

2019-03-28 AAI Developers Meeting

Attachments (including meeting recording)

File Modified









No files shared here yet.

Chat Server

private group aai-dev on Rocketchat server: <http://onap-integration.eastus.cloudapp.azure.com:3000/group/aai-dev>

Agenda Items

START RECORDING



	Title	Responsible	Status	Last discussed	Notes
1	Integration Blockers	Tian Lee Steven Blimkie	DONE	21st Mar 2019	 AAI-2267 - Wrong model parsing CLOSED Already got a gerrit submitted to fix it.
2	New AAF Certificates at startup	Jimmy Forsyth	IN PROGRESS	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/gerrit/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>
3	Shared Cassandra Database	Mahendra Raghuvanshi Roger Maitland	IN PROGRESS	14th March 2019	<p>Regarding  OOM-4652 - Upgrade AAI to use common shared Cassandra instance CLOSED /</p> <p> OOM-4676 - Upgrade AAI to use helm hooks for rolling upgrade CLOSED /</p> <p> OOM-4486 - Migrate AAI Cassandra charts to OOM common CLOSED /</p> <p> OOM-4487 - Upgrade AAI to use common Cassandra charts CLOSED /</p> <p> OOM-4494 - Create Shared Instance of Cassandra CLOSED /</p> <p> OOM-4499 - Upgrade AAI to use common Cassandra instance CLOSED :</p> <ul style="list-style-type: none">• Is this something AAI team needs to be aware of? Or is it OK for OOM team to just switch it around?• Will this introduce unexpected dependencies between AAI, AAF, OOF, Portal and SDC that will create difficulties for upgrade/downgrade/backup/restore/maintenance/schema change?• Is it going to exacerbate the performance problems already noted? <p>A number of gerrit review issues raised: https://gerrit.onap.org/r/#/c/79425/</p> <p>The "rolling upgrade" change has been combined with the "shared cassandra" change.</p> <p>The "shared cassandra" change has been combined with the "AAF shared cassandra" change, which means it's also combined with the "AAI shared cassandra" change.</p> <p>This sounds like a recipe for disaster.</p> <p>18th Mar: New patchsets to address our concerns https://gerrit.onap.org/r/#/c/82418/</p>
4	ETSI SOL 005	Bharath Thiruvedula	IN PROGRESS		<p>RE: Support ETSI NFV-SOL 005 (Os-Ma-Nfvo ref point) between SO & VF-C/NFVO</p> <p>Hi AAI team,</p> <p>Can you please review the patches regarding the SOL 005 requirements?</p> <p>Here are the gerrit links</p> <ol style="list-style-type: none">1. https://gerrit.onap.org/r/#/c/82892/2. https://gerrit.onap.org/r/#/c/82935/3. https://gerrit.onap.org/r/#/c/82948/ <p>Best Regards</p> <p>Bharath T</p> <p>26th Mar: Also needs  AAI-2279 - Register v16 to MSB while bringing aai-resources up in Dublin CLOSED</p>

