

2019-01-24 AAI Developers Meeting

Attachments (including meeting recording)

File	Modified
Multimedia File zoom_0.mp4	Feb 11, 2019 by James Forsyth
Text File chat.txt	Feb 11, 2019 by James Forsyth


[Download All](#)

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	Schema-service roadmap	James Forsyth Pavel Paroulek	ON HOLD	24th 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. 31 Jan 2019</p>
2	Graceful handling of AAF run-time errors	Keong Lim	OPEN		<p>Raised concerns in</p> <div> AAAF-722 - aaf continues to be available to aai-resources even though aaf database appears to be down</div> <p>CLOSED</p> <p>, but it was closed as being "by design".</p> <p>Does AAI have the ability to handle this run-time error more gracefully?</p> <p>Can AAI bubble up notifications to clients and operators?</p>
3	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>
4	AAI Backup and Restore	Keong Lim	ON HOLD	10th Jan 2019	<p>Brian Freeman asked on Re: Backup and Restore Solution: ONAP-OOM :</p> <div><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></div> <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>

5	Dublin use-cases	Chandra Cinthala	DONE	24th Jan 2019	<p>5G Use Case (Dublin)</p> <p>BBS Broadband Service Use Case (Dublin)</p> <p>CCVPN Use Case (Dublin)</p> <p>Change Management Dublin Extensions</p> <p>Edge Automation Functional Requirements for Dublin</p> <p>OpenSource Access Manager (OSAM) Use Case</p> <p>K8S based Cloud Region Support</p> <p><input checked="" type="checkbox"/> Chandra Cinthala will go over all these use cases and will write a 1 page report per use-case where he lays out the impact and needed high-level description of changes in order to support the use-cases in A&AI. He will send this reports to James Forsyth so he can plan accordingly. The 1 pagers should be discussed with the modelling comitee before sending them to James Forsyth . Chandra Cinthala should start here https://wiki.onap.org/display/DW/Joining+the+ONAP+Technical+Community when integrating to the ONAP community and the mailing list https://lists.onap.org/g/main. The deadline for Chandra Cinthala to send the reports to James Forsyth is 15 Jan 2019</p> <p>The Dublin use-cases for review are Release 4 (Dublin) Use Cases and functional requirements</p> <p>On AAI weekly call on Wednesday Chandra Cinthala will report on his progress in analyzing the use-cases.</p>
6	CCVPN Extension Dublin	Davide Cherubini Susana Sabater emmanuel sarris Christopher Sheng Tony Noori	DONE	24th Jan 2019	<p>During the VF2F we have presented the CCVPN Use Case extension for Dublin Release.</p> <p>In particular, we have defined the E-LAN Service (EP-LAN, EVP-LAN) sub-Use Case. More details can be found at this page (the presentation given to the Use Case Subcommittee and during the VF2F can also be found in that page).</p> <p>We have also created 2 additional pages where we describe the:</p> <ol style="list-style-type: none"> 1. Implementation Options 2. Initial Impact assessment <p>In the former page, we have identified Option 2 (for Service Creation) and Option A (for Service Modification) as the "ideal" implementation. The Impacts in the latter page are based on such an assumption.</p> <p>We would be grateful if you could help us with the assessment for A&AI and if you could confirm the project commitment to CCVPN.</p> <p>14th Dec: Regarding the Options and Impact assessment linked above:</p> <div> <p>AAI</p> <ol style="list-style-type: none"> 1. <i>Support Composite Service Instances (multi-operator)</i> </div> <p>As far as I know, the ServiceComposite pattern hasn't even exited the Modelling team's discussions for R3 or R4 workshops Hui Deng :</p> <ul style="list-style-type: none"> ■ Composite vs Recursive Discussion Page ■ Composite Pattern UML diagram <p>This would seem like a stretch goal given the impacts on all the downstream clients of AAI, if it would require reworking any existing schema elements.</p> <p>However, if you are looking for a zero-impact method of implementation, the EdgeRules already contain service-instance-to-service-instance relationship with label "org.onap.relationships.inventory.ComposedOf", so AAI can do the composite data with this relationship.</p> <p><input checked="" type="checkbox"/> James Forsythwill send a reminder e-mail to Tony and Christopher to engage them on this task, the plan is to discuss this issues and modeling in general in the first 2 weeks in January</p> <p>Update 3rd January 2019 from CFL comment:</p> <div> <p>LIN MENG</p> <p><i>Hil think there might be some confusions in CCVPN use case and the current service IM in the sheet.</i></p> <p><i>I don't think CCVPN has extension demands for A&AI in Dublin release.</i></p> <p><i>What we want to do in D release is to provide a CCVPN service by using one service template to includes the SOTN and SD-WAN resources, which will be a requirement for SDC and modeling DM (In C release, we have to design each resource as a service based on SDC's ability)</i></p> <p><i>We've actually talked about composite pattern at the beginning of R3 and the modeling team accepts this idea but have no implementation in this concepts.</i></p> <p><i>Also, I think the composite pattern is the same as nested service from service IM perspective. And I think it would be better if we unify them as one concept in ONAP for less confusion, which I will propose to next week's service IM call. Back to the point, we're in redesigning the new service IM of CCVPN and currently, CCVPN doesn't have this need.</i></p> <p><i>Pavel Paroulek James Forsyth Hui Deng Davide Cherubini jianguo zeng Chuanyu Chen</i></p> </div>

