Telemetry and Analytics for the NFV World: IOAM, PNDA, DCAE

Frank Brockners, Cisco

December 11, 2017

What the industry had to invent

REACTIVE, COMPUTE NETWORK, STORAGE AND SECURITY – MODULAR, COMPOSABLE ARCHITECTURES

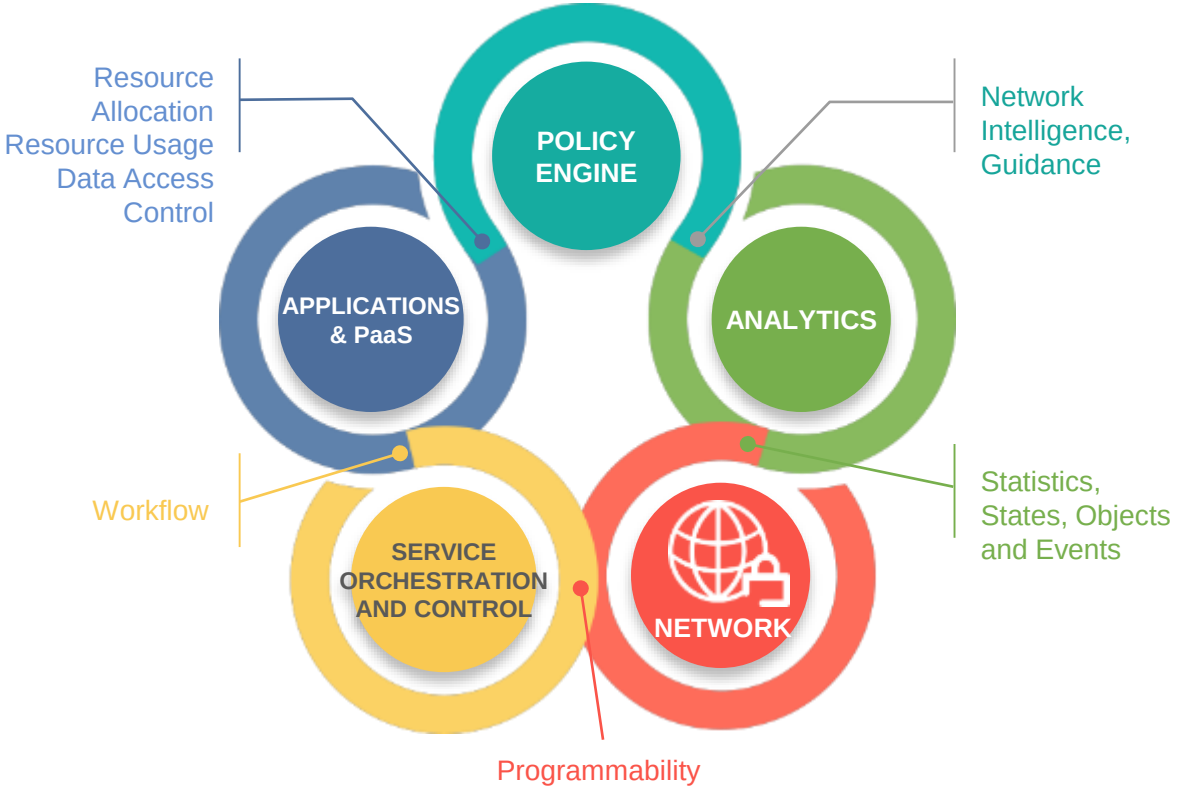

Let’s assume you’re interested in the behavior of your live user-data traffic.

What is the best source of information?

Well… probably the live user-data traffic itself.

Let’s add meta-data to the customer traffic,
so that we can observe the customer traffic itself

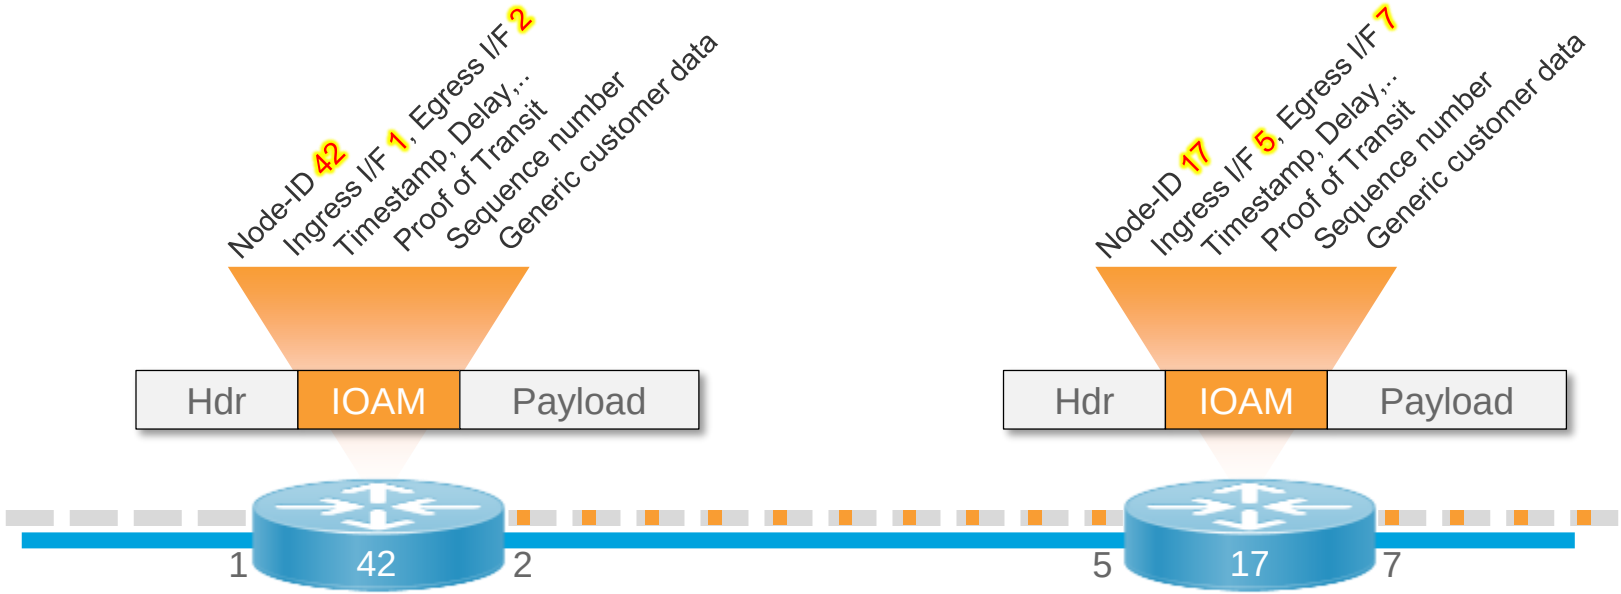

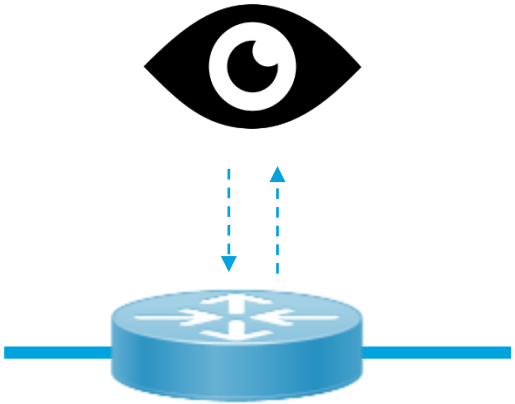

Observe
(SNMP, NetConf)

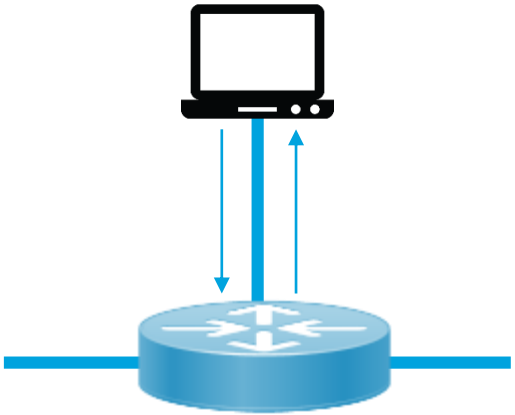

Probe
(ping, traceroute)

In-Band OAM
(per packet telemetry)

