

2019-02-14 AAI Developers Meeting

Attachments (including meeting recording)









File	Modified
Text File chat.txt	Feb 15, 2019 by James Forsyth

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	Shared Cassandra Database	Mahendra Raghuvanshi Roger Maitland	OPEN		Regarding  OOM-1194 - Create Shared Instance of Cassandra  /  OOM-1499 - Upgrade AAI to use common Cassandra instance  : <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?
2	Change PNF to use pnf-id as unique key	Benjamin Cheung Chesla Wechsler David Perez Caparros user-063b3 Matthieu Geerebaert	OPEN		Potentially breaking change:  AAI-2096 - PNF Plug and Play  See also: <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 Questions: <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? <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) <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
3	AAI-BBS for Dublin	Keong Lim	OPEN		Started collating proposed changes in wiki AAI-BBS Proposals for Dublin Release Added  AAI-2454 - Add EdgeRules to cp/generic-vnf/pnf to support BBS use case  : discussions about options. Perhaps some of this should go into the "AAI data model principles document"?
4	AAI REST API	THIRILOSHINI KRISHNAKUMAR	OPEN		1. I see APIs to ADD/DELETE/UPDATE a single PNF to/from AAI. Is it possible to do a bulk ADD/DELETE /UPDATE ? That is, ADD/DELETE/UPDATE multiple PNFs in one API call instead of doing multiple calls for each PNF. The reason behind the question is that, I'll have to get Nodes from SDNC (third party) and add it to AAI. When there are 100s of nodes, making more than 100 API calls individually to add the nodes becomes expensive. If the bulk operation is not supported, Can you please help me understand if there is any specific reason for not supporting it ? 2. I added some p-interfaces as part of adding a pnf. Please find below a sample json that I used. I use POSTMAN for all REST operations. The API to add the PNF: PUT https://10.1.8.105:8443/aai/v14/network/pnfs/pnf/T310-NE4 INPUT: {

```
"pnf-name": "T310-NE4",
"in-maint": false,
"p-interfaces": {
  "p-interface": [
    {
      "interface-name": "otu-1/1/0/E1.1",
      "in-maint": false
    },
    {
      "interface-name": "otu-1/2/0/E1.1",
      "in-maint": false
    }
  ]
}
```

The pnf addition was successful. When I try to fetch the pnf, I expected to see the p-interfaces as well under the pnf element as that's how I see in the pnf model. But I had to make a separate GET API call to see the interfaces.

The API to get the pnf: **GET** <https://10.1.8.105:8443/aai/v14/network/pnfs/pnf/T310-NE4>

OUTPUT:

```
<?xml version="1.0" encoding="UTF-8"?>
<pnf xmlns="http://org.onap.aai.inventory/v14">
  <pnf-name>T310-NE4</pnf-name>
  <pnf-id>059b7753-0c29-3407-a855-11174a1c77bc</pnf-id>
  <in-maint>false</in-maint>
  <resource-version>1550012822891</resource-version>
</pnf>
```

----- I expected to see the p-interfaces in the above output

The API to get the p-interfaces: **GET** <https://10.1.8.105:8443/aai/v14/network/pnfs/pnf/T310-NE4/p-interfaces>

OUTPUT:

```
<?xml version="1.0" encoding="UTF-8"?>
<p-interfaces xmlns="http://org.onap.aai.inventory/v14">
  <p-interface><interface-name>otu-1/2/0/E1.1</interface-name><resource-version>1550012823703</resource-version><in-maint>false</in-maint></p-interface>
  <p-interface><interface-name>otu-1/1/0/E1.1</interface-name><resource-version>1550012824476</resource-version><in-maint>false</in-maint></p-interface>
</p-interfaces>
```

<https://docs.onap.org/en/casablanca/submodules/aai/aai-common.git/docs/AAI%20REST%20API%20Documentation/nodesQuery.html#depth>




☒


James Forreth will fix the missing link on the bulk APIs documentation.

