

2018-11-29 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
AAI Tutorial updates	Keong Lim	IN PROGRESS	29th Nov	<p>Following updates for AAI Developer Environment Setup - Casablanca need to update the link in step 14:</p> <p>Tutorial: Making and Testing a Schema Change in A&AI</p> <ul style="list-style-type: none">Need separate pages for Beijing, Casablanca, Dublin, etc. AAI-1973 - Schema update wiki is out of date of Casablanca CLOSED <p>Harish will update the documentation - the section Tutorial: Making and Testing a Schema Change in A&AI</p> <p>Need to distinguish on the wiki tutorials for:</p> <ul style="list-style-type: none">running A&AI (just installing for non-devs)setting up a dev environment to develop AAI ("Developer Onboarding" section on wiki) <p>James Forsyth will check if we still need to support HEAT deployment for Dublin</p>
aai-cassandra performance issues	Keong Lim	OPEN		<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>
Schema Service	Manisha Aggarwal	IN PROGRESS	29th Nov 2018	<p>Discuss about the Schema Microservice</p> <p>AAI Schema Service</p> <p>11th Oct: Suggested Use Case Proposals for Dynamic AAI Schema Changes based on CCVPN usecase experience</p> <p>1st Nov: William Reehil Robby Maharajh Venkata Harish Kajur will review requirement updates and research the open questions so that a final draft can be prepared for implementation</p> <p>8th Nov: Added AAI Schema Service Use Case Proposals for discussion and planning</p> <p>15th Nov: Reviewed the Requirements section in AAI Schema Service again.</p>

AAI GraphGraph	Keong Lim	IN PROGRESS	29th Nov 2018	<p>William Reehil wrote introduction to (proposal for?) A&AI GraphGraph</p> <p><i>What must the solution provide?</i></p> <ul style="list-style-type: none"> • a generic interface to interact with the schema and edge information, so it can be accessed via end users and microservices • ability to be configured with any schema given in the set formats (JSON for edge rules, OXM for schema) • ability to easily communicate to an end user a node type's attributes and edge rules when provided with a node type as input <p>Looks like an API to reflect on the schema from an instance in the database.</p> <p>Is there some overlap with the AAI Schema Services?</p> <p>Is this leveraging AAI Schema Services e.g. as client? proxy? facade? implementation detail?</p> <p>Is it an alternative to AAI Schema Services?</p> <p>Update: Keong Lim I do see an overlap with the schema service mainly for the retrieval of the data, but GraphGaph will offer more on top of that(UI, NLP), during the POC code for GraphGaph there was no schema service. Implementation details can be discussed on our call for how exactly to leverage the schema service, the poc code, and where the core logic resides for this functionality.</p> <p>15th Nov: Reviewed A&AI GraphGraph and agreed that there is now overlap with AAI Schema Service that should be addressed in updating the POC version.</p>
Spring profiles	Pavel Paroulek	DONE	29th Nov	<p>Do we need the "production" profile?</p> <p>Pavel Paroulek - creating JIRA task to remove the profile (resources, traversal, cachier, graphadmin, oom, testconfig, schema-service)</p>
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>
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>
Return codes and messages for WS	Pavel Paroulek	OPEN		<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>
Copyright license header restricted to 1 company	Pavel Paroulek	OPEN		<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>
AAI Traversal queries	Pavel Paroulek	OPEN		<p>The AAI Traversal API is fairly complex, are there examples and descriptions of what the inputs and outputs of the rest api are, or alternative way to gain know-how? Code is quite difficult to read.</p>

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> <hr/> <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) <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> <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 TEST-133 - Demo.sh init step is not adding availability-zone information in AAI CLOSED and INT-705 - Add availability zone call to AAI in Robot framework CLOSED.</p> <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?
AAI to slow for OOF /HAS	Keong Lim	OPEN	<p>Under OOF Homing and Allocation Service (HAS) section, Dileep Ranganathan wrote about Project Specific enhancements:</p> <p><i>Optimize - AAI cache</i></p> <ul style="list-style-type: none"> • Use MUSIC or any other alternative in memory caching like Redis etc? • Optimize flavor retrieval from A&AI and Cache the information if necessary <p>See also OPTFRA-268 - OOF - project specific enhancements CLOSED / OPTFRA-204 - 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>

<p>Relationship between POMBA Common Model and AAI schema?</p>	<p>Keong Lim</p>	<p>OPEN</p>	<p>Under POMBA Common Model, Geora Barsky and Sharon Chisholm discuss objects that seem to overlap the AAI schema:</p> <ul style="list-style-type: none">• Network has direct relationships with PNFs and physical links• PNF has direct relationship with p-interfaces and l-interfaces <p>Curious to know what is the relationship between the POMBA Common Model and the AAI schema? There seems to be an overlap in these object definitions and relationships.</p> <p>Is POMBA a potential client for AAI schema services?</p> <p>15th Nov: As per Geora's and Sharon's comments:</p> <p><i>POMBA is a client of AAI APIs. It retrieves certain objects from AAI and transforms it into POMBA COMMON model which is aimed to represent flat structure of service instance representation.</i></p> <p>and</p> <p><i>The POMBA model enables us to normalize data from different data sources to facilitate auditing. To save time and work, we often use A&AI as our starting point, but conceptually, this model can have more than is currently or makes sense to have in A&AI.</i></p> <p>So, POMBA needs to evolve independently of AAI, even though it has common ancestry.</p> <p>As we are having discussions about the AAI Schema Services with a view to future dynamic schema updates, e.g. via SDC modelling, we will need to be aware of the downstream impacts of such a change.</p> <p>I think this needs to be added to the use cases for AAI Schema Services (and GraphGraph?) behaviours.</p>
--	------------------	-------------	---

OPEN

While helping UII team to debug a family of related issues such as

 [USECASEUI-483](#) - create ext-aa-network api can not accept json request CLOSED, found that their JSON request bodies had inserted extra lines that were not acceptable to AAI, e.g.

