

Table of Contents

Table of Contents	1
Attendees	1
Donna's roadmap review	2
Blaine McDonnell vOLTHA 2.0 / Vanilla Architecture	3
Parking lot	5

Attendees

Ito Tetsuro NTT-East.

Donna Reineck-Whelan vOLTHA scrum master.

Shawn Ying. Product Owner

Chip Boling, Adtran developer. Been working for 6 months on vOLTHA.

Julie Lorentzen, AT&T. Reports to Bill Hurst. Editor for OMCI in FSAN.

Kent McCammon, AT&T San Ramon. Moved from Wireless back to PON, working for Mounier with Shawn.

Sisir Chowdhury, Radisys SW Architect.

Kim Kempf, System Architect at Radisys. Open white-box solns for vOLTHA.

Frank DiMayo, Broadband Adaptation layer, BAL

Sam Chen, BRCM PON MGR, OMCI

Jim McKeon, Brcm, marketing

Dave Baron, BRCM

Sireesha Kora, Nokia

Sergio Slobodrian Ciena, Dir Dev, MDO CORD, ONOS

David Bainbridge, Ciena.

Tom Anschutz, DMTS AT&T

Ali Al-Shabibi. Cerra.

Bill Snow ON.Lab engineering.

Jeff Gatlan Edge-Core.

Steve Wallace, Broadcom Maple engineer.

On the phone:

Mounier, AT&T.

Mark Shotack. AT&T design.

Blaine McDonnell

Michel, Ciena

Steve Crooks, Ciena

Donna's roadmap review

Objective for today: Agree on what we'll include on a time -based release. We'll populate the user stories and story points during our Tuesday meetings.

What is the definition of done?

Saurav: There is an OLT simulator and emulator. If it runs against OFTest and the simulator and emulator, that is done.

PONSim wont generate PMs and alarms, but does support traffic.

The emulator will generate PMs and alarms.

Both need to be used for testing. You'll get 90% coverage.

Each vendor is responsible for testing with their own hardware.

July 28th is the target feature freeze for vOLTHA 1.0 release..

Phase 3 is pushed out by 3 weeks to start August 14th.

See the notes section of the presentation here for more notes:

Blaine McDonnell vOLTHA 2.0 / Vanilla Architecture

YANG Explorer: YANG GUI from Cisco.

External and federated data store: Could be consul, zookeeper

- Local in-memory store manages only the instances of OLTs that vOLTHA has under control.
- Federated datastore contains everything.
- What manages the vOLTHA cluster? I.e., what manages the federated datastore? TBD.

Performance Management

- There is an existing model that specifies the PM data structures.
- If there is a list of required PMs. The adapter would register the PMs it will support. Once registered, the device can report those which it has informed vOLTHA it will report.
- It will be possible to prune PMs that are not desired. The adapter won't report pruned PMs.
- There doesn't necessarily even need to be a pre-ordained list of PMs hard-coded into vOLTHA. When the adapter registers, it registers with the capable PMs.
- Up to the adapter vendor to determine how to mute the PMs. Could be muted at the device or the Kafka bus.
- Hardware all reports KPIs differently. Should the adapter normalize those, or should that be done in vOLTHA core?
 - The issue with doing it in the adapter is that you lose resolution. Some of the translations are lossy.
 - Sergio's opinion is that this should not be done in the vanilla core. It's post-processing / mathematical translations / pattern matching should be done outside vOLTHA. Elastic data store, for example. Don't pre-cook data. It's irreversibly lossy. You could add containers that cook the data. There are already other tools that do this.
- Are PMs object model driven, or API driven? Today, it's model-driven, protobuf based at compile-time.
 - Auto-generating interfaces produces more change, but the benefit is leaner, more flexible code.

High-level architecture

- Saurav Das indicates the vOLTHA 2.0 / Vanilla architecture looks more like ODL. Are we doing a rip-and-replace and changing the architecture for 2.0?
 - No. Unintentional change. Will stick with gRPC.