5	Return codes and messages for WS	Pavel Paroulek	IN PROGRESS	28th Feb 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> <p><input 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</p>
6	Image footprint reduction as part of CIA project	Paul-Ionut Vaduva	IN PROGRESS	7th Feb 2019	<p>Reduction in size is mostly onf aai-common image as that based on ubuntu.</p> <p>2/7 - Move the base image to be a part of ONAP Build, maybe aai-common repo</p> <p><input type="checkbox"/> Venkata Harish Kajur will create a Jira for it in Dublin Release</p> <p>Move the aai-common Dockerfile RUN into the resources, traversal, graphadmin, cacher, schema-service microservice</p> <p>26th Mar: Also AAI-2235 - last build of aaionap/haproxy is 9 months old CLOSED</p>
7	Change PNF to use pnf-id as unique key	Benjamin Cheung Chesla Wechsler David Perez Caparros user-063b3 Matthieu Geerebaert	IN PROGRESS	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 type="checkbox"/> William LaMont will check for existing migration utility that handles this use case (changing the key from one existing attribute to another)</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>
8	OOM Artifacts	Pavel Paroulek	OPEN		<p>Some of our top level OOM deployment artifacts are not unique (i.e. don't take namespace into account as all other deployables), is that intentional?</p>
9	AAI too slow for OOF/HAS	Keong Lim	OPEN		<p>Under OOF Homing and Allocation Service (HAS) section, Dileep Ranganathan wrote about Project Specific enhancements:</p> <div> <p><i>Optimize - AAI cache</i></p> <ul style="list-style-type: none"> • <i>Use MUSIC or any other alternative in memory caching like Redis etc?</i> • <i>Optimize flavor retrieval from A&AI and Cache the information if necessary</i> </div> <p>See also OPTFRA-268 - OOF - project specific enhancements CLOSED /</p> <p> OPTFRA-294 - 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>What can we do about this?</p> <p>Would the AAI Cacher AAI-1337 - AAI Cacher CLOSED help to improve performance?</p>

10	MultiCloud usage of AAI for HPA telemetry /time-series data to OOF	Keong Lim	OPEN	<div><div>Bin Yang and Lianhao Lu wrote in</div><div><div>MULTICLOUD-274 - HPA telemetry data collection and persistence</div><div>CLOSED</div></div></div> <div>:</div> <div><div></div><div>HPA telemetry data collection and make it persistent in A&AI, from which OOF can leverage during its decision making process.</div></div> <div>and</div> <div><div></div><div>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.</div></div> <div>and</div> <div><div></div><div>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.</div></div> <div>Some concerns:</div> <div><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</div>
----	--	-----------	------	--