Analyze when a packet enters and exists the network…

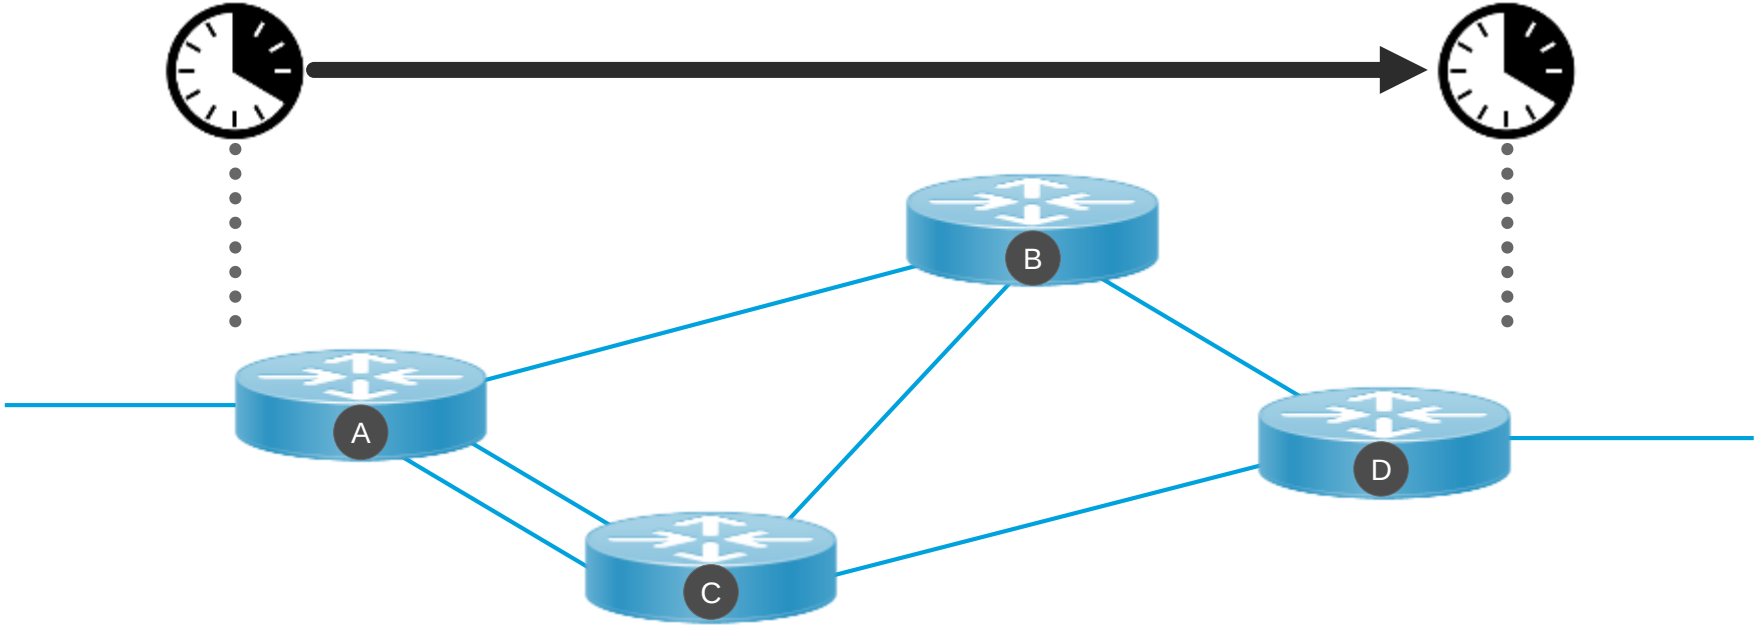

…at what rate packets arrive at a particular hop…

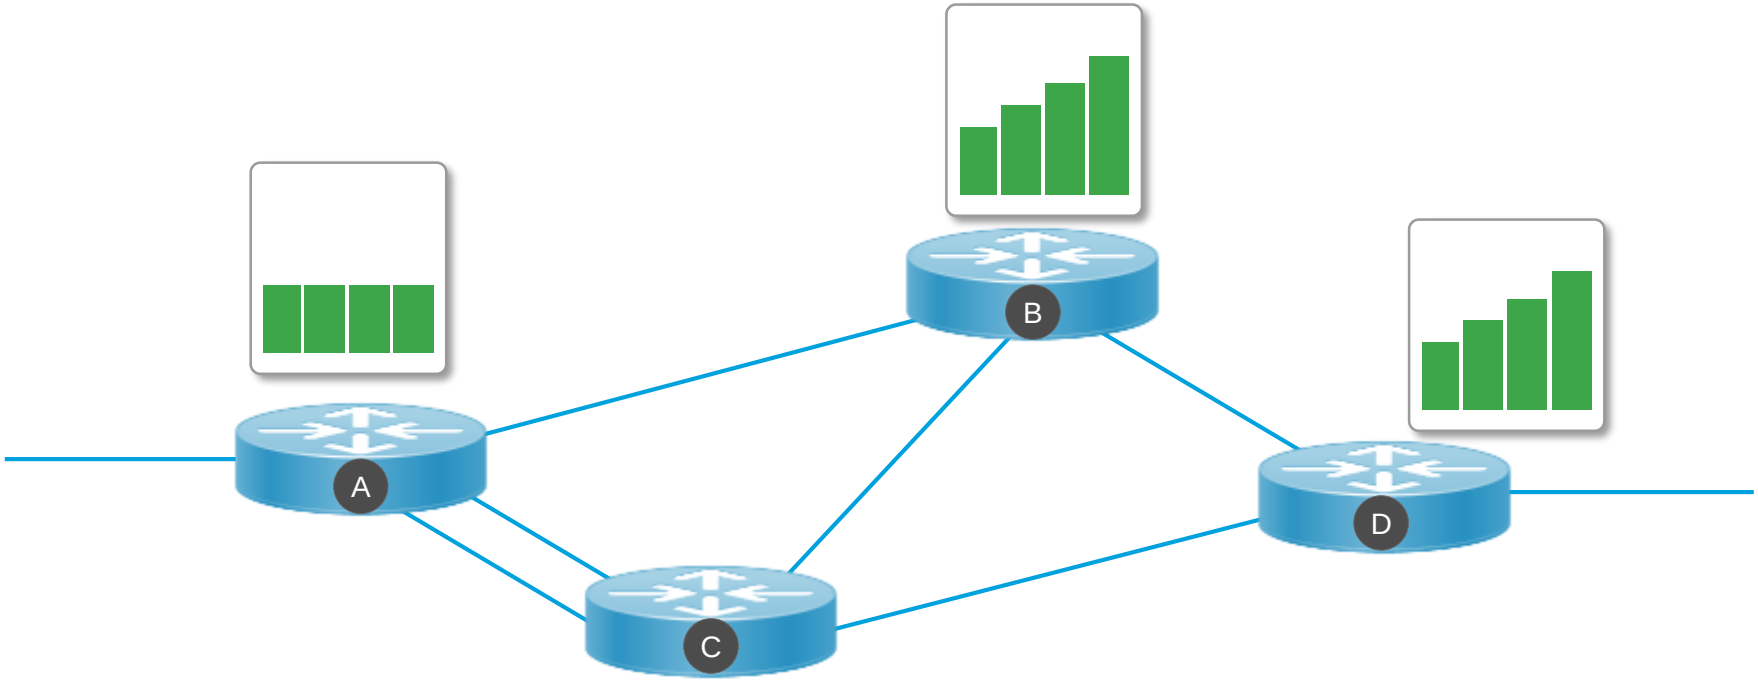

… what path a packet takes …

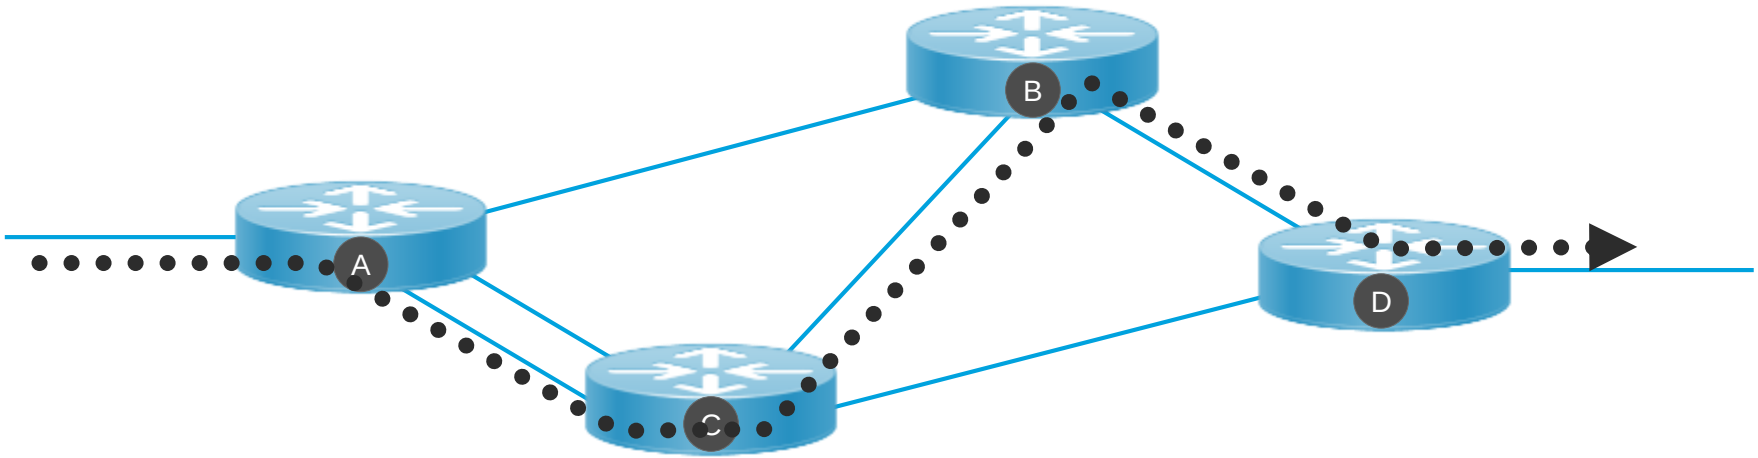

… whether the packet takes the path it is supposed to take (because of TE, NSH, SR)…

