

Q: What is BBCA?

A: BBCA is BMC BladeLogic Client Automation, client discovery product. RE allows for you to merge multiple different discovery sources together in an effective way.

Q: Will the recon engine identify a server by using the simple CI Name and the fully qualified CI Name as and merge them as a single CI into the golden dataset?

A: One of the operations of reconciliation is to identify things that are like. You can write rules to find a match or the ones we ship. It's a choice that is left up the RE Admin. You may match on a specific attribute with the same value. You could write rules that do LIKE comparisons to try and match say "abc" and "abc.myco". It is up to the rules you write. Simple CI name and FQ CI name can be identified

Q: In this Architecture the Reconciliation Engine looks like a single processing method. How is an outage on the Reconciliation Engine impact the updating of the Golden Data Set?

A: Server group is load balanced. All of the RE engines of the CMDB follow the failover rules of server groups. So, if a server goes down, it will fail over to the backup system and continue functioning. Then, when the main system comes back, things will move back to the primary system. Impact of Recon Engine can be handled by ranking of Recon on multiple servers. In case of failure or server is down, Recon Engine on other server will pick up. You can set up any number of failover systems up to all the servers in your server group. So, it can stay up as long as there is at least one server running in your group.

Q: We have three sources and we reconcile CIs with the Name attribute but still reconciliation engine creates duplicates even when we give NAME as priority 1 in identifying.

A: If CIs are getting duplicated, then will have to re-visit the ID rules and precedence rules. If the Name is not unique then multiple match found errors should be seen in logs. The information in logs is specific enough to identify which CIs have multiple match.

Q: How you prevent overwriting of attributes in BMC.ASSET dataset

A: In best practice, you would never have something in BMC.ASSET that has not come from a source dataset through reconciliation, unless you're using the Sandbox. That way, there is never an issue with overwriting anything in BMC.ASSET as it always follows the attribute precedence from source.

Q: Are there use cases of real world usage of the Sandbox?

A: Sandbox is used very heavily by Asset Management operators to prevent overwriting of golden dataset. They use of precedence settings where the CI data can be controlled to have proper data updates. It can also be normalized in the Sandbox if needed. The concept of Sandbox is to avoid direct change in golden dataset. Sandbox will establish the process to have changes reviewed compare to other datasets and will not overwrite/modify the CI accidentally. An example is the Service Model Editor where you create/update models in a "sandbox" (an overlay dataset) and when you are happy with it and have reviewed it, you then "commit it" and it joins the other data.

Q: Are the rules scripted or written by interacting with an interface within RE?

A: Rules have been shipped OOTB. You can use the RE Standard Rules which is the recommended way to start with. You may need to make some changes depending upon your data.

Q: Is there a recommended field to store the key id value from the original datasource... I.E. the primary key id or guid from the datasource solution that the data originated from?

A: The ReconciliationIdentify field will keep track of the id of the original datasource. Best practice is to have a dataset for each source. So, you can tell EXACTLY what came from a given source by looking at its dataset. Then, that source merges data into the production dataset. So, you always know where data came from. The ReconciliationIdentity is that tracking attribute. Also the AttributeDataSourceList keeps track of where the CI came from and which attributes won the precedence gates.

Q: We too have duplicates in our CMDB

A: Duplicate CI problems could be mainly due to improper RE rules. You might want to look at the Class of duplicate CIs and compare with the Class ID Rules. Then correct the RE Identification rule for that class. The reconciliation engine prevents from having two same ReconciliationIdentities within the same dataset now. Should a duplication be attempted, then a log entry is added to reflect that and the CI will remain without an ID. Although this change was only introduced in 7604 SP3 and later.

Q: We use just ID rule for "Name" attribute, but still we see issues with duplication as we can't use any other ID rules apart from NAME

A: You might want to look complete CI value and if there is any other attribute can be added with Name. Even the Standard Rule Editor will allow for AND or OR type rules. Name is not the only attribute that can be used for identification. Several other attributes like TokenId, SerialNumber and others can be used for identification. They can be combined and even sequenced, each tier just must be able to find a unique match. Once a CI is identified, then that CI will not need to be identified again. It will just be merged to the matching record in the destination dataset.

Q: is there a performance hit for having more data sets in a precedence set? Should precedence sets be kept as small as possible?

A: Precedence set is kind of checklist in cache during RE merge operation. This cache will be visited before merging the CI into Golden dataset. No performance impact has been observed even with more datasets.

