

2019-05-16 AAI Developers Meeting



Attachments (including meeting recording)

File	Modified
No files shared here yet.	

Status	OPEN	IN PROGRESS	ON HOLD	DONE	CANCELLED
--------	------	-------------	---------	------	-----------

Agenda Items


START RECORDING

	Title	Responsible	Status	Last discussed	Notes
1	Unreleased docker images in helm chart	Team	OPEN	2019-05-16	Merge change if artifacts are released in timely fashion, need to add release notes (or documentation somewhere) on the various options in the helm charts, things like replica count for cassandra local/shared, enabling/disabling sidecar, aaf, etc.
2	Optimise ad-hoc retrieval of any resource with any of its cousin relationships	Keong Lim	OPEN		<div>Raised as JIRA case</div> <div> AAI-2437 - Optimise ad-hoc retrieval of any resource with any of its cousin relationships CLOSED</div> <p>but holding back the gerrit review until after the Dublin branch is cut so it doesn't count against release criteria.</p> <p>Would like to gather opinions on the implementation (see the comments in the JIRA case). The query has been tested in my dev environment but unsure if the query is done the best way.</p> <div><input type="checkbox"/> SoumyaNaik is invited to a future session to discuss further</div>
3	Dublin Branching	James Forsyth	IN PROGRESS	2019-05-09	<p>El Alto runs Jun - Sep - technical debt and S3P/deployability release</p> <p>Frankfurt starts in Sep</p> <p>As of now, master will accept feature development for Frankfurt</p> <p>We tag dublin at release time (early June)</p> <p>The Dublin branch becomes El Alto (similar to a Dublin Maintenance Release)</p> <p>El Alto changes will be cherry picked back to master</p>
4	Review El Alto proposals				AAI Epics in JIRA: 11 issues
5	Cached vs realtime connection	Keong Lim	DONE	9th May 2019	<p>Have been testing with Traversal custom queries and Resources bulkadd API.</p> <p>It appears that the caching causes a delay of approx. 1 to 10 seconds between creating new objects in Resources API and observing them in Traversal API (seems to have 2 or 3 modes which are fairly consistent).</p> <p>What is the reasoning for using cached rather than realtime connection mode for some clients but not others? Cache being out-of-date would seem to be a cause of problems.</p> <p>There is some discrepancy in the default configs for AAI microservices:</p> <ul style="list-style-type: none">■ In graphadmin/src/main/resources/etc/appprops/aaiconfig.properties: Line 47: aai.realtime.clients=RO,SDNC,MSO,SO■ In resources/aaai-resources/src/main/resources/etc/appprops/aaiconfig.properties: Line 92: aai.realtime.clients=RO,SDNC,MSO,SO■ In traversal/aaai-traversal/src/main/resources/etc/appprops/aaiconfig.properties: Line 69: aai.realtime.clients=SDNC,SO,MSO <p>Apparently the HTTP header "Real-Time: true" can be used to force a realtime connection. Would this refresh the Traversal cache so that all clients can immediately observe new objects? (Further testing shows the cache is not refreshed; should this be raised in a JIRA bug?)</p> <p>Is there some way to know how far behind a cached connection is from realtime? e.g. by observing the latest resource-version value? Would want to avoid polling the microservice.</p> <p>Assuming that the "Real-Time" header has not been specified, is there any part of the response that indicates whether a "Real-Time" connection was used by default or not?</p> <div>9th May:  AAI-2446 - [Code Cleanup] - Deprecate Cached mode connections in AAI microservices CLOSED</div>