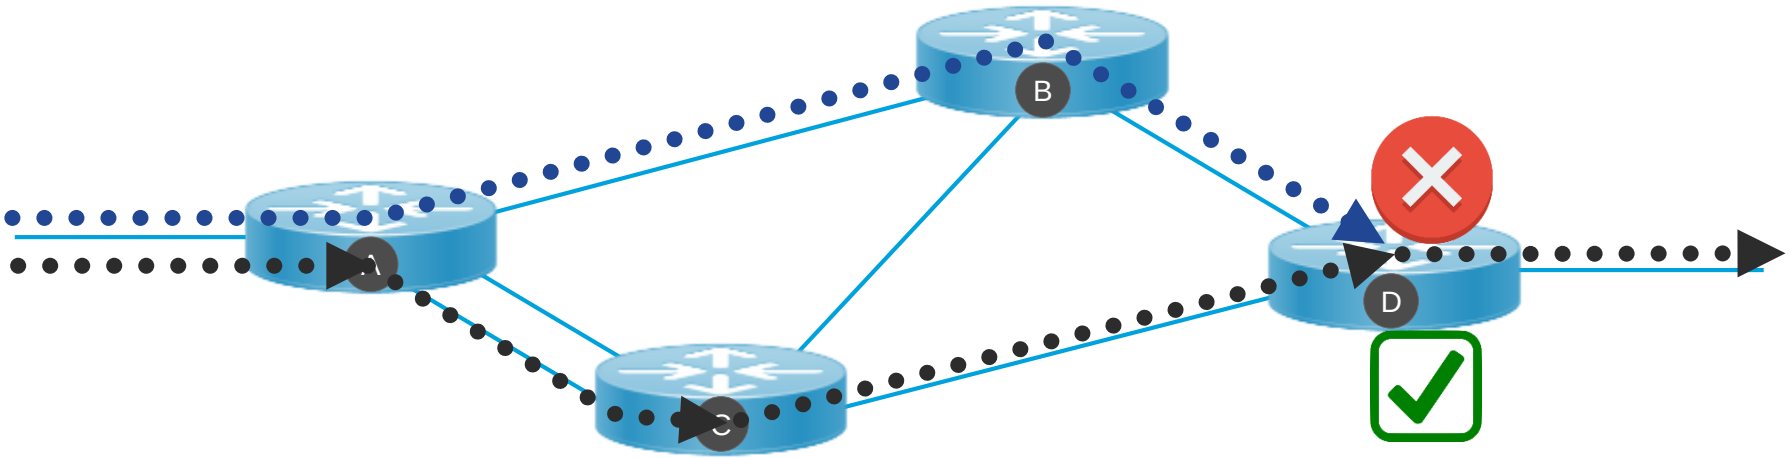

…how long the packet spends at each hop, and which node is experiencing congestion…

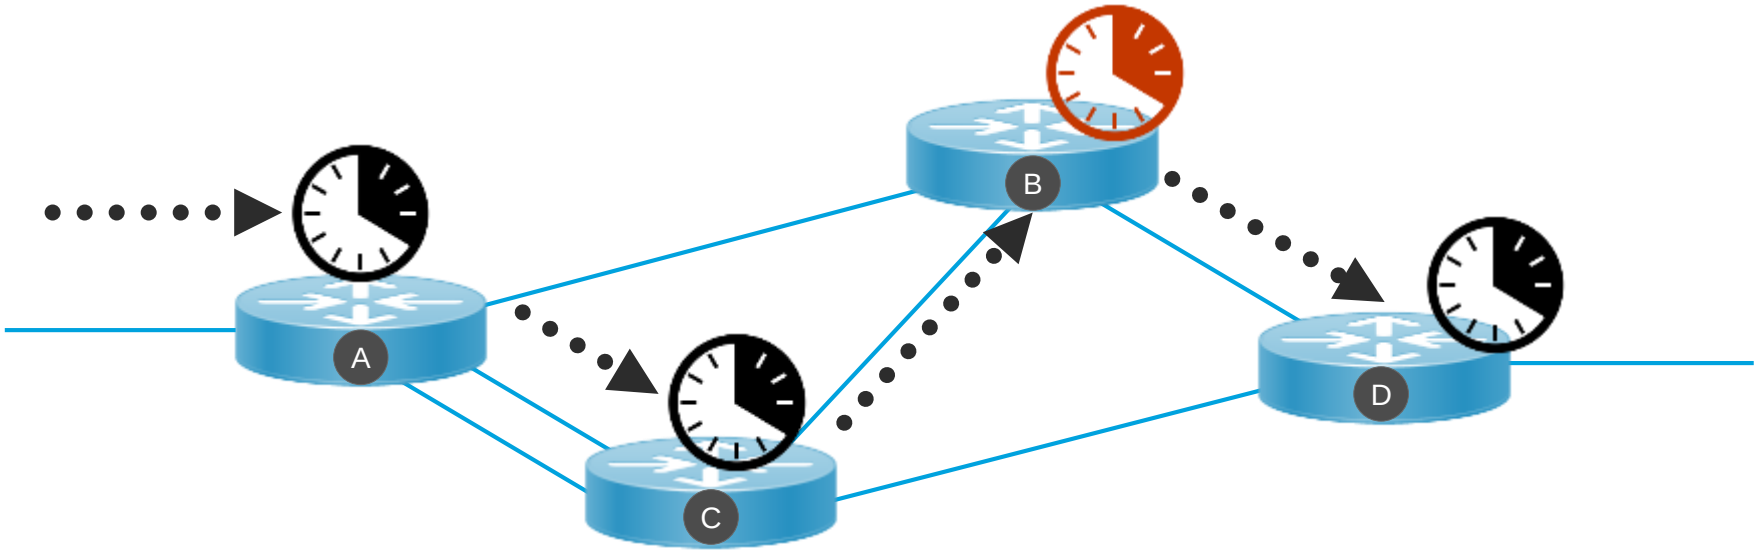

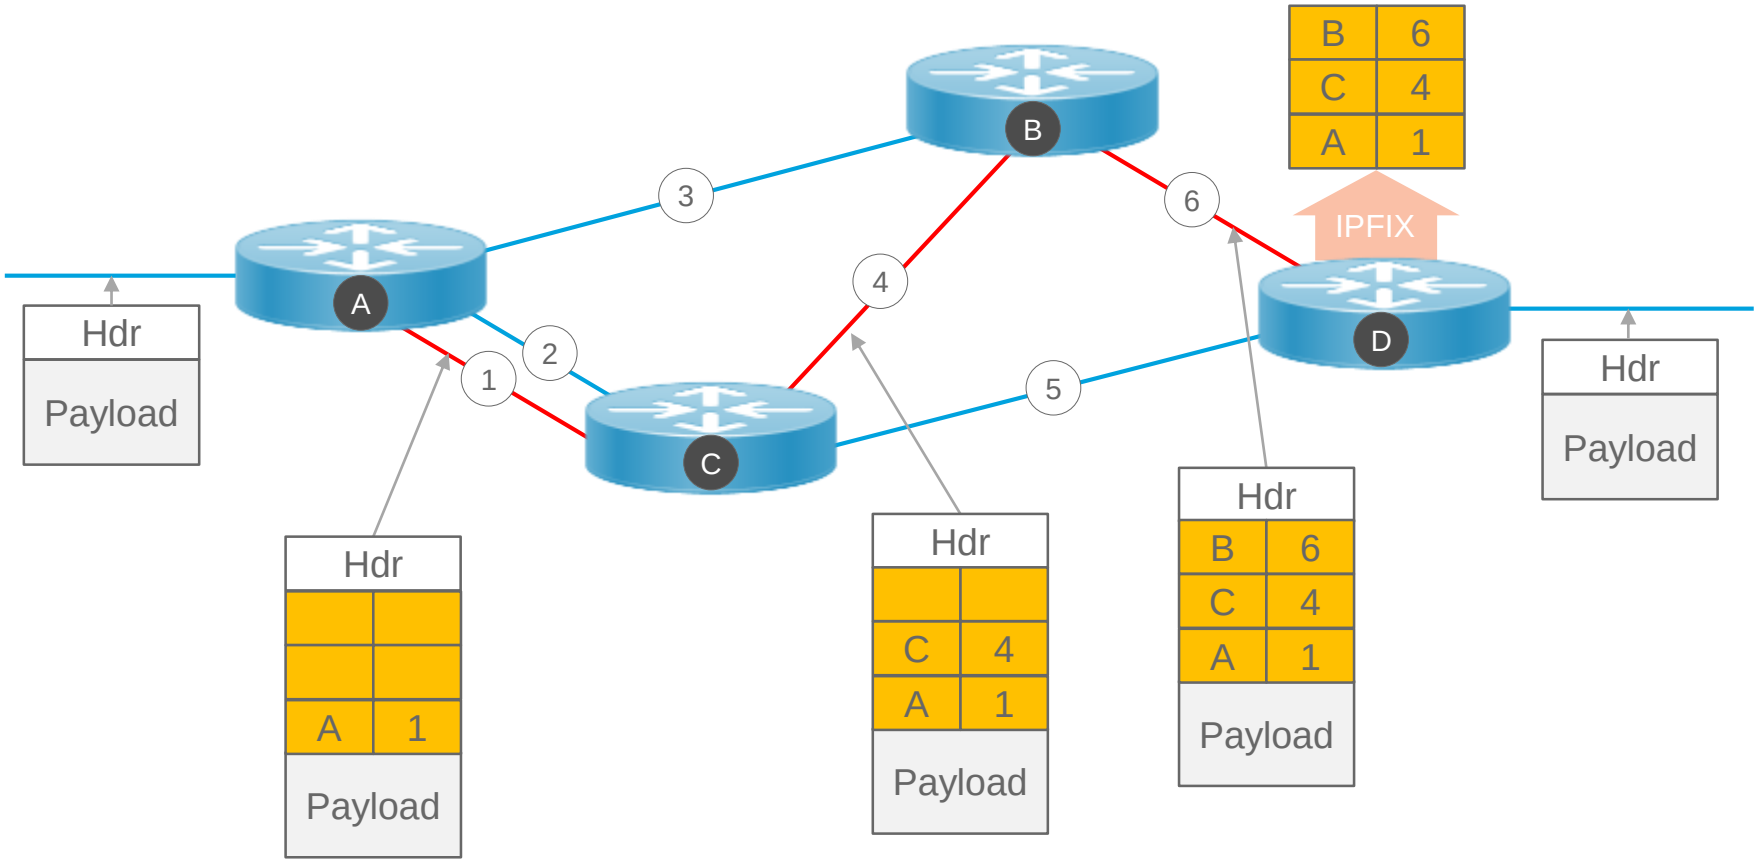

IOAM data fields are defined in a protocol independent way

IOAM data fields can be carried in various protocols

IPv6
VXLAN-GPE
NSH
Segment-Routing v6
GRE
…

IOAM is standardized by the IETF...