☒ **AAI-2457** - Add missing bulkApi file to AAI documentation


CLOSED

5	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>
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, cache, schema-service microservice</p>
7	AAI Questions	THIRILOSHINI KRISHNAKUMAR	DONE	7th Feb 2019	<p>We are trying to bring in transport network into ONAP as part of CCVPN usecase. We have some questions on AAI for the same.</p> <ul style="list-style-type: none"> - We are using a SDN controller which exposes TAPI Rest APIs. I'll have to map the TAPI model to ONAP AAI Model to add the topology to AAI. We are planning to use a mediation which does the model mapping and push the topology to AAI. Is this approach valid or Is there any other standard approach of doing it ? Is there any example that i can refer to where somebody has done mapping from one model to AAI model. - Are there any examples of AAI Rest API to add/delete/update nodes,its interfaces and links to and from AAI inventory that i can refer to ? - The AAI model, may i know if it follows any standards ? - In the model under physical link and logical link, i do not see any endpoint information. May i know how the links are formed when it doesnt know what the endpoints are - We have a onap cluster running Beijing version. Is there any standard approach of upgrading it to Casablanca. <p>Casablanca RESTAPI Documentation</p> <p>https://onap.readthedocs.io/en/casablanca/submodules/aai/aai-common.git/docs/AAI%20REST%20API%20Documentation/AAIRESTAPI_CASABLANCA.html</p> <p>https://nexus.onap.org/service/local/repositories/releases/archive/org/onap/aai/aai-common/aai-schema/1.3.4/aai-schema-1.3.4.jar!/onap/aai_swagger_html/aai_swagger_v14.html</p> <p>2019-02-06 AAI Meeting Notes Amit Sinha discussed an approach for upgrading version</p> <p>A&AI Data Restore for OOM - Casablanca</p>
8	Schema-service roadmap	James Forsyth Pavel Paroulek	ON HOLD	31th Jan 2019	<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>Open question:</p> <ol style="list-style-type: none"> 1. What schema abstractions are planned in schema-service? 2. When will the abstractions (from point 1.) be implemented? <p><input 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>Postponed until 21st February 2019 when the position of ecomp regarding schema-service functionality should be clear</p>
9	New AAF Certificates at startup	Jimmy Forsyth	IN PROGRESS	24th Jan 2019	<p>AAF will generate certificates to the 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 wether it is documented 31 Jan 2019</p>

10	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>
11	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> <div>  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 </div> <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 ls to saturate - Prudence Au and Sanjay Agharam 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> <div>  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 </div> <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>
12	get notified of AAI Cassandra issues automatically	Keong Lim	IN PROGRESS	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>Question: Keong Lim should this be part of a larger A&AI monitoring and failure prevention initiative</p>

13	Purpose of fields in AAI	Keong Lim	OPEN		<p>Dénes Németh wrote in  AAI-4104 - Invalid handling of external system certificates CLOSED :</p> <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> <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	Copyright license header restricted to 1 company	Pavel Paroulek	OPEN	24th Jan 2019	<p>In the copyright</p> <pre>/* * =====LICENSE_START===== * * org.onap.aai * * ===== * Copyright © 2017-2018 AT&T Intellectual Property. All rights reserved.</pre> <p>If there is a company other than AT&T the build fails saying the license header is wrong</p> <p>New cases found in gerrit reviews on aai/traversal:</p> <ul style="list-style-type: none"> ▪ https://gerrit.onap.org/r/#/c/78286/ ▪ https://gerrit.onap.org/r/#/c/78284/ <p>Keong Lim to try a license change in aai/resources to see if it also fails the same verify job.</p>
15	Return codes and messages for WS	Pavel Paroulek	OPEN		Is there a guide for the description of the error message and the error codes? How are new error states (message + code) added?
16	OOM Artifacts	Pavel Paroulek	OPEN		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?

17	AAI test data bootstrap	Keong Lim	OPEN	<div>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?</div> <div><h3>Generating AAI data</h3><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.json2. To add/remove HPA Capabilities edit the flavors section in the body of <code>PUT Cloud-Region{x}</code>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> <div>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:</div> <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> <div>Similarly:</div> <div><div><div>TEST-133</div><div>Demo.sh init step is not adding availability-zone information in AAI</div><div>CLOSED</div></div><div><div>INT-705</div><div>Add availability zone call to AAI in Robot framework</div><div>CLOSED</div></div><div><div></div><div>vCPE Use Case Tutorial: Design and Deploy based on ONAP</div><div></div></div></div> <div>Related <div><div>AAI-1948</div><div>Correct broken links</div><div>CLOSED</div></div> on the brittleness of the ReadTheDocs links to data files.</div> <div>One for VIM: How-To: Register a VIM/Cloud Instance to ONAP and <div><div>AAI-1928</div><div>AAI-Internal server Error:500</div><div>CLOSED</div></div></div> <div>Potential issues:</div> <div><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?</div>
18	AAI too slow for OOF/HAS	Keong Lim	OPEN	<div>Under OOF Homing and Allocation Service (HAS) section, Dileep Ranganathan wrote about Project Specific enhancements:</div> <div><div><h4>Optimize - AAI cache</h4><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</div></div> <div>See also <div><div>OPTFRA-268</div><div>OOF - project specific enhancements</div><div>CLOSED</div></div> / <div><div>OPTFRA-294</div><div>Optimize the use of AAI cache in HAS Data component</div><div>CLOSED</div></div></div> <div>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.</div> <div>What can we do about this?</div> <div>Would the AAI Cacher <div><div>AAI-1337</div><div>AAI Cacher</div><div>CLOSED</div></div> help to improve performance?</div>