11	Orchestration Scenarios for VNFS	Keong Lim	OPEN	<div>Comments on Orchestration Scenarios related to AAI:</div> <div>Viswanath Kumar Skand Priya / Viswanath KSP said:</div> <div><div><div>Thank you Ranny Haiby & Fernando Oliveira . I agree partly, but I still have following queries.</div><div><ul style="list-style-type: none">I agree & acknowledge that atleast for a foreseeable future, we would need a way to specify the VNFM / NFVO as part of "Design Decision", which I believe can be reflected as part of VNFD/NSD (using some special attribute) or as part of Internal Model that SDC might build before distributing the same. SO can then use this hint to select relevant actors. My only question is, why this has to be maintained in AAI which is exclusively for runtime record? All AAI cares about is what is running in the network irrespective of how that got orchestrated. Isn't it ?</div><div>On a broader note, I would like to understand what's the original intent of AAI (atleast in ECOMP world) ? Are we simply assuming that, just because AAI has "available inventory" in its name, we are expecting it to keep track of cloud inventory realtime ? Because our entire story (including the new G-FPS proposal) is based on this assumption. Can anyone from AAI team or ATT clarify on this ?</div><div>Because AFAIK, AAI neither has the schema to host such available inventory, nor the MC has the pub/sub or polling mechanism (today) to refresh the cloud inventory inplace. Ofcourse those can be scoped for further releases, but my original question is, was that the original intent behind AAI or are we now including it in the scope?</div></div></div> <div>and Fernando Oliveira replied:</div> <div><div>For the first question: I think that A&AI needs to maintain the VNF instance VNFM instance and the NS instance NFVO instance relationship for subsequent life cycle operations, i.e. a scale or heal operation. The path would be something like Event (VNF Instance, Busy) DCAE (policy for VNF instance) Policy Evaluation (VNF instance, Scale-out) SO (VNF instance, Scale-out) A&AI (find VNFM instance for the VNF instance) SO (VNF instance, VNFM instance, Scale-out) SOL003 Adapter (VNFM instance, VNF instance, Scale-out) VNFM instance (VNF Instance, Scale-out).</div></div> <div>As I understand, ESR has "esr-vnfm-list", which has an "esr-vnfm", which has "esr-system-info-list", which has "esr-system-info", which has a "relationship-list" that can contain relationships to "generic-vn" and other AAI objects.</div> <div>The "generic-vn" object also contains "self-link", "ipv4OamAddress", "ipv4OamGatewayAddress", etc, which links the AAI object back to its "source-of-truth" external-system.</div> <div>Is there some new data, new schema or new API that is required on top of this?</div> <div>Fernando Oliveira; Apologies for my lack of knowledge, but a few comments:</div> <div><div><div>1. For the VNF/VF VNFM case, I think that there needs to be a reference from a VNF/VF instance record to the specific instance of the VNFM that was used to deploy the VNF/VF. If there is already such a reference from the VNF/VF through the ESR to the specific item on the esr-vnfm-list, then I think that would be sufficient. If not, I think that would be a new requirement.</div><div>2. For the Service NFVO case, Is there an equivalent NFVO/Orchestrator list in the ESR? The esr-nfvo-list would need the same set of info as the VNFM case. If the esr-nfvo-list does exist, I think that there needs to be a reference from the Service Instance record to the specific NFVO instance that deployed the Service. Is there such a reference? If not, I think that would be a new requirement.</div></div></div> <div>Bo Lv can comment more on the current ESR capabilities, but I believe there are only 3 kinds of systems so far: EMS, VNFM and third-party SDNC.</div> <div>ESR could be extended to handle VNFO as another kind of system.</div> <div>Fernando Oliveira : I created JIRA stories:</div> <div><div><div><div><div><div>▪</div><div><div><input checked="" type="checkbox"/></div><div>REQ-422 - ONAP tracking of available external (ETSI compliant NFVO) orchestrators</div><div>DONE</div></div></div></div><div><div><input checked="" type="checkbox"/></div><div>REQ-423 - ONAP tracking of Service dependency on an external orchestrator</div><div>DONE</div></div><div><div><input checked="" type="checkbox"/></div><div>ONAPARC-399 - ONAP tracking of VNF dependency on an external ETSI compliant VNF Manager</div><div>CLOSED</div></div><div><div><input checked="" type="checkbox"/></div><div>(VNFM)</div><div>CLOSED</div></div></div></div></div> <div>for various parts of the scenario.</div> <div>Bharath Thiruveedula is this item related to your question for Support ETSI NFV-SOL 005 (Os-Ma-Nfvo ref point) between SO & VF-C/NFVO?</div> <div>Keong Lim it is related to the question</div>
----	----------------------------------	-----------	------	---



12	AAI test data bootstrap	Keong Lim	OPEN	<p>Looking at AAI usage in OOF - HPA guide for integration testing by Dileep Ranganathan, wondering whether there is a better way to bootstrap AAI test data?</p> <div> <h2>Generating AAI data</h2> <p><i>Note: Required only if the Multicloud has no real cloud-regions and HPA discovery cannot happen.</i></p> <p><i>If Multicloud team has data for creating the Cloud-region and doesn't have the HPA, then please update the existing data with the flavors with HPA.</i></p> <ol style="list-style-type: none"> 1. Import the postman collection CASABLANCA_AAI_postman.json 2. To add/remove HPA Capabilities edit the flavors section in the body of PUT Cloud-Region{x} 3. Once all the necessary Use postman to add the complex and cloud regions in the order specified below (snip screenshot of specific sequence) 4. Use the GET requests to verify the data. (snip screenshot of specific sequence) </div> <p>Similarly, Scott Seabolt and J Armstrong / Joss Armstrong wrote for APPC Sample A&AI Data Setup for vLB/vDNS for APPC Consumption and Script to load vLB into AAI:</p> <div> <p><i>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.</i></p> </div> <p>Similarly:</p> <div> <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 </div> <p>Related AAI-1048 - 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 AAI-1028 - AAI-Internal server Error:500 CLOSED</p> <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?
13	Purpose of fields in AAI	Keong Lim	OPEN	<p>Dénés Németh wrote in AAI-1104 - Invalid handling of external system certificates CLOSED :</p> <div> <p><i>In think it would be good to answer what is the meaning of the field (collection of PEMs of the CA xor URL)</i></p> </div> <p>Questions:</p> <ol style="list-style-type: none"> 1. Is AAI intended to strictly prescribe how the fields are used and what contents are in the values? 2. 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>

