

Migration & Transitioning from TM1 10.2.2 to IBM Planning Analytics:

Guidance and Recommendations

V2.1


Andreas Kugelmeier

Executive Consultant, FOPM

IBM Data and AI Expert Labs

Kugelmeier@us.ibm.com

+1-215-384-7302



Migration & Transitioning to Planning Analytics: Strategy, Approach, & Immediate Benefits

Migration and Transitioning to Planning Analytics: Strategy & Approach

1. **1:1 Migration:** PA is fully compatible with TM1 10.2.2 / Fundamental TM1 Architecture and TM1 Database Objects have not changed. => Install PA, copy databases, and test...
2. **Conversion of Perspectives Web Sheets to PAX: Planning Analytics TM1 Web understands Perspectives sheets just like TM1 10.2.2 TM1 Web.** For Web Deployed workbooks, Perspectives to PAX conversions can be done on an as-needed/as-desired basis.
3. Start to leverage Cognos Command Center, Cognos Integration Server for Data-Integration and Process-Automation and Orchestration



Migration and Transitioning to Planning Analytics: Immediate Benefits

1. Better Performance and Scalability:

- Planning Analytics TM1 Web is significantly faster (rendering)
- Planning Analytics Database Server = TM1 Server V11
 - Is more scalable (significantly less locking)
 - Re-processes Feeders much faster (with multi-threaded feeder processing)
 - Starts up much faster (with optional multi-threaded cube load and multi-threaded feeder processing)

2. Improved End-User Experience:

- Self-Service Capabilities with [Planning Analytics Workspace \(PAW\), Cognos Analytics \(CA\), and Planning Analytics for Excel \(PAX\)](#)
- Support of Attribute-based Hierarchies as 'Virtual' Dimensions
- Support of Alternate Rollups as 'Virtual' Dimensions



Planning Analytics **Cloud** Upgrade & Transitioning Considerations

- Hierarchies
 - When to integrate
- Data- & Systems-Integration
- Authentication and Security
- Perspectives to PAX transitioning:
 - Conversion necessary? (if end-users do not use Perspectives, you do not need to convert Perspectives Web Sheets: TM1 Web fully supports Perspectives and PAX sheets)
 - Inventory - How many really?
 - Triage
 - Prioritize
 - Performance of existing Workbooks (time to rebuild/redeploy)
 - Installing and Updating
 - PAX and Perspectives Co-Existing?
 - [Perspectives to PAX migration: proven practices](#)
 - Differences between PAX and Perspectives: See [Planning Analytics Local Upgrade Guide \(on-premises\)](#) & [IBM TM1 to IBM Planning Analytics Transitioning Guide \(Cloud\)](#)
- Testing
 - Unit
 - Perspectives to PAX Worksheets
 - System Integration
 - User Acceptance
 - Parallel
- Workshops
 - PAW
 - PAX
 - Administration
- PAW & PAX Integration
 - Self Service
- COMMUNICATE, COMMUNICATE, COMMUNICATE!



Planning Analytics **Local** Upgrade & Transitioning Considerations

- Hierarchies
 - When to integrate
- Perspectives to PAX transitioning:
 - Initially, continue using Perspectives with PA, just like you did with TM1 10.2.2
 - Gradually transition from Perspectives to PAX
 - => PAX and Perspectives Co-Existing
 - Differences between PAX and Perspectives: See [Planning Analytics Local Upgrade Guide \(on-premises\)](#) & [IBM TM1 to IBM Planning Analytics Transitioning Guide \(Cloud\)](#)
- Testing
 - Unit
 - Perspectives to PAX Worksheets where needed
 - System Integration
 - User Acceptance
 - Parallel
- Workshops
 - PAW
 - PAX
 - Administration
- PAW & PAX Integration
 - Self Service
- COMMUNICATE, COMMUNICATE, COMMUNICATE!



TM1 to IBM Planning Analytics **Local** Upgrade Practices and Guidance

- [IBM TM1 to IBM Planning Analytics Upgrade Guide](#)
- [Perspectives to PAX Migration Proven Practices](#)
- [Cognos Analytics with Planning Analytics - Introductory Guidelines and Proven Practices](#)
- [Installing Planning Analytics Workspace \(PAW\): Overview](#)
- [PAW Architecture and advanced Administration](#)



TM1 On-Premises to IBM Planning Analytics **Cloud** Migration and Transitioning Practices and Guidance

- [IBM TM1 to IBM Planning Analytics Transitioning Guide](#)
- [Perspectives to PAX Migration Proven Practices](#)
- [Cognos Analytics with Planning Analytics - Introductory Guidelines and Proven Practices](#)



Planning Analytics Cloud FAQ:

https://www.ibm.com/support/knowledgecenter/en/SSD29G_2.0.0/com.ibm.swg.ba.cognos.tm1_cloud_mg.2.0.0.doc/c_tm1_cloud_faq.html

Planning Analytics Cloud FAQ

Can you tell me about the security and privacy of my data in the Cloud?

- TM1 on Cloud SSAE16 SOC 1 third party audit available upon request.
- SoftLayer data center SOC 2 audit report available upon request.
- ISO27001 compliant and certification currently in process.
- Sellers are now able to sell to FFIEC regulated customers.
- IBM Data Security and Privacy Principles available online: <http://www.ibm.com/cloud-computing/built-on-cloud/saas-security>
- IBM Software Products and Software-as-a-Service Privacy Statement: <http://www.ibm.com/software/info/product-privacy/>
- A Letter to Our Clients about Government Access to Data: <http://asmarterplanet.com/blog/2014/03/open-letter-data.html>
- IBM Business Conduct Guidelines: <http://www.ibm.com/investor/governance/business-conduct-guidelines.html>
- GDPR Ready



Planning Analytics Cloud FAQ

Is my Cloud data encrypted?

- **IBM Planning Analytics Cloud is using IBM Cloud Block Storage to implement Encryption at Rest at the storage device level.** Click [here](#) for details. Encryption at Rest is applied
 - for the shared folder (where the TM1 database directory and transaction log resides), and
 - for all backups are encrypted at rest.
- Furthermore, the content store for Planning Analytics Workspace and the content store for Cognos Analytics reporting both leverage database level encryption at rest.
- Because IBM Planning Analytics Cloud leverages Encryption at Rest at the file storage level, the implementation of Encryption at Rest on IBM Planning Analytics Cloud is transparent to PA Cloud processes and application. It follows that encryption/decryption is managed implicitly in the PA Cloud, i.e. customers do not need to (nor can) manage encryption/decryption themselves.



Planning Analytics Cloud FAQ

Can you tell me about High Availability, Backup, & Restore?

• High Availability:

- PA Workspace Application tier built for the cloud with multi-tenant infrastructure.
- TM1 data and transaction logs are stored on a RAID SAN.
- Monitoring for issues on all tiers as per previous slide
- In the event of an issue, tier will be restored to another instance by CloudOps

• Backup & Restore

- Backup and Restore is to/from SAN and application databases.
- Restore is requested through a support ticket
- Restore for TM1 database can replace existing data or in a separate tm1 database for comparison with current data
- Backup and Restore is also to/from another data center for disaster recovery.
- IBM will retain the last 7 daily backups and an additional four weekly backups.



Planning Analytics Cloud FAQ

How do I restore data from a file backup?

- To restore data from a backup, contact the Cloud Operations team. The Cloud Operations team performs daily backups of the data in your shared folder. They retain the last seven daily backups and an additional four weekly backups.
- **Important:** Ensure that your data resides in your shared folder. Any files that are stored in a location other than the shared folder will not be backed up.
- You can of course also make your own backups and use those to restore your data



Planning Analytics Cloud FAQ

Can I change the tm1s.cfg file?

- Yes, you can modify the tm1s.cfg file. However, you must use the version that is provided with the Cloud deployment. This Cloud version has settings specific to Planning Analytics, such as the TM1 Server port number. Do not use a customer's existing tm1s.cfg file if you are moving their TM1 model to the cloud. Some of the settings require a server restart before they become effective. Use the IBM Planning Analytics Administration to [start or stop the TM1 Server](#).
- Do not modify the following parameters in the tm1s.cfg file:
 - UseSSL
 - CAMUseSSL
 - SSLCertificateID
 - SSLCertAuthority
 - ClientCAMURI
 - ServerCAMURI
 - DatabaseDirectory
 - LoggingDirectory
 - ServerName
 - PortNumber
 - MessagePortNumber
 - HTTPPortNumber



Planning Analytics Cloud FAQ

Have there been any changes to TM1 configuration parameters since 10.2.2?

- Yes, parameters have been added to the TM1 Server configuration file (tm1s.cfg), the TM1 Web configuration file (tm1web_config.xml), and the TM1 client configuration file (tm1p.ini) since 10.2.2.
- Please refer to https://www.ibm.com/support/knowledgecenter/en/SSD29G_2.0.0/com.ibm.swg.ba.cognos.tm1_cloud_mg.2.0.0.doc/c_tm1_cloud_faq.html for changes that were made after the 10.2.2.3 version of the *TM1 Installation and Configuration Guide* was published.



Planning Analytics Cloud FAQ

What TM1 Servers are set up initially?


- A single blank TM1 Server, with the name TM1, is set up initially when the system is provisioned. Currently, this default TM1 Server cannot be renamed or removed.
- Other TM1 Servers (names chose by customer) are set up per request. See ['how do I set up new TM Servers'](#) on the following slide:



Planning Analytics Cloud FAQ

How do I set up new TM1 Servers?

- To set up additional TM1 Servers, send a request to the Cloud Operations team by creating a PMR. Provide the name of the TM1 Server that you require.
- After the TM1 Server is set up by the Cloud Operations team, they will send you an email with a new welcome kit. The welcome kit contains the details of the new server, such as the server name and the native TM1 login and password.
- The customer can now use the server as is or change the directory structure to organize the TM1 objects in a specific way. By default, all objects will reside in the base directory.
- The best practice is to create a folder structure to that shown in the following figure.



The following conditions must all be met:

- The tm1s.cfg file must exist in the base directory.
- Any folders that you create must be below the base directory.
- The location or the name of the base directory cannot be changed.

- In the above example, the name of the tm1 application/server in this case is smartco. Note that the default server was named TM1. Therefore, for this TM1 Server, the base directory is: \\data\s\prod\smartco



Planning Analytics Cloud FAQ

How can I migrate my existing TM1 Server content and settings to the Cloud?

- You can migrate your existing TM1 Server content and settings to the Cloud by following these steps:
- Create a PMR for the Cloud Operations team to set up a new server. See [How do I set up new TM1 Servers?](#).
- **Note:** Use the tm1s.cfg file that is provided by IBM. Do not use your existing tm1s.cfg file.
- Delete the default TM1 objects from the new TM1 Server.
- Copy the TM1 objects from your existing TM1 data directory to the data directory in the new TM1 Server.
- Add an existing Cognos user to the Cognos TM1 ADMIN group to act as the administrator. For more information, see [Defining a Cognos user to function as a Cognos TM1 administrator](#)
- **Tip:** All Planning Analytics user accounts are already configured by default with Cognos security mode 5.
- For more information, see [Loading and migrating data in to IBM Planning Analytics](#).



Planning Analytics Cloud FAQ

How do I access text files in Architect, Turbo-Integrator, or Performance Modeler?

Name of text box on Data Source tab	Description
Data Source Name	This is the path used by Architect and Performance Modeler when you run remote desktop. The full path is required. For example, type \\data\s\prod\tm1\data\Integration\OracleGLSubAccountDim.csv
Data Source Name on Server	This is the path used by TurboIntegrator processes on the TM1 Server. The path should be relative to the TM1 data directory, such as .\Integration\OracleGLSubAccountDim.csv, although absolute paths are possible using S:\... => If you enter \\data\s as the Data Source Name on Server value, the TI process will fail.

All files must reside on the Shared Drive because:

- The TM1 Server cannot see the Remote Desktop file system.
- Only the shared drive is backed up. Therefore, you risk losing your data if you store files on the remote desktop disk drive.