19	MultiCloud usage of AAI for HPA telemetry/time-series data to OOF	Keong Lim	OPEN	<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>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>
----	---	-----------	------	---

20	Orchestration Scenarios for VNFs	Keong Lim	OPEN		<p>Comments on Orchestration Scenarios related to AAI:</p> <p>Viswanath Kumar Skand Priya / Viswanath KSP said:</p> <div> <p>Thank you Ranny Haiby & Fernando Oliveira . I agree partly, but I still have following queries.</p> <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 ? <p>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 ?</p> <p>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?</p> </div> <p>and Fernando Oliveira replied:</p> <div> <p>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).</p> </div> <p>As I understand, ESR has "esr-vnfm-list", which has an "esr-vnfm", which has "esr-system-info-list", which has a "relationship-list" that can contain relationships to "generic-vnf" and other AAI objects.</p> <p>The "generic-vnf" object also contains "self-link", "ipv4OamAddress", "ipv4OamGatewayAddress", etc, which links the AAI object back to its "source-of-truth" external-system.</p> <p>Is there some new data, new schema or new API that is required on top of this?</p> <p>Fernando Oliveira: Apologies for my lack of knowledge, but a few comments:</p> <ol style="list-style-type: none"> 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. 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. <p>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.</p> <p>ESR could be extended to handle VNFO as another kind of system.</p> <p>Fernando Oliveira : I created JIRA stories:</p> <div> <ul style="list-style-type: none"> <input checked="" type="checkbox"/> REQ-422 - ONAP tracking of available external (ETSI compliant NFVO) orchestrators DONE <input checked="" type="checkbox"/> REQ-423 - ONAP tracking of Service dependency on an external orchestrator DONE <input checked="" type="checkbox"/> ONAPARC-399 - ONAP tracking of VNF dependency on an external ETSI compliant VNF Manager (VNFM) CLOSED </div> <p>for various parts of the scenario.</p>
21	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 A AI Schema Service Use Case Proposals and AAI Schema Service.</p>

22	AAI too slow for Holmes	Keong Lim CT Paterson	IN PROGRESS	1st Nov 2018	<p>Guangrong Fu mentioned AAI in Baseline Measurements based on Testing Results:</p> <p>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.</p> <p>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.</p> <p>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?</p> <p>31st Oct: https://lists.onap.org/g/onap-discuss/topic/27805753</p> <p>1st Nov:</p> <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) HOLMES-186 - Improve the performance on AAI queries. CLOSED <p>Would the AAI Cacher AAI-1337 - AAI Cacher CLOSED help to improve performance?</p>
23	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 CLOSED</p> <p>?</p>
24	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> <p><input checked="" type="checkbox"/> James Forsyth creates a JIRA ticket to define the list of AAI subprojects and create the categories (essential, full "experience") for the OOM deployment AAI-2025 - Classify aai docker images CLOSED</p>
25	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?
26	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 Identifiedandsupportedpatternsandfeatures</p> <p><input checked="" 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>

27	named-query replacements	James Forsyth	IN PROGRESS	31th Jan 2019	<div> AAI-1989 - AAI Named Query Deprecation OPEN </div> <div> APPC-1245 - AAI named-queries are being deprecated and should be replaced with custom-queries CLOSED </div> <p>: Takamune Cho pointed to inherited CCSDK functions, so added a new case to follow-up there too</p> <div> POLICY-1278 - AAI named-queries are being deprecated and should be replaced with custom-queries CLOSED </div> <div> VID-355 - AAI named-queries are being deprecated and should be replaced with custom-queries OPEN : Itta y Stern has provided an update with 2 known queries used. </div> <div> CCSDK-962 - AAI named-queries are being deprecated and should be replaced with custom-queries CLOSED </div> <p>: Raised to Dan Timoney based on info about APPC component.</p> <p>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</p> <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 <p>Christopher Shang is leading the effort to deprecate named queries in e-Comp.</p> <p>Next steps: get data from teams in order to prepare for testing the change from named queries to custom queries.</p> <p>See also AAI Named Queries</p> <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-queries <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)
28					