14	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>
15	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><i>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 ?</i></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 Forsyth will 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>
16	aai-cassandra performance issues	Keong Lim	ON HOLD	10th Jan 2019	<p>Michael O'Brien has documented performance issues in aai-cassandra:</p> <ul style="list-style-type: none"> Cloud Native Deployment (search for LOG-376 to find the specific references to AAI) https://jira.onap.org/browse/LOG-376?focusedCommentId=29358&page=com.atlassian.jira.plugin.system.issuetabpanels:comment-tabpanel#comment-29358 <p><i>hector has discovered that the stress test jar (liveness probe?) in aai-cassandra is hammering the cpu/ram/hd on the vm that aai is on - this breaks the etcd cluster (not the latency/network issues we suspected that may cause pod rescheduling)</i></p> <p>Is there something that should be tweaked in AAI config? Or documentation on the recommended setup to run the VM?</p> <p>I'll come to the next AAI meet (conflicts with pomba meet) -</p> <p> LOG-376 - Logstash full saturation of 8 cores with AAI deployed on one of the quad 8 vCore vms for 30 logs/sec - up replicaSet 1 to 3 or use DaemonSet CLOSED</p> <p> 20190108 work continues to find the cause - I see 7 vCore spikes on cassandra as well as a saturated logstash on that particular vm - we are no longer a DaemonSet (13 instances on a 13+1 cluster) - I will reduce the current ReplicaSet from 5 to 2 or 1 until I can label the nodes and/or find out what is causing is to saturate - Prudence Au and Sanjay Agraharam mentioned cassandra - I have seen cs high on several "top" sessions - will post screen caps - bottom line is correlation - I have a 2nd cluster where I can just run aai,dmaap and log</p> <p>LOG Meeting Minutes 2019-01-15</p> <p><input checked="" type="checkbox"/> LOG-915 - Reduce Logstash core limit to 1 from 3 until LOG-LS and AAI-CS perf issue on the same VM is determined OPEN</p> <p><input type="checkbox"/> ask Michael O'Brien about performance problems - if they persist and what the problem exactly is.</p> <p><input type="checkbox"/> Venkata Harish Kajur will inform Michael about the schema performance fix - he should test with the casablanca maintenance release.</p> <p>On-hold for 3 weeks (end of January) - if until then no performance issues reported agenda item will be closed</p>
17	get notified of AAI Cassandra issues automatically	Keong Lim	IN PROGRESS	24th Jan 2019	<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>Question: Keong Lim should this be part of a larger A&AI monitoring and failure prevention initiative</p>