Latest IETF Drafts

- In-band OAM Authors: Cisco, Comcast, Facebook, JPMC, Bell Canada, Mellanox, Marvell, Barefoot, rtBrick
  - In-band OAM data types: draft-ietf-ippm-ioam-data-01
  - Encapsulations:
    - VXLAN-GPE: draft-brockners-ioam-vxlan-gpe-00
    - NSH: draft-brockners-sfc-ioam-nsh-00
    - Geneve: draft-brockners-nvo3-ioam-geneve-00
    - Additional encaps defined in: draft-brockners-inband-oam-transport-05 (will evolve into protocol specific drafts over time)
  - In-band OAM transport: Proof-of-transit: draft-brockners-proof-of-transit-04
  - In-band OAM requirements (no longer maintained): draft-brockners-inband-oam-requirements-03
- In-band OAM manageability – YANG models and methods defined in IETF LIME WG
  - draft-ietf-lime-yang-connectionless-oam-03
  - draft-ietf-lime-yang-connectionless-oam-methods-00

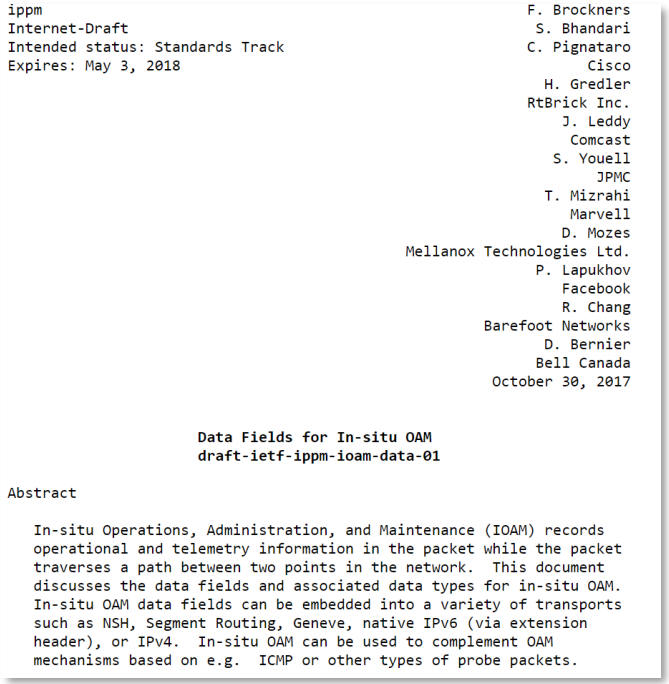

ippm
Internet-Draft
Intended status: Standards Track
Expires: May 3, 2018

F. Brockners
S. Bhandari
C. Pignataro
Cisco
H. Gredler
RtBrick Inc.
J. Leddy
Comcast
S. Youell
JPMC
T. Mizrahi
Marvell
D. Mozes
Mellanox Technologies Ltd.
P. Lapukhov
Facebook
R. Chang
Barefoot Networks
D. Bernier
Bell Canada
October 30, 2017

Data Fields for In-situ OAM
draft-ietf-ippm-ioam-data-01

Abstract

In-situ Operations, Administration, and Maintenance (IOAM) records operational and telemetry information in the packet while the packet traverses a path between two points in the network. This document discusses the data fields and associated data types for in-situ OAM. In-situ OAM data fields can be embedded into a variety of transports such as NSH, Segment Routing, Geneve, native IPv6 (via extension header), or IPv4. In-situ OAM can be used to complement OAM mechanisms based on e.g. ICMP or other types of probe packets.

... and running IOAM code exists.

IOAM Implementation

- Dataplane Implementation:
  - Open-Source: FD.io/VPP (see fd.io)
  - IOS (ISR-G2) – PI31 (CCO: End of July/16)
- Silicon vendors supporting IOAM
  - Broadcom (Trident T3), Netronome (Agilo), Barefoot (Tofino), Mellanox
  - Cisco (planning)
- Controller Implementation:
  - OpenDaylight (Carbon release)

IOAM Use Cases

- Service/Quality Assurance – Fabric OAM
  - Prove traffic SLAs, as opposed to probe-traffic SLAs; Overlay/Underlay
  - Service/Path Verification (Proof of Transit) – prove that traffic follows a pre-defined path
- Micro-Service/NFV deployments
  - Smart service selection based on network criteria (intelligent Anycast server/service selection)
- Operations Support – Fabric Visibility
  - Network Fault Detection and Fault Isolation through efficient network probing
  - Path Tracing – debug ECMP, brown-outs, network delays
  - Derive Traffic Matrix
  - Custom/Service Level Telemetry

Now that you have all this data...

How do you make use of it?

Observe
(SNMP, NetConf)

Probe
(ping, traceroute)

In-Band OAM
(per packet telemetry)

Network data is becoming a big data problem …

3-fold increase in total IP Traffic

>60% increase in devices and connections

Telemetry data streamed in near real-time

Today's silo'ed analytics pipelines

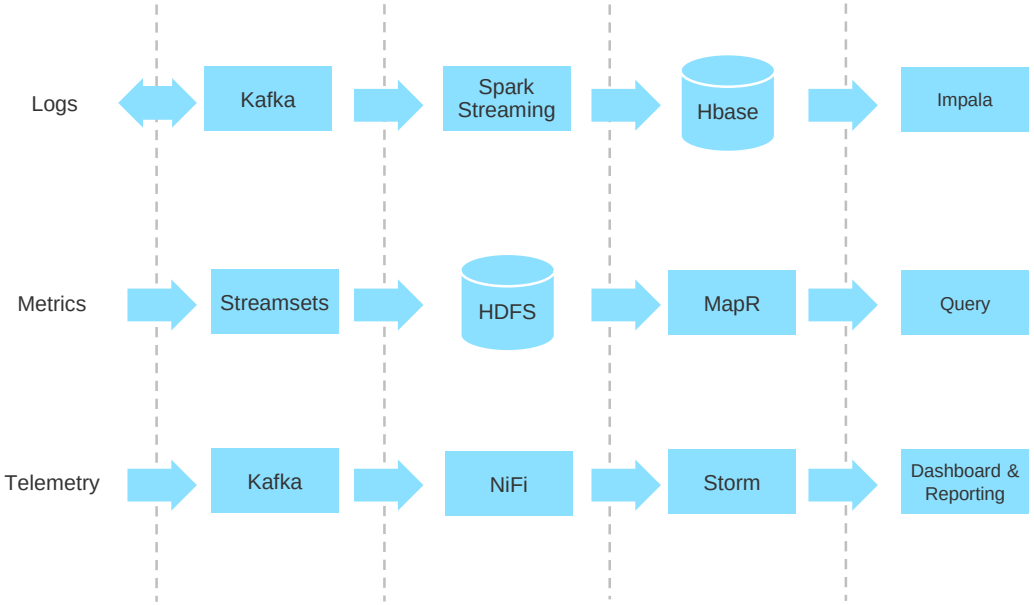

- Tight coupling of data aggregation/store/ analysis
- Multiple analytics pipelines implemented from open source components
- Common design patterns ~75% of effort wasted / duplicated
- Siloes limit the potential of big data analytics and lead to industry divergence

Gather all data into one domain...

... so that you can correlate it.

PNDA brings together a number of open source technologies to provide a simple, scalable open big data analytics Platform for Network Data Analytics

Linux Foundation Collaborative Project based on the Apache ecosystem

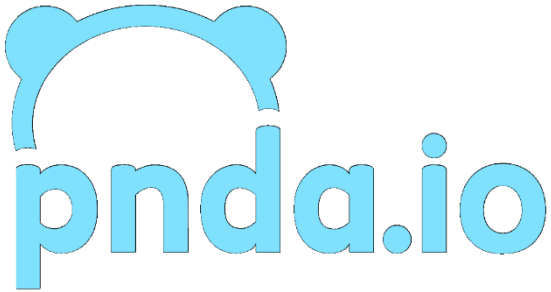

Why PNDA?

There are a bewildering number of big data technologies out there, so how do you decide what to use?

The PNDA project has evaluated and chosen the best tools, based on technical capability and community support.

PNDA combines them to streamline the process of developing data processing applications.