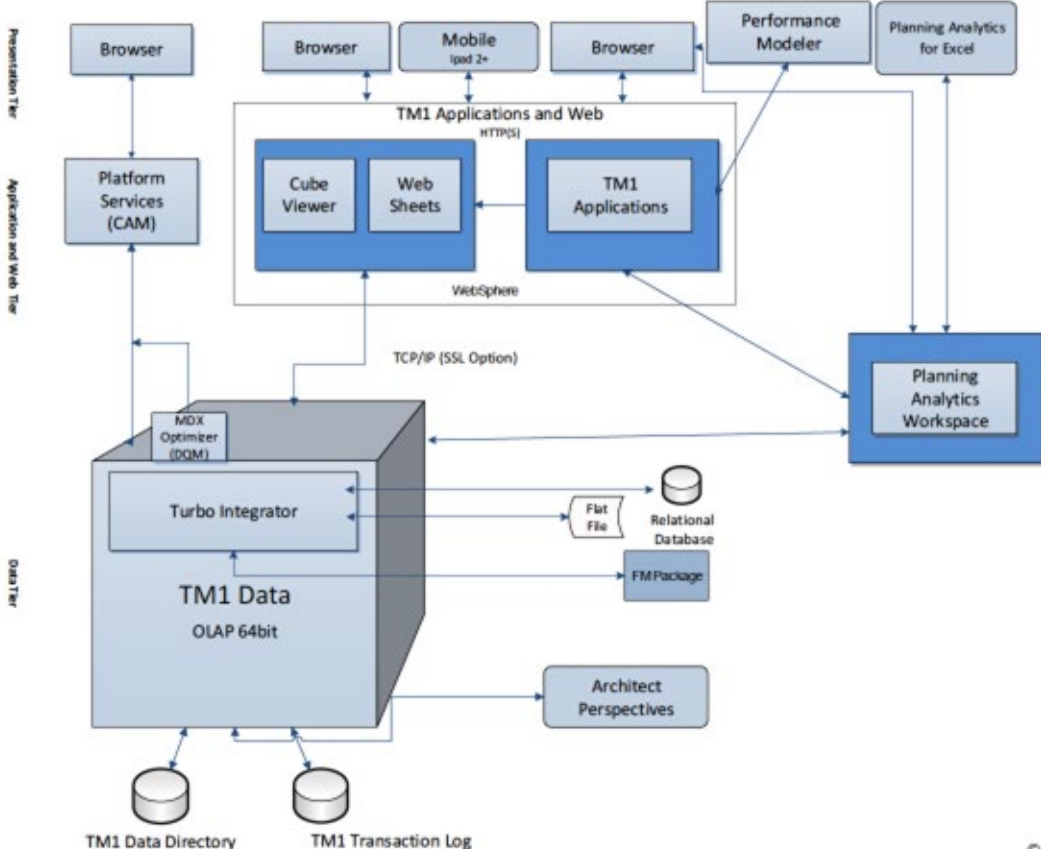
This video shows you how to access text files: <https://youtu.be/qoFW--HVm0w>

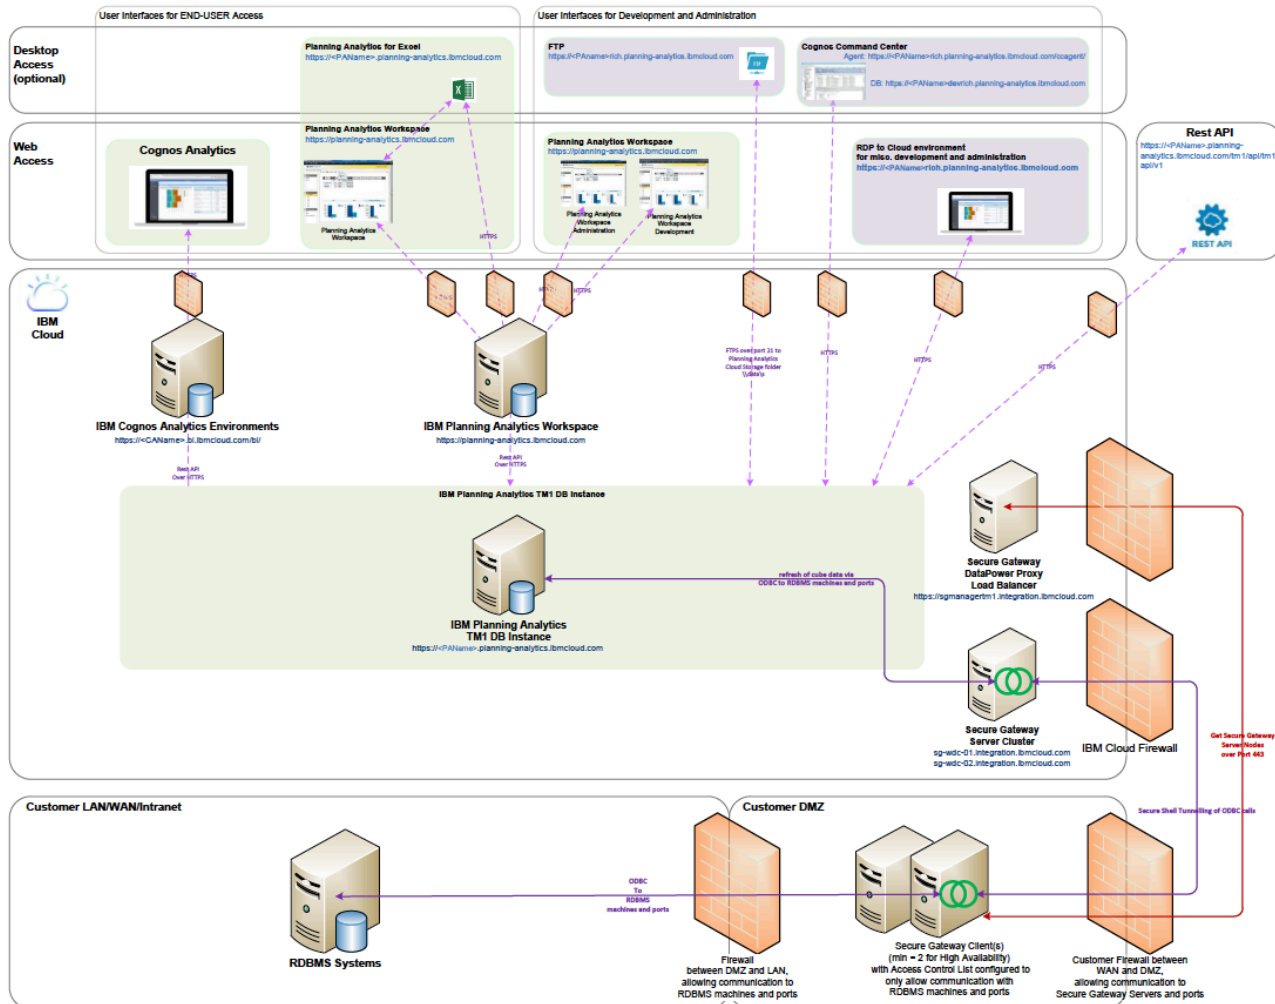


Setup & Configuration

- **TM1 Server (TM1s.cfg) and other TM1 Database Configuration Settings**
- **TM1 Web (WebSphere Liberty Configuration)**
- **Planning Analytics Workspace**

Planning Analytics Architecture





TM1 Database Server Configuration: TM1s.cfg

- Click [here](#) for documentation on Planning Analytics tm1s.cfg parameters
- Apart from parameters needed to start and operate the database (such as security authentication settings, port numbers, paths etc.), the TM1s.cfg should ONLY contain parameters that are
 - Proven to be needed or beneficial for the installed PA release version
 - Used to achieve a certain user experience / enable a feature
 - Non-default
- Any other parameter changes (such as changes prompted by IBM support to circumvent a bug or other issues) should ONLY be temporary, because:
 - Parameter changes are often of a type that will ‘throttle’ TM1
 - If temporary parameter changes are kept in the tm1s.cfg, the database will over time be configured to run in an ‘outdated’ operating mode.



TM1 Database Server Configuration: TM1s.cfg

- On MTQ:
 - What is it? -> MTQ stands for 'Multi-Threaded-Queries'. It allows TM1 to split a query into multiple chunks that then are processed in parallel CPU threads. Without MTQ, every TM1 query runs in a single thread, with MTQ, multiple threads can be leveraged, hence significantly improving performance.
 - **MTQ significantly improves performance and should typically be set to leverage all available cores.** Caveat: While MTQ=ALL ensures that hardware is leveraged at its maximum, it is only a recommended setting provided that the HW-sizing is adequately matched with the demands of the TM1 Database(s) that are running on the machine. Operating TM1 databases with MTQ=ALL yet on machines with insufficient HW capabilities may over-tax the HW and lead to performance degradation.
 - TM1 10.2 Default: MTQ OFF => Recommended setting in TM1 10.2.2 is MTQ=ALL
 - **PA / TM1 11 Default: MTQ=-1 / MTQ=ALL (all available cores).**
 - If your HW supports Hyperthreading, ensure you enable it on the Bios and/or OS. A Hyperthreaded Core provides about 80% of the performance of a 'real' CPU core. Therefore, enabling Hyperthreading can improve performance by about 80%.



TM1 Database Server Configuration: TM1s.cfg

- On MTQ (continued):
 - The MTQ value (MTQ=X, like MTQ=4) is the total thread pool for all users that the TM1 Server makes available for Multi-Threaded-Queries.
 - ⇒ Example: If you set MTQ=4 on an 8-CPU machine, only 4 threads will be used in Total for MTQ . If two users run a query that leverages 4 MTQ cores, one user will get 1 non-MTQ thread, and the other user will get 1 non-MTQ Thread plus 4 MTQ Worker Threads.
 - ⇒ MTQ is not the MTQ pool per user, but for the TM1 Database server in total.
 - An MTQ pool thread becomes 'available' to other users once it has finished.
 - ⇒ Example: If a query initially gets assigned 4 MTQ worker threads, this does not mean it will keep all those threads. One thread may finish and the query will continue with 3 worker threads. Or 2 threads will finish, but the TM1 MTQ engine will determine that the next 'part' of the query is best served with an additional 8 threads, resulting in 10 threads being used...
 - ⇒ As a particular MTQ worker threads is finished (right after 'commit' in TM1Top), the thread becomes available again, and it may be used by the same query (for a new worker thread), or by a different user query.



TM1 Database Server Configuration: TM1s.cfg

On MTQ (continued):

- **MTQ is NOT dependent/reliant on the # of CPUs on the machine!**
 - ⇒ **MTQ will leverage multi-core processing capabilities by splitting and resolving queries via parallel worker *threads***
 - ⇒ If you set MTQ to ALL or leave it at the default (ALL), TM1 will set an MTQ thread pool equal to the number of cores on the machine.
 - ⇒ You can set MTQ to a value higher than the number of cores on the machine. You will see more worker threads than CPU cores. This is acceptable, even though on most hardware, it is typically not going to result in performance gains, yet at lower concurrency typically will also not lead to performance degradation.
 - The CPU time and resources given to a MTQ worker thread are entirely handled by the Operating System.
 - An MTQ worker thread will ONLY occupy an entire CPU core if the core is not busy otherwise.
 - TM1 simply initiates a parallel thread. Where and how this thread is being processed is handled by the Operating System.
- ⇒ it is typically best to leave MTQ at its default and hence have it leverage as many worker threads as there are CPU cores.

Caveat: While MTQ=ALL ensures that hardware is leveraged at its maximum, it is only a recommended setting provided that the HW-sizing is adequately matched with the demands of the TM1 Database(s) that are running on the machine. Operating TM1 databases with MTQ=ALL yet on machines with insufficient HW capabilities may over-tax the HW and lead to performance degradation.