Protobuf vs. REST vs. NETCONF

- Today Chameleon re-packages REST messages into gRPC

- You don't need Chameleon if you don't want a REST interface and want to use gRPC directly instead.
- In protobuf version 3, dynamic definition not supported. Only static. Reason is that it's meant for lower-level devices.
 - SSH is heavier than gRPC and not as high performance.
- In Ali's opinion: NETCONF / YANG is antiquated (from the 70's) with a limited toolset and therefore should not be used. There's no other industries using NETCONF and YANG for configuration outside of networking. Protobuf is much better supported. Google uses YANG only as an artifact. All devices are gRPC and protobuf.

Standalone adapters

- There will be a gRPC / protobuf interface between the individual adapters and the vOLTHA core in 2.0. This replaces the current 1.0 API.

"Generic" Adapter for the Broadcom B0

- Saurav targeting to get the generic adapter done for the phase 3 trial. Without additional help, it won't get done.
- Radisys Broadcom adapter yet could be used as a starting point.

Scale

- vOLTHA manages on the order of 10k devices.

G.fast Adapter

- We'll use the same API for both PON and G.fast.
- Broadcom adapter could start using the BAL to communicate with Broadcom HW. The adapter will adapt BAL to gRPC. BAL is generic enough to be used for PON or G.fast. It communicates using the concept of flows.

Parking lot

Netconf YANG vs. gRPC / protobuf

Persistence

- It's console / lmdb today. Changing this for 1.0 will delay the release
- Hooks are there to change that to a cluster / HA and we're changing it. It's an in-progress user story.
- By 7/28, we plan to have a persistent datastore with information in the cluster of consuls. Each vOLTHA pushes data into it. If any vOLTHA fails & restarts. The restarted vOLTHA will query consul and pick up where it left off.

Service discovery

Auto-discovery

Auto-generate protobufs

Disaggregation in version 2.0. Break out the three pieces of vOLTHA core:

- Managing devices
- Proxy / load balancer between gRPC clients & vOLTHA instances
 - Every time a new OLT comes online, this assigns the OLT to a vOLTHA and pre-provisions it. From then on, that vOLTHA is responsible for that OLT.. Round robin today, but in the future we can weight it based on load.
- Pre-provisioning system would check with consul the vOLTHA managing a particular ONT / OLT, then pre-provision on that vOLTHA. Assign OLT / vOLTHA on pre-provisioning because there's a lot of data during pre-provisioning that needs to live somewhere.
 - Each device would be named, and consul would map the name to device ID.
- Why have an enable step? Just pre-provision.
- Control: If an instance dies, a new one comes up. There's currently an election between all vOLTHA cores.

Container Management

- Container management will be either Swarm or Kubernetes. Sergio recommends Docker swarm for 1.0 because it's simpler. Won't need Registrar in swarm mode because swarm covers DNS.
- Better to use labels than tags for Dockers because tags get overwritten. Labels will be part of the 1.0 release.

Documentation

- There's no documentation today. It's the code. We need to change that going forward.

ONOS tenant applications

- Applications need to be updated.
- Won't be using segment routing (no leaf)
- Need to agree on a version # for ONOS.

Phase 3 will use Ubuntu 16.04 for compute platforms.

Requirement to LAG between OLT and TOR, and TOR <-> 7450.

- There's a single TOR in the POC.
- OLT should have LACP to TOR. High priority from AT&T..
 - Edge-Core not planning on implementing on the Broadcom OLT or the leaf.
 - Needs to be added to release plan of vOLTHA.
 - Requirement needs to be documented.
 - Number of interfaces
 - Interface speed
 - Active / Active

Discovery (According to Tom A.)

- Idea is you onboard HW, follows ONIE model (uboot), comes up, reaches out & loads firmware.
- Data collection via IPMI & BMC.
- Trigger to spin up instance to manage it as necessary.
- Not part of vOLTHA. Part of a broader infrastructure management.
- ONT SW shouldn't be on the OLT.
- Ideally there's a separate system responsible for SW management on ONTs and OLTs. vOLTHA gets involved if the image comes over OMCI. OLT / vOLTHA is data path only.
- Today in AT&T, currently ONUs maintain their image locally. Other systems don't store their image and boot from a repository.

Image Management for

- This is a currently requirement for Phase 3 POC for the vOLTHA.
- Mark S. doesn't think this is required for phase 3 for the OLT and ONT. There's a line item in section 5, medium priority.

Phase 4 field trial will be in Atlanta.