PNDA

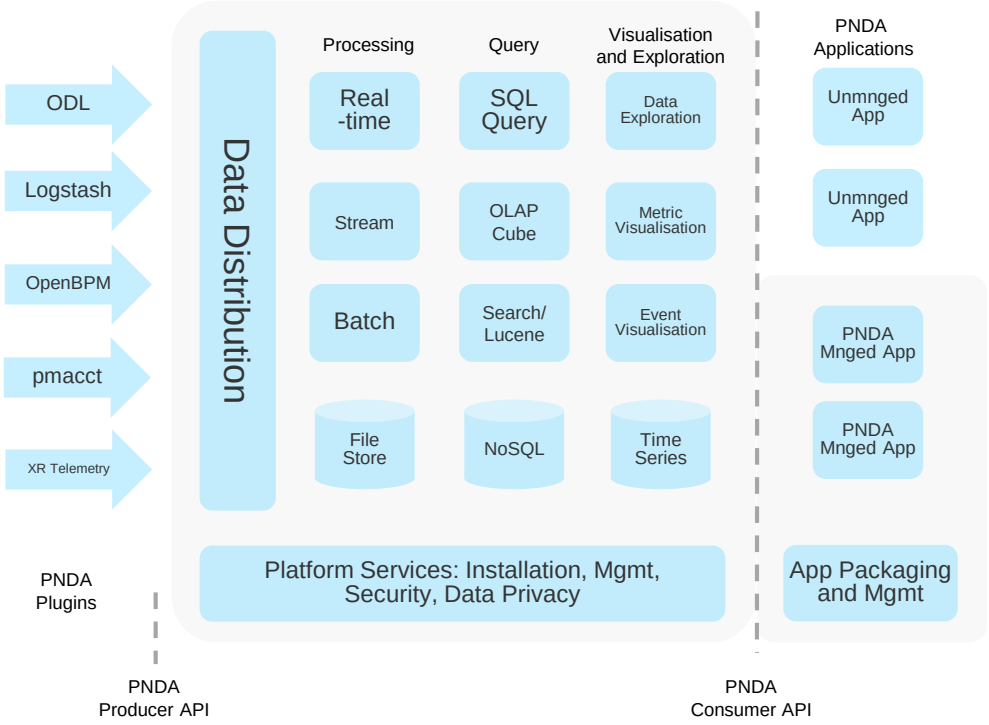

- Simple, scalable open data platform
- Provides a common set of services for developing analytics applications
- Accelerates the process of developing big data analytics applications whilst significantly reducing the TCO
- PNDA provides a platform for convergence of network data analytics

PNDA

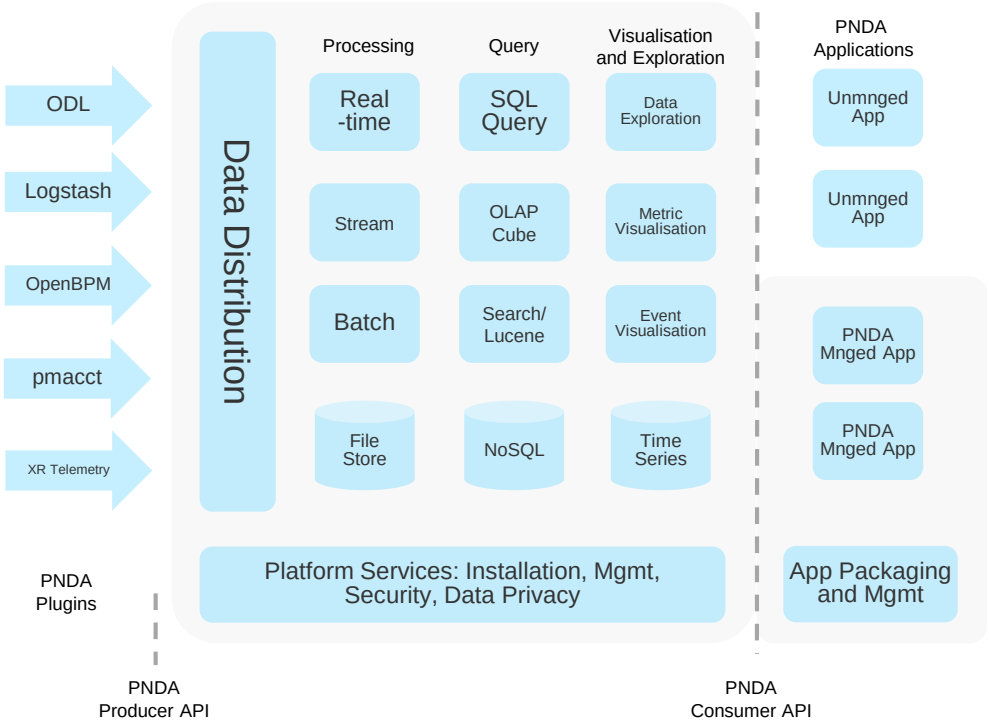

- Horizontally scalable platform for analytics and data processing applications
- Support for near-real-time stream processing and in-depth batch analysis on massive datasets
- PNDA decouples data aggregation from data analysis
- Consuming applications can be either platform apps developed for or client apps integrated with PNDA
- Client apps can use one of several structured query interfaces or consume streams directly.
- Leverages best current practise in big data analytics

PNDA 3.5 Capabilities

- Platform for data aggregation, distribution, processing and storage
- Automated installation, creation, and configuration
  - Openstack, AWS and baremetal
  - Ubuntu and RHEL
  - Typical install ~1hr
  - Online and offline install; modular install
- Open producer and consumer APIs
  - Avro platform schema
- Plugins for Logstash, pmacct, OpenBMP, OpenDaylight, Cisco XR-telemetry, bulk ingest …
- Data distribution – Apache Kafka
- Data store:
  - Automated data partitioning and storage (HDFS)
  - OpenTSDB – time series analysis
  - Hbase - NoSQL
- Support for batch and stream processing:
  - Apache Spark and Spark Streaming
- Jupyter notebook server for app prototyping and data exploration
- Impala-based SQL query support
- Grafana for time series visualisation
- PNDA application packaging
- PNDA management and dashboard

PNDA Console

- The PNDA console provides a dashboard across all components in a cluster
- Inbuilt platform test agents verify the operation of all components
- Active platform testing verifies the end-to-end data pipeline

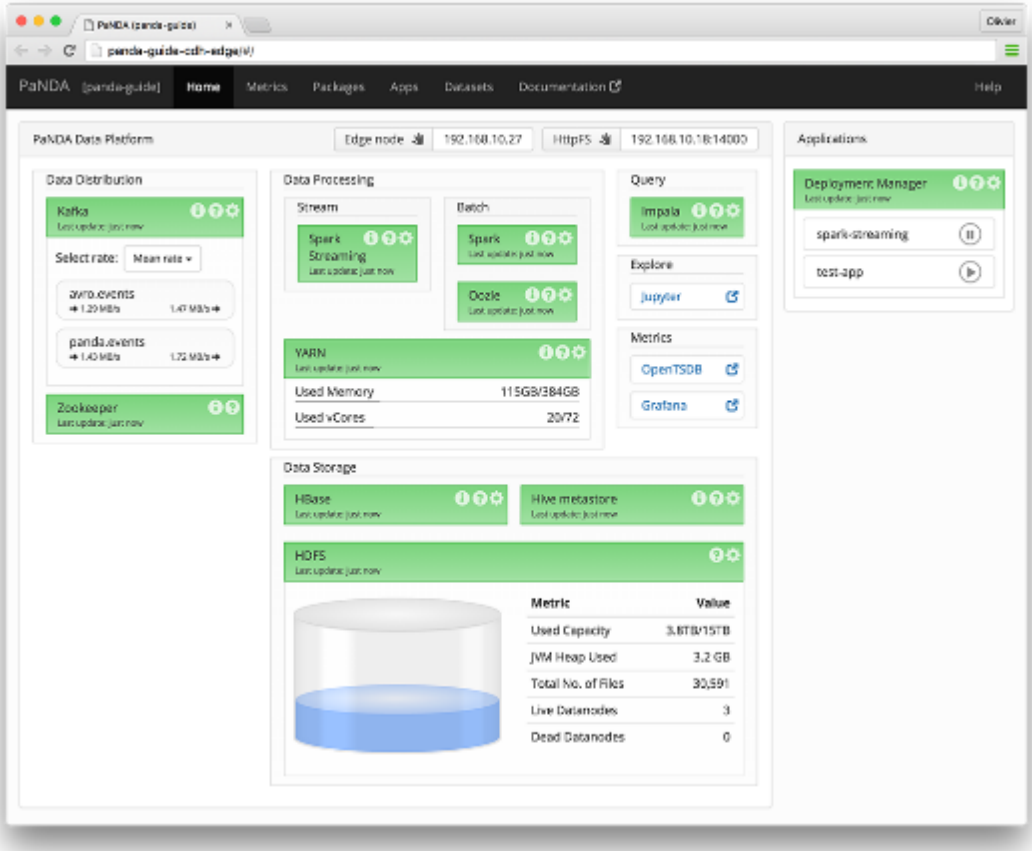

PNDA Flavors span the Data Science Lifecycle

pnda standard

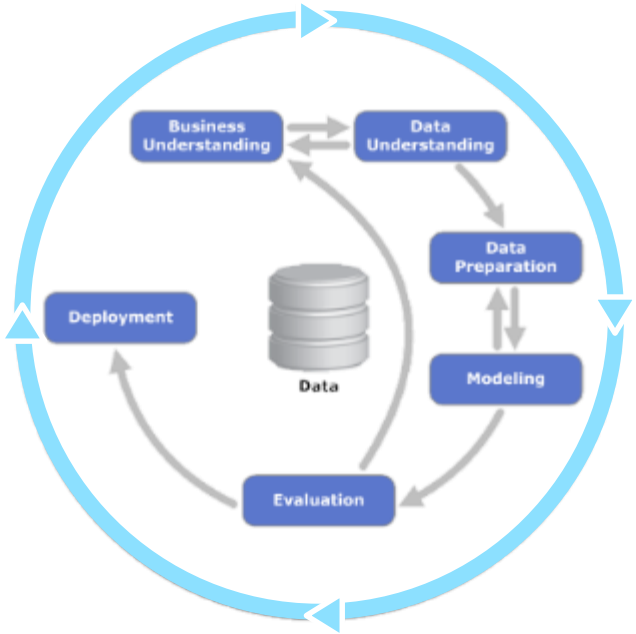

pnda pico

Red pnda

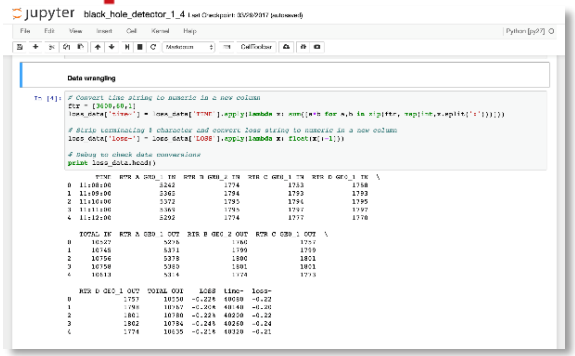

Red PNDA

- Smaller, simpler subset of PNDA designed for development, demonstration and education
- Can run on your laptop
- Consistent technologies, including:
  - PNDA data-ingest (Kakfa/AVRO)
  - Data-exploration tools: Jupyter, OpenTSDB and Grafana
  - Apache Spark and Hbase
- Doesn’t include HDFS and other Hadoop infrastructure for distributed processing.