6	debugging Traversal custom query problem	Keong Lim	DONE	10th May 2019	<p>Calling Traversal custom query with curl:</p> <pre>curl --verbose --silent --insecure --user AAI:AAI --header 'Accept: application/json' --header 'X-TransactionId: testaai' --header 'Content-Type: application/json' -- header 'X-FromAppId: AAI' --header 'Real-Time: true' -- upload-file data-customquery-count-vnf.json https://<IPADDR>:30233/aai/v14/query</pre> <p>Request body for PUT is:</p> <pre>{ "start": ["network/generic-vnfs/generic-vnf/test-bulkadd-gvnf-4"], "query": "query/count-vnf-byVnfType" }</pre> <p>Response returns error:</p> <pre>{ "requestError": { "serviceException": { "messageId": "SVC3002", "text": "Error writing output performing %1 on %2 (msg=%3) (ec=%4)", "variables": ["GET", "v14/query", "Internal Error:java.lang.ClassCastException: java. util.HashMap\$Node cannot be cast to org.apache.tinkerpop. gremlin.structure.Vertex", "ERR.5.4.4000"] } } }</pre> <p>Stored-query definition appears to be:</p> <pre>"count-vnf-byVnfType":{ "stored-query":"builder.getVerticesByProperty('aai- node-type', 'generic-vnf').groupCount().by('vnf-type'). store('x').unfold()" }</pre> <p>How do I debug/fix this query?</p> <p>9th May: Venkata Harish Kajur could you take a quick look at this and email me?</p> <p>Keong Lim This custom query seems to only support one format which is the console. The reason is the output of this query just return a list of counts grouped by vnf type. It doesn't make sense to be able to support any other format. The default format I believe is the graphson so thats why its failing. Add the query parameter format=console and it should fix this issue. We should probably update the documentation to say that this custom query will only be working with format=console</p> <p>10th May: added to</p> <div>  AAI-2408 - [DOCUMENTATION] - Documentation on the AAI Traversal custom query API is incomplete CLOSED </div> <p>. Also raised</p> <div>  AAI-2447 - Examples API for AAI Traversal custom query API CLOSED </div>
7	GraphGraph demo	Pavel Paroulek	IN PROGRESS	2nd May 2019	<p>A 5-10 minute demo of GraphGraph</p> <div>  AAI-534 - Tooling and utilities in support of providing visualization for the AAI model, edge rules, and schema CLOSED </div> <p>.</p> <p>Feedback needed!</p>

8	AAI Modeling Multi-part key for schema elements	Keong Lim	IN PROGRESS	25th Apr 2019	<p>From https://lists.onap.org/g/onap-discuss/topic/31317665</p> <p>Discussion:</p> <ul style="list-style-type: none"> <input type="checkbox"/> The cloud-region schema element is unusual in that it has a two-part key i.e. "cloud-owner" and "cloud-region-id". There are not many other usages of it ("ctag-pool", "service-capability" and "route-target" are three others, out of over 100 other schema elements) <input type="checkbox"/> Is it possible to enhance the error message to indicate that part of the key value is missing from the relationship-data? ❌ AAI-2394 - Improper behavior when relationship-data is missing key attribute(s) CLOSED <input type="checkbox"/> Is it time to deprecate the relationship-data and switch over to using the related-link only? <input type="checkbox"/> Is there any modeling guidance that would steer new designs away from using multi-part key for schema elements? <input type="checkbox"/> Are there other caveats to using the multi-part key design for schema elements? <input checked="" type="checkbox"/> Can we get feedback from Chandra Cinthalaon the key design for multi-part keys and whether this will be more common going forward? <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p><i>From:</i> CINTHALA, CHANDRA [mailto:cc1196@att.com] <i>Sent:</i> Tuesday, April 30, 2019 12:16 AM <i>To:</i> Keong Lim <Keong.Lim@huawei.com>; FORSYTH, JAMES <jf2512@att.com> <i>Cc:</i> AGGARWAL, MANISHA <amanisha@att.com> <i>Subject:</i> Re: [confluence] Keong Lim has assigned tasks to you in "2019-05-02 AAI Developers Meeting"</p> <p>Keong,</p> <p>I think we have no plans to deprecate the relation-data option the in the A&AI relationship payload.</p> <p>It's another option for the client to specify the relationship.</p> <p>Thanks</p> <p>Chandra</p> </div> <p><input type="checkbox"/> See also email from Marco Platania in https://lists.onap.org/g/onap-discuss/topic/31385256</p>
9	Return codes and messages for WS	Pavel Paroulek	IN PROGRESS	25th Apr 2019	<p>Is there a guide for the description of the error message and the error codes? How are new error states (message + code) added?</p> <ul style="list-style-type: none"> <input checked="" type="checkbox"/> William LaMont will send James Forsyth the output of a script that formats the error.properties file to make a wiki page and readthedocs <input type="checkbox"/> James Forsyth should commit that script and create a wiki for the error properties

