

2020 Q2 Hyperledger Avalon Project Update

Prepared by Eugene (Yevgeniy) Yarmosh

Project

[Hyperledger Avalon](#)

Project Health

The Avalon team has been working according to its 2020 plan and is on track to release 0.6 version in July. This release includes baseline implementation for the worker pools (a foundation for the TEE scalability), separation of the key management from the work order execution, initial K8S support and improved cryptographic support. The team has also completed evaluation and prototyping for LibOS runtime integration (Graphene, Occlum, SGX-LKL). The team participated in a number of the industry events even though Covid19 has a high negative impact in this area. The upcoming release 0.7 shall focus on auto-scaling (by utilizing K8S), LibOS runtime integration and test automation integrated with CI. The team continued to focus on expanding its community and contributor's diversity.

Questions/Issues for the TSC

There are no issues for the TSC currently.

Releases

July 2020 – Release 0.6 shall address the following

- Worker pools and key management separation from the the work order execution
- Initial K8S support
- Improved cryptographic support
- Bug fixes bugs and corner cases.

The next plan is to have 0.7 release at the end of the year with 2-3 intermediate stable releases (0.6.1, 0.6.2, etc.).

Overall Activity in the Past Quarter

Avalon team was focused on scalability and key management isolation along with overall code improvements and prototyping for upcoming releases.

- Delivered initial worker (aka enclave pool) implementation
- Implemented KME (Key Management Enclave) and WPE (Worker order Processing Enclave) that isolated key management from the work order processing
- Integrated Avalon with K8S
- Improved cryptographic support in python modules by utilizing "native" Python libraries instead of SWIG wrappers for C++ code
- Added support MbedTLS library in addition to OpenSSL. It provides a base for a broader TTE runtime support
- Updated connectors for Hyperledger Fabric and Hyperledger Besu in the Avalon proxy model
- Dockerised front-end load balancer (Nginx based) to distribute transactions to multiple listeners
- Prototyped Graphene runtime integration
- Refactored overall code repository structure and improved build process
- Improved CI support
- Improved test coverage
- Started work on the test automation framework
- Substantially improved project documentation and added Doxygen auto API spec generation
- Expanded Wiki documentation

Avalon team participated in the following activities:

- Continued regular Avalon Technical Forum calls every other week with a good community participation
- Relied on the GITHUB issues for the bug tracking and feature request
- Published several video tutorials
- Presented Hyperledger Avalon at Consensus 2020
- Submitted Avalon tech talk proposal to Grace Hopper Celebration India (GHCI) - the largest gathering of women technologists in Asia
- Continuously utilized email and rocket chat for community support

There are at least 7 active maintainers. Covid19 slowed community building activities.

Current Plans

The Avalon team planned to have quarterly release but realized that this cadence is not optimal. Primary reason is complex nature of the tasks under development. Additional reason is a three-week sprint pattern utilized by the team. Overall, this pattern works well for the team, but having only four sprints between releases makes it hard to address unforeseen changes. Based on the above the plan was changed to have

- A major project release 0.7 at the end of the 2020
 - 2-3 intermediate stable (and tagged) releases - 0.6.1., 0.6.2, etc.
- 0.8 release in Q2 of 2021 focused on the code stability and robustness (aka "active" stage candidate)

Next release 0.7 will focus on

- Finalization of worker pool implementation
- 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 3rd party attestation (aka DCAP)
- Integration of high-level LibOS runtimes (Graphene, Occlum)
- Adding attested oracles end-to-end use case (to be contributed by Chainlink)
- Adding test automation into CI
- Refining Avalon Architecture with focus on auto-scaling, multi-tenancy and TEE runtime integration

The team plans to regularly review GITHUB issues and consider utilizing JIRA for the project management. Avalon Developer Forum calls (every other weeks) generally proved to be frequent enough, but the team still considers starting a new series of Avalon Architecture Forum calls.

Integrated to building up Avalon community and to mitigate Covid19 impact, the team will look for the online opportunities and will publish technical blogs and user case studies along with additional video tutorials.

Maintainer Diversity

Formal Avalon maintainer list has not changed during this quarter, but we started ramping up a new maintainer candidate from Wipro and in the process of recruiting maintainer candidates from three more organizations. Based on the recent activity and our current plans we anticipate that the Avalon maintainers list will be extended in the second half of 2020.

Contributor Diversity

Avalon project gets contributions and participation from multiple companies – Intel, WiPro, IBM, iExec, Kaleido and Santander. AntFin Occlum team modified Occlum project according to the Avalon requirements in preparation for Occlum runtime integration with Avalon. The project team sees a lot of interest from and anticipates additional contributions from ConcenSys Health and Chainlink.

Avalon Developer Forum calls happens every other week well attended by members from different organizations. Building up Avalon community was less efficient during Q2 of 2020 primarily due to bigger than expected Covid19 impact. In Q3 and Q3 of 2020 the team will work on expanding Avalon community by more actively utilizing on-line industry events, video tutorials and publishing 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

