Project Alvarium Annual Review 2023

Prepared by: Trevor Conn
trevor.conn@dell.com

4-Oct-2023

DELL Technologies

Agenda

- Short refresh on what Alvarium is
- Review of work done this past year
- Goals for Coming Year
- Challenges

Executive Summary

A description of the Alvarium concept and its relevance for modern use cases.

- Modern applications are extensively distributed
- Data is no longer a fixed asset stored in a silo.
- Data traverses the network and can be transformed along the way.
- Create metadata that attests verifiable authority at the origin of data
- Create metadata describing how data was handled as it traverses the eco-system
- Metadata is created at “trust insertion points”
- A measure of trust is calculated at each insertion point and can be weighted
- Trust can be rolled up into an overall confidence score for a piece of data
- A trust score may be used to govern system behavior or alert operators to an attack

Example: Data annotated in traversal

The Alvarium code base is a lightweight SDK that annotates data streams (e.g., sensor data) with trust metadata and confidence scores, forming a Data Confidence Fabric (DCF)

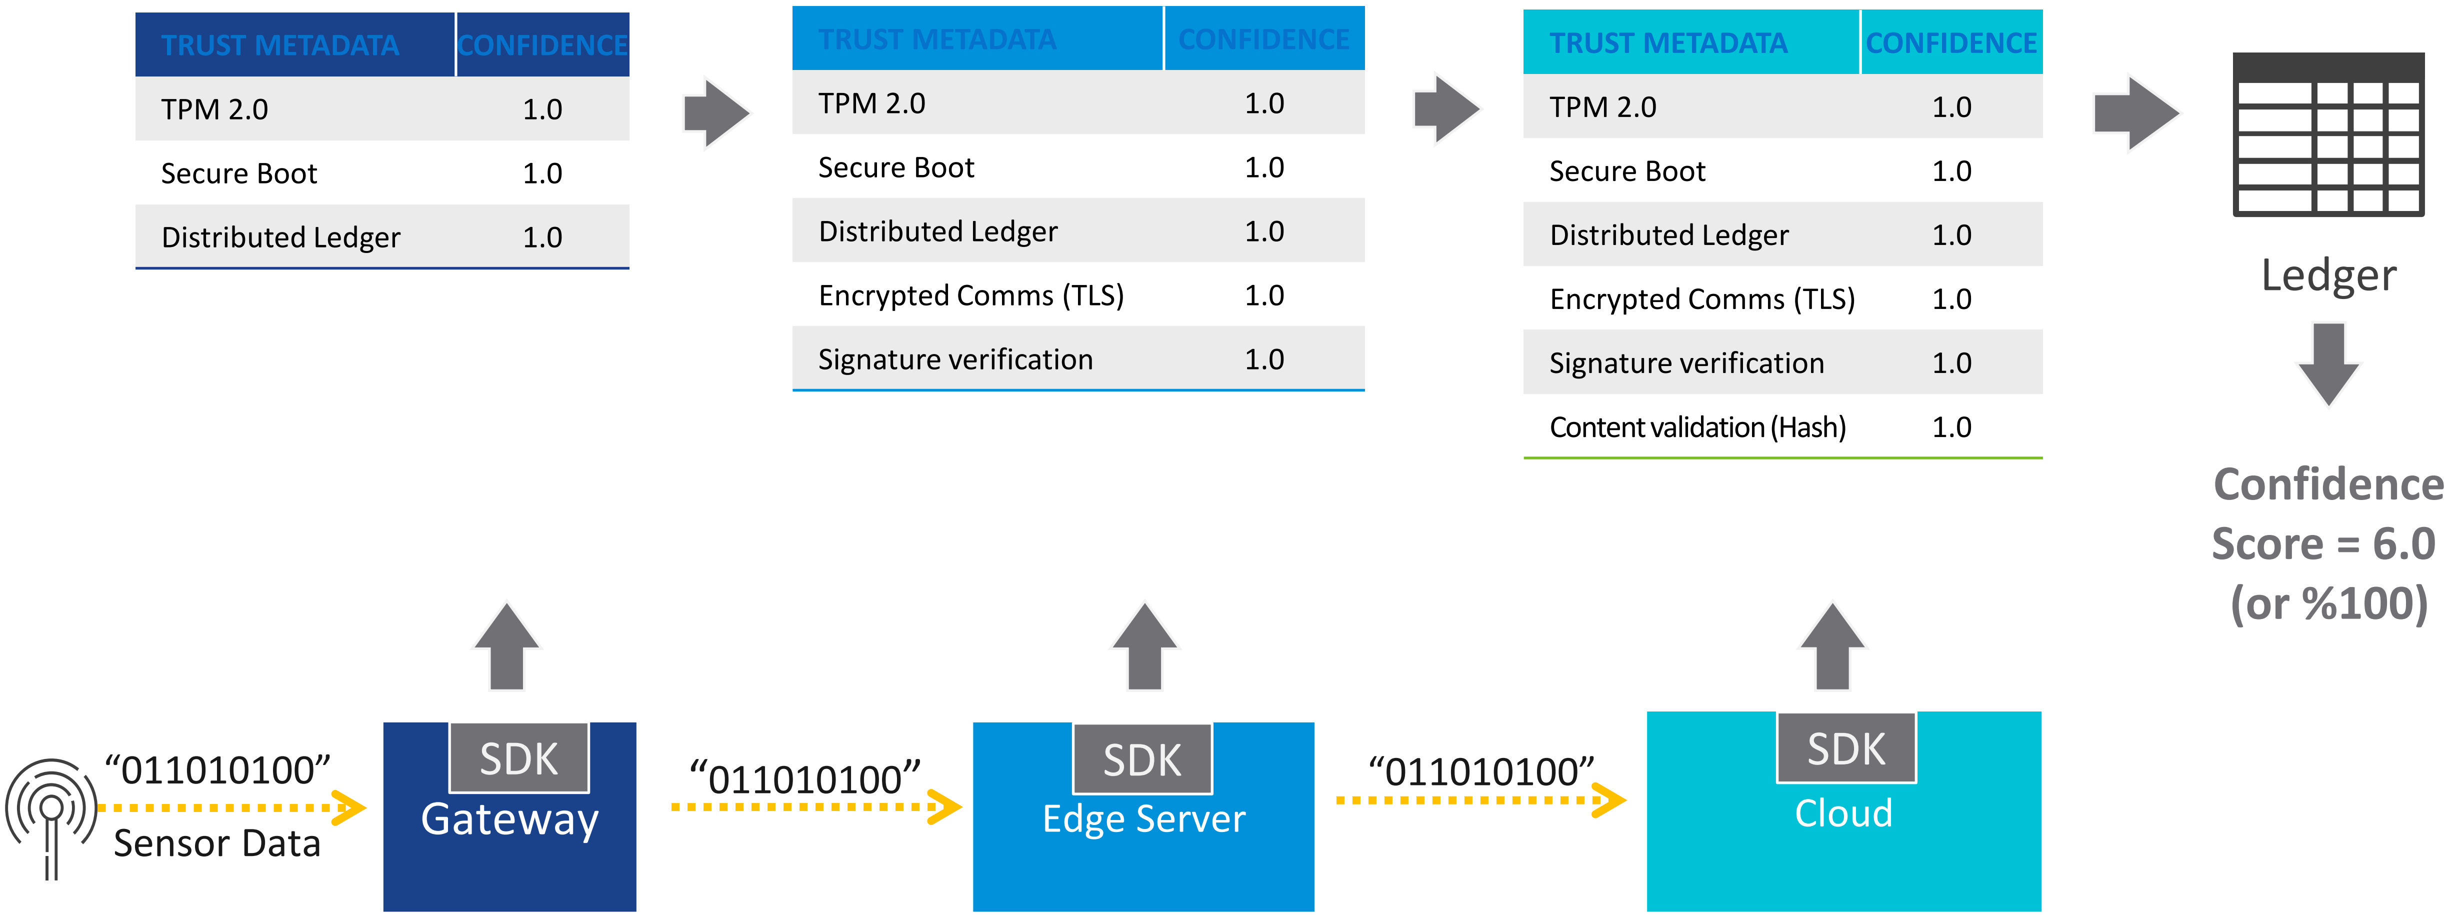

TRUST METADATA	CONFIDENCE
TPM 2.0	1.0
Secure Boot	1.0
Distributed Ledger	1.0

TRUST METADATA	CONFIDENCE
TPM 2.0	1.0
Secure Boot	1.0
Distributed Ledger	1.0
Encrypted Comms (TLS)	1.0
Signature verification	1.0

TRUST METADATA	CONFIDENCE
TPM 2.0	1.0
Secure Boot	1.0
Distributed Ledger	1.0
Encrypted Comms (TLS)	1.0
Signature verification	1.0
Content validation (Hash)	1.0

Data Confidence Graph (Conceptual Diagram)