18	Modeling team R4 discussion, including extra AAI attributes in a model-driven way	Keong Lim	IN PROGRESS	15th Nov 2018	<p>Modelling team having Service Instance thoughts by Chesla Wechsler, which will affect AAI schema.</p> <p>Also referred from comments on ONAP R4+ Service Modeling Discussion Calls</p> <div> <p>9) "vhn-portal-url"? "Bandwidth", "QoS", "SLA", etc, attributes that not all the services need but still need to be stored in certain service instance: stored as a schemaless field on the service-instance vertex (Chesla will follow up) (my concerns: according to the call, is that ok if we set a "global-type of service" and a "customized-type of service", then mapped it with internal descriptor, and A&AI's model only stores global type in service instance's schema, but stores the customer-faced attributes of service in a schemaless way? Chesla Wechsler Kevin Scaggs Andy Mayer)</p> </div> <p>See also Modeling 2018-11-13</p> <p>The service-instance already uses a "metadata" relationship, which can store an arbitrary list of key-value pairs, but perhaps AAI should extend the use of the "properties" element, which is also an arbitrary list of name-value pairs or the "extra-properties" element, which is also an arbitrary list of name-value pairs.</p> <p>15th Nov: Having seen Chesla's presentation, it should be called "Model-driven schema" rather than "schemaless" behaviour, since the idea is that the changes are controlled by SDC modelling. Seems aligned to the eventual goal in AAI Schema Service Use Case Proposals and AAI Schema Service.</p>
19	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> <div> <p> LOG-877 - S3P: Logging streaming/format alignment for dublin - China Telecom, Deutsche Telekom, Vodafone</p> <p>CLOSED</p> </div> <p>?</p>
20	Disable unused web services (see also Helm chart requested values)	Pavel Paroulek	IN PROGRESS	20th Dec 2018	<p>Could we disable unused (i.e. not integrated) A&AI web services, so that the deployment is faster and the resource footprint is smaller? e.g. Champ (any other ws?)</p> <p>Motivation: Decrease the resource footprint for A&AI (ONAP) deployments</p> <p>Idea: we could support 2 different deployments 1. full (normal) deployment and 2. barebones deployment. The point of the "barebone" deployment would be to deploy only the essential services necessary for proper functioning of A&AI (leaving out services like cacher, sparky, graphadmin, having 1 cassandra node instead of 3 or 5 etc).</p> <p>In order to reduce hardware/cloud costs (mainly the memory footprint) it could be beneficial to support a minimalistic A&AI deployment.</p> <p>1st Nov:</p> <p>Venkata Harish Kajur and Pavel Paroulek - investigate how to disable/enable charts in A&AI so we can create a core group of pods which handle the use-cases and then extended group will all the services. Consider a group of unused/unintegrated services (like Champ). Consider other possible groups (like GUI?)</p> <div> <p><input checked="" type="checkbox"/> James Forsyth creates a JIRA ticket to define the list of A&AI subprojects and create the categories (essential, full "experience") for the OOM deployment</p> <p> AAI-2025 - Classify aai docker images</p> <p>CLOSED</p> </div>
21	AAI Champ	Pavel Paroulek	IN PROGRESS	1st Nov 2018	<ol style="list-style-type: none"> Who is responsible for the project? What is the roadmap for the project? Who will do the integration?
22	AAI HAProxy and 2-way-TLS	Pavel Paroulek	IN PROGRESS	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 Identified and supported patterns and features</p> <div> <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 whether they support criteria from the Dublin S3P requirements 31 Jan 2019</p> </div>