Multiple Feedback Loops

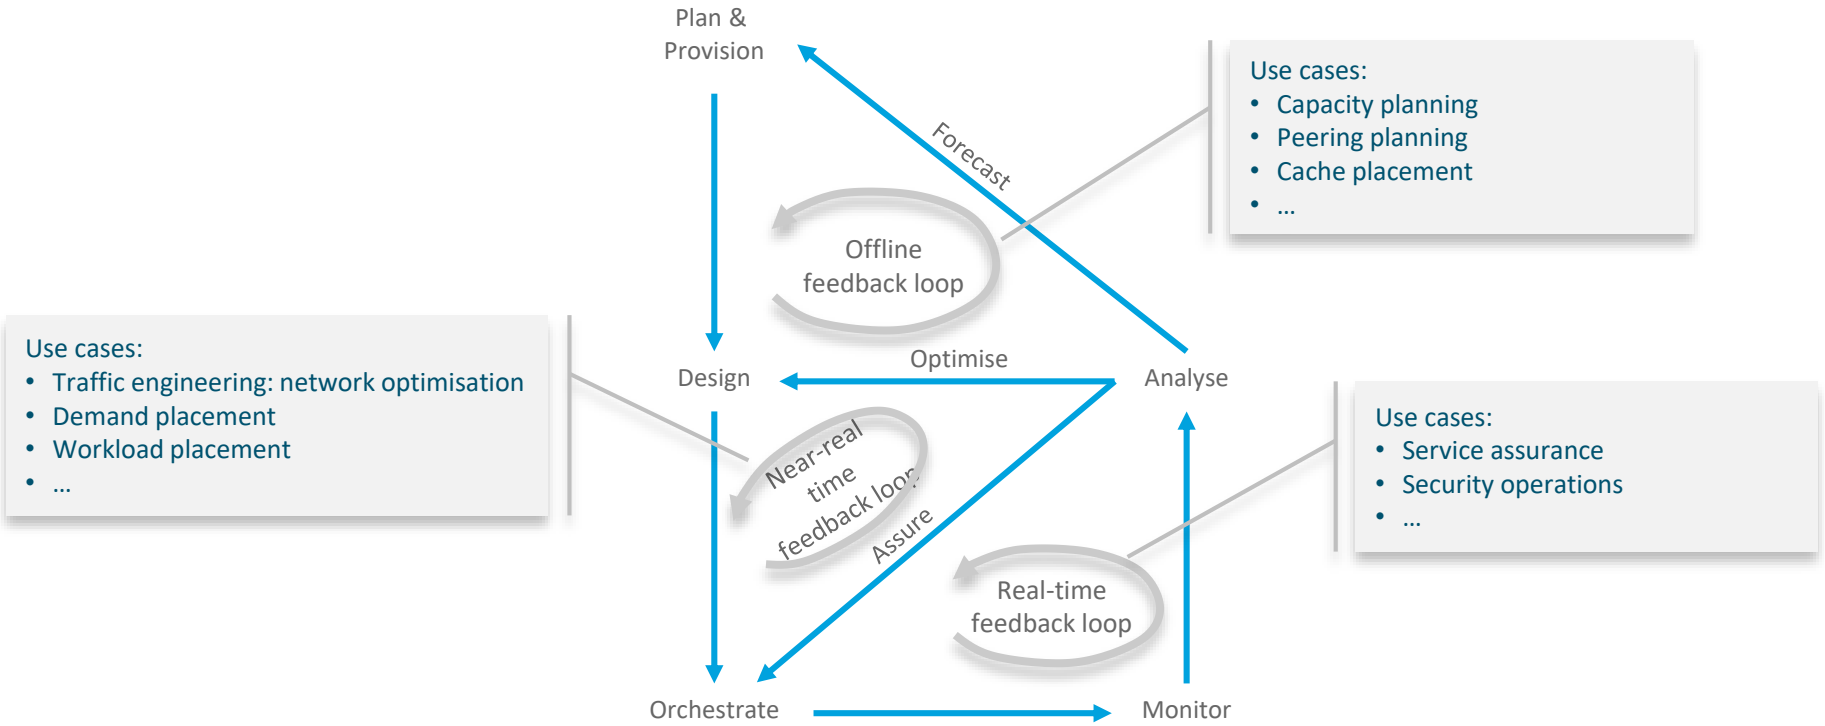

Anomaly detected:
Traffic fails proof of transit for configured path

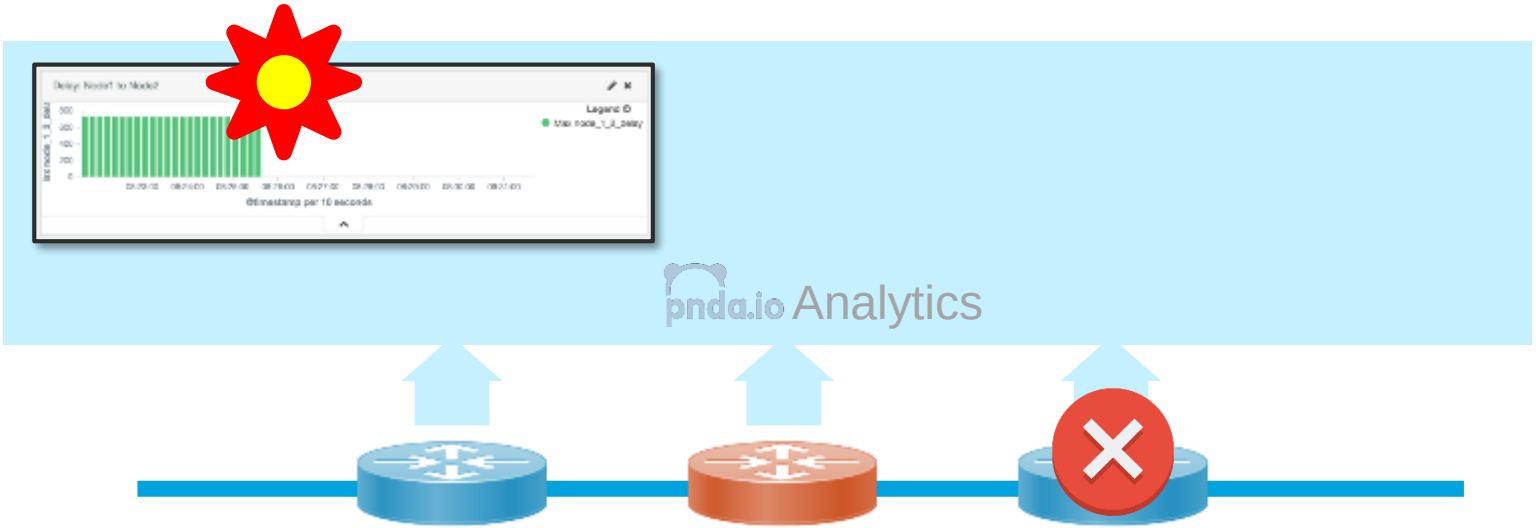

There is a new low-bandwidth, high-delay link…

