

Q&A Session for Connect with Remedy Webinar- BMC Discovery CMDB Synchronization Performance Tuning webinar

Date: Wednesday, February 22, 2017

Q: What is a typical sync throughput?

A: There is no "typical" throughput as it is very contextual. On a high-performing AR system, with 8 concurrent workers set for the Sync in Discovery 11.1, we have seen 1000-2000 discovered Devices/hour during an initial synchronization to an empty dataset.

Q: Regarding incremental resync, how does it work exactly?

A: See <https://docs.bmc.com/docs/display/DISCO111/Resynchronizing+a+CMDB+connection>

In the first phase of a resync, BMC Discovery works out all the changes to CIs and relationships required to bring the CMDB dataset into synchronization with Discovery. In the second phase, it makes all the required changes. In a Complete Resync, it makes all the changes immediately; in an Incremental Resync, it remembers the required changes but does not immediately apply them. During normal synchronization activity it then applies the remembered resync changes as it considers the CIs and relationships for that sync.

Q: Can batch size be increased to improve performance?

A: It can be increased, but we would not normally expect it to improve performance. All environments are different, however, so sometimes larger batches may lead to a small improvement. You should experiment to find a good batch size. Normally, if changing the batch size is required at all, it is to reduce the batch size to prevent timeouts when using a slow CMDB.

Q: Is the number of concurrent workers limited by AR Server specs/configuration?

A: Discovery permits any number between 1 and 20. It does not read the AR Server configuration to determine the limits. The best number does depend on the AR server specification and configuration, but also on the network between Discovery and the CMDB, and other factors. If you have a high latency network, for example, using more workers in Discovery means that multiple requests can be in flight across the network simultaneously, meaning you obtain higher throughput even if the AR configuration means it only processes one call at a time.

Q: What version of Discovery (ADDM) is being shown in the slides?

A: It's 11.1 patch 2 (11.1.0.2). We always recommend using the latest version, which is currently 11.1 patch 3.

Q: In the filter we just have radio button and it doesn't ask for save when a filter is set and radio button is selected.. would have been nice if we could have a prompt or confirmation message when a radio button is changed....

A: This is covered by an existing defect; DRUD1-16219

Q: is it possible to extract a report csv or any format for the devices and relationships updated in the CMDB in the last sync or synced for a particular time period?

A: Yes, you can create a report through midtier using BMC.CORE:BMC_BaseElement or BMC.CORE:BMC_BaseRelationship where the dataset is BMC.ADDM and the Modified date is greater or equal to today and the Modified by is NOT Remedy Application Service. In Discovery, you can query the root devices based on when they were last successfully synchronized. To find all the CIs and relationships, you would have to look at the shadow copy store in the Discovery data store.

Q: Is "concurrent workers" in 11.1.x the same as 'resync threads' in 10.2?

A: During a resync, concurrent workers is the same as the previous resync threads option. During normal sync activities, concurrent workers allows multiple root nodes to be transformed and synchronized simultaneously. Previous releases always processed one root node at a time.

Q: I have to perform a resync after I additionally set a filter?

A: No, unless you want to see the previously synchronized CIs removed by this filter immediately, you don't have too. If you don't resync after changing filters, then previously-created CIs and relationships will be marked as deleted the next time the corresponding root node (e.g. Host) is synchronized. The filter applies to each and every nodes that are then synchronized. You can also chose to manually synchronize all nodes with the new filter.

Q: Does "using private queue" mean "sharing reconciliation/normalization queue" or can we have additional queues exclusive to CMDB sync?

A: At present, we have only two CMDB private queue ports available for CMDB client I.e. 390398 and 390699. In future release, we have considered broadening this range so CMDB clients can utilize separate private queue port.

I would suggest look at the load on ADDM CMDB Sync, Normalization and Reconciliation and access which components can be shared on same private queue port.

Usually, I have seen Recon Engine configured on 390698 port and ADDM CMDB Sync & Normalization Engine configured on 390699 scheduled without operation overlap.

More information is shared on community:

https://communities.bmc.com/community/bmcdn/bmc_atrium_and_foundation_technologies/bmc_atrium_cmdb/blog/2016/10/17/what-to-look-for-when-synching-addm-data-to-cmdb-in-ar-server-group

Q: Is BMC Planning to modify product for writing queries which work better for performance? For e.g. queries for normalization uses where clause as NormalizationStatus NOT IN (50, 60), which can be made better with equal to or OR clauses.

A: There are only 7 possibilities:

Other;10

Not Normalized;20

Not Applicable For Normalization;30

Normalization Failed;40

Normalized but Not Approved;50

Normalized and Approved;60

Modified After Last Normalization;70

The NOT IN or IN clauses

Typically, I have seen IN clause performs better than NOT IN. Again, it depends on the record volume in the table too. I would pull out the SQL execution plan to take next action. If you need BMC's assistance, please raise support ticket.