23	named-query replacements	James Forsyth	IN PROGRESS	31th Jan 2019	<div> <div> AAI-1989 - AAI Named Query Deprecation OPEN </div> </div> <div> <div> APPC-1245 - AAI named-queries are being deprecated and should be replaced with custom-queries CLOSED </div> <div> Takamune Cho pointed to inherited CCSDK functions, so added a new case to follow-up there too </div> </div> <div> <div> POLICY-1278 - AAI named-queries are being deprecated and should be replaced with custom-queries CLOSED </div> </div> <div> <div> VID-355 - AAI named-queries are being deprecated and should be replaced with custom-queries OPEN </div> <div> : Ittay Stern has provided an update with 2 known queries used. </div> </div> <div> <div> CCSDK-962 - AAI named-queries are being deprecated and should be replaced with custom-queries CLOSED </div> <div> : Raised to Dan Timoney based on info about APPC component. </div> <div> Need to replace custom queries currently in use by these systems (and others?) in Dublin toward the retirement of the named-query API in Dublin </div> <div> <ol style="list-style-type: none"> 1. Coordination with the ROBOT team needed for data population 2. Coordinate with each team (POLICY, VID...) to have the specific data for each named-query </div> <div> Christopher Shang is leading the effort to deprecate named queries in e-Comp. </div> <div> Next steps: get data from teams in order to prepare for testing the change from named queries to custom queries. </div> <div> See also AAI Named Queries </div> <div> <ul style="list-style-type: none"> <input checked="" type="checkbox"/> Keong Lim will ask the teams about the data mentioned above. <input checked="" type="checkbox"/> James Forsyth will remind Christopher to create a LF account 24 Jan 2019 <input checked="" type="checkbox"/> James Forsyth will ask for a Fitness test which tests the correctness of conversion from name-query to custom-quesry <input checked="" type="checkbox"/> James Forsyth will try to find out if any documentation can be provided from ecomp regarding the named-queries (how they work and how to read the named-query files) </div> </div>
24	AAI too slow for Holmes	Keong Lim CT Paterson	IN PROGRESS	7th March 2019	<div> Guangrong Fu mentioned AAI in Baseline Measurements based on Testing Results: </div> <div> <ol style="list-style-type: none"> 1. Cache the AAI data and refresh them periodically so that Holmes won't have to make an HTTP call to AAI every time it tries to correlate one alarm to another. </div> <div> The problem for caching is how to know when to update the cached data. Even though the access time may be fast for Holmes, the risk is using out-of-date data, so the correlations will be wrong anyway. Also, duplicating the AAI data outside of AAI is probably a bad architectural decision. Making AAI faster for these use cases would be better. </div> <div> Has there been a performance analysis of where the time is spent? Could it help to use Elasticsearch (e.g. as in sparky)? Should Holmes have a batch interface to get more AAI data in fewer calls? Or a better correlation API that results in fewer calls? </div> <div> 31st Oct: https://lists.onap.org/g/onap-discuss/topic/27805753 </div> <div> 1st Nov: <ul style="list-style-type: none"> • Guangrong Fu will try custom queries for queries that took too long to return • The hardware (mainly storage) influences the query speed - need to find out what hardware was the speed test conducted on (Guangrong Fu will provide HW specs) </div> <div> <div> HOLMES-196 - Improve the performance on AAI queries. OPEN </div> </div> <div> Would the AAI Cacher <div> AAI-1337 - AAI Cacher CLOSED </div> help to improve performance? </div> <div> 5th Mar: Guangrong Fu </div> <div> Hi, </div> <div> Sorry for my late response. It took me a long time to set up AAI in my own env. For Item 10, here's some information: </div> <div> Main APIs invoked in Holmes for different use cases: </div> <div> VoLTE <ul style="list-style-type: none"> • Getting the VM query URL via: /search/nodes-query?search-node-type=vserver&filter=vserver-name:EQUALS: - once • Getting VM info via: the URL returned by the query above - once • Getting the VNF data via: network/generic-vnfs/generic-vnf - once </div> <div> CCVPN <ul style="list-style-type: none"> • Updating terminal point via: /network/pnfs/pnf/{pnfName}/p-interfaces/p-interface/nodetd-{pnfName}-ltpId-{ifName} - once • Getting logical links via: /network/pnfs/pnf/{pnfName}/p-interfaces/p-interface/nodetd-{pnfName}-ltpId-{ifName} - 3 times • Getting VPN binding info via: /network/pnfs/pnf/{pnfName}/p-interfaces/p-interface/nodetd-{pnfName}-ltpId-{ifName} - once • Getting connectivity info via: /network/vpn-bindings/vpn-binding/{vpndId} - once • Getting service instance info via: /network/connectivities/connectivity/{connectivityId} - once </div> <div> Performance </div>