TM1 Database Server Configuration: TM1s.cfg

- On MTQ (continued):
 - If you have two or more TM1 databases on the same machine, it is still recommended to set MTQ to ALL, because: An MTQ worker thread does **not** occupy a CPU core. MTQ worker threads will share CPU cores where needed and where applicable. It is the Operating System that handles CPU time and resources.
 - ⇒ If you have two databases on an 8 core machine, each set to MTQ=8, and on each a user runs a query leveraging 8 MTQ threads, the 16 MTQ threads are balanced among the 8 cores on the machine, utilizing the HW in an optimal way.
 - ⇒ If only one user runs a query leveraging 8 MTQ threads, the HW is still used in an optimal way.
 - ⇒ But if you set MTQ=4 and just one user runs a query, the HW is utilized at only 50%.
 - ⇒ If we extend these simple examples to many users and dozens to hundreds of threads (like in a typical, large TM1 environment), we can argue that every TM1 database should be configured to run queries as efficiently as possible, leveraging the available HW as much as possible.
 - ⇒ Balancing HW utilization is the job of the Operating System



TM1 Database Server Configuration: TM1s.cfg

- On MTQ (continued):
 - On very powerful HW (high # of CPU cores) on smaller to mid-sized databases, it may not make much of a difference to set MTQ to a value that is lower than the total number of cores.
 - On very large Databases, a high MTQ # typically does make a big difference
- ⇒ It is a good practice to start with MTQ=ALL (= default).
- ⇒ To evaluate if different MTQ settings improve performance, adopt a holistic testing approach:
 - **MTQ Worker Thread \neq CPU Core** => An MTQ Worker Thread is 'just' a TM1 Thread.
 - What is the peak load on each database?
 - What is the Avg CPU utilization at Peak?
 - Do we currently max-out our HW capabilities? (HW should be used as much as possible; HW idle time is not good, because it means 'things' could be running faster)
 - How do the # of cores affect our end-user queries? This depends on the database/cubes: are they very large, do they leverage rules, ...
- ⇒ Analyze!
- ⇒ Is there an MTQ setting where using additional worker threads do not matter much anymore? If Yes, then this should be your max MTQ setting (i.e. do not go higher than where you see a tangible difference)



TM1 Database Server Configuration: TM1s.cfg

- On MTQ (continued):

Some additional considerations when evaluating MTQ Settings:

- One could argue that setting MTQ to a slightly lower value than the # of cores will allow more CPU time to be available for TI-processing (for example). But: this is **only** true IF the total concurrency on the database(s) is sufficiently low (compared to the # of cores). Example: The Operating System is balancing 200 concurrent TM1 users (and their TM1 threads) among its 36 cores. In such a context, setting MTQ to 36 or 34 or 30 will not really make a difference, because the CPU will be operating at maximum capacity regardless of the MTQ setting.
- ⇒ When considering to lower MTQ, take into account the overall concurrency on the database at peak time and the data volume that you are querying.
- ⇒ If concurrency is high and the data volume high, a higher MTQ value typically results in better performance in that end-user queries will be faster.
- ⇒ When you operate multiple Databases on one server and each database has a high concurrency, lowering MTQ will not necessarily free up CPU time for processing etc., because the overall load on the environment is already high.
- ⇒ Lowering MTQ in environments where multiple TM1 Database Instances share the same hardware will only have a positive impact on performance if the concurrency on each database is sufficiently low such that HW resources are kept available for other TM1 databases (and if those other databases need the free CPU time).
- ⇒ When setting MTQ, do not optimize it for low concurrency times. Optimize it for when the database is used most.



TM1 Database Server Configuration: TM1s.cfg

- Parameters that are to be set to their default values are not to be used explicitly. Not required parameters should only be used if their default value will result in a slower or more 'contentious' database operation
- Temporary parameters (such as parameter changes prompted by support) should be commented as such and should include an explanation of why they were added (logical reason, not just 'as told by IBM support')

	Tm1s.cfg parameter	explanation
Recommended non-default parameters:	AllowSeparateNandCRules=T ViewConsolidationOptimizationMethod=Tree LoadPrivateSubsetsOnStartup=T ReduceCubeLockingOnDimensionUpdate=T PersistentFeeders=F vs. T	<ul style="list-style-type: none"> Default is F. It is the proven practice to separate N- and C-level rules Array is default (legacy). Tree is typically faster Default is F. T will improve performance for end users with private subsets Default is F. T will significantly improve scalability and reduce contention Default is T, consider/test setting to F and using MTFeeders (see below). For new TM1 databases, set to F and only turn on if/when needed.
Optional recommended non-default parameters	MTCubeLoad=T MTFeeders=T MTFeeders.AtStartup=T	<ul style="list-style-type: none"> Default is F. Allows TM1 to use multiple threads/cores for Cube Load on Startup. Will likely allow for faster database startup. Click here for more information Default is F. Allows TM1 to use multiple threads/cores for feeder processing. Will likely result in significantly faster 'refresh' & startup of database instances. Click here and here for more information. Note: use of MTFeeders functions is typically OK with conditional feeders that are based on attributes or regular lookup cubes.
New default to be handled with Caution	MTQQuery=T (default) vs. MTQQuery=F	<ul style="list-style-type: none"> Default=T. MTQQuery will use MTQ to query and 'pre-'cache C-level views used by TI-processes. While this typically leads to performance improvements, it can also significantly increase RAM consumption for larger views. That is because TI with MTQQuery=F / without MTQQuery will not load a view into RAM, but instead 'cycle' through a view in smaller blocks, not really increasing RAM usage much. Yet with MTQQuery=T, TM1 will attempt to query and load the entire C-level view into RAM prior to processing it via TI ⇒ MTQQuery=T (default) can increase RAM use significantly ⇒ only use IF (i) size of views is managed (known to be small enough) and (ii) sufficient RAM is available, otherwise, use MTQQuery=F

TM1 Database Server Configuration: TM1s.cfg

- **Recommendations:**

- Remove Duplicate TM1s.cfg entries and entries that are not applicable
- Apply TM1s.cfg settings as per prior slides on TM1s.cfg
- For as long as IT manages the HW & SW environment and infrastructure / unless the Business Owns and Manages the HW & SW environment and infrastructure:
 - ⇒ Performance-related parameters should be configured identically for all databases (i.e. TM1 shall 'operate' identically in all environments)
 - ⇒ Allow database-specific parameters only if they are related to enabling/disabling/modifying user experience & features, such as Time-Outs, EnableNewHierarchyCreation, EnableSandboxDimension, EnableTIDebugging, ...
- Always check what the default parameter values are



Planning Analytics: Configuring Time-Outs

- TimeOuts: [Cognos Analytics timeouts should be set to be smaller than TM1 Web and TM1 Database Timeouts](#),
 - ⇒ CA Session timeout
 - < PAW Session Time-Out
 - < TM1 Web & TM1 Web Applications timeout
 - < TM1 Database timeouts (for TM1 Database Servers, keep in mind that there is a 'legacy' client IdleConnectionTimeoutSeconds, and a Rest-API connection HttpSessionTimeout.
- [Recommended TM1 Web Session Time-Out when using PAW](#): 60 minutes
- [Planning Analytics Time-Outs Overview \(TM1 DB Server, Web, Applications, PAW, ...\)](#)
- New as of 2.0.6: [TM1 Web HttpSessionTimeout](#)

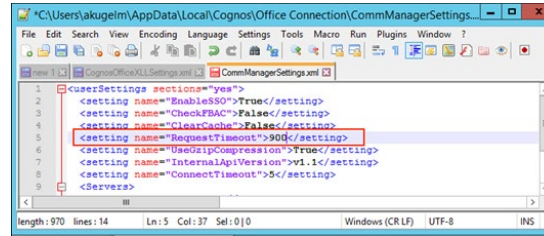


Planning Analytics for Excel: Time-Outs

- In **network environments with high latency (high ping rate and/or many network hops)**, it can be beneficial to increase the Request Timeouts and Execution Timeout for PAX in the PAX xml configuration files:

- Request Timeout in CommManagerSettings.xml:**

```
<setting name="RequestTimeout">300</setting>
```

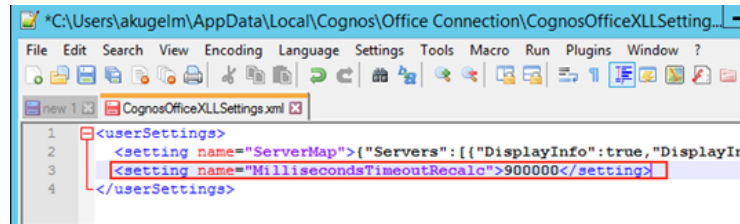


```
*C:\Users\akugelm\AppData\Local\Cognos\Office Connection\CommManagerSettings.xml
File Edit Search View Encoding Language Settings Tools Macro Run Plugins Window ?
new 1 x1 CognosOfficeXLLSettings.xml CommManagerSettings.xml
1 <userSettings section="yes">
2 <setting name="EnableSSO">True</setting>
3 <setting name="CheckFBAC">False</setting>
4 <setting name="ClearCache">False</setting>
5 <setting name="RequestTimeout">300</setting>
6 <setting name="UseZipCompression">True</setting>
7 <setting name="InternalApiVersion">v1.1</setting>
8 <setting name="ConnectTimeout">5</setting>
9 </Server>
length: 970 lines: 14 Ln: 5 Col: 37 Sel: 0|0 Windows (CR LF) UTF-8 INS
```

- RestAPI Execution Timeout in CognosOfficeXLLSettings.xml**

```
<setting name="MillisecondsTimeoutRecalc">300000</setting>
```

The default for the execution timeout is 300000, and by default, there is no entry for this setting. Note that this setting is in milliseconds, not seconds




```
*C:\Users\akugelm\AppData\Local\Cognos\Office Connection\CognosOfficeXLLSetting...
File Edit Search View Encoding Language Settings Tools Macro Run Plugins Window ?
new 1 x1 CognosOfficeXLLSettings.xml
1 <userSettings>
2 <setting name="ServerMap">{"Servers":[{"DisplayInfo":true,"DisplayIn...
3 <setting name="MillisecondsTimeoutRecalc">900000</setting>
4 </userSettings>
```



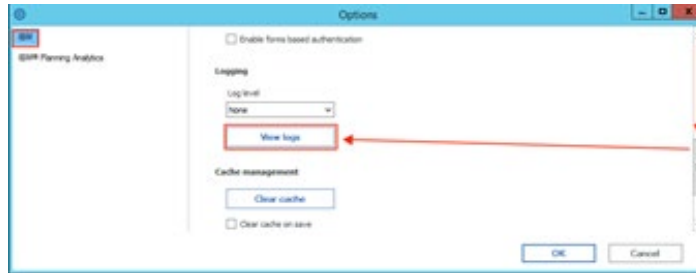
Planning Analytics for Excel: Time-Outs

How to find the location of the PAX xml config files:

i. In PAX, click on 'Options':



ii. In the top left corner, click on 'IBM', then scroll down and click on 'View Logs':



iii. In the window that opens, go up one folder



Logging: Overview

- System and Performance Monitoring is an essential tools for the Analysis and Remediation of problems (actively and retroactively)
- Click [here](#) for an overview of available tools, methods, etc.



