

# Presentation References

## Duplicate Data

<https://docs.bmc.com/docs/display/public/ac91/Investigating+issues+with+duplicate+data+in+CMDB>

<https://docs.bmc.com/docs/display/public/ac81/Removing+duplicate+computer+systems+from+the+CMDB>

<https://docs.bmc.com/docs/display/public/ac81/Modifying+reconciliation+identification+rules+to+avoid+duplicates>

## AR Log Analyzer tool

<https://communities.bmc.com/docs/DOC-25473>

## Installing

<https://docs.bmc.com/docs/display/public/ars91/Installing+AR+System+Log+Analyzer>

## Running

<https://docs.bmc.com/docs/display/public/ars91/Running+AR+System+Log+Analyzer>

## Analyzing

<https://docs.bmc.com/docs/display/public/ars91/Viewing+AR+System+Log+Analyzer+output>

## Investigating

<https://docs.bmc.com/docs/display/public/ars91/Investigating+slow+BMC+Remedy+AR+System+API+calls>

## Recommended Hotfixes

<ftp://ftp.bmc.com/pub/ARRecommendedFixes/>

---

Q: Can we write an AI job to load the BMC.CORE part, run reconciliation, check reconciliation status and once this has completed read the ReconId and update AST:Attributes record with correct data?

A: Yes, you can do that. In fact ITSM already ships an OOTB job to do this. "CI-CMDB".

---

Q: I'm currently on 7.6.04 (We are going to 9.X in a year +/- after engineering processes). I need to change all current Names and other data to lower case will the reconciliation engine be able to do this? Slides indicates this is possible in 9.1

A: You can use Recon Engine for this change, however, if you know what all fields to be changed and it is for all CI type, then you can either use Atrium Integrator/Pentaho based job or Direct SQL approach which will be more efficient. Normalization Custom Rule is also more appropriate for this kind of rule enforcement.

---

---

Q: I recall I can start reconciliation job by writing to some AR form but how should I check if the job has completed?

A: You can use Reconciliation Job Runs form for the purpose of checking a given job status.

---

Q: What is best practice for modifying the standard identification rules? Is it recommended?

A: As mentioned in the webinar modifying a standard rule set is a decision that you need to consider. There are clear exceptions to the problem, for example if your data source is just one dataset, then you can go ahead and do modifications to the STD ID Rule set, but keep in mind that rule you modify as part of that STD Rule -will be modified across all jobs that use standard rules. It depends on your business need. Yes, we have provided std rules, but it is okay to make change to fit your business requirement. As long as the change in the standard rule is considered a change for all standard of your business needs.

---

Q: About Precedence: First review Dataset Merge Precedence and then Precedence value in Precedence Group, don't it?

A: It tries to indicate that Merge Precedence and Precedence Group is very important part of Job configuration and it is important to configure right. Otherwise, you may see data overlap problem.

---

Q: So I can use SQL to change information in the database without causing problems on other forms?

A: This depends on what you're modifying. Each change to data should be "taken action" on by the workflow. So if you're making a change directly through SQL then you're bypassing that workflow. You can do that. Just to be cautious with this approach. In the case of data cleanup that is we do not recommend deleting with a SQL. You should mark your CIs with a SQL but delete with AR Escalation/RE Job/Etc as mentioned in the webinar so as to run the workflow.

---

Q: Can I force re-identification after ReconId has been assigned but attributes used to determine it changed?

A: To force Re-identification, you will have to reset ReconId to 0 for all CIs and associated relationships. Make a note - Relationship holds ReconId for End points which have to be reset as well. You can use CMDBDiag tool (cmdb\server\bin) for this purpose.

---

---

Q: What forms can we use in order to clear out old reconciliation data? Our last upgrade a lot of old data was migrated and we would like to clean this up.

A: Great question, there many ways to do it, as Gustavo is explaining now. The form can be BMC.CORE:BMC\_BaseElement and you can use the ModifiedDate as your scope. Basically anything that is older than a <DATE>. You can mark it soft deleted first and Purge. You can also use Delete activity with a ModifiedDate <= "1/1/2014" qualification for BaseElement and BaseRelationship

---

Q: Sometime we see Identification Job could not identify the CIs and have Recon ID stays are 0

A: In this case, you may have to review RE job debug logs and RE ID rules to make sure your ID rules are aligned with Data.

---

Q: In Asset I can use SQL to modify the data in the database without causing problems on other forms as I would if I made a change on the People form?