10	AAI test data bootstrap	Keong Lim	IN PROGRESS	25th Apr 2019	<p>3rd May: See also AAI Simulator in Policy: https://gerrit.onap.org/r/gitweb?p=policy/models.git;a=blob;f=models-interactions/model-simulators/src/main/java/org/onap/policy/simulators/AaiSimulator.java;hb=HEAD . Was recently updated for custom query https://gerrit.onap.org/r/#/q/pramod+custom and contains hard-coded AAI data /schema. Maybe discuss with Pamela Dragosh about how this data is maintained?</p> <p>25th Apr: Wiki page OOF - HPA guide for integration testing updated for https://gerrit.onap.org/r/#/c/86588/</p> <p>Similarly, Scott Seabolt and J Armstrong / Joss Armstrong wrote for APCC Sample A&AI Data Setup for vLB/vDNS for APCC Consumption and Script to load vLB into AAI:</p> <div> <p>The below put_vLB.sh script can be used to submit the vLB data to A&AI in order to run ConfigScaleOut use case. This script and referenced JSON files are used on an AAI instance where the cloud-region and tenant are already defined.</p> </div> <p>Similarly:</p> <ul style="list-style-type: none"> TEST-133 - Demo.sh init step is not adding availability-zone information in AAI CLOSED INT-705 - Add availability zone call to AAI in Robot framework CLOSED VCPE Use Case Tutorial: Design and Deploy based on ONAP <p>Related AAI-1948 - Correct broken links CLOSED on the brittleness of the ReadTheDocs links to data files.</p> <p>One for VIM: How-To: Register a VIM/Cloud Instance to ONAP and</p> <div> AAI-1928 - AAI-Internal server Error:500 CLOSED </div> <p>Potential issues:</p> <ul style="list-style-type: none"> fragility of static import data file w.r.t. schema changes and version upgrades for each ONAP release? how "common" is this knowledge, i.e. what to load, where to get it, who else should be using it, etc? should it be automated/scripted, rather than manual steps to bootstrap? should it be a simulator program or test harness, rather than a static data file? should it reside within AAI CI/CD jobs for maintenance and upgrade of schema versions? who maintains the data itself? Is there a "data repository" which can be delegated to other teams, e.g. like documentation repository links in git? how many other teams have similar private stashes of AAI bootstrap data? does it need to be published at a stable URL to avoid linkrot? <p>Possible solution/action:</p> <ul style="list-style-type: none"> <input type="checkbox"/> Look at the examples API and possibly enhance it to get the desired behavior <input type="checkbox"/> Collect all the known data samples, commit to test-config repo, update the teams/wiki to point to test-config repo instead of keeping private stash of AAI data <div> AAI-2399 - Collate test data from wiki into AAI repository CLOSED </div>
11	Purpose of fields in AAI	Keong Lim	IN PROGRESS	18th April 2019	<p>Dénes Németh wrote in AAI-1484 - Invalid handling of external system certificates CLOSED :</p> <div> <p>In think it would be good to answer what is the meaning of the field (collection of PEMs of the CA xor URL)</p> </div> <p>Questions:</p> <ol style="list-style-type: none"> Is AAI intended to strictly prescribe how the fields are used and what contents are in the values? Or does AAI simply reflect the wishes of all the client projects that use it to store and retrieve data? <p>Even if (1) is true, AAI is not really in any position to enforce how clients use the data, so really (2) is always true and we need to consult the original producers of the data and the ultimate consumers of the data to document their intended meanings.</p> <p>How do we push to have documentation on the purpose and meaning of the fields in AAI?</p> <p>Where does all this documentation go?</p> <p>Should the documentation be backed up by validation code?</p> <p>See also discussion about AAI in 2018-11-28 ExtAPI Meeting notes</p> <p>29th Nov: Started on new wiki page AAI Schema Producer-Consumer Pairings</p> <p>18th Apr: Can we have this documentation go into ONAP in a generic way?</p> <p>24th Apr: See also questions about "sw-version" in https://wiki.onap.org/display/DW/5G++PNF+Plug+and+Play?focusedCommentId=63996633#comment-63996633</p>