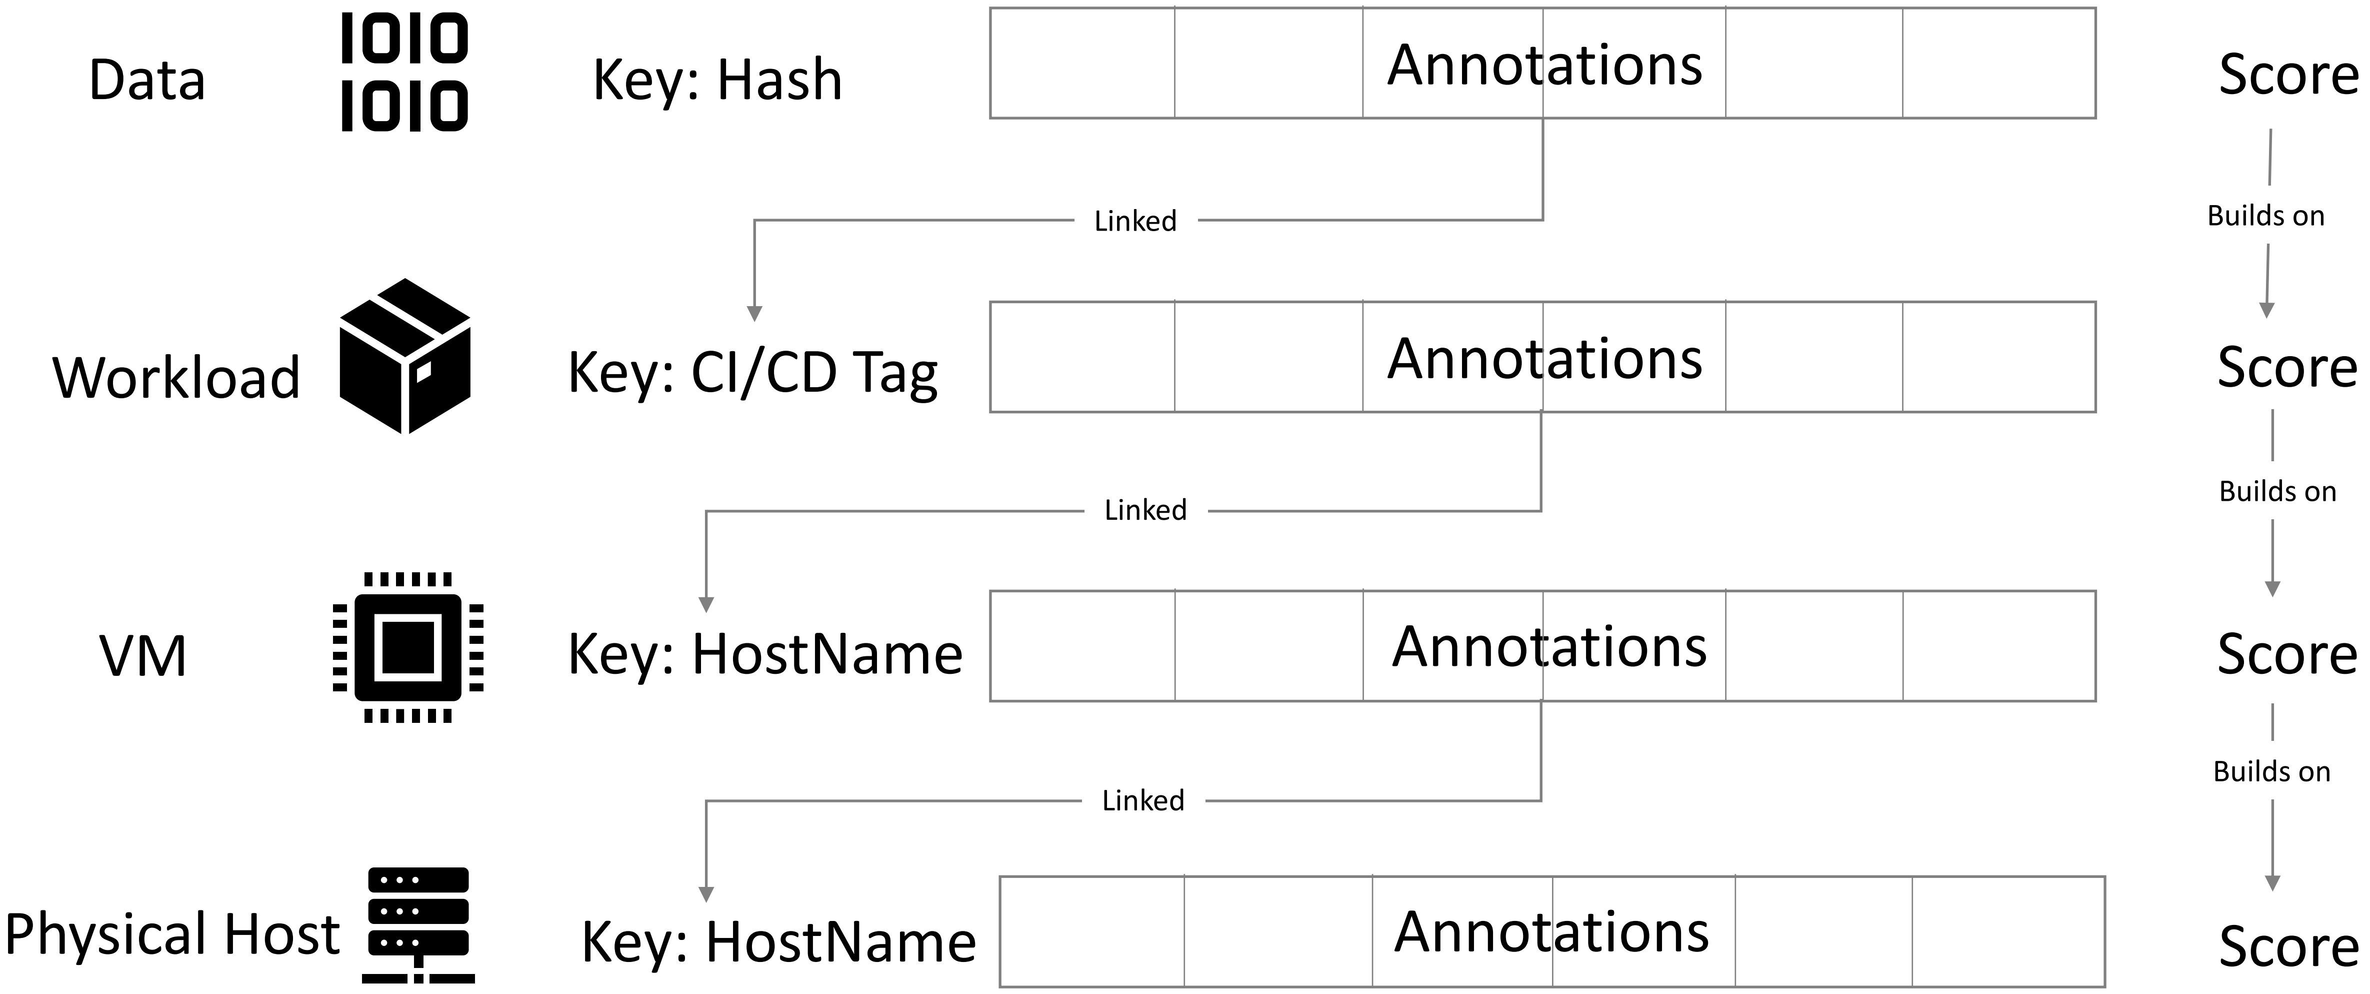

Internal Use - Confidential

Prior Year’s Work Review

DELL Technologies

Work Review

- All code contributions during past year have come from Dell team members
- Focus of our effort has been on attempting to realize a prototype of the “Data Confidence Graph” vision shown in the previous section
  - Integrated Alvarium Go SDK with EVE-OS Adam Controller
    - Work performed on a team member’s fork to annotate the presence of a TPM when device is onboarded
  - Integration of Alvarium Java SDK with Jenkins pipeline
    - Wrapped Alvarium SDK in a shared library for use by multiple pipeline steps
    - Three new annotators demonstrating auditability of CI pipeline, scoring of published artifact
      - SourceCodeAnnotator – verifies integrity of cloned code prior to build
      - VulnerabilityAnnotator – scans dependency files (such as Maven or go.mod) for known vulnerabilities
      - ChecksumAnnotator – compares the checksum value of a build artifact against a known good value
    - Using a “forked” copy of the EdgeX Foundry pipeline for integration
  - Integration of Alvarium Go SDK with EdgeX Foundry services
    - In progress, discussion ongoing with EdgeX TSC/Architects
    - Initial work to conduct integration has begun in team member forks based on Minneapolis release
      - Device-sdk-go, device-virtual-go

