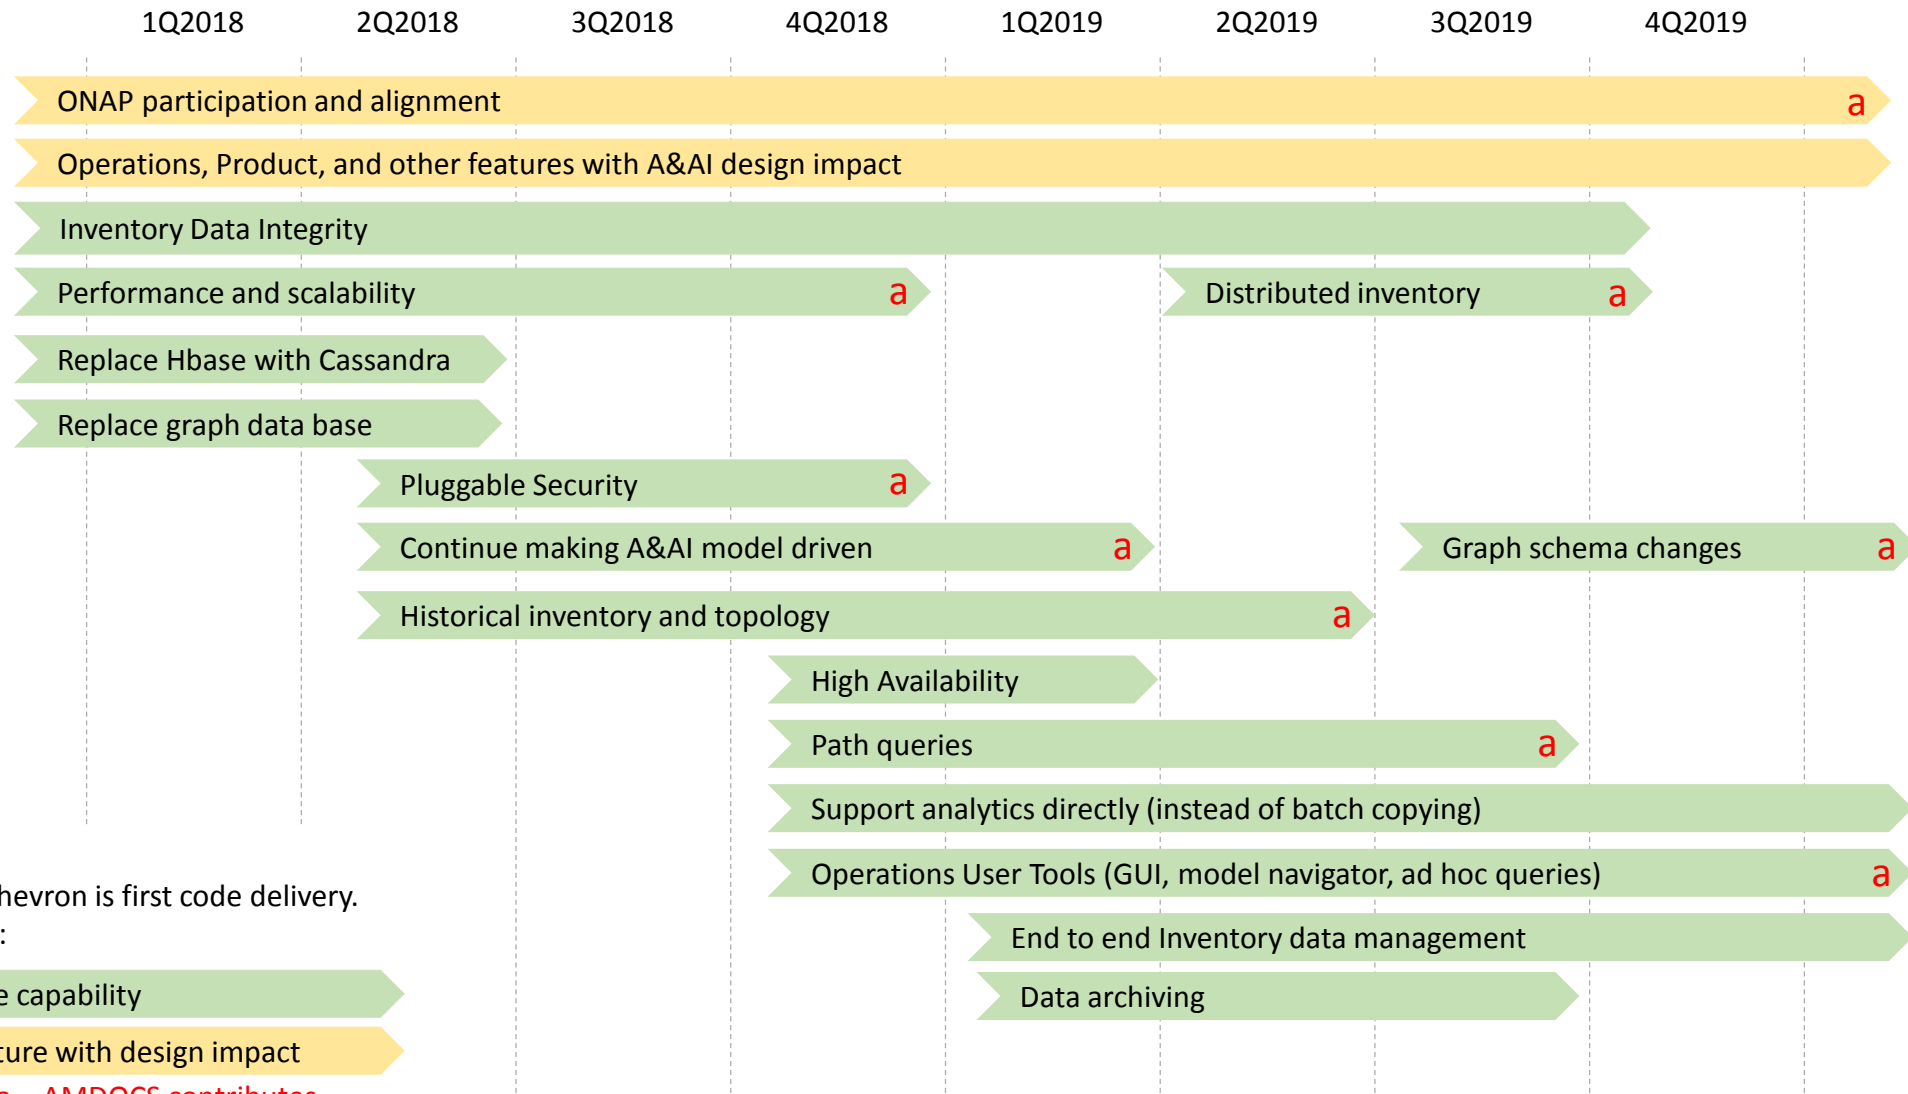


Inventory (input to ECOMP and ONAP Roadmaps)



Left end of chevron is first code delivery.
Color coding:

Core capability

Feature with design impact

a = AMDOCS contributes

Next steps (alphabetical order)

Continue making A&AI model driven – socialize in ONAP and work to select a standard TOSCA parser within ONAP
Data Archiving – document use cases, corporate retention requirements
Distributed inventory – align with DCAE plans
End to End Inventory Data Management – define target and roadmap
Graph schema changes – socialize in ONAP
High Availability – align with other proposals to make ONAP and ECOMP components high availability
Historical Inventory and topology – review architecture recommendations, evaluate current time stamps available
Inventory data integrity – establish governance process, proposed roadmap for fixing known problems
ONAP participation and alignment - socialize A&AI roadmap and recruit collaborators
Operations, Product, and AIC features with A&AI design impact – review Ops requirements for inventory
Path Queries – Define or select a DSL – Domain Specific Language that will simplify queries
Performance and scalability – continue with assessments, modifications (1806), clarify use cases for DCAE analytics
Replace graph data base – finish evaluation, deliver JanusGraph
Replace Hbase with Cassandra – finish evaluation, deliver feature
Security – Socialize in ONAP
Support analytics directly – deliver queries for DCAE enrichment and vertical correlation from VNF

Continue making A&AI model driven

Goal: derive A&AI graph schema and APIs from SDC models

A&AI currently uses Java reflection techniques to read an OXM text file and create graph node data structures without having to restart the data base.