A: Yes, you can do that. People data is tightly integrated with BMC\_Person CIs. If you make change in People form, make sure it is matching with BMC\_Person CI. Otherwise, it requires change in BMC\_Person as well.

Note that CTM:People is managed by Asset Management. CMDB would only reconcile the BMC\_Person records

---

Q: We use spoon to bring LDAP people into the system. The data is in a staging form. When every the data is updated it creates multiple updates in BMC Person. Many times the people are duplicate CI's. For the most part we have to delete people from BMC Person

A: CTM:People form has OOTB integration with BMC\_Person form. This integration can be configured to reconcile in bulk. You will have to change this configuration and then run RE Sandbox Bulk job to reconcile all BMC\_Person CIs together.

---

Q: What is the recommended and minimum intervals that can be set for continuous jobs?

A: The minimum interval for continuous job is 120 seconds

The recommended number depends how frequently you want to have your data updated in CMDB. If you are looking for very real time data, you can keep it very low interval.

Doc reference -

<https://docs.bmc.com/docs/display/public/ac81/Creating+a+continuous+reconciliation+job>

---

---

Q: RE ADDM and duplicates: I can see a lot of them when ADDM created a new key for slightly changed Host and then started to age out the old one. Before it aged out and was deleted both are synced to CMDB.

A: Yes, in continuous sync, you will see old CI gets soft deleted.

---

Q: If we don't want new assets created in the golden dataset what are all the settings we need to configure? Our customer would like to manually review these before they are created.

A: To get some time to review data before gets reconciled, you may want to have a RE job with Identification and Merge separately. After ID, you can have data reviewed and then merge them. Mainly, the feature you want to avoid is using the "Generate IDs" check box in the identification activity. With this box disabled RE will only merge updates (CIs both in source AND target). New CIs will not be given a Reconciliation Identity. Which means you will have to figure out some other way to give them an identification.

---

Q: When I manually (or via AI for instance) "zero out" the ReconID do I also need to zero out all relationships for this CI? In 8.1 SPO I have to, otherwise merging of the previously identified relationships fails (invalid ReconID of the parent/child)

A: Yes, it is a good idea to do that when re-identification is required. The cmdbdiag tool can "zero" them out as well with limited qualification i.e. you can select class/dataset, but complex qualification.

---

Q: What does gsql mean?

A: Get List SQL, it's an AR driver command that allows to run direct SQL from AR into the DB and print the results on screen.

---

Q: Does the cmdbdiag tool, when attempting to delete orphaned CIs, allow you to use a qualification? Such as, dataset id?

A: CMDBDiag allows you to qualify your option for DatasetId, Classid, but it doesn't support complex qualification.

---

---

Q: In the presentation you are referring to both ac81 and ac91 documentation. With a 81 installation does that mean that the ac91 documentation it not related to version 81?

A: ac81 is documentation for 8.1 version and ac91 is for 9.1 version

For Recon Engine THEORY perspective, ac81 and ac91 just represent different version, there is not major product change between these two versions except product improvement. Meaning the 9.1 documentation applies for 9.1 RE as well as for older versions. Sometimes if there is a new feature in the product you will only see it documented in the latest documentation only.

---

Q: Are there plans to include recon logs in some AR form?

A: I would suggest you submit Idea on BMC Community for this.

---

Q: One or 2 jobs in our environment are never showing completed status rather shows SGPaused when checked at RE:Job\_Runs format at Atrium Core console it has only start time and no End time and kind of greyed color instead of green or red.

A: That is usually related to your Server Group configuration via the AR System Ranking form where RE is set to failover to another system because AR Server stopped responding to RE. SGPaused status is showed when server has failed over in server group while the recon job is running. You should make sure your system is configured properly and system has enough resources to be able to process all actions. There are Knowledge Modules on this as well as community posts. Look for Fine tuning keywords.

---

Q: Are there plans to make recon be a failover service like NE in a future release?

A: Reconciliation has always been setup for failover. NE failover feature was recently added because it was missing. Ranking for and suspended flags are taken into account when reconciliation starts or checks its metadata with AR Server

---

Q: Do we need to give the thread count based on the CPU only? What will happen if we give more than the CPU count?

A: Not necessary. Use of CPU in thread configuration is to simplify the formula but that is not always an answer. Too many threads at AR server level would cause context switching and overall slowness. I would suggest you to review overall thread count on AR server and make an adjustment.

---

---

Q: Regarding DB performance, does BMC have an expected # transaction per second that is deemed as a good or normal range?