Q: what is the significance of sandbox?

A: Nair, Sandbox allows for a staging area where a CI or a Bulk load of CIs gets to be updated without an impact to the golden (BMC.ASSET) data and then (and only then) promoted into the golden data. There you can also track the changes by Modified User as each login ID is associated with the sandbox and user ID. Using the Sandbox prevents the records from being corrupted while the record is in edit mode.

Q: I'd imagine that we'd be modifying these rules or creating new ones to meet our specific needs. Specifically correlating FQDN and non-FQDN CI's. Would those rules be created/modified through a tool?

A: There is a front end where you create the rules (RE Standard Rules Editor). The CMDB has a set of Administrative forms where you create the RE rules and run the jobs and look at progress of jobs and the like.

Q: when reconciling multiple sources to BMC.ASSET; is it better to merge 1 source at a time to BMC.ASSET or merge all the sources into another dataset before finally merging to BMC.ASSET?

A: to merge 3 datasets, do I use one job to merge all three or do I merge 2 and then merge the result with the third -- merge 3 is best

If you are going to merge multiple (more than 2) datasets together, you should do them at ONE time rather than two together and then those two to a third. You can do it piecewise, but it is most effective to do it together so you can best manage the rules and precedence and avoid confusion about who wins. The only time to do it piecemeal is if you need to resolve issues between two to create something clean as a "single source" that then merges with others.

Q: you can tick each dataset to merge is it better to merge one dataset at a time or multiple?

A: Is it a good idea to merge into an intermediate dataset and then merge that with production?

When you have multiple sources, first of all, it is important to understand how much data overlap between different sources. If there is more overlap between different sources, then better to have one consolidated dataset before merging into Asset. There is value in having an intermediate set so you can confirm all is ready to go and do comparison of new vs. old and such.

Q: how many dataset at the max can we create in Atrium?

A: You can create as many datasets as you have data sources. It is basically logical data partition and gets tracked as Attribute value. There is no artificial limit but there is a realistic number of datasets that you can have.

Q: Is sandbox common for all or configurable for some set of users?

A: It is only common for BMC Asset Management users. Basically if you're making edits to existing golden set data (BMC.ASSET) then all edits go through the Sandbox first. But you also can import into the Sandbox Bulk dataset to import larger set of data (rather than doing single edits). Bulk load is usually done with People (Person) records. In a use case example you may have a computer that needs to be associated with a person by a relationship. This "Used By" relationship would be done in the Sandbox and promoted. This would be transparent to the operator of Asset Management data

Q: What technologies are used to apply the rules to the datasets?

A: Rules are applied by using underlying database search capabilities for each dataset that is being reconciled.

Q: how does reconciliation treat numeric values such as market cost that default to 0? Are these treated as NULLs and ignored if ignore NULLs option is used?

A: Numeric value will be considered as value. Defer if NULL option won't be considered here. The value of field should be really NULL.

Q: How does the reconciliation job switch in server group if main server is down?

A: This is performed by the standard server-group feature within the AR Server where there is a 'ranking' given to each server to dictate which should take over when the current primary goes down/fails. RE activities can start where it left. It means there won't be reprocessing of CIs. In server group environment, when job stops on one server in SG, the job status goes into SGPause mode. Same job operation will be picked up by RE on other server in SG.

Q: BMC.ASSET is an intermediate data set right?

A: BMC.ASSET is the production dataset and consumed by applications to drive key CMDB use-cases.

Q: How product catalog communicates with reconciliation engine

A: Product catalog is not directly used in reconciliation engine. However, PC is used in the normalization engine to "cleanup" and "normalize" software titles in Product CIs. Doing this before RE allows RE to do better matching of products in RE.

Q: Is there any performance impact to using explicit precedence rules? Versus blanket rules for an entire class?

A: Similar question came earlier. We haven't seen any performance impact when the precedence rule at class level vs. attribute level.

Q: Is there a document that lists the expect values and formats of the tokenId field for each class? Why do standard rules use token Id for all classes when it is only supplied by ADDM for Computer System?

A: There is no such document. There is really no rule for what tokenId is. At this time, it is just a location where a value can be placed that is an "ID". You have to determine what to put there for your environment/sources.

Token id has been used by ADDM to provide a fast and quick way to reconcile computers since the rules for tokenId for computers are known and defined and followed by the products within BMC. Please refer CMDB7.6.04_DataModelingGuide for common practice on attribute value.