This core capability will be extended to ingest SDC models and apply them (e.g. convert SDC formats to a format readable by Java reflection).

Other areas of enhancement include:

- making graph edges model-driven
- separating out the data presentation (APIs) from data storage (graph schema)
- Support model-driven backward compatibility

ONAP: this work should be of great interest to ONAP participants.

Critical path dependency: we need to settle on a single TOSCA parser within ONAP.

Next steps:

- Socialize within ONAP, and secure agreement on single TOSCA parser

Data archiving

Goal: archive A&AI inventory and topology data for long-term use

ONAP: assume other ONAP users will archive the data in their own environment; retain a clean separation of the archive mechanism vs. transactional data

Status:

- Rules have been defined for coordinating with historical inventory and topology, including use cases for changing the retention interval

Next steps:

- document use cases and corporate data retention requirements

Distributed Inventory

Support low-latency transactions for DCAE by placing inventory data co-resident with DCAE (which has multiple locations).

ONAP: review with ONAP team

For use by DCAE, co-locate A&AI with DCAE and store inventory data for the portion of the network that the local instance of DCAE will be getting network events from. For simplicity, this scope of data can be defined as a set of locations.

Status:

- DCAE is currently copying inventory data instead of querying A&AI directly.

Next steps:

- Align with DCAE
- Define data flows for inventory updates

End to end Inventory data management

Goal: apply data federation, replication, and other data management techniques to the end to end inventory data flows.

ONAP: this work will be pursued under the existing ONAP reference architecture (e.g. A&AI and DCAE)

Includes hybrid services, transactions, etc.

Status:

- Some issues are already documented (e.g. many separately-managed copies of the same inventory and topology data)
- List of key questions about federation is documented (Phil Blackwood)

Next steps:

- Define target and roadmap
- Align with DCAE and controllers
- Review current documented principles for inventory data sourcing
- Define when to use enrichment feature vs. other federation techniques

Graph schema changes

When a new graph schema is introduced, existing data must be converted and clients do not want to experience A&AI downtime.

Goals:

- Convert existing data (this is not always done today, e.g. when set of valid values changes)
- Provide backward compatibility per ECOMP/ONAP standards
- Streamline the process
- Minimize or eliminate downtime for clients

ONAP: collaborate with other ONAP contributors

Status:

- AMDOCS provided a proposal that includes queuing updates and ability to replay as needed to support migration
- Whenever SDC models are changed, translation from the old model to new model should also be provided

Next steps:

- Socialize in ONAP

High Availability

Currently, A&AI high availability is limited.

ONAP: coordinate closely with ONAP participants.

Status:

- Current A&AI support for high availability comprises:
 - Local Hbase/HDFS clusters within a Site
 - One active site + customized data replication to standby site
 - Manually invoked failover scripts (requires about 20 minutes)
- Finish the migration to JanusGraph and Cassandra before implementing this feature

Next steps:

- Align with other proposals for high availability of all ONAP/ECOMP components

Historical inventory and topology

Goal: Provide historical views of inventory and topology including past views that have been corrected

ONAP: plan to contribute to ONAP.

Status:

- Evaluated 5 options. Recommend basing the feature on Gremlin time queries.
- Architecture recommendations are available (Phil Blackwood).

Next steps:

- Review architecture recommendations
- Take to ONAP A&AI sub-team
- Evaluate currently available timestamps
- Take recommendation for timestamps from sources to ONAP architecture sub-team

Inventory Data Integrity

Work areas:

- Comparison audits that can sync data to its source
- Customer identifiers
- Data provenance
- Data quality metrics
- Data validations
- Fix root causes of known problems
- Improve data handoffs (guaranteed delivery, retry strategy, etc.)
- Migrate older data to latest version of schema, populate defaults
- Notification and process for correcting of errors
- Set standards for data sources
- Set standards for data migration
- Track sources of data (data provenance)

Status:

- List of known problems exists

Next steps:

- Establish roadmap for fixing root causes of known problems
- Establish governance process

draft

ONAP participation and alignment

Goal: ONAP first. Core capabilities should be socialized with ONAP team and built in ONAP environment.