A: Our best measured records/second measurement in a controlled environment was around 90 CI's/sec. System had 32 Gigs of RAM and 8 CPU Cores. The worst we've seen was 1 CI/sec. That was a VM machine with low resources. I would say, if you experience 20-30 CIs/Second, that's good performance for RE. Consider there is Identification as well as Merge performance. Both can be quite different.

---

Q: What is the best method to find the errors in the Recon Job run instead of running through 100s of log file generated for a single job?

A: You should keep the job log level to ERROR. With that setting, every error counts. If you have too many errors with similar message then you can use regular expression to "count" how many times the error appears and then see if they have a common threads. Then you may use the method thought in this webinar to resolve them.

---

Q: is this possible to concurrently run multiple recon jobs with different source datasets but the same target dataset but with not overlapping CIs?

A: Yes, it is. There is a setting in AR.CFG that allows "Parallel RE Jobs" to run. However, as you said, the data must not overlap in any way. There are some concurrency issues running parallel jobs which are complex. Like for ex. Two CIs, one in each source, but with no CI in target to identify against. If you run ID job for both datasets, probably both CI will receive an ID, but not the same ID. Later when you run merge, both CIs will be pushed to the target as DIFFERENT CIs. That is why sometimes it's better to use an intermediate dataset or to run RE jobs in serial mode and not in parallel.

---

Q: Do we need to follow the same count for VM and Physical system CPU cores ?

A: Typically there is up to 20% performance difference between VM's and real systems. Real system will obviously perform better because do not have to share its I/O or disk space allotments

---

Q: Is there a way to normalize two different classes (CI-Type) like BMC\_Product and BMC\_Softwareserver, when ADDM discover a SW-Product with the same "Model" and "Manufacturer" ?

A: Not currently, but it is under consideration.

---

---

Q: With queue 390698 and 390699 setup, you recommend 3 and 5 threads per cpu for each queue?

A: Private queue configuration goes 2-3 thread per CPU, however it is not always the case. It is important to see total thread count on AR server and make necessary adjustment.

---

Q: What is the recommended Recon log file size which would help in analyzing in case of multiple files generated for a job? Default value seems to be 300KB

A: With ERROR level, 300 KB is a lot of errors and if you use DEBUG level and you're creating multiple logs then you'll have hard time decoding what the log is trying to tell you. 3000 KB would be the max I'd recommend and keep it to ERROR.

---

Q: Can data be pushed directly to the gold dataset using Spoon?

A: You can, but that is a terrible practice because you cannot validate the quality of the data in staging dataset. Always use a staging dataset to push data to.

---

Q: Is AR Log Analyzer a paid/Licensed tool from BMC or free?

A: It is not a licensed tool and available for free. The utility and support is provided through BMC Communities - <https://communities.bmc.com/docs/DOC-25473>

---

Q: In 9.0 the AST:Attributes is populated when the CI is created by AI jobs in staging datasets, compared with 8.0 when the AST:Attributes entries were created after the recon

A: That's correct. In 9.x all CIs are added to AST:Attributes and use instanceid for InstanceID and ReconciliationIdentity until the CI is identified and merged into BMC.ASSET

---

Q: What is the name of the OOTB UDM job that loads CIs and Reconciles them?

A: You can find the list here:  
<https://docs.bmc.com/docs/display/public/itsm91/Transactional+data+mapping>

---

---

Q: After I merged any attributes from Dataset "A" they are kept in AttributeDatasourceList and cause Errors in recon if that "A" dataset is not part of the merge precedence set. May I remove entries from AttibuteDatasourceList manually?

A: Again, you can do such thing, but the question really is, why has the data changed? Why is the same CI coming from different dataset? Reconciliation by definition means to compare the weight of the data and make the result 100% reliable (in BMC.ASSET). Best practice here is to add DatasetA to the merge precedence set so that the weight does not get ignored.

---

Q: Can we delete Data from a regular form using AI job directly?

A: We don't have a plugin to delete data directly from a regular form. If a form representing CMDB class, then you can use logic of AI (via Spoon) to soft delete i.e. mark data with MarkAsDeleted = Yes which will be hard deleted later in the life cycle.

---

Q: With the mandatory version on normalization does that mean you need an unknown for a version for all things coming from Discovery

A: Version is not mandatory. Version starts being mandatory for a specific product when you start populating versions for that product. So if for product A you populate version X. Then CIs of product A need to either have version X or will be rejected. If they do not have a version and you want to ALLOW those CIs in, you can create another version for product A, called UNKNOWN.

---