Standard rules are addressing generic data source and considering different discovery sources and not only ADDM.

Q: Why are the sequence of activities in Job in hundreds always?

A: They don't have to be 100's. Any number will do, but with 100's it gives room for inserting sequences in between.

Q: if one of your attributes for example in computer systems (Attributedatasourcelist) gets overwritten when merge is done, by the discovery dataset attribute, how would you prevent

A: AttributeDataSourceList is internal attribute used as part of RE merge operation. This attribute shouldn't be considered as part of your CI data.

Q: We have ADDM of 8.3.2 version and planned to migrate to ver 9.0, do we need to reconfigure Data map set again to populate data in CMDB

A: No, you do not have to reconfigure your CDM Mapping. It will be compatible as far there is no change on CMDB CDM side about attribute. Generally, there won't be any change in CDM side.

Q: is there a way to download the standard identification rules?

A: Yes, there is an indirect way to get the standard rule set. It is imported with the CMDB installation and the "ARS Transfer File a.k.a arx" has them if they need to be restored. However there is no way to get a "download" of it without the product

Q: I have about 200 000 CI. What time for merge is normal for this quantity? How to improve performance of merge activity?

A: the performance numbers will depend upon your system configuration. Just to ballpark... IF correctly configured and multi-threaded, you should get rates of 30 or more CIs per second. So, an initial load of 200K should do the job in the hours range (a couple for identify and a couple for merge).You can refer our performance white paper to apply the recommendations.

Q: What is BMC Atrium Explorer DataSet for? When we use it?

A: The Atrium Explorer datasets are sandboxes for the users. Promotion is a reconciliation job executed once the user has completed all his/her changes to push complete model changes to the BMC.ASSET dataset a.k.a Golden Dataset

Q: What infrastructure components are required to spin up the RE stack?

A: RE Stack is an application that can be run on one of the servers in a server group and requires the AR/CMDB engine running on that server.

Q: What happens if attribute precedence is 200 for class 100 and attribute precedence is 100 for class 200?

Which attribute would be selected for final asset?

A: If attribute precedence is different than class level, then attribute precedence will overwrite the class precedence even if it's lower.

In all cases, the attribute value would be from whatever source has the highest precedence. If the precedence value is the same for more than one dataset/class, then the 'tie-breaker' value will dictate which 'wins'. If the tie-breaker matches too... then the last source to be merged will produce the current value in the BMC.ASSET CI.

Q: when merge is done, it populates the import config list and when users try to make a change to CI, it will not update

A: This sounds like a precedence configuration problem.

Q: is the use of consolidated dataset in any BMC documentation?

A: Page 111 (Identify and Merge Activities) in RE/NE 7604 Guide.

Q: what field is set when a CI is reconciled? Is the field that gets set to signify when a CI is reconciled set when using each type of merge order? I have seen in my env, only one merge order, separately by class, updates the merge status field.

A: On Source, ReconciliationMergeStats will be set to "Merge Done" (enum 50) when the merge activity has completed

Q: Can we run multiple jobs or activities for the same dataset?

A: yes, all these options are configurable. Although we recommend caution before doing that.

Q: should the use of the 'cascade delete' option be standard (removing orphan child CIs) - OOTB it is not set

A: I would avoid using cascade delete option which hides the delete operation even if orphaned CIs. You might be losing some information. Let discovery tool convey which CI to delete.

Q: how would you prevent the attributedatasourcelist from getting overwritten when merge is done?

A: AttributeDataSourceList is internal field used for RE merge activity, so it will be always overwritten. It shouldn't be considered as part of your CI data.

Q: I mean can we run the jobs separately. Based on the sources?

A: Yes, each activity can be done separately either from the job configuration (by setting only one activity to Active and others to Inactive), or by creating a separate job with just one activity.

Absolutely! Data sources rule the world of Reconciliation :).

Q: Can you demo how to merge multiple datasets at the same time into a consolidated dataset?

A: you can refer the Default merge precedence set (OOTB) which has multiple datasets covered.

Q: how much important is reconciliation engine. Means if we do all these jobs in AIE itself...

A: These are radically different processes. AIE/AI LOAD data from a source. RE figures out if two things from different sources are the same and it has rules about which attributes from each are the "best". Very different processes for different purposes.