Logging: TM1 Server Logging

- Location: TM1s-log.properties file, in the same directory as the tm1s.cfg
- Only use DEBUG loggers if
 - prompted by support or
 - for temporary troubleshooting/testing or
 - if the particular debug logger has a low or insignificant performance impact (such as the Lock.Exception debugger below)
- Recommended loggers:
 - log4j.logger.TM1.Lock.Exception=DEBUG
(useful for identifying locking issues and analyzing means to remediate)
- [Available loggers](#)



Logging: TM1 Database Activity Logging

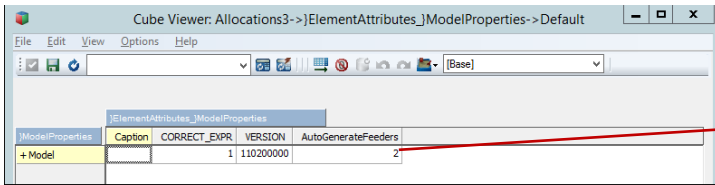
- Log Activity of Users to disk!
- TM1Top and/or Operations Console can both be configured to write/log TM1 database / user activity to disk
- Those activity logs are essential in telling a story of ‘what was going on / what happened’, and are the best tool in determining how to approach and remediate a problem.
- View the document on [Operations Console](#) for details or – if using TM1Top – consult the [TM1top.ini documentation](#) (see parameters LogFile and LogPeriod)
- [TM1Top logging on PA Cloud](#)



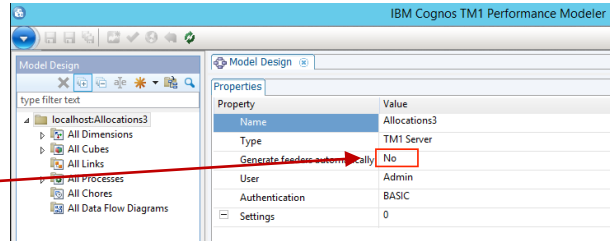
Performance Modeler: recommended default settings

(to prevent changes to TM1 Database if/when an admin logs into TM1 via Performance Modeler)

- Performance Modeler, upon login by an admin will attempt to modify (and recompile) rules right after login unless the following settings are set such that Performance Modeler is prevented from preventing the behavior:
- In }ElementAttributes_}ModelProperties, ensure that the value for 'AutoGenerateFeeders' is 2. This setting will set the Performance Modeler setting 'Generate Feeders Automatically' to No

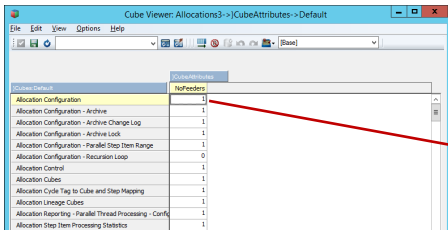


}ElementAttributes_}ModelProperties			
ModelProperties	Caption	CORRECT_EXPR	VERSION
+ Model		1	110200000
			AutoGenerateFeeders
			2

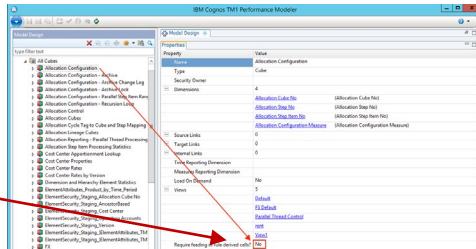


Property	Value
Name	Allocations3
Type	TM1 Server
Generate feeders automatically	No
User	Admin
Authentication	BASIC
Settings	0

- To prevent Performance Modeler from generating feeders when modeling/creating new rules, Disable feeder generation explicitly per cube in }CubeAttributes by ensuring the setting for NoFeeders is 1 for every cube:



}CubeAttributes	
Attribute	NoFeeders
Allocation Configuration	1
Allocation Configuration - Archive	1
Allocation Configuration - Archive Change Log	1
Allocation Configuration - Archive Change Log 1	1
Allocation Configuration - Parallel Step Item Range	1
Allocation Configuration - Recursion Lock	0
Allocation Control	1
Allocation Cubes	1
Allocation Cycle Tag to Cube and Step Mapping	1
Allocation Lineage Cubes	1
Allocation Reporting - Parallel Thread Processing - Conf	1
Allocation Step Item Processing Statistics	1



Property	Value
Administration Configuration	Administration Configuration
Type	Cube
Security Owner	4
Dimension	Allocation Cube No
Allocation Cube No	Allocation Cube No
Allocation Step No	Allocation Step No
Allocation Step Name No	Allocation Step Name No
Allocation Configuration Measure	Allocation Configuration Measure
Separate Links	0
Target Links	0
Hypercube Links	0
Missing Reporting Dimension	No
Used On Demand	Default
Default	Default
Default	Default
Health Check Control	Health Check Control
Health	Health
Require Feeder Generation	No



Planning Analytics Local: TM1Web on Planning Analytics (from Apache to WebSphere Liberty)

- [TM1 Web installation](#)
- [TM1 Web Integrated Login \(configuration\)](#)
- [TM1 Web Configuration Parameters](#)
- [TM1 Web & WebSphere Liberty Application Server configuration and performance tuning](#)
- [Increasing Heap Size in TM1 Web \(if increasing Heap Size generates an error on startup\)](#)
- Important Note on Heap Size / Maximum Memory for WebSphere Liberty:
 - ⇒ especially in environments where use of TM1 Web is not managed by IT but the TM1 Web environment is managed by IT, it is a good practice to set the Maximum WebSphere Liberty Memory to a significantly higher value than the default of 1.5 GB.
 - ⇒ When upgrading from TM1 Web 10.2.2 to TM1 Web PA, set the Maximum memory in WebSphere Liberty to the same value as Apache Tomcat was set to in TM1 Web 10.2.2. Note the [recommended limits for PA 2.0.5](#) or lower and upgrade to PA 2.0.6 or higher if needed.



**Self-Service with
IBM Planning Analytics Workspace (PAW),
IBM Planning Analytics for Excel (PAx), &
Cognos Analytics (CA)**

Self-Service Contribution and Analysis with IBM Planning Analytics Workspace

The New Face of TM1: Visual, Intuitive, Insightful, Social, Mobile

Rich interactive workspace

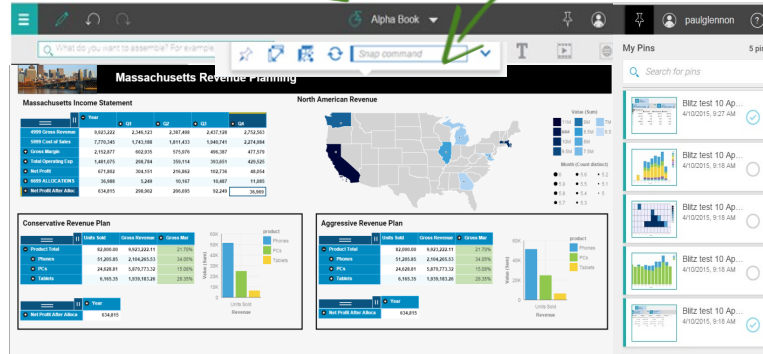
Natural language discovery & manipulation

Pin to share & reuse

Collaborate with Colleagues via Chat

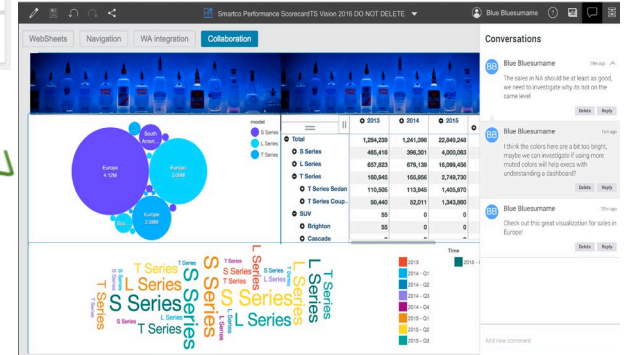
Organize into multi-tabbed workbooks

Self-Service & Collaborative Analysis and Contribution



Streamlined Navigation between Books

Synchronize Views, Charts and Websheets within a Book and Book-to-Book when you navigate



Self-Service Contribution and Analysis with IBM Planning Analytics for Excel

The ultimate Excel UI for Financial Analytics & Performance Management:

- Ultra-fast self-serve analysis & contribution, even over WAN
- Interoperability with Planning Analytics Workspace

Self-Service & Collaborative Analysis and Contribution

Content Store

Uniform UX

Pop Up Viewers

Pop Up Set Editor

The screenshot displays the IBM Cognos Product Forecast application within an Excel environment. The main interface shows a grid of data for various beer styles across different time periods. Several pop-up windows are overlaid on the main view:

- Active Form:** A window showing a list of beer styles (Craft, Shiner, Dogfish, Elysian, Purple Haze, Sierra Nevada, New Belgium, Fat Tire, All Day IPA) with columns for region, units, and version.
- Set Editor (ProductMix):** A window for editing the current set of members, showing a search for available members and a list of selected members (Shiner, Dogfish, Elysian, Purple Haze, Sierra Nevada, New Belgium, Fat Tire, All Day IPA, Long Trail, Heady Topper).
- Available Members:** A window showing a search for available members and a list of selected members (Total, Beer, Spirits, Wine, Specialty Soda, Mixers).
- Open Report in Cube Viewer:** A button to open the report in the Cube Viewer.

The main data table shows sales and gross margin data for 2014 and 2015:



	2014		2015	
	Sales	Gross margi..	Sales	Gross margi..
Beer	138,460,421	72.2%	74,275,122	85.9%
Craft	46,541,572	61.4%	33,354,690	85.4%
Shiner	0	#N/A	4,814,391	85.5%
Dogfish	7,872,305	67.4%	2,815,110	76.3%
Elysian	6,169,272	58.4%	5,249,763	86.7%
Purple Haze	0	#N/A	4,819,280	85.6%
Sierra Nevada	15,481,154	83.4%	6,525,146	89.3%
New Belgium	8,620,309	70.3%	4,301,926	83.8%
Fat Tire	8,398,532	69.5%	4,829,084	85.6%
All Day IPA	0	#N/A	0	#N/A

Self-Service Contribution and Analysis with IBM Planning Analytics with Cognos Analytics

Modern, Self-Service Business Intelligence

Reporting and Visualization Features:

- Contextualized smart search
- Intuitive interface helps all users quickly author content
- Dashboards can be created using drag and drop on mobile device or desktop
- Automatic recommendation of the best visualizations for particular data
Templates and styles let you format reports instantly
- On-demand menus for access to full capabilities over a clean workspace
- Single interface to create ad hoc or pixel perfect reports, frees up IT




Data(Analysis) Modeling Features:

- A complete web-based experience
- Easy upload of personal and external data
- Direct reporting from data sources
- Effortlessly combine data sources
- Automatic data model generation based on keywords

Data Integrity Features:

- Data protected with layers of permissions, authentication, and history
- Report integrity maintained regardless of range of inputs across business
- Controls to protect data whether you're creating one report for many or many are creating one report
- Scheduling and alerts



DevOps Optimization and Cloud-Readiness: Considerations for

- **Operations Management & Orchestration**
-> Cognos Command Center
- **Data Integration**
-> Cognos Integration Server
- **Performance and Regression Testing**

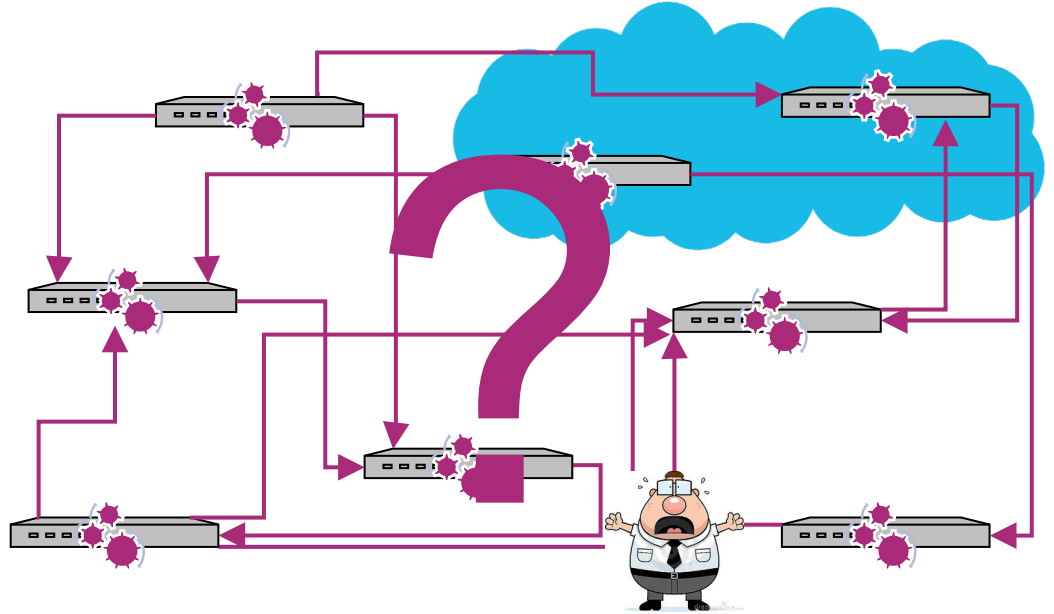
Cognos Integration Server

In the context of MDM, Master- & Meta-data Governance, Data-Integration & Transformation, IBM Cognos Integration Server provides out-of-the-box hybrid-cloud integration capabilities, hence minimizing compliance risk and improving IT-process efficiency