Status:

- AT&T/AMDOCs inventory roadmap was shared with ONAP A&AI team.
- Orange expressed interest in working together on historical inventory, data integrity, and methods for integration with legacy inventory

Next steps:

- Continue working with ONAP A&AI team

Operations, Product, and AIC features with A&AI design impact

Identify any architecture issues or enhancements to core capabilities that are needed to support features requested by Operations and Product Management, or needed to support AIC changes.

Status:

- Hybrid services raise issues about data federation and data replication
- Some services needs A&AI to provide views of service that are presented to external customer during service configuration

Next steps:

- Identify any design or architecture impacts in 2018 feature list
- Identify core capabilities that are of interest to multiple ONAP participants

Operations User Tools

Support near-term operations needs for access to inventory data.

Status:

- Operations users have highlighted the need for near-term support
- Prototype of graph graph concept (from ONAP community) lets user navigate graph schema and see definitions of nodes, edges, and attributes

Next steps:

- Set regular meetings with stakeholders to identify high-value near-term deliverables (quick hits)

Path queries

Path queries include paths along service chains and more dynamic paths taken by packets through the network.

ONAP: we expect to contribute this feature to ONAP

Status:

- Complex service modeling includes some path modeling
- Proof of concept was built including NEPAL query language

Next steps:

- Define or select a DSL – Domain Specific Language that will simplify queries

Performance and Scalability

The current A&AI implementation needs to be enhanced to meet performance and scalability goals.

Good performance is especially important to enable DCAE to start using A&AI queries instead of creating separate copies of the data.

ONAP: findings from evaluations should be shared with ONAP team

Status:

- AMDOCS provided software architecture recommendations (e.g. efficient processing, move some functions to lower level of processing)

Next steps:

- Document non-functional requirements
- Get examples of DCAE analytics that can be used to test performance. (cases that are available in ONAP community)

Pluggable Security

Goal: improve security of transactions in A&AI consistent with ONAP and other ECOMP components

ONAP: coordinate closely with ONAP participants.

Status:

- Draft proposal is available and has been reviewed with key stakeholders in AT&T

Next steps:

- Socialize in ONAP

Replace Graph Data Base

Currently, A&AI uses the open source Titan graph data base, which no longer has an active open source community supporting and improving it. Since Titan is essentially defunct, it is necessary for A&AI to migrate to a different graph data base.

ONAP: coordinate closely with ONAP participants.

Status:

- DSE and JanusGraph are the two remaining candidates.
- Issues with transaction handling by DSE graph
 - DSE-Graph does not support TinkerPop transaction API.
DSE transactions are implicit, and do not provide Read Consistency
- Prefer to move to JanusGraph
 - Must understand level of engagement by Janus community, and availability of commercial maintenance services (Horton works, IBM, etc.)
- AMDOCS has done evaluation of JanusGraph
- Plans have been socialized in ONAP and there is alignment on migration to JanusGraph

Next steps:

- Migrate to JanusGraph.

Replace Hbase with Cassandra

Cassandra has been identified as a better data persistence solution than Hbase.

ONAP: coordinate closely with ONAP participants.

Status:

- Evaluations are complete.

Next steps:

- Migrate from Hbase to Cassandra

Support Analytics Directly

Currently, A&AI and DCAE support separately-managed copies of network inventory data.

ONAP: socialize with ONAP community and partner with those who want their analytics to directly query A&AI (they might provide use cases for performance studies and tuning). DCAE ONAP work should steer ONAP in this direction and not in the direction of copying A&AI data.

Goal:

We want DCAE analytics to use A&AI inventory directly (instead of using a separately-managed copy of the inventory data).

ONAP: prepare to accommodate a variety of scenarios. Many operators will have legacy applications and want a variety of ways to replicate data for their use.

Status:

- DCAE and A&AI teams have been working together to document existing flows and direction

Next steps:

- Get requirements from DCAE for enrichment and vertical correlation of VNF