Q: How my IP Range(s) impact performance?

A: If you scan more IP addresses, it will take longer to scan. If those scans find more devices, they will take more time to synchronize to the CMDB than if you had not scanned them.

Q: Does Custom Patterns run during Resync process?

A: yes, if they are custom sync mapping files

Q: What does it mean when the AVG time of the API is high.. but the SQL statements are low AVGs?

A: It depends. APIs may have multiple SQLs. The addition of all SQLs will summarize the API time. Sometimes there is a delay in server itself, in those cases API takes longer than SQL, so it normally means the server is overloaded (maybe you need RAM/CPU).

Q: in API/SQL logs how to recognize a particular API/SQL statement is part of CMDB sync (I understand the part about private queues one slide back)?

A: In AR API, SQL log, you can look the user and Client-RPC to filter out the logs.

e.g. here, ADDM CMDB Sync is not using CMDB Private queue port. So, all CMDB API from CMDB Sync operation will be executed using 390696 and user name. In case of private queue configuration, you will see 390698/390699.

<Client-RPC: 390696 > <USER: Demo >

Q: Is there any such things as a performance dashboard for CMDB Sync on the AR Server Side?

A: At present, we don't have Performance dashboard for CMDB sync activity.

Q: how about versions? What's the recommended optimal version (ADDM and AR server)

A: We always recommend using the latest available versions and service packs of all our products.

Q: are there any levels of performance you guys see as red flags?

A: As long as you are able to sync at least as fast as you scan, performance is ok. If you have restricted sync windows and are not able to empty the sync queue during those windows, that's a red flag. Occurrences of timeout errors of type ARERR(91) are also a red flag.

Q: Can we query the ADDM data? for third party application to use to minimize querying CMDB

A: Yes. Using the REST API, for example. <https://docs.bmc.com/docs/display/DISCO111/Using+the+REST+API>

Q: Since at the CMDB DB level, the same tables are being used for Recon, Norm, and Syncs, is the general guideline to avoid two of these operations running at the same time? What if you have continuous Normalization enabled?

A: You will have continuous enabled once you did the initial data loading. Deltas should be small so you should not have a problem. If you re-sync ADDM or during initial load you should NOT have NE set to continuous.

Q: Do you have any recommendation to improve performance / turnaround time while we open the CAM / SAM / CI in CMDB Atrium Explorer (synced with CMDB) it takes too much of time to open all the CI in explorer

A: It is hard to say. We should analyze logs during Atrium Explorer activities to see where the time is being spent.

Q: Is the AR Server Log Analyzer available for version 90? What is the link?

A: <https://communities.bmc.com/docs/DOC-2973>

Q: Does using NAT between BMC Discovery and CMDB has any noticeable performance impact? Has this been tested before?

A: No noticeable impact.

Q: What Database ADDM use?

A: A proprietary NoSQL graph database that we implemented.

Q: Is it more efficient to Filter on the ADDM or CMDB side of the CMDB sync filtering configuration? If within ADDM, will there eventually be more filters (i.e. SoftwareInstance nodes)

A: The filters do different things. In the "Discovery" part of the filter, you get to choose which Hosts and other root devices to synchronize at all. E.g. you can use it to sync just Hosts running Windows. In the "CMDB" part of the filter, you can turn on and off particular CIS, based on class and data values. E.g. you could choose to only sync BMC_SoftwareServer CIs where the Manufacturer is BMC.

Q: In the continuous sync in ADDM , if the connection breaks for sometime & come back again then will pick up the pending discovery items?

A: Yes. In all versions, during normal sync operation, Discovery maintains a queue of work to do, and pauses the queue if the CMDB connection is temporarily broken.

During a resync -operation, if the connection breaks for more than a very short period of time, the resync stops. An enhancement in version 11.1 means that you can resume the resync where it left off. In previous versions, you had to restart the resync from the beginning. See <https://docs.bmc.com/docs/display/DISCO111/11.1+Enhancements#id-11.1Enhancements-CMDBresynchronizationimprovements>

Q: In addm 10.2 you have the synchronization with impact relationships. why we don't have this data model in ADDM 11.1? or the synchronization with impact attributes is that the same function?

A: Impact Relationships were deprecated in the CMDB. They are converted to the newer Impact as Attributes. The data model is still present, just hidden. If you have a truly ancient CMDB version, you could still use it, but really you should not be using the deprecated BMC_Impact relationships.

Q: If using Remedy on Demand, can we simply ask BMC Support to monitor our CMDB Sync activities and ensure our AR Server/CMDB is configured correctly for best CMDB Sync performance?

A: Sure, you can certainly do that. RoD Support team should be able to help here.