12	range query	Keong Lim	IN PROGRESS	7th Feb 2019	<ul style="list-style-type: none"> if I had some AAI data with attributes that are strings but nominally contain date/timestamps, is there a way to query for a particular range of values? is there a way to do partial match? regex? PUT /aai/v13/query?format=raw <pre>{ "gremlin": "g.V().has('aai-ts', org.janusgraph.core.attribute.Text.textContains('URL'))" }</pre> <p>https://github.com/JanusGraph/janusgraph/blob/master/janusgraph-core/src/main/java/org/janusgraph/core/attribute/Text.java</p> <p><input type="checkbox"/> Invite Arul Nambi and CT Paterson to next week's dev call to talk about how sparky/elastic does partial/range matches</p>
13	2 Types of logging in A&AI WS	Pavel Paroulek	ON HOLD	1st Nov 2018	<p>There are 2 types of logging in the services</p> <ul style="list-style-type: none"> one read from EELFManager the other Logger log = Logger.getLogger(... <p>Is that correct? Shouldn't there be just 1 type?</p> <p>1st Nov:</p> <p>After Casablanca release investigate logging guidelines and figure out what library to use in order to unify logging within A&AI</p> <p>26th Nov: See also ONAP Application Logging Specification - Post Dublin</p> <p>29th Nov: how does this fit with</p> <p> LOG-877 - S3P: Logging streaming/format alignment for dublin - China Telecom, Deutsche Telekom, Vodafone</p> <p><input type="button" value="CLOSED"/></p> <p>?</p>
14	Schema-service roadmap	James Forsyth Pavel Paroulek	ON HOLD	21st Mar 2019	<p>31st Jan 2019:</p> <p>The schema-service is ready. Currently it provides file-sharing capabilities in terms of schema/edgerule files.</p> <p>In order for GraphGraph to take advantage of the schema parsing/processing in schema-service additional abstractions have to be implemented on top of the crude file2string functionality currently in schema-service.</p> <p><input checked="" type="checkbox"/> Venkata Harish Kajur will ask Manisha Aggarwal if the current functionality of the schema-service is the final version for Dublin and if there will be further enhancements in next releases.</p> <p><u>GraphGraph needs the following functionality:</u></p> <p>Venkata Harish Kajur and Manisha Aggarwal What is missing in schema service that is needed in graphgraph is the following:</p> <ul style="list-style-type: none"> rest call to get available schemas list of all schema nodes/items (like vserver, tenant, p-interfaces..) for example on a REST path /schemas/{schema}/nodes all relevant attributes of a given node/item for example on REST path /schemas/{schema}/nodes/{node} edges/relationships with their attributes between schema nodes/items (for example on REST path /schemas/{schema}/edges where you specify a "from" "to" schema items as query params) subgraph of the schema, where you specify 1. initial (root) items/node (like tenant or vserver) 2. schema version and 3. number of parent/cousin/child hops from the initial item/node all paths in a given schema graph between 2 items/nodes (like vserver and tenant) for a given schema version edges in the schema graph should be composed of edges in the schema file + edges created from the edgerules file edges should contain basic attributes when delivered via the subgraph call (like parent/child relationship and important properties from edgerules) and have additional (or all) attributes when queries via /schemas/{schema}/edges REST endpoint. <p>20. Mar 2019:</p> <p>Open questions for schema-service:</p> <ol style="list-style-type: none"> what is the current implemented functionality? what are the business use-cases in ONAP for schema-service? Description of functionality in relation to other services/projects is needed. In other words who needs it and why? if <u>no business</u> use-cases can be formulated we should consider removing schema-service from A&AI and replacing it with standard file-sharing mechanisms. <p>21st Mar 2019:</p> <p>Based on William Reehil comments</p> <p>https://wiki.onap.org/display/DW/AAI+Schema+Service?focusedCommentId=45299262#comment-45299262 what is "our future proposed functionality"?</p>
15	New AAF Certificates at startup	Jimmy Forsyth	ON HOLD	24th Jan 2019	<p>AAF will generate certificates to be used by the containers at startup; AAI services should use the run-time generated certs instead of the ones that are in the repos or oom charts.</p> <p>In dublin the services will mount a volume with certificates. This is on the roadmap for Dublin as a feature.</p> <ul style="list-style-type: none"> is this for all service and/or HAProxy? Where are the certificates coming from (OOM/gerit/generated by AAF) <p><input checked="" type="checkbox"/> James Forsyth will ask Jonathan Gatham when the certificate init image is going to be available in ONAP and whether it is documented 31 Jan 2019</p>