7	AAI behaviour in error conditions	Keong Lim	OPEN	24th Jan 2019	<p>Just raised</p> <div>  AAI-2882 - aai-resources gives incorrect output when aai-cassandra has shutdown with failure CLOSED </div> <p>apart from getting a fix, if possible, may need some help with recovery actions.</p> <p>Moved from AAI Weekly Status Meeting.</p> <p>Question: Keong Lim can this be closed?</p>
8	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 <div> <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> </div> <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 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> <div> <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 </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>
9	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> <div> <p><i>f. AAI team wanted to get notified of AAI Cassandra issues automatically</i></p> <p><i>i. Can we setup a Nagios or equivalent to monitor both rancher/k8 and the applications for rancher/k8 issues ?</i></p> </div> <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>
10	Purpose of fields in AAI	Keong Lim	OPEN		<p>Dénes Németh wrote in <div> AAI-4484 - Invalid handling of external system certificates CLOSED</div> :</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"> 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>

11	Postel's Law / Robustness Principle for AAI data inputs	Keong Lim	DONE	24th Jan 2019	<p>While helping UII team to debug a family of related issues such as  USECASEUI-183 - create ext-aai-network api can not accept json request  , found that their JSON request bodies had inserted extra lines that were not acceptable to AAI, e.g.</p> <div><h3>Not Accepted Input</h3><pre>{ ext-aai-network: { "aai-id": "VDF", "schema-version": "version-1", "resource-version": "1542082337501" } }</pre></div> <p>vs</p> <div><h3>Expected Input Format</h3><pre>{ "aai-id": "VDF", "schema-version": "version-1", "resource-version": "1542082337501" }</pre></div> <p>Initially, I thought this could be due to XML-to-JSON translation error, but it could also have been a copy-paste from the output of a GET, e.g.</p> <pre>> GET /aai/v14/network/ext-aai-networks HTTP/1.1 { "ext-aai-network": [{ "aai-id": "createAndDelete", "schema-version": "version-1", "resource-version": "1542247826990" }, { "aai-id": "aaiId-2", "schema-version": "version-2", "resource-version": "1542029867153" }] }</pre> <p>In the spirit of Postel's Law ("be liberal in what you accept"), could/should AAI accept both variations of the input data above?</p> <p>There could be additional validation that the element name inside the request body matches the element name in the URL of the API (similar validation is already performed for the key ID value).</p> <p>It would also allow for simple methods of data transfer/migration, where AAI output is directly accepted as AAI input.</p> <p>7th Dec: As per Jimmy's comment below, the difference between the "Not Accepted Input" and "Expected Input Format" is the behaviour controlled by "SerializationFeature.WRAP_ROOT_VALUE" as described in:</p> <ul style="list-style-type: none">▪ https://stackoverflow.com/questions/37882478/serializationfeature-wrap-root-value-as-annotation-in-jackson-json▪ https://fasterxml.github.io/jackson-databind/javadoc/2.6/com/fasterxml/jackson/databind/SerializationFeature.html <p>Result: A&AI has to pick one or the other way to accept JS objects. It was decided that we go for the non-wrapped version (in what we send and receive). Changing this setting would break the code of existing clients.</p>
12	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>