Data Confidence Graph Toolchain

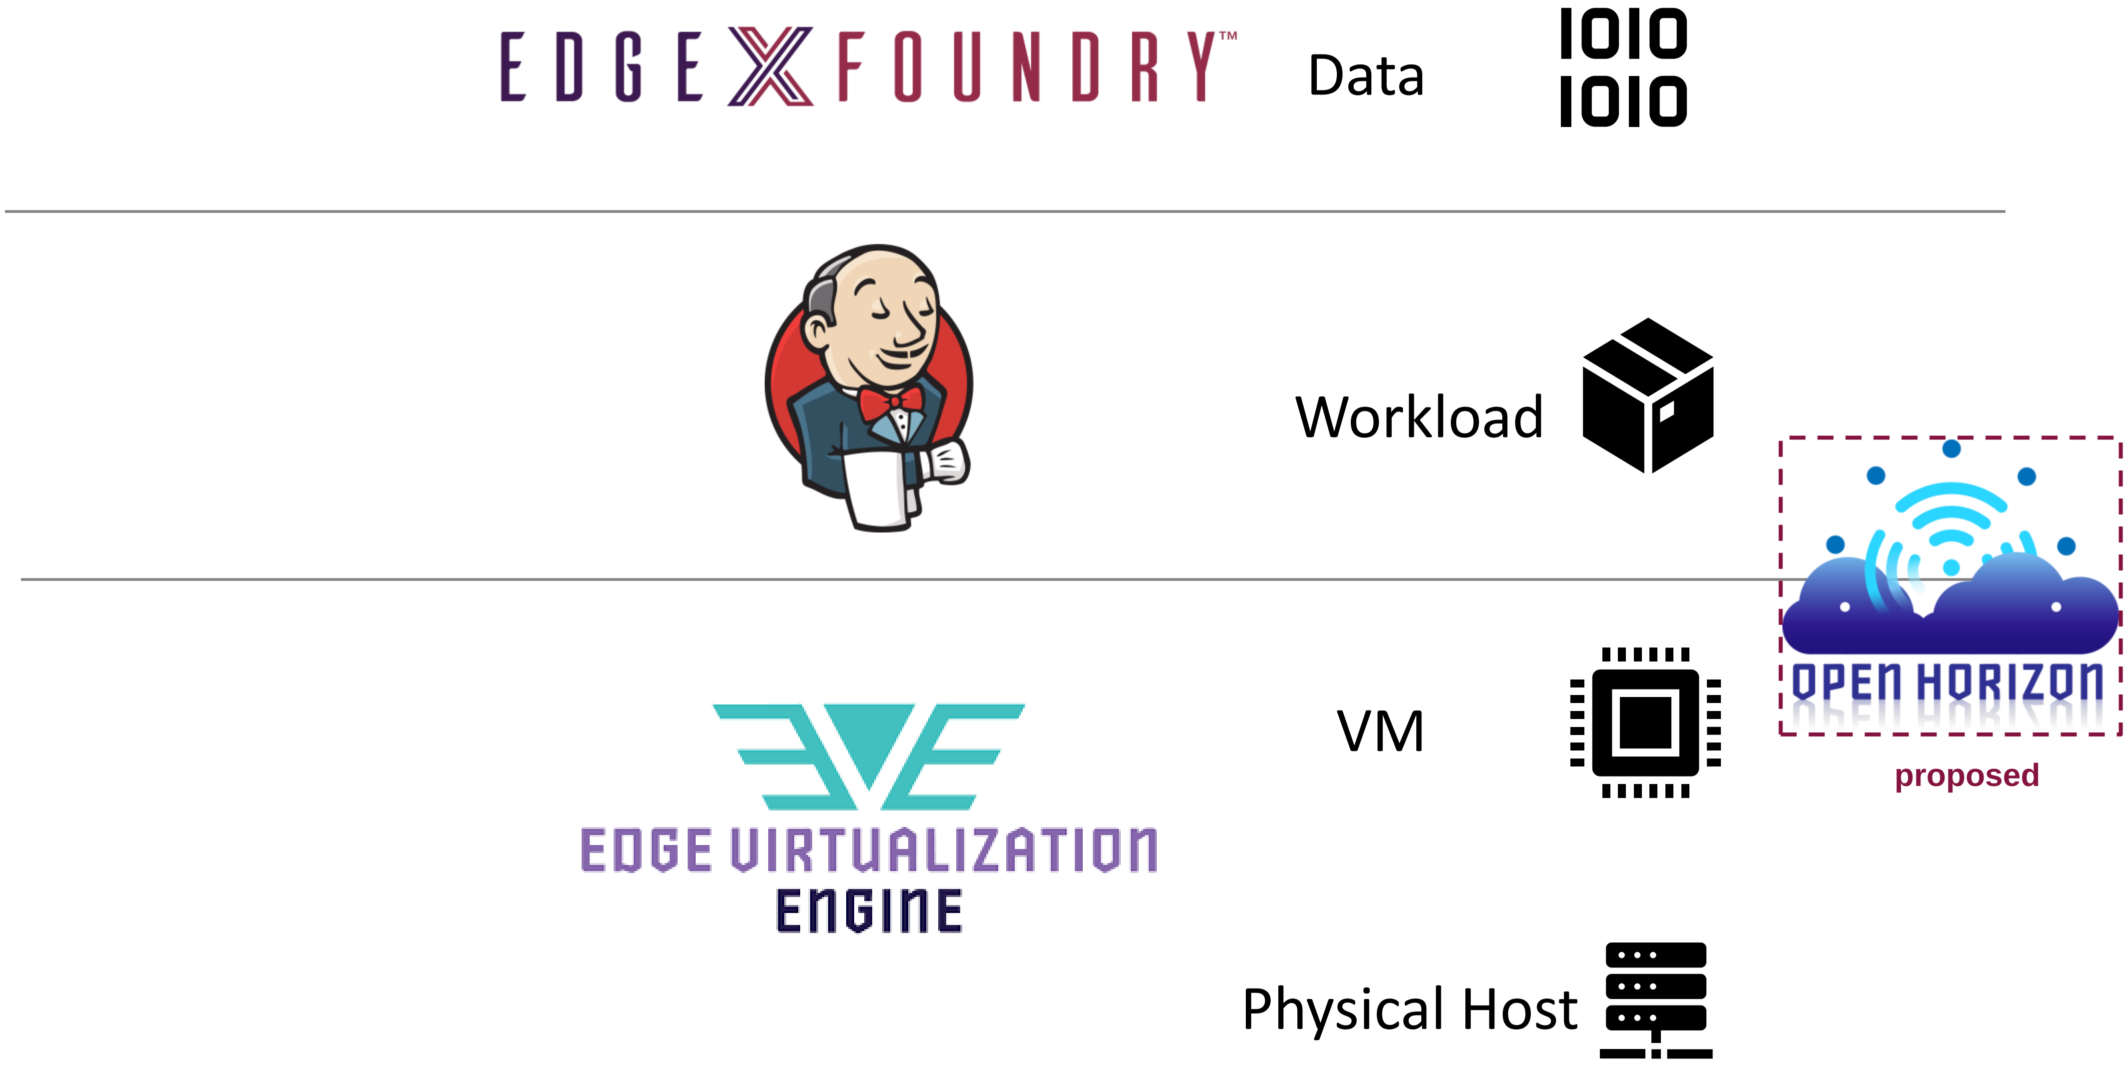

Data

Workload

VM

Physical Host

proposed

Data Confidence Graph Annotations

- Data Signature
- TLS
- Content Checksum

EDGEXFOUNDRY™

Data

- Source code integrity
- SBoM / Dependency Check
- Artifact Checksum

Workload

OPEN HORIZON

proposed

- TPM Validation

EDGE VIRTUALIZATION ENGINE

VM

Physical Host

Data Confidence Graph Orchestration

- Data Signature
- TLS
- Content Checksum

EDGE X FOUNDRY™

Data

- Source code integrity
- SBoM / Dependency Check
- Artifact Checksum

Workload

- TPM Validation

EDGE VIRTUALIZATION ENGINE

VM

Physical Host

OPEN HORIZON

proposed

Open Horizon orchestrates deployment according to workload confidence threshold configured via policy

DELLTechnologies

Goals for Coming Year

Goals

- We hope to complete the aforementioned proposed integration with other LF-Edge projects