					<p>We set up an AAI env on a VM (8 cores, 16GB memory, 160GB storage) following the guidance https://wiki.onap.org/display/DW/How+to+Docker+setup+on+Single+VM+HEAT+Deployment and tried to run a VNF query using "/aai/v11/cloud-infrastructure/cloud-regions/cloud-region/example-cloud-owner-val-45051/example-cloud-region-id-val-56689/tenants/tenant/example-tenant-id-val-51834/vservers/vserver/example-vserver-id-val-51834" (which is returned by "/search/nodes-query?search-node-type=vserver&filter=vserver-name:EQUALS:") for 1000 times. It took ~95ms per query. Also, we tried to query a VNF for 1000 times via "/aai/v11/network/generic-vnfs/generic-vnf/example-vnf-id-val-92494" and the average time is ~86ms.</p> <p>From the result, we know that even for a single request, the time cost reaches around 100ms. Let alone there will be several requests sent to AAI when an alarm is processed by Holmes. Taking CCVPN for example, for each alarm, there are up to 7 requests made. That means it'll take around 600-700 ms for Holmes to interact with AAI. In case of alarm storms, it is hard for AAI to support such intensive queries.</p> <p>6th March: Guangrong Fu</p> <p>In my opinion, the performance of AAI queries is not only impacted by the computation inside AAI, but also impacted by the HTTP request itself.</p> <p>I've done another test. I tried to send requests to the health check API (which does nothing but return immediately after it receives a request) of Holmes. The average time cost is also ~ 70ms. So it seems to be a problem with the time cost caused by setting up and releasing HTTP connections.</p> <p>6th March: Keong</p> <p>Regarding these queries:</p> <div> <ul style="list-style-type: none"> Getting logical links via: <code>/network/pnfs/pnf/{pnfName}/p-interfaces/p-interface/{nodeId-{pnfName}-ltpId-{ifName}}</code> - 3 times Getting VPN binding info via: <code>/network/pnfs/pnf/{pnfName}/p-interfaces/p-interface/{nodeId-{pnfName}-ltpId-{ifName}}</code> - once </div> <p>What depth is used on these GET calls? If the defaulting to depth=0, then perhaps some improvement can be made by using "depth=1" or "depth=2"? Fewer calls returning more data could improve overall performance.</p> <p>Same could be achieved by changing to Nodes query, e.g.</p> <p>GET /aai/v14/nodes/p-interfaces?interface-name=nodeId-{pnfName}-ltpId-{ifName}</p> <p>Question1: Can the Bulk API be used with GET calls? Documentation does not show any examples of GET actions. https://onap.readthedocs.io/en/casablanca/submodules/aai/aai-common.git/docs/AAI%20REST%20API%20Documentation/bulkApi.html</p> <p>Question2: Would it help to have the Holmes pod co-located with the AAI haproxy and AAI resources pods? Reduced network latency could improve overall performance.</p> <p>Guangrong: Holmes is acutally deployed by DCAE. I'm not sure whether your proposal is feasible. What's more, the performance data I got was based on the fact that Holmes and AAI were deployed on the same VM, sharing the same docker env.</p>
25	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 Kajurwill ask Manisha Aggarwalif 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 no business 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>

26	CSIT / CSIT on OOM	James Forsyth	IN PROGRESS	21st March 2019	<p>New requirement to have our csit run as part of oom test environment in windriver</p> <p>Long term goal: one robot test case per endpoint that can run as part of e2e test</p> <p>AAI R4 Integration Sanity Test Plans</p> <p>See also Contributing To AAI Best Practise</p> <p>4th March:  AAI-2208 - Add robot test cases to support CCVPN use case CLOSED / https://gerrit.onap.org/r/#/c/79583/</p> <p>13th March: Success! First commit merged, more to follow!</p> <p>20th March: Second commit merged https://gerrit.onap.org/r/#/c/82623/</p>
27	Maven plugin for consistent code style	Keong Lim	DONE	21st Mar 2019	<p>Raised JIRA  AAI-2496 - Consistent code style in AAI for maven-checkstyle-plugin (onap-java-style) CLOSED</p> <p>with commits submitted for aai-common, cacher, gizmo, and more to come.</p> <p>Also added to Contributing To AAI Best Practise</p> <p>Please test and review.</p>
28					