16	Change PNF to use pnf-id as unique key	Benjamin Cheung Chesla Wechsler David Perez Caparros user-063b3 Matthieu Geerebaert	ON HOLD	7th March 2019	<p>Potentially breaking change: AAI-2096 - PNF Plug and Play CLOSED / blocked by</p> <p>ONAPARC-409 - Dublin Architecture Review: AAI CLOSED</p> <p>See also:</p> <ul style="list-style-type: none"> Proposal to Change AAI PNF Entity to use PNF-ID as key AAI R4 Use Case and Functional Requirements Impacts Suggestion in comment Re: AAI REST API Documentation - Dublin <p>Questions:</p> <ul style="list-style-type: none"> how to minimise impact of the transition from pnf-name as unique to pnf-id as unique key? <ul style="list-style-type: none"> would the v14 URL be different from the v15 URL? would both paths be equally supported for GET/PUT /etc? what forwards-compatibility or backwards-compatibility will be supported? how to migrate forwards or backwards database versions, ONAP versions, etc, across this transition? who is going to implement it? Test it? what is the impact of this not going ahead? <p><input checked="" type="checkbox"/> William LaMont will check for existing migration utility that handles this use case (changing the key from one existing attribute to another). Changes to pnf object in all oxm versions would be needed, and a migration similar to what was done in UriMigration but limited to pnf node-type to update the aai-uri, and a schema mode to add an index on pnf-id.</p> <p><input checked="" type="checkbox"/> James Forsyth will socialize the breaking change on the PNF in the next PTL call so clients can prepare to do a search for ?pnf-name=\${pnf-name} instead of /pnfs/pnf/\${pnf-name}. They also need to handle doing the PUT operation differently - Added to PTL agenda PTL 2019-02-19</p>
17	A&AI monitoring and failure prevention initiative	Keong Lim	ON HOLD	24th Jan 2019	<p>Mike Elliott wrote in OOM Meeting Notes - 2018-12-5</p> <p><i>f. AAI team wanted to get notified of AAI Cassandra issues automatically i. Can we setup a Nagios or equivalent to monitor both rancher/k8 and the applications for rancher/k8 issues ?</i></p> <p>Keep an eye out for new issues!</p> <p>This should be part of a larger A&AI monitoring and failure prevention initiative!</p>
18	AAI HAProxy and 2-way-TLS	Pavel Paroulek	ON HOLD	29th Nov	<p>Technical solution to either decommission the proxy or make design changes to AAF to enable client side certificates.</p> <p>After VF2F we will know if this is a requirement in Dublin. We discuss after this date.</p> <p>question raised: MSB - would client authentication be supported?</p> <p>15th Dec: https://wiki.onap.org/display/DW/Pluggable+Security#PluggableSecurity-7.10 Identifiedandsupportedpatternsandfeatures</p> <p><input type="checkbox"/> James Forsyth creates a task for encryption of communication between A&AI services and Cassandra 24 Jan 2019</p> <p><input type="checkbox"/> Tian Lee and Steven Blimkie to report on the Amdocs managed A&AI microservice wheter they support criteria from the Dublin S3P requirements 31 Jan 2019</p>
19	AAI too slow for OOF/HAS	Keong Lim	ON HOLD	25th Apr 2019	<p>Under OOF Homing and Allocation Service (HAS) section, Dileep Ranganathan wrote about Project Specific enhancements:</p> <p><i>Optimize - AAI cache</i></p> <ul style="list-style-type: none"> Use MUSIC or any other alternative in memory caching like Redis etc? Optimize flavor retrieval from A&AI and Cache the information if necessary <p>See also OPTERA-268 - OOF - project specific enhancements CLOSED /</p> <p>OPTERA-394 - Optimize the use of AAI cache in HAS Data component CLOSED</p> <p>Similarly to the "AAI too slow for Holmes" item below, this introduction of extra caching of AAI data is a worrisome development and sad indictment of the performance of the system architecture.</p> <p>For holmes, we could possibly create a custom query to address it.</p> <p>What can we do about this?</p> <p>Would the AAI Cacher AAI-4337 - AAI Cacher CLOSED help to improve performance?</p>