System re-surrects required state

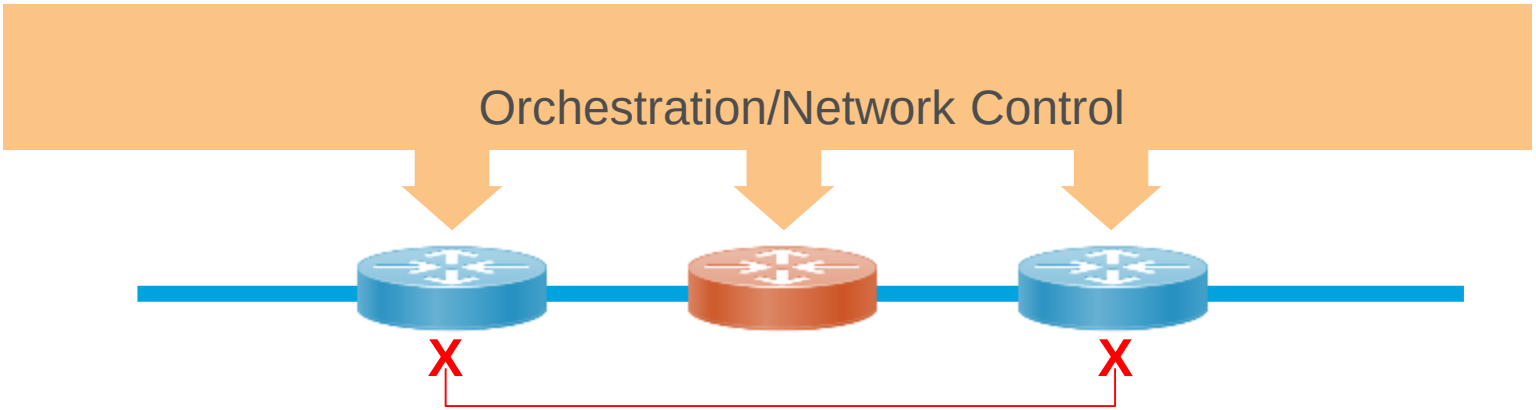

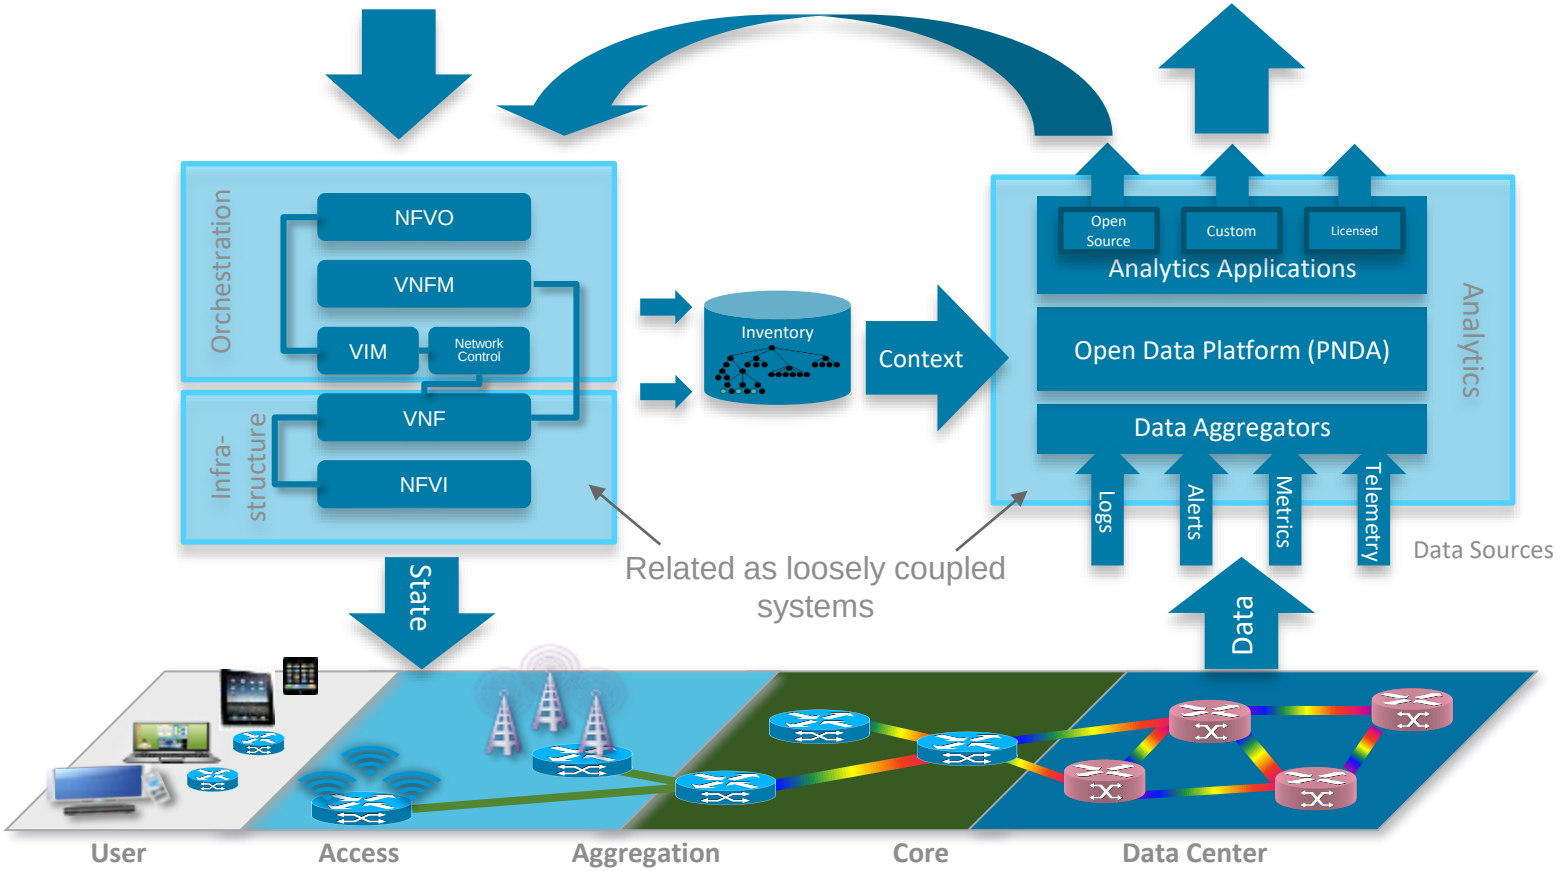

ONAP R1 Architecture

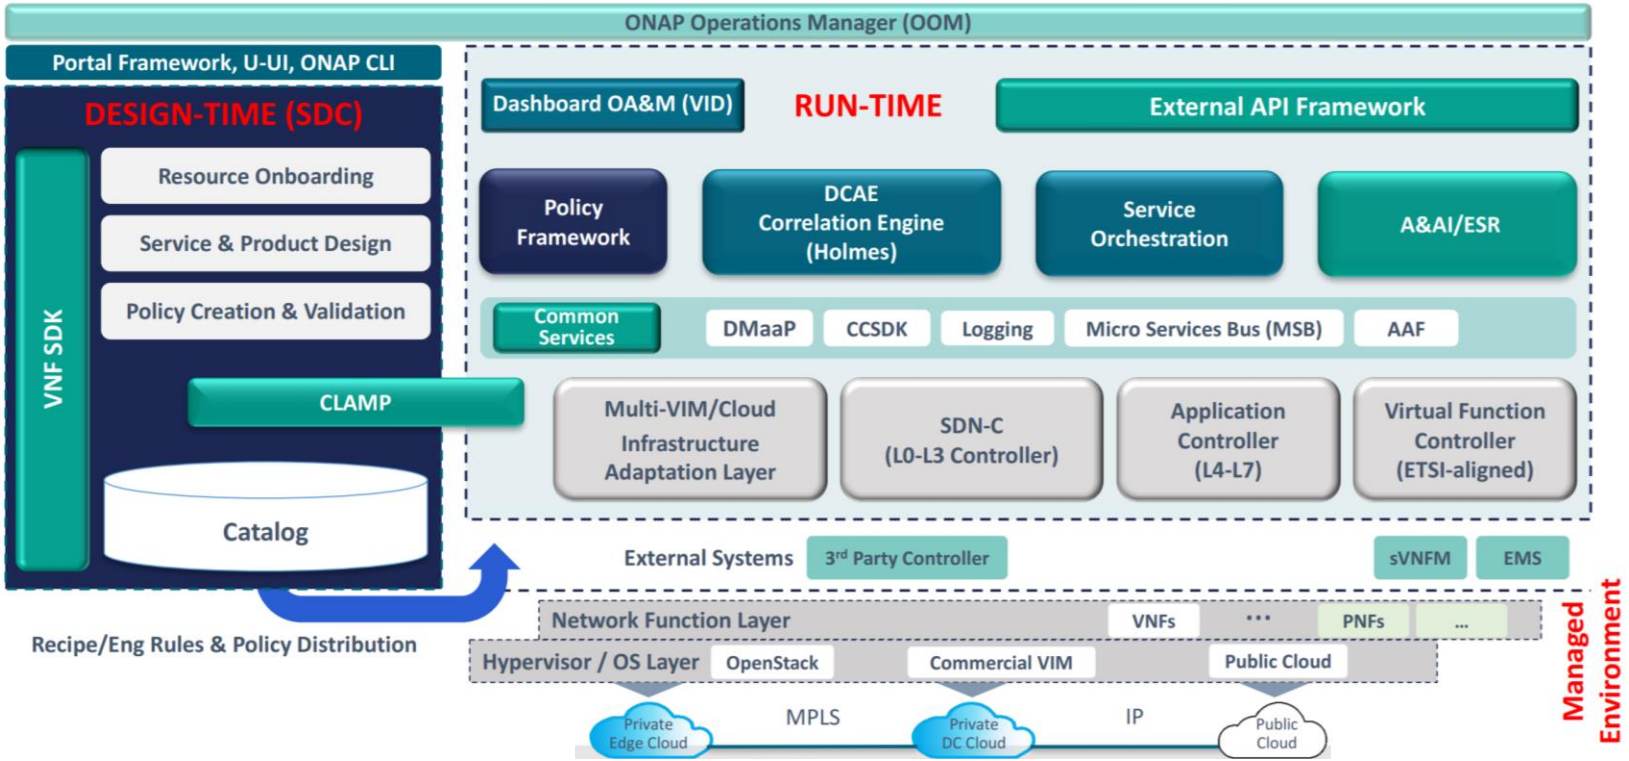