  - We believe this will provide a compelling, tangible demonstration of what we perceive to be Alvarium's value add
  - This work will also help us to establish more connection and communication with other LF-Edge projects
- Subsequent to completion, we would be happy to promote the work through conferences, written content or podcasts on behalf of LF-Edge.
- Once complete, we hope this workstream will provide some clarity as to the future direction of the project
  - See "Challenges" section below.

Challenges

Challenges

- Once again, project participation remains the key challenge
  - All contributions at this point come from Dell team members
  - Internal resources are not full-time dedicated and there can be organizational fluidity/churn
- “Digital Confidence Graph” is great ideation but is widening the aperture of the project outside of its original scope
  - For example
    - To what extent is full stack provenance already provided in an implicit way through attestation?
    - If Alvarium plans to provide confidence regarding automated CI/CD practices, does it plan to align with other projects in this space such as SigStore and SLSA?
      - “800 pound gorilla” problem
    - Alvarium enablement throughout the stack requires different programming models, divergent SDKs. Seems complex.
- A topic that has repeatedly come up but which we can’t seem to nail down is whether or not the annotation schema rather than the SDK implementation should be the focus.
  - Annotation schema should support explicit confidence measurement at multiple layers of the stack
  - Extend current arbitrary schema to support specialized messages through abstraction
  - Turn the project towards more of a “standards body” approach rather than implementation.

DELL Technologies