Q: Are Normalization and Recompilation dependent on each other?

A: Yes and no. They run independently and do different jobs. Normalization CLEANS the data. Reconciliation IDENTIFIES and MERGES the data. Cleaning the data, helps with identifying the data. So, they cooperate but are not "dependent" on each other.

Q: Can AI and AIE processes run on the same server?

A: Atrium Integrator and Atrium integration engine can run on the same server

Q: what type of through put is expected from the merge process?

A: The performance varies by environment. My personal experience, I have seen RE merge performance in range of 30-50 CIs/sec. For more information, please refer performance benchmarking white paper.

Q: Why is it that after merge when users try to update any CIs the updates will not work

A: RE merge activity will qualify the CI based on REMergestatus attribute status with By Class in separate Transaction merge order. The CI should have ReconciliationMergeStatus = "Ready to Merge". Whenever source CI will be updated by any means, CMDB engine updates the ReconciliationMergeStatus to "Ready to merge" to qualify CI for merge operation.

Q: The RE Console performance is so slow that it is almost impossible to create and test rules in a reasonable amount of time. Any tips for improving the Console performance?

A: Generically the performance of the RE UI is influenced by the performance of MidTier/AR. There are some web-specific tuning recommendations in the AR Performance tuning guide for BSM which should help with RE UI slowness.

Q: Is Normalization mandatory for every Reconciliation job?

A: Best practice would be to normalize CIs and then reconcile it. However it is not mandatory.

Q: if I update the attributedatasource list, then the updates will work

A: AttributeDataourceList is internal field used during Reconciliation Merge operation. User shouldn't change the value of it.

Q: Regarding only processing normalized CIs, it seems to be an unfair expectation that a customer would want to normalize every class. For example, why normalize networkcard if they are not considered via request/incident separate from the host computer?

A: this is configurable. You are in control what class is normalized.

Q: Will federation be discussed in a later session?

A: This is of course a VERY different topic and is at least an entire session. This is definitely a candidate for a future session.

Q: Here I have one issue in my reconciliation job's Merge activity; my identifies the CI & during merge activity, the jobs completes but the record count is 0 only

A: If there is a problem with the merge then the RE Logs need to be reviewed for info.

Q: Sorry the question is: How can we trigger Promotion from sandbox users to BMC.ASSET Dataset automatically when a Change request had finished (state=closed).Is it possible? Or promotion is only manually?

A: Actually, I want to step back. The best practice is to let DISCOVERY find the changes and bring them in and NOT to manually (through promote of a sandbox) update the CMDB. Discovery tells you what REALLY happened not what you hoped happen.

You want to let your discovery sources bring in the updated things as they find what is REAL. You only update things manually that cannot be discovered (or when you find something your discovery is getting wrong).

For Automatic promotion of sandbox changes, you can have Workflow in place which gets triggered when Change request state change to "Closed". For workflow, please refer Reconciliation user documentation. (Page 82 - 7.6.04 Normalization and Reconciliation guide)

Q: How can I schedule a job to run monthly?

A: Reconciliation only allows for weekly job schedules. Running on monthly basis would need to become an enhancement request. Here is enhancement request - SW00449598 for your reference.

Q: Private queue setting is depending on Number of CPUs??

A: Private queue can be set according to CPU# / Core #; For CMDB client, 390698 and 390699 are available. For RE, you can specify any of this number with min and max threads. You can set no. of threads = 2*(No of CPU*No of cores). Make a note, thread configuration varies by environment as well as other AR thread configuration.

Q: If I have ADDM (discovery job runs several time in the day for different segments) +Normalization+Reconciliation Standard job what sequence type is the better for use: Schedule type or Continue?

A: Schedule type OR continuous will be based on your data load. If it's initial load, make sure you use scheduled type. With ADDM runs happening several times a day, you could set this up in a continuous mode since the schedule might not be known in advance. If you know the schedule accurately, then scheduled jobs can be done.

How to trigger RE job on completion on NE job? How to integrate basically?

- 1) On Completion on NE job, one filter needs to write on Normalization Job Runs form.
- 2) Filter implementation should consider following points.
 1. Keep monitoring specific NEJobRunId which has status "**Started**".
 2. When NEJobRunId status transits to "**Completed**", filter will trigger RE job. Refer Reconciliation user documentation. (page 82 - 7.6.04 Normalization and Reconciliation guide)

Q: Does AI support integration to MySQL database?