Feature	Advantage	Out-of-the-Box Benefit
Integrated and Automated Extraction	<ul style="list-style-type: none"> Automates data extraction Extract TM1, Oracle Hyperion Essbase, HFM, SAP BW Fact-, Meta-, & Master-Data Dynamically and Automatically generates <ul style="list-style-type: none"> RDBMS Target Schemas Files Target Dimensions and Cubes TM1 TI-scripts 	<ul style="list-style-type: none"> Establish and maintain a Dimension Meta- & Master-data backup & repository Automatically Integrate different (Meta)-Data formats Automate Meta-, Master- & Fact-Data export/import Integration and export/import from <ul style="list-style-type: none"> Cube Model to RDBMS (CIS Star Schema), Cube Model to Text File, Cube Model to Cube Model, CIS Star Schema to Star Schema CIS Star Schema to Text File CIS Star Schema to Cube Model
Hybrid-Cloud Support	Native, RestAPI-based support of TM1, Planning Analytics Local, Planning Analytics Cloud	<ul style="list-style-type: none"> Out-of-the-box integration of Hybrid-Cloud TM1/PA Environments (on premise to on premise, on premise to cloud, cloud to on premise, cloud to cloud) Transfer/Replicate/Copy data between TM1/Planning Analytics instances
Performance Optimized	supports continuous, near real-time extraction & import	Support for continuous planning & analysis in a Hybrid-Cloud environment

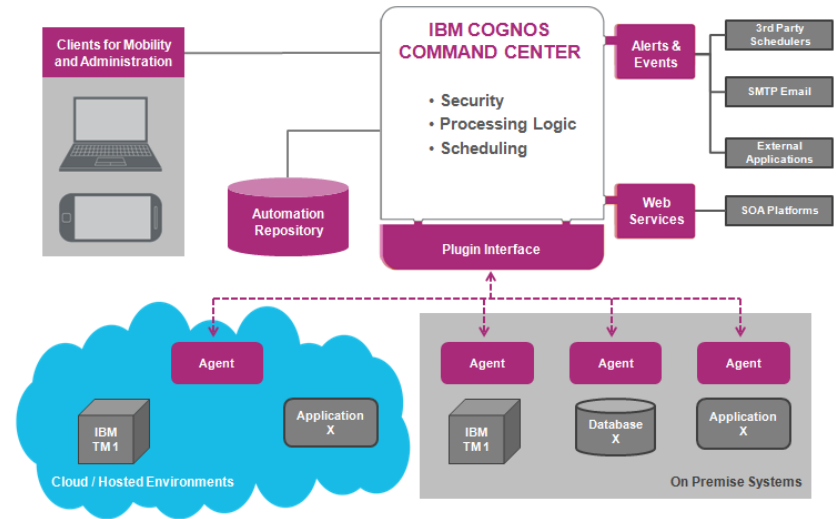
Cognos Command Center

Cognos Command Center provides a solution for orchestrating recurring maintenance & administration tasks around Master-/Meta- & Fact-Data Management for TM1, hence unlocking significant DevOps productivity gains while establishing full procedural auditability



Cognos Command Center

- Orchestrates applications, either on-premise or in the cloud
- Libraries of business functions eliminating custom code
- Business friendly interface that supports mobility
- Central repository for compliance and audit
- Handshake with enterprise schedulers and web Services
- Each application with a command-line interface



Features:

- ⑩ Email automation
- ⑩ File automation (copy, move, delete,...)
- ⑩ FTP, SFTP, FTPS automation
- ⑩ Systems & Monitoring utilities
- ⑩ IBM Cognos Integration Server integration (run CIS tasks)
- ⑩ RDBMS:
 - ⑩ SQL Server,
 - ⑩ DB2,
 - ⑩ Oracle DB
- ⑩ IBM Cognos TM1:
 - Check TM1 Server
 - Delete Server Object
 - View & Filter Server Log, Audit Log, Transaction Log
 - List Server Objects
 - List Users
 - List/Delete Views
 - Manage BLOBs
 - Read/Write Cell
 - Run TI Chore, TI Process
 - Start/Stop TM1 Server (Windows)

Cognos Command Center : How does it work?

✓ Process Complete



IBM Cognos® Command Center

Process Overview

Home Process Overview

Process List Filters

Ecosystem: Corporate Planning
 Environment: (All)
 Process: (All)

Status: Running Succeeded
 Paused Failed
 Never Run Cancelled

Status	Process	Ecosystem	Environment	Last Started
❗	Consolidation	Corporate Planning	PROD	05/01/13 02:23
🔄	File ComparisonCondit	Corporate Planning	PROD	05/01/13 02:18
🕒	Compare File Values	Corporate Planning	PROD	05/01/13 02:17
✅	Email Test Process	Corporate Planning	PROD	04/30/13 08:42
✅	BasicRunTask	Corporate Planning	PROD	04/30/13 08:41

Task Detail - Consolidate

Task type: ⚙️

Home Process Overview > Process Details - Consolidation

Failed
Exit Code: -6

Started: 05/01/13 02:23:46 PM
 Completed: 05/01/13 02:23:47 PM
 Result: Invalid Login

Task Output

```
Attempt Connection To:
  HFM Server Domain:
  HFM Server Name: sa-demo
  HFM Application Name: CommaCL
  HFM UserID: admin
  HFM Password: *****

HFM version: 11.1.2.2.000.3609

Failed to Connect

Error #-2147217085 was generated

Elapsed Time: 00:00:00

Exit Code: -6
```

On Premise Applications

Unix Agent

Windows Agent

- Automation Binary & resources
- Unmanaged & no version
- Hardcoded values & locations



Automation Repository

- Audit, version control and life cycle management
- Leverage dynamic variables

Cloud Applications

Remote Agent

Cognos Integration Server & Cognos Command Center Guides, Training, & Tutorial

- Click [here](#) to access a Cognos Command Center (CCC) deck which includes a thorough introduction and tutorial
- Click [here](#) to access a folder with CCC & CIS training courses (ppt decks and recordings)
- Click [here](#) for a document on how to use the CIS command line utility (highly recommended)



Performance and Regression Testing, TM1 Rest API

FOPM Performance and Regression Testing: Proven Practices

Tools and methodologies for testing via [TM1 Rest API](#):

- [Cucumber](#) (click [here](#) for some guidance)
- [IBM DeveloperWorks API Testing Framework Using BDD Approach Through Cucumber and Rest- Assured \(involving Jenkins\)](#)

TM1 Rest API Materials:

- [Introduction and Tutorial on the TM1 Rest API](#)
- [TM1 Rest API Developer Guide](#)
- [TM1 Rest API wrapped into Python](#)
- [Mastering the Rest API with Postman](#)

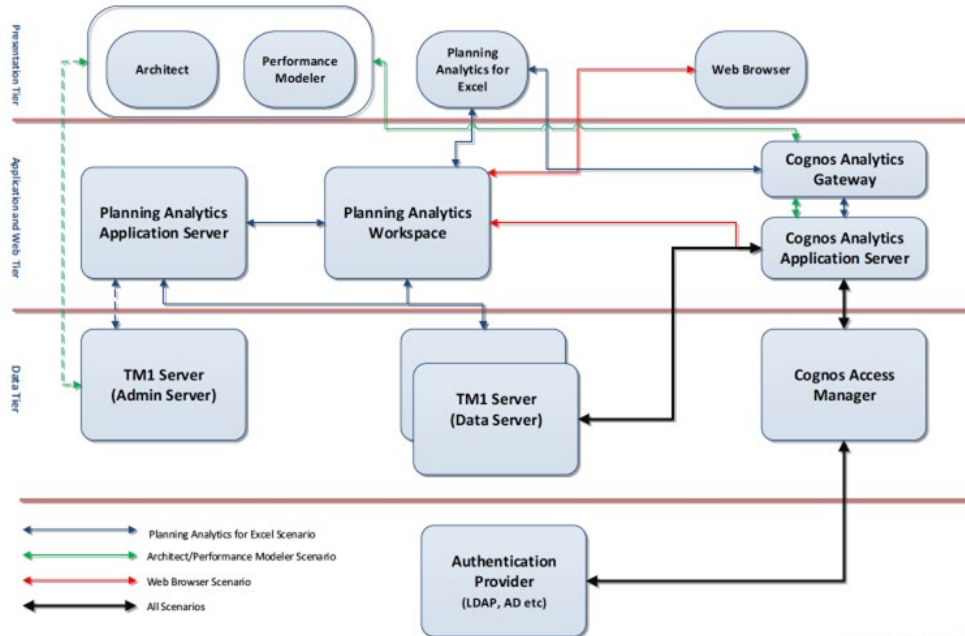


Configuring Planning Analytics **Local** Security Authentication

Security Authentication

Configuring CAM Security for

- [TM1, TM1 Web, PMHUB, PMP SVC](#)
- [Panning Analytics Workspace \(PAx via PAW\)](#)
- [Detailed Walkthrough and Guidance on CAM Security with PA](#)



Encryption

- **Data Transmission Encryption**
- **Encryption at Rest**

SSL/TLS encryption of TM1/PA Data Transmissions

- SSL encryption via TLS 1.2 is enabled with TM1s.cfg parameter UseSSL=T
- [Overview of TLS Data Transmission Security](#)
- To configure custom (WBA-issued) SSL certificates, please refer to the following links:
 - [Using Custom TLS](#)
 - [Planning Analytics 2.0 Local - Configure Custom SSL Certificates for TM1 Database Server and TM1 Admin Server](#)
 - [Planning Analytics 2.0 Local - Configure Custom SSL Certificates for TM1Web, pmpsvc, and pmhub](#)



New with Planning Analytics: Encryption at Rest

- Why 'Encryption at Rest'?

In its default state (and as per a common default practices for database objects), the data objects in a TM1 database data directory and the TM1 database log files are not encrypted. In the unencrypted state, the restoration of a TM1 database is possible just based on data directory files and - where applicable - transaction logs. Additionally, unencrypted TM1 objects can be read using a text editor.

=> Ideally, we would want to have the TM1 Database Objects be encrypted. That is what Encryption at Rest does:

- What is 'Encryption at Rest'?

Encryption at Rest provides the capability to encrypt TM1 objects and log files such that their content cannot be read and such that the objects themselves are not sufficient to restore a TM1 database without having access to the encryption keys used by the TM1 server.

[Click here for a comprehensive guide on how to configure, enable and manage TM1 Database Encryption at Rest](#)



General Recommendations after migrating from TM1 10.2.2 to Planning Analytics

TM1 Database Design & Development Practices: General Recommendations when moving from TM1 10.2.2 to Planning Analytics

Optimize User Experience by using new User Interfaces

- Use Planning Analytics Workspace (PAW), i.e.
 - i) Transition from TM1 Web Cube Views to native PAW Books.
 - ii) Transition from using TM1 Web/Perspectives Active Forms /PAX Dynamic Reports to native PAW Books

Advantages: higher functionality/features, ability to build dashboard-like analysis, collaboration, self-service capabilities, faster performance than Active Forms / Web Sheets)

Note: TM1 Web Excel Workbooks can be embedded in PAW, i.e. PAW replaced TM1 Web Cube Views. TM1 Excel Workbooks embedded in PAW still use TM1 Web

- Leverage PAX to provide users with improved self-service desktop analysis capabilities
- Immediately switch to using PAX for Web-Sheet development. Convert Perspectives Sheets to PAX on an as-needed/applicable basis, for example when editing a web sheet to add a new feature or when fixing it. Because TM1 Web understands both Perspectives and PAX sheets, a conversion of all Perspectives Sheets is not necessary, but sheets should be converted and tested as they are changed in the course of regular DevOps cycles. Note that Perspectives formulas are identical to PAX formulas. Differences between PAX and Perspectives: See [Planning Analytics Local Upgrade Guide](#)



TM1 Database Design & Development Practices: General Recommendations when moving from TM1 10.2.2 to Planning Analytics

Leverage the new Planning Analytics 'Hierarchies' functionalities for Analysis

- In cases where dimension member attributes are dimensionalized within a cube because end-users would like to slice&dice/analyze by those attributes: In Planning Analytics, such dimensionalization is not necessary, because attributes can be leveraged in a 'virtual' dimension, hence reducing cube size, simplifying cube maintenance, and improving functionality.
- Also, alternative hierarchies/rollups can in PA V2.0 be put in a different 'hierarchy container', allowing customers to have true alternate rollups. In PA V2.0, a consolidation like 'Sweden' for example can have different descendants (children) in one hierarchy container vs the other. In TM1 10.2.2, a consolidation can only have one set of descendants, in PA V2.0, a consolidation can have different descendants in different PA V2.0 hierarchy containers.



