

Cloud Native Fabric Task Force

- Assumption for Fabric consumers is “I just want to deploy a network, write some chaincode, and integrate with an app tier.” Reality != assumption.
- Assumption is “kubernetes native.” (Ahoy! The ship has sailed.)
- 26+ different LABS initiatives with overlapping feature sets. The fragmentation is a barrier to general adoption.
- State of the art for “cloud native” has evolved radically, like clockwork, every 18 months for the past decade. Fabric has not kept up.
- Learning curve for “Building applications with Fabric” is assumed to be: “Contributing to core Fabric.” This is a major barrier to new adoption.
- Learning curve for “Kubernetes” is steep. (Another barrier.)
- Fragmentation is extreme and unhealthy for the continued relevance of Fabric. (Yet another barrier).

The TF will identify key contributors, design patterns, standards, and technologies to guide Fabric towards “Cloud Native” outcomes:

- Encapsulate necessary but minimal value-add configuration tasks to Kubernetes Operators
- Define common formats and APIs to encourage interoperability and collaboration between competitive platforms.
- New start development efforts take minutes, not weeks/months/years/never.
- Applications run equivalently in dev, test, dev, QA, staging, and production environments. (local dev = hybrid = cloud)
- Applications run predictably, reliably, and with minimal intervention in production environments.
- Corporate contributions are aligned for community collaboration, not parallel efforts.
- Fabric ... Just works.

Tasks to be Completed

- Articulate “Cloud Native” terminology and TF objectives.
- Author and publish TF recommendation guidelines, positioning, and best practices in a position paper, public forum, or *manifesto*.
- Organize HL Workshop(s) to promote best practices, TG recommendations, and provide community outreach/training.
- Identify feature sets, functional gaps, and areas of improvement necessary to realize “cloud ready” state. Key areas include, but are not limited to:
 - Container-based Application Development
 - Operational practices : “Day 0”
 - Operational practices : “Day 1+”
 - Visibility, Observability, and Monitoring
 - Service mesh integration practices
 - Kubernetes operator patterns
- Identify functions that may require fabric-RFP updates and community review. Candidate topics:
 - Cloud-native x509 CA alternatives (e.g. fabric-ca → [cert-manager.io](#))
 - [FAB-10734 Fabric CLI Redesign](#)
 - [0000-opentelemetry-tracing.md](#)
 - [fabric-admin-sdk](#)
 - fabric-samples : upgrade Docker-compose → Kube orchestration
 - fabric-docs : clear guidance on “production” deployments → Kube
 - TBD
- Align open-source corporate contributions towards shared, collaborative outcomes.

Time to Complete

- 4 months + option to extend with targeted focus
- monthly cadence

Organizer

@Josh Kneubuhl

Initial Participant List

@Antoine Toulme - Splunk

@Arnaud Le Hors - IBM

@Daniel Szegö - Hyperledger meetup organizer

@David Boswell - Linux Foundation

@Hart Montgomery - Linux Foundation

@Josh Kneubuhl - IBM

@Sownak Roy - Accenture

Discord Channel

#cloud-native-fabric-task-force ?

References

- Fabric Operations Labs: <https://wiki.hyperledger.org/display/CA/Fabric+Operations+Labs>