A: Yes. AI support many a number of integration features. Much more than what was supported by AIE.

Q: So Normalization and Reconciliation go hand in hand?

A: Not necessarily. There is a flag in every Reconciliation Job to make the job run on Normalized instances or not. Best practice recommendations dictate that you try to use that option as much as possible.

Q: On initial data reconciliation, should you set the auto-identify to on?

A: You use "autoidentify" when you have a SINGLE source for a given type of data or where you consider a specific source as the MASTER source of data. Then, you would have merge identify rules find things in other datasets to get the identity copied to. Be careful not to set "autoidentify" to multiple datasets discovering the same things or you will get the same object from different sources with different reconciliation IDs and end up with duplicate data

Q: Is there any possibility to promote automatically outside Atrium Explorer?

A: Yes, by scheduler.

Q: I find when I use the merge order of related (commit separately or together), the merge job looks at all source CIs regardless if they have been changed or not

A: Correct. This is because of relationships related to the reconciled CI

Q: Earlier there was some CI examples that you could import into CMDB ver 7.5. They were accessible on developer community. Any plans for an 8.1 version of the same?

A: I believe you're referring to sample data. That is included with Asset Management installation.

Q: The sample I refer to is: <https://communities.bmc.com/communities/docs/DOC-6189>. Is this supported in CMDB 8.x?

A: No, that package is outdated.

Q: We have CI that are in the explorer view, if relationship is deleted, the explorer still shows the flow...Do I need to run a RE Job to clear?

A: Yes, if that is in a sandbox view then changes need to be promoted

Q: I was just trying to get a feel for what additional infrastructure components would be required for us to leverage RE. So this just sits on top of the existing CMDB?

A: Yes, RE is a part of the existing CMDB.

Q: Hi, it's possible if mal configuration server group, cmdb don't working and classes stayed in pending

A: Yes. Classes' modification can only happen in Admin server.

Q: When you recon only normalized CIs how can you review CIs that were not normalized?

A: You can go into the import dataset and search for all CIs with "Not normalized" or "Normalization Failed" normalization status flag

Q: Any Plans to enhance the Compare Activity report to be configurable or (more) user friendly? The report shows differences by InstanceId of CI rather than CI Name or Asset Id

A: We can accept an RFE to enhance the results. Compare Activity creates a comprehensible report that is easy to use but we need to understand your specific needs.

Q: is Asset dataset included in the Merge Precedence or do you have to add it? (To prevent manual updates)

A: Asset is included in the Default Precedence set of dataset. (Default Merge Precedence Set). If you setup a custom Precedence Set you need to add it.

Atrium CMDB – Understanding the Atrium Reconciliation Engine Webinar – Q&A

March 14, 2013

Q: We have duplicates in our Production Dataset, is there a way to reconcile production CIs

A: You don't want to be reconciling production CIs. Things Reconciled merge by matching Reconciliation IDs. What you need to do is to find out why things are duplicating and correct the source of the problem. Then, you can remove the duplicates and let the next reconciliation job now merge the data that has been correctly identified and not duplicated within the source.

Q: if you choose to only merge normalized data - and there isn't any values in the product catalog for it - e.g. for IP Endpoint etc., what is the approach to take to ensure that the data is made available to the Asset dataset?

A: You can have different reconcile jobs for different classes. So, if there are classes you have where you have no normalization rules for them, you can have an RE merge job for those CIs that allows merge of un-normalized data and have another RE job for those items where you do want to normalize that only merges normalized.

Q: Both the REMergeStatus and Normalization Status attribute mentioned in the Q&A are not found in the BMC_CORE forms in my 7.6.04 out of the box install.

A: REMergeStatus and NormalizationStatus attributes are internal fields and hidden to avoid the any manual update which can influence the RE and NE job operation. These field contain help internal mechanism to function and they should not be edited manually unless recommended by BMC.

Q: When and where does arrecond.log get generated?

A: By default in the <AtriumCore home>/Logs directory.

The arrecond.log is log for Recon Engine and will be generated as soon as RE process starts.

You can change location and log details in the RE Configuration Editor.

Q: Is there is any way to restrict the log size?

A: Yes, log size is configurable in the RE Configuration Editor

Q: Follow up on the Norm and RE status fields. If they are not in the view how customers are supposed to reference them?

A: you can run a report even if they are hidden. Or make them visible from Dev. Studio and use them for reference.