Q: With v11.1, does clustering improve CMDB sync performance? So, is one possible way to improve performance is to add additional appliances to the consolidation cluster? Is one big appliance better to do CMDB sync, or 2 smaller appliances in a cluster?

A: Clustering won't really help the CMDB Sync performance. But, it helps scanning performance. (CMDB Sync is only performed by the "coordinator" member of the Cluster).

Q: About CMDB Sync Blackout window...is it recommended to blackout during the CMDB Normalization and/or Reconciliation?

A: If you run NE and RE as batch jobs (and not continuously/inline), then the answer is yes. Better not to Sync ADDM during this time.

Q: Is there any recommendation for normalization during the synch? inline v/s batch v/s continuous?

A: For Initial load, we recommend Batch Normalization. For incremental load, we recommend continuous mode to process data. If system has some resource constraint where NE can't be run in continuous mode, Batch NE can be used to normalize CIs in Blackout window. We don't recommend Inline normalization as the normalization failure would block CI to create.

Q: We cannot use 390696 to ADDM for synch and sort operations and leave the others for recon and Normalization?

A: See

https://communities.bmc.com/community/bmcdn/bmc_atrium_and_foundation_technologies/bmc_atrium_cmdb/blog/2016/10/17/what-to-look-for-when-synching-addm-data-to-cmdb-in-ar-server-group

Q: Is there worker threads configuration best practices documented for Discovery 11.x? Ex. = xCPUs = suggested X threads.

A: It is too dependent on the CMDB performance and configuration for us to give any fixed guidelines. There are many factors including the network latency between Discovery and CMDB that can affect the best number of workers.

Q: if I set 'inline normalization' on the dataset which Discovery syncs to I suppose it is not a good idea to sync on the normalization private queue (98)?

A: Private queues are not always the best solution. You need private queues to not interfere with user queues. If you have an integration server you may not need private queues for NE/RE.

Q: Sometime we see that the connection with the CMDB is interrupted from ADDM side: "Error retrieving server details" in our prod env. How to check why? Servers are communicating with each other network wise.

A: You can check exporter logs to look for specific error code reported by CMDB/AR server. Also see this post on communities in case it applies, as there is a diagram in the replies that shows how to manage the queues https://communities.bmc.com/community/bmcdn/bmc_atrium_and_foundation_technologies/bmc_atrium_cmdb/blog/2016/10/17/what-to-look-for-when-synching-addm-data-to-cmdb-in-ar-server-group

Q: Case: Migration to a new ADDM environment. What should one do: drop the BMC.ADDM dataset data and start with the fresh env OR keep the dataset, export+import root_nodes_key into new env and start the new env from that? (Background: which "sync" is faster?)

A: The root node key export / import happens with Discovery data, not CMDB data. If you have the old Discovery environment, then, yes, exporting the keys would allow you to keep the old dataset in the CMDB. It would probably end up being faster.

Q: if you want the ADDM dataset to reconcile with the ASSET dataset is it best practice to create the ASSET first and then run ADDM or let ADDM create the ASSET and edit it after?

A: OOTB ASSET dataset is created in CMDB. Irrespective of source, it is always best practice to get external store data into CMDB in the working dataset. Depending upon your use-case, the dataset needs to run through normalization. Then, the dataset has to be merged into BMC Asset. You should not load data directly into BMC Asset i.e. Production dataset which may cause data problems.-

Q: Will there be CMDB sync performance issues if you have a multimember (6) consolidator cluster? If the coordinator is the only one performing the sync and the datastore is spread across all members wont sync be more resource intensive?

A: A cluster consolidator will share the load of everything else within members. The synchronization won't be slower overall because the main time spent is not on discovery side but on AR side.

Q: I read somewhere to avoid "Resync" - which would mean not root_node_key export /import, better to just start with the fresh ADDM CMDB synchronizing into an empty CMDB dataset

A: You should not "avoid resync" just for the sake of avoiding it. It takes some time, so you shouldn't intentionally get yourself into a situation of needing it on a regular basis, but as a one-off operation, it's fine.

Q: Hello, is it possible to have more Discovery appliances configured to one CMDB Dataset? Let's say we have 2 test appliances and we want them to sync data to same dataset - is it possible?

A: No. Each Discovery instance thinks it totally owns the contents of its dataset. They will fight with each other if you share the dataset between multiple Discovery instances. Just create more datasets.

Q: similar onDemand question - if a DB Copy/Refresh or Upgrade were occurring, would we expect onDemand support to carry out Root Node Key Export or is this an ADDM administrator activity?

A: Root node key export / import is a purely Discovery activity. It is only for the situation that you are replacing one BMC Discovery system with another one and want to retain the data keys. It has nothing at all to do with situations in which the Remedy / CMDB database is changed, and is therefore out of scope for the Remedy On Demand team.