Planning Analytics & TM1 Application and Database Design & Development Practices:

Overview of proven design & development practices

IBM Planning Analytics Proven Practices, Assets, Utilities, and Templates

Proven Practices, Assets, Utilities, Configuration & Design Guidelines for TM1 Modeling and Configuration

- [Introductory Guidelines and Proven practices for TM1 Configuration, Architecture, and Design](#) (or “An introductory architecture and design guide for building very large yet fast performing TM1 models”)
- [Proven Practices for Version/Scenario Management, TM1 Multi-Threaded-Queries: Configuration Guidelines](#)
- [Guidelines & Proven Practices for IBM Planning Analytics and IBM Cognos TM1 Integration and Performance Optimization with Cognos BI and Cognos Analytics](#)
- [Use Case, Benefits & Functionalities of a Cube, Dimension, Master-, Meta-, and Fact-Data Management & Maintenance Framework for IBM Planning Analytics](#)
- [TM1 Turbo Integrator Utilities for View- & Subset-Generation](#) (download objects [here](#))
- [TM1 Turbo-Integrator Utilities for Dimension and Hierarchy Maintenance](#) (download objects [here](#)) (utilities include TI processes to build PA V2.0 Hierarchies, using new PA V2.0 TI functions)
- and [more](#)

Security Management and Automation

- [Proven Practices for Security Management and Auditability](#)
- [Management of TM1 Security: An Introduction](#)
- [IBM Planning Analytics Security Management Framework](#)

Workflow

- [Workflow with TM1 and Planning Analytics: Workflow Types, Design Options & Guidelines](#)

Authentication

- [IBM Planning Analytics Local: Security Authentication Modes](#)



Query Caching: VMM & VMT – Where do I start?

- Query caching behaviour is configured per cube via the VMM & VMT value in the }CubeProperties cube, where
 - The **VMM** value defines the maximum amount of memory to be used for caching of Stargate views per cube (i.e. the memory pool in RAM that is made available for caching). If no VMM value is specified the default value is 128KB. The valid range is 16 – 42,934,943,296 kb (16-2³²kb).
 - The **VMT value defines the** query time threshold (in seconds) as of which Stargate Views will be cached (default = 5 seconds).
- What's a Stargate View? A Stargate view is a calculated and stored subsection of a TM1 cube that is created when querying a TM1 cube, with the purpose of allowing quicker access to the cube data. A Stargate view is different from a TM1 view object in that it contains only the data for a defined section of a cube (current title elements and row and column subsets), and does not contain the formatting information and browser settings that are in a view object.



Query Caching: VMM & VMT – Where do I start?

VMM:

- For applicable cubes, it is a good practice to increase the VMM value to a new 'default' value of between 1MB and 5MB (or higher). This will allow caching of relatively large as well as many views and improve performance accordingly.
- Running and monitoring representative sample queries via TM1 Operations Console will indicate when the Query Cache is leveraged (short query time and no MTQ activity) vs. if the cache is not being used (long query times & MTQ activity). One can also log stargate creation with

```
log4j.logger.TM1.Cube.Stargate=DEBUG
log4j.logger.TM1.Cube.Stargate.ViewStorage=DEBUG
```

ONLY use such loggers in DEV or Test environments & for testing/evaluation purposes
- If the VMM value(s) are too low, A query may not be cached. In such a scenario, running the same query repeatedly will cause repeated query execution & calculation vs. retrieval from the cache. If such scenarios are common, i.e. if certain common queries will not be cached due to low cache size, increase the cache until the query result is retrieved from cache.
- Check the logs for stargate size(s) and increase VMM if needed. For very large cubes, a cache size in the high MB to GB range can be advantageous!



Query Caching: VMM & VMT – Where do I start?

VMT:

- For smaller environments, keeping VMT at its default of 5 seconds is typically a good idea.
- For large environments, where MTQ is leveraged heavily and where many processors may get engaged, an decrease of the VMT threshold may be beneficial in that CPU intensive queries, even if short, will be cached and hence will employ only one CPUs once cached rather than many.



Query Caching: VMM With MTQ




The use of MTQ typically will require an increase in VMM size. If VMM cache is set too low, even queries that were cached without MTQ use may not be cached anymore once MTQ is enabled. In such cases – to avoid unnecessary re-execution of MTQs - increase the VMM value until repeated query execution will not trigger MTQ activity anymore (indicating the cache is used).

- ⇒ Running and monitoring representative sample queries via TM1 Operations Console will indicate when the Query Cache is leveraged (short query time and no MTQ activity) vs. if the cache is not being used (long query times & MTQ activity).
- ⇒ If the VMM value(s) are too low, a query may not be cached. In such a scenario, running the same query repeatedly will cause repeated query execution & calculation vs. retrieval from the cache.
- ⇒ If such scenarios are common, i.e. if certain common queries will not be cached due to low cache size, increase the cache until the query result is retrieved from cache.



Cube Logging

- TM1 Data that is loaded or entered into a cube is not automatically committed to disk (the *.cub file) but retained in-memory only.
- For Cubes with Cube Logging enabled (default) data changes will be logged in the TM1s.log file.
- Cube Logging may be disabled for a particular cube by setting the cube's 'Logging' value in }CubeProperties to 'No':



- Use the functions
 - CubeSaveData() to commit one cube's data and/or
 - SaveDataAll() to commit all data to disk.

When CubeSaveData() or SaveDataAll() are called, all corresponding log entries will be removed from the TM1s.log file and the logs will be 'archived' in new log files named tm1s<DateTimeStamp>.log such as tm1s20140210172014.log.



Cube Logging

- `SaveDataAll()` and `CubeSaveData()` will acquire a write lock on the cube/database. It is hence not recommended to use `SaveDataAll()` and `CubeSaveData` at the end of each TI load process (because that will cause lock-contention with other load processes that load data to any cube (`SaveDataAll()`) or also load data to the same cube (`CubeSaveData()`) in a parallel load scenario (see below). There should be only one TI process that calls the `SaveDataAll()` function. Use a stand-alone, single, distinct chore to execute the `SaveDataAll` operation.
- Commit to disk prior to backups and/or at the end of a business day / at night. Do not commit to disk during regular working hours, as this may affect the user experience. There is no need to commit to disk more than once a day
- Undo operations will evaluate all applicable TM1 logs in the TM1 Log directory. If a large number of log files can be found in the Log directory, the UNDO operation may take a longer time (and may lead to a lock on the corresponding cube).
- For **larger read-only cubes it is a good practice to disable cube logging** as this will speed up data load performance. Because a read-only cube cannot be accidentally overwritten by users, cube logging is not needed for audit purposes.
- For Plan/Budget and other **input cubes, cube logging should however be enabled** to allow for proper auditing to occur and to allow for data-recovery in case of outages.



Cube Logging

- Option: To build/change a TI process that automatically & temporarily will disable cube logging during data load one can implement the following logic:

Prolog:

```
nLogFlag = 0;  
sOldCubeLogChanges = CUBEGETLOGCHANGES(sTargetCube);  
CUBESETLOGCHANGES(sTargetCube, nLogFlag);
```

Epilog:

```
CUBESETLOGCHANGES(sTargetCube, sOldCubeLogChanges);
```

Caveats:

- If such a TI process does not complete the epilog (due to processquit for example or because processing was cancelled by an admin via TM1Top or Operations Console) or if the TM1 server is shut down while such a TI process is running, logging will be **DISABLED** for the cube.
- In the case of unplanned server outages or crashes, the absence of logging could then lead to a data loss.
- To prevent such a scenario from occurring (where logging is turned off inadvertently), it can be a good practice to set up a chore that will periodically check if logging for crucial cubes is enabled, and if not, re-enable cube logging.



Optimization of internal Cube Dimension Order (Cube Optimization)



The order of dimensions significantly affects cube memory requirements and query performance

- The affect on query performance tends to be more significant than the affect on RAM utilization, i.e. the relationship between memory requirements and query performance is not linear.

For example, a 20% reduction in memory requirements *may* be indicative of a >1000% improvement in query speed.

⇒ **Look at Cube Optimization as an essential tool and process**



Optimization of internal Cube Dimension Order (Cube Optimization)

So how do I optimize the cube dimension order to maximize performance?

- The TM1 'Cube Optimizer' lets you re-arrange & evaluate a different, internal order of dimensions in a cube and apply the new sort order to the cube if desired.
- When you optimize the order of dimensions in a cube, TM1 does not change the actual order of dimensions in the cube structure. TM1 does change the way dimensions are ordered internally on the server, and because the cube structure is not changed, any rules, functions, or applications referencing the cube remain valid.
- You can have TM1 itself evaluate the cube and dimension structures and determine a system-generated optimal dimension sort order
- Or you can manually optimize dimension sort order, allowing for further optimization beyond the system-generated optimal sort order:



Optimization of internal Cube Dimension Order (Cube Optimization)

➤ To manually optimize dimension sort order:

1. Divide the dimensions into two groups: sparse and dense dimensions.
2. Order the dimensions as follows: smallest sparse to largest sparse, followed by smallest dense to largest dense.
3. Exceptions: It typically is better to put a very small, dense dimension before a very large but sparse dimension. For example, a dimension such as Version/Scenario (Act, FCST, Plan) that has only two or three elements is better positioned before a very large but sparse dimension, such as Product, which might have thousands of elements. => be flexible and experiment with different configurations.
4. The TM1 manual states that one “should optimize the order of dimensions in a cube only in a development environment while you are trying to determine optimal cube configuration”. This should be interpreted as follows: Determine the optimal cube order via tests in a production-like development environment. Then – at a time of zero-concurrency – implement the tested new dimension order in your production environment.
5. If a dimension (such as a measures dimension) contains string elements, it needs to be the last dimension.
6. Very often, moving the measures (or last) dimension (if it does not include string elements) upwards (from the last position) can result in significant memory savings and performance gains. The TM1 system-generated cube optimization will not consider moving the measures/last dimension. The impact of moving the measures/last will have to be evaluated manually.
7. If you find that moving the measures dimension upwards will render significant performance gains but you would like to use string elements too (for comments etc.), you may consider building a separate cube just for string measures.



General Architecture & Design: Dimensions, Hierarchies, Consolidations

The Number of Dimensions Matter:

An additional dimension (regardless of sparsity) can increase total memory consumption of the cube by up to 50% due to the additional storage required for each populated cell's pointer or index (like an address). Additional Dimensions will also affect the speed at which data can be retrieved by affecting the trie (the ordered tree data structure that is employed for storing and retrieving values from a TM1 cube)

Therefore:

- Try to limit the # of dimensions in the cube to the dimensions that are really needed.
 - Do not create a specific 'measures' dimension unless it really is needed.
- ⇒ When using Cognos BI / Cognos Analytics, evaluate if a 'regular' dimension can fulfill the technical purpose of the measures dimension (for example the currency dimension). For more details on PA/TM1 with BI/CA, go to [Guidelines & Proven Practices for IBM Planning Analytics and IBM Cognos TM1 Integration and Performance Optimization with Cognos BI and Cognos Analytics](#)




General Architecture & Design: Dimensions, Hierarchies, Consolidations



Leverage TM1's natural consolidation algorithm to its fullest potential

Example 1 - **Version/Scenario Dimension**: For Variances and other simple calculations and rollups, use weighted hierarchy consolidations to conduct variance calculations

1. Use a dedicated 'set' of Version dimension elements for which you will define leaf level rules. Those are your 'open' Versions:




Then, in the context of Version Management,

2. copy/submit data from those 'open' FCST & Plan versions to 'historical' Versions:

Version	Version Identifier	Hierarchy Suffix	Actuals Through Date
-- Open Versions			0
Actual	Actual 201503		201503
Forecast	Forecast 201410	FCST201410	201410
Plan	Plan 201501	Plan201501	201501