Not Accepted Input

```
{
  ext-aa-network: {
    "aai-id": "VDF",
    "schema-version": "version-1",
    "resource-version": "1542082337501"
  }
}
```

vs

Expected Input Format

```
{
  "aai-id": "VDF",
  "schema-version": "version-1",
  "resource-version": "1542082337501"
}
```

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.

```
> GET /aai/v14/network/ext-aa-networks HTTP/1.1
{
  "ext-aa-network": [
    {
      "aai-id": "createAndDelete",
      "schema-version": "version-1",
      "resource-version": "1542247826990"
    },
    {
      "aai-id": "aaiId-2",
      "schema-version": "version-2",
      "resource-version": "1542029867153"
    }
  ]
}
```

In the spirit of Postel's Law ("be liberal in what you accept"), could/should AAI accept both variations of the input data above?

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).

It would also allow for simple methods of data transfer/migration, where AAI output is directly accepted as AAI input.

MultiCloud usage of AAI for HPA telemetry /time-series data to OOF	Keong Lim	OPEN		<p>Bin Yang and Lianhao Lu wrote in MULTICLOUD-374 - HPA telemetry data collection and persistence CLOSED :</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)?
Purpose of fields in AAI	Keong Lim	OPEN		<p>Dénes Németh wrote in AAI-4404 - Invalid handling of external system certificates CLOSED :</p> <p>In think it would be good to answer what is the meaning of the field (collection of PEMs of the CA xor URL)</p> <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>
HAproxy in OOM	Pavel Paroulek	IN PROGRESS	1st Nov 2018	<p>Why is the pod for HAproxy not named (hard to figure out that there is a proxy), unsure how it is logging and where</p> <p>James Forsyth creates JIRA tasks to 1. to have the pod named 2. add logging to the proxy</p> <p><input checked="" type="checkbox"/> AAI-4807 - Change the aai pod name to aai-haproxy CLOSED</p> <p><input checked="" type="checkbox"/> AAI-4840 - Add option to enable logging for haproxy docker CLOSED</p>
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> <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> <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>

Jackson Replacement	Team Keong Lim Tian Lee	ON HOLD	1st Nov 2018	<p>Security subcommittee has recommended teams move away from jackson, and will be presenting alternatives and asking for an assessment from each project. Our team will need to do an analysis - this would not be trivial, especially given how many of our repos are impacted. As of now, this would be a very high LOE for the team, we need to understand what the recommendation from the SECCOM is before we can provide better details on what the LOE would be.</p> <p>Updated: Using Google gson vs FasterXML Jackson</p> <p>10th Oct: Present to Seccom meeting</p> <p>15th Oct: Present to PTL meeting</p> <p>31st Oct: Debatable whether the cost of swapping Cassandra and changing code is worth the benefit of removing Jackson from the vulnerabilities list.</p> <p>On-Hold until James Forsyth consults with other PTLs: PTL 2018-11-05</p>
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> <ol style="list-style-type: none"> 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>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) <p>Would the AAI Cacher AAI-1337 - AAI Cacher CLOSED help to improve performance?</p>
2 Types of logging in A&AI WS	Pavel Paroulek	ON HOLD	1st Nov 2018	<p>There are 2 types of logging in the services</p> <ul style="list-style-type: none"> one read from EELFManager the other Logger log = Logger.getLogger(... <p>Is that correct? Shouldn't there be just 1 type?</p> <p>1st Nov:</p> <p>After Casablanca release investigate logging guidelines and figure out what library to use in order to unify logging within A&AI</p> <p>26th Nov: See also ONAP Application Logging Specification - Post Dublin</p> <p>29th Nov: how does this fit with</p> <p>LOG-877 - S3P: Logging streaming/format alignment for dublin - China Telecom, Deutsche Telekom, Vodafone</p> <p>CLOSED</p> <p>?</p>
Disable unused web services	Pavel Paroulek	IN PROGRESS	1st Nov 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 than extended group will all the services. Consider a group of unused/unintegrated services (like Champ). Consider other possible groups (like GUI?)</p>
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?
Dublin 5G Use Case	Benjamin Cheung	IN PROGRESS	25th Oct 2018	<p>Dublin AAI changes in support of 5g use cases.</p> <p>Link for presentation: 5G - PNF Plug and Play (Casablanca carry-over items)</p> <ol style="list-style-type: none"> NF A&AI: S/W Image Repository Info NF A&AI: Cloud Home Info NF A&AI: External Manager IP Address NF A&AI: S/W Version Object Association