13	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?
14	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?
15	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> <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.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> <p>Similarly:</p> <ul style="list-style-type: none"> ▪ TEST-433 - 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 AAI-1928 - 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? </div>
16	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"> • 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> <p>See also OPTFRA-268 - OOF - project specific enhancements CLOSED / 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>

17	MultiCloud usage of AAI for HPA telemetry/time-series data to OOF	Keong Lim	OPEN	<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
18	AAI-BBS for Dublin	Keong Lim	OPEN	<p>Started collating proposed changes in wiki AAI-BBS Proposals for Dublin Release</p> <p>Investigating code proposal 1:</p> <p>In AAI, when the "orchestration-status", "service-instance-location-id", "input-params" for a Service Instance changes, notify ExtAPI so it does not need to poll AAI</p> <p>Is it already covered by spike? If so, how to consume the events? If not, how to make AAI publish new events to DMaaP?</p>

19	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><i>Thank you Ranny Haiby & Fernando Oliveira . I agree partly, but I still have following queries.</i></p> <ul style="list-style-type: none"> <i>I agree & acknowledge that atleast for a foreseeable future, we would need a way to specify the VNF / 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 ?</i> <p><i>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 ?</i></p> <p><i>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?</i></p> </div> <p>and Fernando Oliveira replied:</p> <div> <p><i>For the first question: I think that A&AI needs to maintain the VNF instance VNF 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 VNF instance for the VNF instance) SO (VNF instance, VNF instance, Scale-out) SOL003 Adapter (VNF instance, VNF instance, Scale-out) VNF instance (VNF Instance, Scale-out).</i></p> </div> <p>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-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 VNF case, I think that there needs to be a reference from a VNF/VF instance record to the specific instance of the VNF 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 VNF 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, VNF 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>
20	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><i>9)"vhn-portal-uri"? "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)</i></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>

21	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> <div> 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> <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>
22	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> LOG-877 - S3P: Logging streaming/format alignment for dublin - China Telecom, Deutsche Telekom, Vodafone CLOSED </div> <p>?</p>
23	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>
24	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?
25	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> <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>

26	named-query replacements	James Forsyth	IN PROGRESS	24th Jan 2019	<div> AAI-1989 - AAI Named Query Deprecation OPEN </div> <div> APPC-4245 - 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-4278 - 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 : Ittay 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
27	Spring Boot Upgrade	James Forsyth	IN PROGRESS	24th Jan 2019	<p>https://spring.io/blog/2018/11/29/spring-boot-1-5-18-available-now - support for 1.5.x spring boot will end in Aug 2019 so we should consider upgrading.</p> <p>We should update to spring 2.X.X.</p> <p>On the meeting on the 10th Jan discuss if anyone has anything against it, if not then James Forsyth creates a JIRA item for Dublin for upgrading the spring in all A&AI repositories.</p> <p>Conclusion:</p> <ul style="list-style-type: none"> <input type="checkbox"/> Venkata Harish Kajur will check with Manisha Aggarwal if the update to 2XX version of spring boot will happen in Dublin 24 Jan 2019