate with the TEPS. You will see something like this in the `*sy_java*.log` file when the S&P is unable to communicate with the TEPS (Prior to ITM 6.2.2FP2):

```
== 70 t=Thread-0 Connecting to the TEP Server at (localhost:1920)
== 71 t=Thread-57 EXCEPTION: connectivityTest request to the TEP Server
failed.
== 72 t=Thread-0 EXCEPTION: getHistoryConfig request to the TEP Server
failed.
== 73 t=Thread-0 EXCEPTION: Failed to init history
== 74 t=Thread-0 EXCEPTION: Failed to setup history config
== 75 t=Thread-0 EXCEPTION: Failed to get parameters
== 76 t=Thread-0 Summarization and pruning agent ended with errors.
```



Not able to view summarized data in TEP and summarized tables do not exist in the TDW problems (*cont.*)

- **Solution 5:** Make sure the S&P agent can communicate with the TDW. You will see something like this in the `*sy_java*.log` file when the S&P is unable to communicate with the TDW with the warehouse URL specified

```
== 6 t=main Attempting connection with URL :
(jdbc:db2://brazil:50001/WAREHOUS) user : (ITMUser)
== 7 t=main com.ibm.db2.jcc.b.DisconnectException:
[ibm][db2][jcc][t4][2043][11550] Exception java.net.ConnectException: Error
opening socket to server brazil/9.27.138.144 on port 50,001 with message:
Connection refused: connect.
```

- **Solution 6:** Make sure the S&P agent can communicate with the TDW. You will see something like this in the `*sy_java*.log` file when the S&P is unable to communicate with the TDW with the warehouse user id or warehouse user password specified

```
== 7 t=main com.ibm.db2.jcc.b.SqlException: [ibm][db2][jcc][t4][2013][11249]
Connection authorization failure occurred. Reason: User ID or Password
invalid.
```

- **Solution 7:** Make sure the S&P agent can communicate with the TDW. You will see something like this in the `*sy_java*.log` file when the S&P is unable to communicate with the TDW with the JDBC driver specified

```
== 4 t=main Connecting to the warehouse database at (jdbc:db2:Warehous)
using (<incorrect driver name here>)
```

Not able to view summarized data in TEP and summarized tables
do exist in the TDW

- **Solution:** Check connection between TEPS and TDW. May need to reconfigure the connection to use the proper ODBC/JDBC driver and/or user id and password. Reconfigure the TEPS to specify the correct configuration parameters.



Typical Questions to Ask

- Does raw or detailed data exist in the TDW?
 - Is summarization and pruning configured correctly in the TEP History Collection Configuration panel?
 - Has the Hub TEMS been seeded with the S&P support files?
 - Has the S&P agent been scheduled correctly and is started?
 - Has the scheduled time elapsed while the agent was in a running state?
 - Does the `*sy_java*.log` report any exceptions or table failures?
 - To recreate the problem, typically you need:
 - ▶ ODI file (install into TEPS base directory on Windows or data on Linux/UNIX)
 - ▶ CAT file (install into TEMS RKDSCATL)
 - ▶ HIS file (install into TEPS SQLLIB)
 - ▶ DDL of table (generate from database administration tool)
 - ▶ Data export in interchangeable format (Oracle exp format or DB2 IXF format, for example)
 - ▶ With all those pieces, create the table from the schema and import the data.
- Restart TEMS and TEPS, configure S&P from TEP and run it.