Version Management TI Process
Will Create Elements & Copy Data



General Architecture & Design: Dimensions, Hierarchies, Consolidations

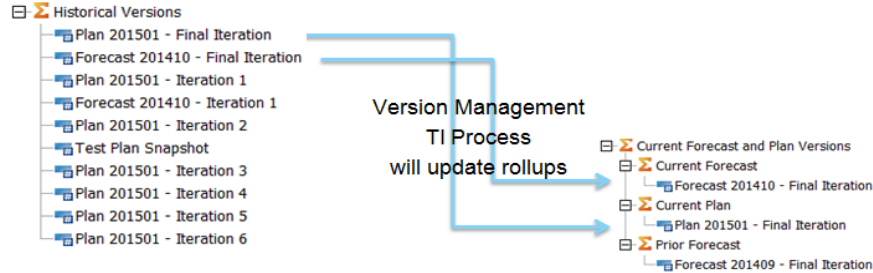


Leverage TM1's natural consolidation algorithm to its fullest potential

- and to create rollups for 'Current Forecast' & 'Current Plan', with 'current' ⇔ last approved FCST/Plan

Current Forecast and Plan Versions

- Current Forecast** (latest official/approved/submitted Forecast)
Forecast YYYYMM – Final Iteration
- Current Plan** (latest official/approved/submitted Plan)
Plan YYYYMM – Final Iteration
- Prior Forecast** (prior official/approved/submitted Plan)
Forecast YYYYMMPrior – Final Iteration
- Current Plan Target** (latest official/approved/submitted Plan target)
Plan Target YYYYMM – Final Iteration



You may now use the elements 'Current Forecast', 'Current Plan', 'Prior Forecast':

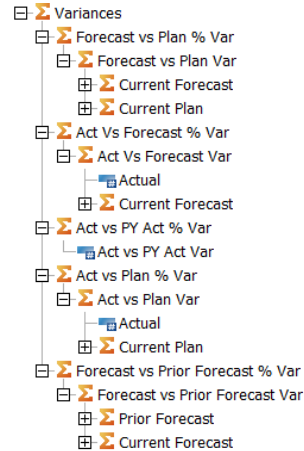


General Architecture & Design: Dimensions, Hierarchies, Consolidations



Leverage TM1's
natural
consolidation
algorithm to its
fullest potential

4. In reports or rules (w/o having to resort to using alias names in rules)
5. For variance calculations:



where

6. Consolidation weights can be used to calculate absolute variances
7. 'Feeder-less' C-Level rules can be put in place to calculate % Variances

$$['Act Vs Forecast \% Var'] = C: (['Actual'] - ['Current Forecast']) \setminus ['Current Forecast'];$$




General Architecture & Design: Dimensions, Hierarchies, Consolidations



Leverage TM1's natural consolidation algorithm to its fullest potential



Example 2 – **Time Period Dimension**: Use a continuous time period dimension with alternate rollups & hierarchies for (relative) time- and variance Analysis

1. Start with a standard 'Fiscal' Time Period hierarchy like




and then,

2. create QTD, YTD & YTD Variance Rollups:



Name	Weight
Σ 201503 YTD	1
Σ 201403 YTD	-1





General Architecture & Design: Dimensions, Hierarchies, Consolidations



Leverage TM1's natural consolidation algorithm to its fullest potential

3. Create rollups for 'Current Month', 'Prior Month', '(current) QTD', '(current) YTD', etc., including Variance rollups

In the below example, 201504 is the current fiscal month, resulting in the following relative rollups:




General Architecture & Design: Dimensions, Hierarchies, Consolidations




Leverage TM1's natural consolidation algorithm to its fullest potential

4. Create a calendar year rollup if applicable.

The following example shows a Calendar Year rollup for a company whose Fiscal Year starts in October:



Member/Element Alias Names have been created to view the Time Period Dimension in Calendar Year Format:



General Architecture & Design: Rule Calculations



TM1 will evaluate a cell for rules if the cell is identified by a rule AREA (the left side of the rule/feeder). This evaluation comes at a cost!

- ⇒ For TM1 rules and feeders, specify the AREA as detailed as possible.
- Example: The following rule will be evaluated against each N-Level data point and the evaluation may determine that the no rule is to be applied. But before TM1 comes to that conclusion it will evaluate each intersection.

```
[ ] = N:  
  IF ( !Currency @= 'Local Currency' &  
      (!Version @= 'FCST' % !Version @= 'Plan'),  
      ...,  
      stet );
```

For small cubes, the duration of this evaluation is insignificant. For medium to larger size cubes, it can add many seconds to even minutes to the query time.

=> Write the rule such that the left side (the AREA) is as specific as possible (even if this means including many elements), like:

```
[{'FCST','Plan'},'Local Currency'] = N: ...;
```

With this rule, there will be no rule evaluation for Currencies other than 'Local Currency' & Versions other than 'FCST' & 'Plan'.



General Architecture & Design: Rules vs. Turbo-Integrator Process (TI-Process)

1. Reporting & Analysis: if the data is available in the source system, use TI
2. If the data 'moves' according to a low-volatility submission or closing cycle, use TI
3. As of TM1 10.2 MTQ, consider rules where you previously did not dare to
But: do not use MTQ to counter a poor rules design. MTQ allows a good design to be more scalable from a concurrency and data volume perspective. It is not a crutch for a poor cube and/or rules design
4. For large & volatile data-sets, TI re-processing may be too time-consuming
5. Are calculations via weighted consolidations/rollups an option (like the Time and Variance calculations in our previous examples)?
6. Are feeder-less C-Level rules an option (like for our % Variance calculations from before)?
7. For small cubes / data volumes, rules are often a better choice than TI, due to
 - insignificant performance difference,
 - higher user-friendliness (real-time calculations & fewer processing steps)
8. Use rules for calculations against highly volatile data (input/planning data)
9. If in doubt & particularly for larger cubes, evaluate the trade-offs between Performance, Usability & Maintenance:

	Performance Impact	Usability & Maintenance Impact
TI-Process	Processing within acceptable timeframe? Y/N?	Can processing be automated? Y/N? Do end-users have to manually process? Y/N?
Cube Rules	Good Calculation Performance? Y/N?	Can cube rules be fed efficiently? Y/N? What is the maintenance effort for the rules & feeders? Low/High?



General Architecture & Design:

Rules vs. Turbo-Integrator Process (TI-Process) - Examples

FX conversion for Input Data	Performance Impact	Usability & Maintenance Impact
TI-Process	Processing within acceptable timeframe? -> TBD	Can processing be automated? -> No Do end-users have to manually process? -> Yes
Cube Rules	Good Calculation Performance? -> Yes	Can cube rules be fed efficiently? -> Yes What is the maintenance effort for the rules & feeders? -> Low

Sales \$ Actuals based on Price & Volume	Performance Impact	Usability & Maintenance Impact
TI-Process	Processing within acceptable timeframe? -> Yes (along with Act update)	Can processing be automated? -> Yes Do end-users have to manually process? -> No
Cube Rules	Good Calculation Performance? -> Yes	Can cube rules be fed efficiently? -> Yes What is the maintenance effort for the rules & feeders? -> Low



General Architecture & Design:

Rules vs. Turbo-Integrator Process (TI-Process) - Examples

FX conversion for Actuals	Performance Impact	Usability & Maintenance Impact
TI-Process	Processing within acceptable timeframe? -> FX conversion data in Source System => Yes	Can processing be automated? -> Yes Do end-users have to manually process? -> No
Cube Rules	Good Calculation Performance? -> Yes	Can cube rules be fed efficiently? -> Yes What is the maintenance effort for the rules & feeders? -> Low

FX conversion for Actuals	Performance Impact	Usability & Maintenance Impact
TI-Process	Processing within acceptable timeframe? -> FX conversion data NOT in Source System, multiple FX conversions per record => No	Can processing be automated? -> Yes Do end-users have to manually process? -> No
Cube Rules	Good Calculation Performance? -> Yes	Can cube rules be fed efficiently? -> Yes What is the maintenance effort for the rules & feeders? -> Low



General Architecture & Design:

Rules vs. Turbo-Integrator Process (TI-Process) - Examples

Driver-based FCST/Planning & Analysis	Performance Impact	Usability & Maintenance Impact
TI-Process	Processing within acceptable timeframe? -> Yes	Can processing be automated? -> No Do end-users have to manually process? -> Yes
Cube Rules	Good Calculation Performance? -> Yes (for large cubes, MTQ may be required)	Can cube rules be fed efficiently? -> Yes What is the maintenance effort for the rules & feeders? -> Low

Profit Center Allocations & Apportionments	Performance Impact	Usability & Maintenance Impact
TI-Process	Processing within acceptable timeframe? -> Yes	Can processing be automated? -> Yes Do end-users have to manually process? -> No
Cube Rules	Good Calculation Performance? -> TBD	Can cube rules be fed efficiently? -> TBD What is the maintenance effort for the rules & feeders? -> High



General Architecture & Design: Rules vs. Turbo-Integrator Process (TI-Process) - Examples

Spreading of FCST/Plan (top-down) by YTD Act	Performance Impact	Usability & Maintenance Impact
TI-Process	Processing within acceptable timeframe? -> Yes	Can processing be automated? -> Yes Do end-users have to manually process? -> No
Cube Rules	Good Calculation Performance? -> Yes (for large cubes, MTQ may be required)	Can cube rules be fed efficiently? -> Yes What is the maintenance effort for the rules & feeders? -> Low

⇒ Outcome depends on particular spreading logic, i.e. what is easier to implement and maintain?



General Architecture & Design: Security



- Do not use rules in }CubeSecurity.cub, }DimensionSecurity.cub, }ProcessSecurity.cub, }ElementSecurity_<Dimension>.cub
- Cube Security ('}CubeSecurity.cub'), Dimension Security ('}DimensionSecurity.cub'), Process Security ('}ProcessSecurity.cub'), and Element Security Data ('}ElementSecurity_<Dimension>.cub') should be processed via TI instead of cube rules.
- If rules used, a security metadata change - for example due to a hierarchy change (with corresponding/resulting security changes for parent and/or child nodes) or due to a new element being added to a hierarchy (like a new archived version for which READ access now to be granted to all applicable groups) – will always require running the 'SecurityRefresh()' command in TM1, effectively rendering all cached security settings invalid and hence renewing/refreshing all security credentials.
- A security refresh on large models will typically lead to a multi- to many minute lock of all user activity due to TM1 refreshing security access credentials for all active users and groups.
- If security is manually entered or processed via TI (and hence directly stored in the corresponding security cube), a security refresh is not necessary for such security changes. The security changes will propagate automatically and with only very short locks.
- Optimize CellSecurity cubes by only using dimensions that are needed for determining cell security access credentials (see next slide for a utility to optimize cell security cubes via TM1 Architect)
- Be aware that CellSecurity by default may overrule ElementSecurity by allowing data access to cells without ElementSecurity access credentials. I.e. if a cell in CellSecurity is READ, but the Element is set to access NONE (empty), a user can still read the value in Excel by typing in the Element Name. Use }CubeSecurityProperties.cub and set CellSecurityMostRestrictive to Yes to only allow Cell Security to restrict access further (and to prevent cell security from being able to 'open' access.
- For details, go to:

[Proven Practices for Security Management and Auditability](#)

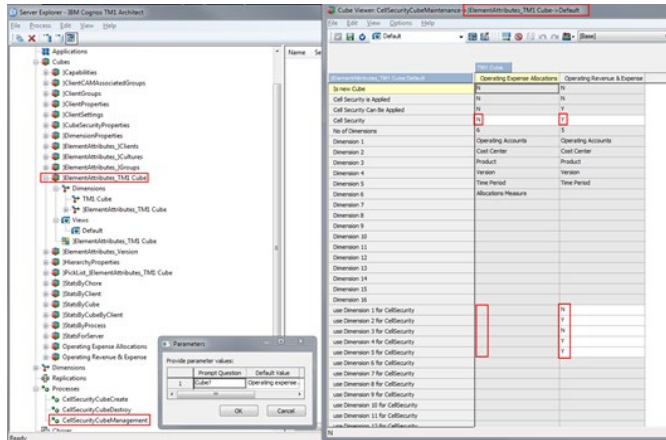
[Workflow with TM1 and Planning Analytics: Workflow Types, Design Options & Guidelines](#)



General Architecture & Design: Security

Click [here](#) for a utility for the creation of optimized CellSecurity Cubes. Note: the utility does not save and re-attach cell security rules. Ensure that a backup of the existing cell security rules is available and attach the rules after the new cell security cube was created. How to use the utility:

- 1) Add the cube name to dimension 'TM1 Cube.dim'
- 2) To delete an existing cell security cube,
 - a) go to }ElementAttributes_TM1 cube.cub and set value for attribute 'Cell Security' to N
 - b) then run process 'CellSecurityCubeManagement.pro'
- 3) To create a new cell security cube,
 - a) go to }ElementAttributes_TM1 cube.cub and set value for attribute 'Cell Security' to Y
 - b) then, in }ElementAttributes_TM1 cube.cub, configure the cell security cube dimensions like in the following screenshot:



- 4) then run process 'CellSecurityCubeManagement.pro'
- 5) to re-create an existing cell security cube with different dimensions, first execute step 2 and then step 3



General Architecture & Design: ODBC vs. File-Based

	ODBC	File
Speed (Throughput)	<p>Low to Very High (depending on DB performance, driver & infrastructure) Very High: >30.000 records / second (against Oracle DB at Fin. Sector Client in Chicago),</p> <p>lower for on-premises to cloud</p>	<p>Medium (for very large files) to Very High</p> <p>Lower for on-premises to cloud b/c file needs to be transferred first</p>
Maintainability	<p>Low Effort (SQL can be parameterized)</p>	<p>Medium to High Effort (DB ETL routines may have to be changed, multiple files may have to be extracted, ETL needs orchestration)</p>
Complexity	Low	Low
Security	High	Medium



General Architecture & Design: Optimizing Processing Speed

1. Use `CubeClearData()` prior to a re-load of all data in a cube
2. Use `ViewZeroOut()` to zero-out sections of a cube
3. Separate TI-processing for
 - a) Dimension Load/Update &
 - b) Data (Fact) Load/Update

Doing so will – via procedural separation of metadata and data-update procedures - allow the load and update of cube data using parallel load threads, hence significantly improving cube load time.

