## 2020 Q3 Hyperledger Avalon Project Update

Prepared by Eugene (Yevgeniy) Yarmosh

## Project

Hyperledger Avalon

# **Project Health**

The Avalon team has been progressing according to its 2020 plan. 0.6 version was released in July. Currently team is working towards 0.7 release that targets end of 2020. Incremental deliverables include support for multiple workers and multiple worker processing enclaves (aka WPE) in the worker pools, improved support for Kubernetes and Docker containers, porting to CentOS (in addition to Ubuntu). The team is on track to deliver Graphene based workers and DCAP attestation support for Intel SGX TEE. The team started Avalon Python SDK project, initially, in the Hyperledger Labs.

### Questions/Issues for the TSC

None at this time

#### Releases

The team has released 0.6 release in July that includes

- · Delivered initial worker (aka enclave pool) implementation with key isolation
- Initial K8S support
- Improved cryptographic support in python modules by utilizing "native" Python libraries instead of SWIG wrappers for C++ code, supprt for MbedTLS
- · Dockerised front-end load balancer (Nginx based) to distribute transactions to multiple listeners
- Prototyped Graphene runtime integration
- Improved CI support and test coverage
- Substantially improved project documentation and added Oxygen auto API spec generation and expanded WiKi documentation

The team currently working towards 0.7 release scheduled at the end of 2020 with an intermediate incremental deliverables with 3 weeks cadence.

#### Overall Activity in the Past Quarter

After completing the 0.6 release in July (see above), the team has been building additional functionality towards 0.7 release that included

- Worker key refresh
- CentOS support
- Improved front-end load balancing that includes support for HAProxy and K8S
- Significant progress towards extending SGX attestation model to support 3<sup>rd</sup> party attestation (aka DCAP)
- Work in progress on the integration of the high-level Graphene LibOS runtime into Avalon
- Started work on Avalon Python SDK

These capabilities simplify Avalon worker development and expand supported use cases, e.g. Graphene integration enables AI, ML, and video analytics use cases.

Due to Covid19 community activity was lower than we would like, but

- The team continued regular Avalon Technical Forums with good attendance
- The project was presented at the Virtual Meetup with Hyperledger Avalon and Chainlink
- The project was also actively presented and discussed at EEA Trusted Compute Work Group

## **Current Plans**

The Avalon team has moved from quarterly releases to less frequent releases (every 6 months) with incremental (tested and operational) incremental functional deliverables at the end of each three-week sprint. Current plan includes

- A major project release 0.7 at the end of the 2020
- 0.8 release in Q2 of 2021 focused on the code stability and robustness (aka "active" stage candidate)

0.7 release targets following functionality (relative to 0.6 release)

Finalization of worker pool implementation, key refresh

- Support for multiple KME (Key Management Enclave) for the worker pools (aka scalable key management)
- Kubernetes integration with elastic compute support
- Multi-tenancy support by utilize SGX KSS (Key Sharing and Separation) to dedicate a TEE enclave for processing workloads from a specific requester only
- Extending SGX attestation model to support 3<sup>rd</sup> party attestation (aka DCAP)
- Integration of high-level LibOS runtimes (Graphene and, optionally, Occlum)
- Improved and expanded load balancing (HAProxy and K8S)
- Porting to CentOS
- Improved test automation into CI
- Refining Avalon Architecture with focus on auto-scaling, multi-tenancy and TEE runtime integration

Many of the capabilities above are already in progress and either has been completed or on the track to be completed in September.

The team will continue to look for the online opportunities and will publish technical blogs and user case studies along with additional video tutorials.

### Maintainer Diversity

The team welcomed Divya Taori as a maintainer for Avalon Python SDK. The SDK is temporarily hosted at a Hyperledger Labs. We expect it become a part of the overall Avalon project before end of the year and Divya to become an Avalon project maintainer.

#### **Contributor Diversity**

Avalon project gets contributions and participation from multiple companies – Intel, WiPro, IBM, iExec, Kaleido and Santander. The project team sees a lot of interest from a number of companies including ASUS, ConcenSys Health and Chainlink. Nevertheless, community growth in Q3 2020 was relatively slow, primarily, due to Covid19 impact. To grow the community in Q4 of 2020 and in the first half 2021 the team will need to be more active in utilizing online industry events and publishing more video tutorials, blogs and case studies.

#### Additional Information

None at this time

#### Reviewed by

- Angelo De Caro
- Arnaud J Le Hors
- Christopher Ferris
- Dan Middleton
- Gari Singh
- Hart Montgomery
- Mark Wagner
- Nathan George
- Swetha Repakula
- Tracy Kuhrt
- Troy Ronda