Q: We have already populated Datasets for BBKA & BBKA, and same data is reconciled to the Golden Dataset. Now how can we normalize the data already present in the BBKA & BBKA Dataset??

A: You can have two NE jobs for BBKA and BBKA to normalize the CIs. Normalized CIs will be updated into Golden dataset as part of next RE merge operation.

Q: Is it mandatory that manual import data must be put into a staging dataset first before we move to Asset dataset, or there is any other way we can manually import data without causing any issue?

A: It is the strongly recommended practice to only let Recon Engine to populate the BMC.ASSET dataset. You really need to know your SOURCE of data so that you can manage it. Otherwise, how do you know whether?

Q: Can I please get instructions in how to schedule a RE and AIE job to run once a month -

A: Regarding AIE and RE schedule configuration, you can refer the documentation. Use AIE User Guide and Reconciliation guide.

Q: Is there future plans to allow the CMDB process run on a non-Primary Admin server within a server group?

A: Ken, CMDB process runs on all servers, primary and not primary. Whether it used there or not depends on the configuration of the server.

Q: Can you please let me know how I can create the NE Jobs for these Datasets??

A: Please look at the Normalization console from the Atrium UI where you can create NE job. Also refer to NE documentation and NE webinar.

Atrium CMDB – Understanding the Atrium Reconciliation Engine Webinar – Q&A

March 14, 2013

Q: Can you post a typical architecture for the use of an Integration server

A: Excellent question. We don't have that ready for the webinar but will consider for future ones.

Q: BMC should index field that are used on the out of the box Identification Rules. Serialnumber, Host Name, Name, etc.

A: Indexing of fields is under constant review. We have KAs on some new changes. <https://kb.bmc.com/infocenter/index?page=content&id=S:KA381681>

Q: These missing indexes are a disservice to customers as all customers that grow their CMDB have to experience the same performance issues.

A: Well received comment. We are always reviewing these conditions for better data processing. The issue for us is that almost all customers have a different idea on what they want indexed or which feature they are actually taking advantage off. We only index things that 80-90% of our customers use.

Q: Can you please help me to configure the normalization jobs for the same??

A: I would suggest to go through' the NE documentation, it's very simple to create NE job for dataset. If you have any specific query, send question on CMDB community.

Q: Do we have any issues with CMDB 7.6.03 delete and purge activity in reconciliation.

A: Not to my knowledge. All CIs that are MarkAsDelete=Yes will be purged. Data that qualifies for delete will be deleted.

Q: If a CI is MarkAsDeleted in the production dataset, and not in the source (so it becomes 'live again' and is picked up by discovery), will the CI then be undeleted in the Asset dataset?

A: The CI will be available in ASSET till Asset Purge job is run. If it is available in source dataset not as marked as Deleted the RE job will reconcile it again, yes, and make it once more available.

Q: Is there a way to enforce cardinality rules within certain CI class to another class relationship 1:1 or 1:5 a definitive number or 1: n which would be indefinite.

A: Good idea, but not currently. The cardinality has to be compliant with the ITIL standard and the current DMTF standard.

Q: Any issues to be aware of in migrating from 7.6 to 8.x regarding the RE?

A: No issues from 7.6.04 to 8.0 for Core CDM definition. If you are keying off of extended Asset attributes from ITSM, those attributes have been removed from the CMDB so you have to adjust there.

Q: In what circumstances the BBKA CIs will be set as Mark as Deleted??

A: "MAD" is set to YES when the device is no longer discovered and absent from the BL database

Q: Are there any BMC Adaptor(s) for SCCM?

A: There is no BMC CMDB adapter for SCCM integration. However, BMC partner provides one such integration (Seamless technology).

Q: We have a desire to off-load the Normalization activity off the Primary server to a non-primary, but can't because the CMDB must remain on the Primary. We've been told if we do that, then we lose the scheduling feature in Normalization.

A: This is considered an RFE at the moment. NE will have some improvements in the next release regarding the way it works in server groups.

Q: What are the BMC Standard Reconciliation Rules for Relationships?

A: Relationships are identified by using its two endpoints. The parent and the child.

Q: KA376464: How to increase Performance in Reconciliation Jobs (with SP3 or higher) - does it mean this is not available in SP2?

A: We have made couple of product fixes for performance reason which are available as part of SP3. I would suggest to go with latest SP (i.e. Sp4) if you plan to upgrade.