20	MultiCloud usage of AAI for HPA telemetry/time-series data to OOF	Keong Lim	ON HOLD	18th April 2019	<p>Bin Yang and Lianhao Lu wrote in</p> <div>  MULTICLOUD-274 - HPA telemetry data collection and persistence CLOSED </div> <p>:</p> <p>HPA telemetry data collection and make it persistent in A&AI, from which OOF can leverage during its decision making process.</p> <p>and</p> <p>1. Multi-cloud to collect the data from time-series data services like Prometheus (http://prometheus.io) or openstack Gnocchi, and push them to A&AI based on the data recording & aggregation rules.</p> <p>and</p> <p>The reason why we propose here is that VES mechanism doesn't store the telemetry data into A&AI. And OOF now can only get those kind of data from A&AI.</p> <p>Some concerns:</p> <ul style="list-style-type: none"> how much additional load will this place on AAI? will AAI cope with this load? is AAI suitable for "time-series data"? is "telemetry data" considered to be "active & available inventory"? should OOF access the telemetry/time-series data via other means (not AAI)? AAI API latency (4-6 second per request as benchmarked in CMCC lab) could be a problem
21	AAI Backup and Restore	Keong Lim	ON HOLD	10th Jan 2019	<p>Brian Freeman asked on Re: Backup and Restore Solution: ONAP-OOM :</p> <p>what would be the approach to backup an entire ONAP instance particularly SDC, AAI, SDNC data ? would it be a script with all the references to the helm deploy releases or something that does a helm list and then for each entry does the ark backup ?</p> <p>What is the AAI strategy for backup and restore?</p> <p>What is the overall ONAP strategy for backup and restore?</p> <p>Should it be unified with the data migration strategy as per "Hbase to Cassandra migration" on 2018-11-14 AAI Meeting Notes?</p> <p><input checked="" type="checkbox"/> James Forsythwill raise the topic of having backups and restore functionality in ONAP - if it is feasible, on the roadmap and what others PTL think</p> <p>Jimmy didn't directly raise the topic but there was movement - Keong Lim asked "if istio service mesh is a no-go, is there a replacement for secure onap communications?is backup/restore/upgradability included in s3p?"</p> <p>Michael O'Brien replied that a reference tool set for backup and restore was introduced in Casablanca: Backup and Restore Solution: ONAP-OOM</p> <p>Mike Elliott said he would look at Brian's question, AAI will provide support as needed.</p>
22					