4. **Where applicable** (i.e. for adding not yet assigned masterdata elements), **use the new TM1 10 functions `DimensionElementInsertDirect` and `DimensionElementComponentAddDirect`** directly in the TM1 data tab along with loading new data (unless the dimension is very large),
5. **Utilize a Parallel Data Load Regime for loading very large amounts of data very rapidly:**



General Architecture & Design: Parallel Data Load & Processing

- **What is it?**

In a parallel data-load regime, data is loaded and processed into **one** or multiple cubes via separate & **parallel processing threads**.

- **What is it good for?**

- With parallel data load / processing, a cube can be loaded using as many CPU threads as available and applicable, hence providing significant performance gains of data-load/processing operations.
- Using a parallel data load regime/framework, TM1 can load upwards of 30.000 records per second per CPU core. => Using 16 CPUs for parallel data load, this can result in an overall data-load/update speed of roughly 480.000 records per second => Using 16 CPUs for parallel data load, you can load roughly 1,000,000,000 records (1 billion) in 35 minutes!



General Architecture & Design: Parallel Data Load & Processing

- In cases where parallel data load is to be used frequently/extensively, the use of RunTI.exe to trigger multiple data load processes built to allow parallel execution and data load is the preferred method to initiate parallel data load.
- To trigger, orchestrate and manage parallel TI process execution using RunTI.exe, the use of IBM Cognos Command Center or - if Command Center is not available – the use of RunTI.exe via scripting language are the preferred methods to initiate parallel data load.
- Proven Practices on how to initiate and manage data load & processing threads using RunTI.exe and Unix, Linux & Windows scripting language are available per request.
- See [Proven Practices for Parallel Data Load and Processing in TM1](#) for details. Templates/Assets can for parallel data load and processing can be downloaded [here](#).
- For information on how to manage TI processes using Cognos Command Center, please refer to the [Cognos Command Center Online Documentation](#) or to http://www-01.ibm.com/support/knowledgecenter/SSPLNP_10.2.1/com.ibm.swg.ba.cognos.ag_ccc.10.2.1.doc/c_ug_ccc_pi_tm1_overview.html for more information on TM1 Integration



General Architecture & Design: Parallel Data Load & Processing

The process(es) used to load/process the data (the processes that are initiated in parallel) should adhere to the following guidelines, conditions & restrictions:

- When performing a parallel data load, the zero-out / clear data operation should ideally be executed before the parallel data load is initiated (but the zero-out can also be part of each parallel load thread).
- The load processes may not update meta- & master-data that is in a dependency-relationship to the cube(s) that are updated. Any data-update-related master- & metadata updates shall occur
 - a) in a single thread per dimension and
 - b) prior to initiating the parallel fact data load,
i.e. via separate TI-processing for Dimension Load/Update & Data (Fact) Load/Update. Doing so will – via procedural separation of metadata and data-update procedures - allow the load and update of cube data using parallel load threads.
- Any subsets and views that are created in the context of the data-load TI processes need to be unique subsets and views (unique within the TM1 instance) to avoid contention during subset and view creation and deletion.



General Architecture & Design: Minimizing Contention

1. Use the [new PA TM1s.cfg parameters](#) to reduce contention
2. Use **SaveDataAll()** & **CubeSaveData()** **only when and where needed**: SaveDataAll() will acquire a write-lock on the database, CubeSaveData() will acquire a write lock on the cube. It is hence not recommended to use SaveDataAll()/CubeSaveData() at the end of each TI load process (because that will cause lock-contention with other load processes that load data to any cube (SaveDataAll()) or also load data to the same cube (CubeSaveData()) in a parallel load scenario (see below).
3. Use **Unique Views and Subsets for TI processing** used in the context of TI-processing. Any subset and view (and the subsets used by that view) that is used by a TI process should have a name that is unique to the user executing the process and the time the process is executed. Using subsets and views with unique names ensures that
 - The Subsets and Views can - if needed – be isolated and debugged without affecting other processing (be it by the same process or by different processes)
 - No contention will occur between different TI processes creating and/or destroying subsets and views at the same time
 - No contention will occur between the same TI processes creating and/or destroying subsets and views at the same time
4. Implement a **synchronization & serialization framework for applicable TI processes** (NOT for all processes – you want to be able to run in parallel what shall run in parallel)



General Architecture & Design: Minimizing Contention

4. **Implement a synchronization & serialization framework for applicable TI processes! Why?**
- The concurrent execution of Turbo Integrator processes that lock the same objects and block each other can lead to contention.
 - TI process contention typically expresses itself successive rollbacks and retries, also known as thrashing. Examples:
 - a) Two or more TI processes may perform an update of one and the same dimension
 - b) The first process to acquire the dimension 'lock' will block the other process from continuing
 - c) The 2nd process will hence do a rollback of its actions before encountering the lock and then will attempt to start anew
 - d) The 2nd process may possibly hit the same lock later should process 1 not have finished or have released the lock (again resulting in a roll-back, i.e. back to the beginning...).
 - e) Or the 2nd process will find that the lock was released, but: the 2nd TI process ,may obtain a lock on a different object that is needed by the first process (but later on, i.e. towards the end of the 1st process running
 - f) Now the 2nd process is locking the 1st process & the 1st process will do a roll-back...
- ⇒ Not only may there be roll-backs and retries, but the rollbacks and retries – depending on the TI procedures in place, may affect one another, resulting in two or more processes locking each other out.



General Architecture & Design: Minimizing Contention - a Lock Contention Example

4. Implement a synchronization & serialization framework for applicable TI processes! Why?

The concurrent execution of Turbo Integrator processes that lock the same objects and block each other can lead to contention. TI process contention typically expresses itself successive rollbacks and retries, also known as thrashing. Examples:

- Two or more TI processes may perform an update of one and the same dimension
 - The first process to acquire the dimension 'lock' will block the other process from continuing
 - The 2nd process will hence do a rollback of its actions before encountering the lock and then will attempt to start anew
 - The 2nd process may possibly hit the same lock later should process 1 not have finished or have released the lock (again resulting in a roll-back, i.e. back to the beginning...).
 - Or the 2nd process will find that the lock was released, but: the 2nd TI process ,may obtain a lock on a different object that is needed by the first process (but later on, i.e. towards the end of the 1st process running
 - Now the 2nd process is locking the 1st process & the 1st process will do a roll-back...
- ⇒ Not only may there be roll-backs and retries, but the rollbacks and retries – depending on the TI procedures in place, may affect one another, resulting in two or more processes locking each other out.

Two Processes start at the same time & both process complete the metadata procedures and start the data tab:

ID	User	Context	State	Function	Type	Object	Info	Time (s)
6800	Admin	Architect	Run:R	ProcessExecuteEx	Process	ProcessB	Data (9)	16
6376	Admin	Architect	Run:R	ProcessExecuteEx	Process	ProcessA	Data (9)	17

When Process B encounters a lock (in the data tab, of dimension 'ContentiontestA') & consequently enters a wait state:

ID	User	Context	State	Function	Type	Object	Info	Time (s)
6800	Admin	Architect	Wait:IXC	ProcessExecuteEx	Dimension	contentiontestA	-	40
6376	Admin	Architect	Run:R	ProcessExecuteEx	Process	ProcessA	Data (9)	41

When Process A has finished and the lock on dimension 'ContentiontestA' has been released, Process B starts anew:

ID	User	Context	State	Function	Type	Object	Info	Time (s)
6800	Admin	Architect	Run:R	ProcessExecuteEx	Process	ProcessB	Metadata (15)	2



General Architecture & Design: Minimizing Contention

So what do we need to do to prevent thrashing?

- ⇒ Applicable processes should be kept from being executed in parallel and instead should only be allowed to be run in serial to avoid thrashing and excessive (or even endless) rollback actions.
- ⇒ The same applies to simultaneous fact updates and dimension updates: if you are running a process that updates fact data you should not attempt to update the dimension master data using a different process as this will cause locking and rollbacks.

So how do we do this?

- ⇒ A semaphore logic needs to be employed to synchronize certain (applicable) TI processes to run in serial execution mode only. The synchronization logic will ensure that a process is told to 'wait' at the very beginning of execution (before it would require a roll-back or data updates etc.) until the 'lock' that would prevent the process from continuing is released.
- ⇒ As of TM1 10, the **SYNCHRONIZED()** function allows for easy serialization of TI processes if needed:
 - For TI processes that are built specifically to updated a certain dimension, add the following code to the beginning of the TI process (TI prolog): **SYNCHRONIZED (<DimensionName>);**
 - Proven Practice: Use **Dimension Names as the lock names** for SYNCHRONIZED()
 - For TI processes that are parameterized such that multiple or varying cubes, dimensions or objects could be accessed and locked, implementation of a configurable TM1 **TI-Synchronization utility** is recommended. The utility is to leverage the SYNCHRONIZED() function in combination with a configurable Lookup model to allow setting, maintaining and testing different serialization configurations in an ad-hoc/dynamic fashion.



Key Takeaways

- TM1 10.2 and Planning Analytics may be much faster than you think or thought
- Planning Analytics (= TM1 11) is even faster than TM1 10.2.2
- Use MTQ. It works. Beautifully.
- Do not count on MTQ to 'offset' the effects of poor rules or cube designs.
- A good cube & model design is paramount. Everything starts there. Therefore:
 - The Version and Time Dimension design matters. A Lot.
 - Don't forget about VMM.
 - Be aware of what may cause TI process contention. Avoid it.
 - Keep it simple. If data belongs together, keep it together (in one cube).
 - Keep it 'real-time' where applicable. Avoid lengthy TI-processing where MTQ-powered rules may fit the bill.
 - But: wherever a time-dependent business process applies, consider TI over rules (moving/copying/promoting data in the context of a closing/submission/approval cycle)
 - Do not forget security & how to maintain it. Avoid having to run SECURITYREFRESH()



Misc.

Misc. Links

- [IBM Planning Analytics YouTube Channel](#)
- [IBM Planning Analytics knowledge center](#)
- [IBM Planning Analytics local software downloads from passport advantage](#)
- [IBM Planning Analytics supported environments](#)
- [IBM Planning Analytics for Excel \(PAX\) and IBM TM1 Perspectives comparison/differences](#)
- [IBM Planning Analytics and IBM Cognos Analytics Integration](#)
- [IBM PA and CA integration technotes](#)
- [IBM Planning Analytics Digital technical engagement \(DTE\): Videos, product tours and hands on labs](#)



Thank You