Q: When creating a dataset - only admin can see the dataset. (The cmdbwritesecurity field on the cfg: dataset form is empty)

A: Dataset creation is an administrative task. There is no restriction on visibility of the data out of the box that's different from the Asset Management guidelines

Q: how do you remove a former staff's sandbox correctly without impacting anything?

A: The sandbox is an overlay so by removing it you're not impacting the underlying data.

Q: If the normalization engine uses the product catalog and say a model is not in the product catalog or alias, what happens? Is there a way to trigger a change or notification process to allow a new model to be added to the product catalog?

A: You can configure things to either "autocreate" an entry into the product catalog for a new item (only desired from a trusted source) or you can find things that are "unnormalized" and review them for handling, in this case creating the product manually.

Q: Is there no longer an AssetLifecycleStatus field in the CMDB? How do we adjust mappings in AI for jobs that set the AssetLifecyCleStatus?

A: Certain non-discoverable attributes have been moved to the AST: Attribures form and only available from Asset forms. If your AI job needs to set that attribute you need to retarget your AI job to point to AST: forms.

Q: How much time will take 4000 endpoints to Reconcile from ADDM to CMDB

A: It depends greatly on several settings and hardware limitations. You can expect around 30 CIs per seconds being identified and then merged into production dataset. 4000/30 gives you around 3 minutes.

Q: Are there any plans on supporting overlays with Atrium CMDB?

A: If you mean supporting overlay datasets as source or target of reconciliation engine... That is a problem that has already been resolved and is available in a patch of 7.6.04 and should be in the 8.1 and later releases.

Overlays in CMDB forms is not expected to be supported.

Q: Is there any plans to create a SQL adapter for UNIX, currently it is only available on windows.

A: I am sorry, what type of SQL adapter? There are already several in place within our products. If you refer to AIE, then no, AIE does not support SQL. You will need to migrate to AI who does support.

Q: In case of a mistaken process (wrong rule definition) and CIs went into the CMDB twice.

Is there a way to identify the duplicates and remove them?

A: Yes, there are several ways to help you identify duplicates. Please submit a ticket, we have a well-defined process that we're working on making public eventually. For now this is done on need per case.

Q: ARERR[100029] Please give values to all required fields.

Merging of record failed Class BMC.CORE BMC_ComputerSystem

Atrium CMDB – Understanding the Atrium Reconciliation Engine Webinar – Q&A

March 14, 2013

A: You must have a field that is left without a value. SQL/Filter/API log will show why.

Q: What is the basic difference in purge and delete?

A: Purge only removes MarkAsDelete =Yes. Delete gets rid of everything in the dataset

Q: What do you typically see around resources used to manage the RE activities? I am wondering if depending on size/load issues occur if multiple people are managing it, or better practice is to have 1 person primarily responsible.

A: It would be very unlikely for one person to be responsible, although possible. There are no issue with multiple people handling Reconciliation.

Q: We removed a person's sandbox but now I am getting this ARERR[101049] Dataset <Sandbox for rlardy> does not exist Activity ended with an error DADS Merge 3

A: You need to check your Merge Precedence Set and remove the dataset reference from there.

Q: In previous versions the log file "arrecond.log" grew and grew and grew (there was no way to limit size this unlike the RE Job logs). How can we can this be configured rather than manually check it every few months (it grows into many GBs)

A: arrecond will grow forever yes. We are working on improving this and providing a backup and archiving mechanism.

Q: If our reconciliation does not give us the results we expected or really screws up our gold dataset, Is there a document on rolling back or best practice for CYA?

A: If you are following best practice and you never manually update the production dataset and you always have your sources come into separate source datasets and those are preserved, you can simply delete all data in the production set and let reconciliation run again with corrected rules. Set the option to reconcile EVERYTHING not just changed things.

Now, if you have not followed good practice and some data exists now only in the production dataset, you can of course not do this and have to be much more careful about cleanup as you have data there that you cannot delete or you lose it.

Q: Are there a Best Practice / White Paper for Relationship classes, I know of the "Data Modeling" chapter in the Data Modeling Guide

A: Yes, there is a document: Look for: Mapping Your Data (Excel sheet)

Q: Answer to Dataset permission question: For non admins to see a dataset the CMDBRowLevelSecurity needs to set to public in the BMC,CORE:BMC_Dataset from

A: correct