PNDA approach and objectives match ONAP DCAE

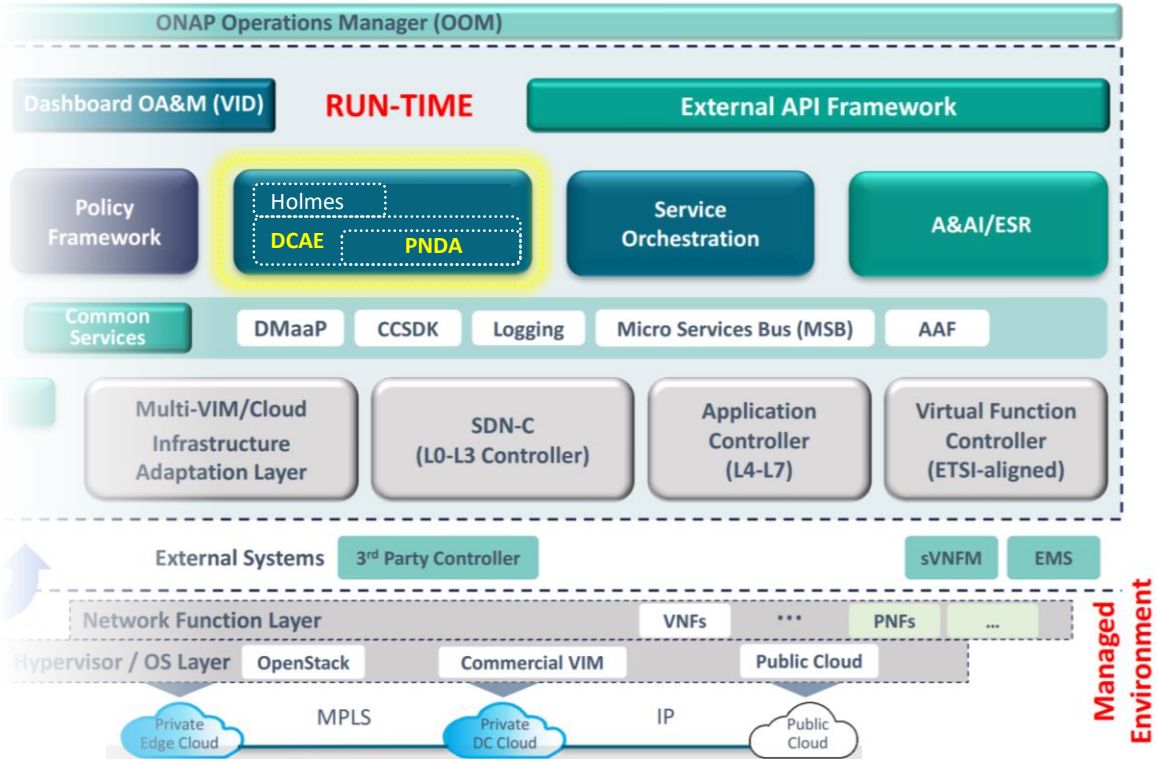

DCAEgen2 – Overview

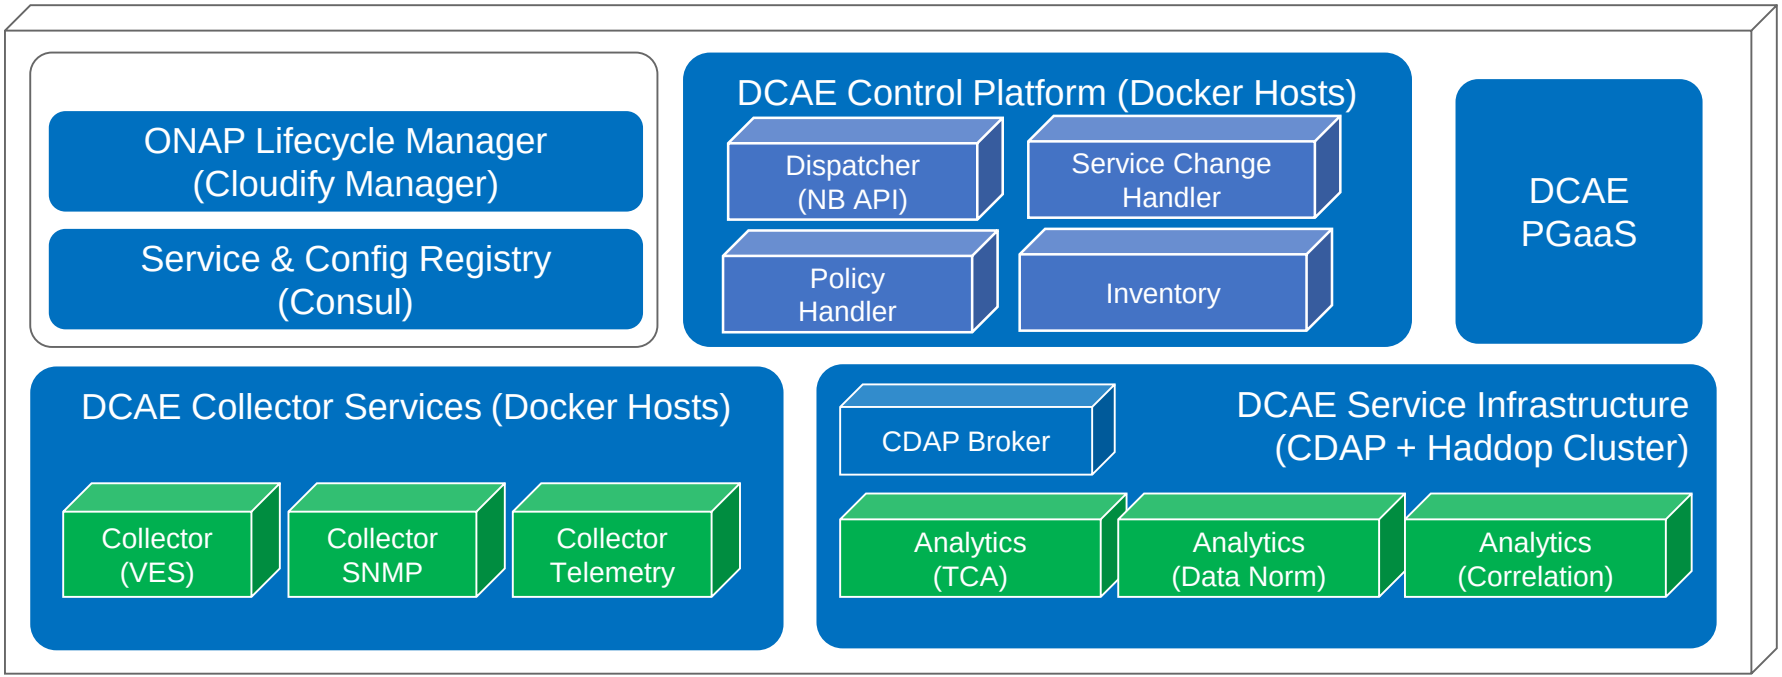

PNDA within DCAEgen2

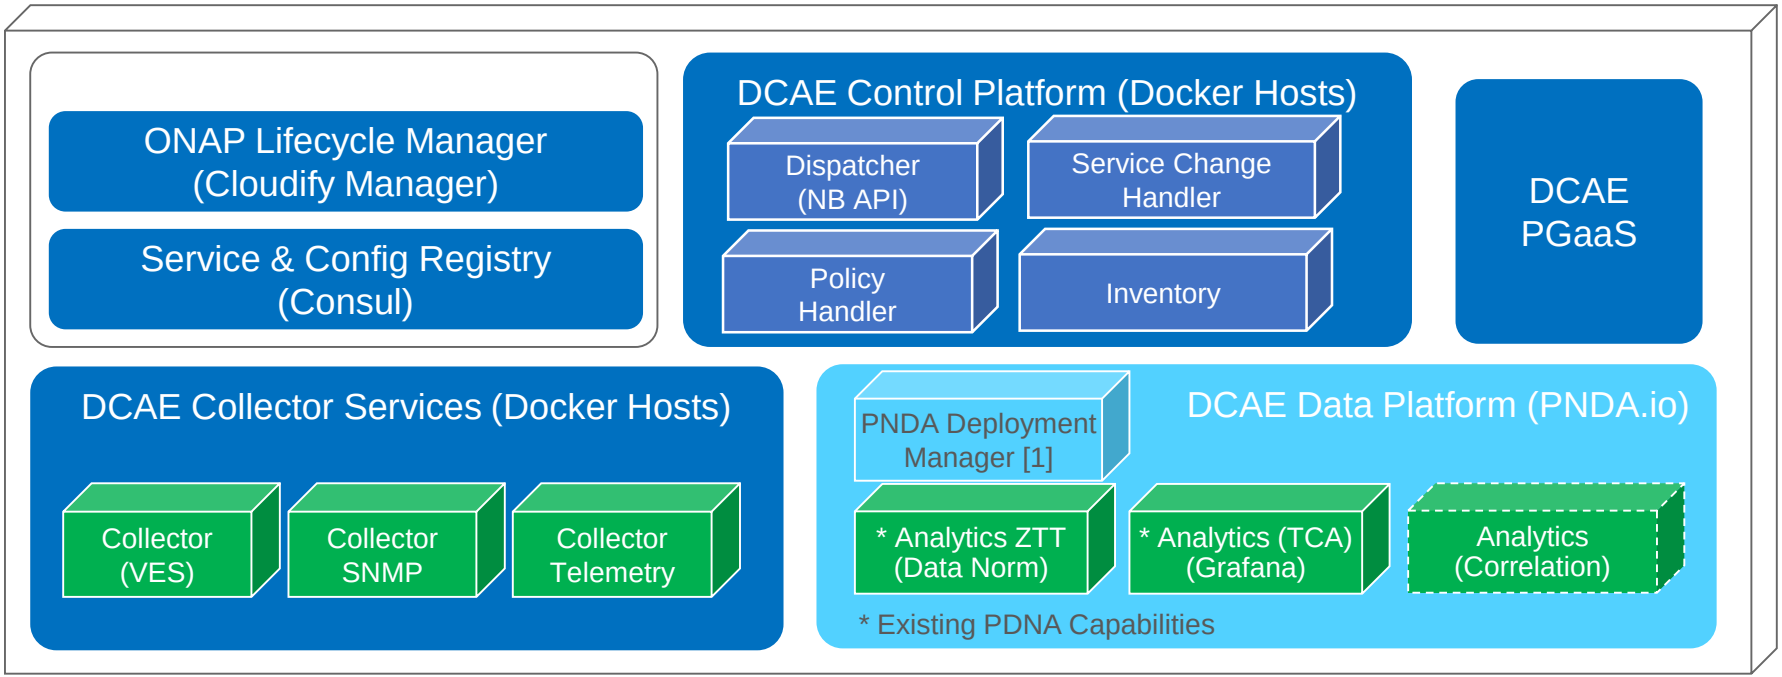

DCAEgen2 with PNDA

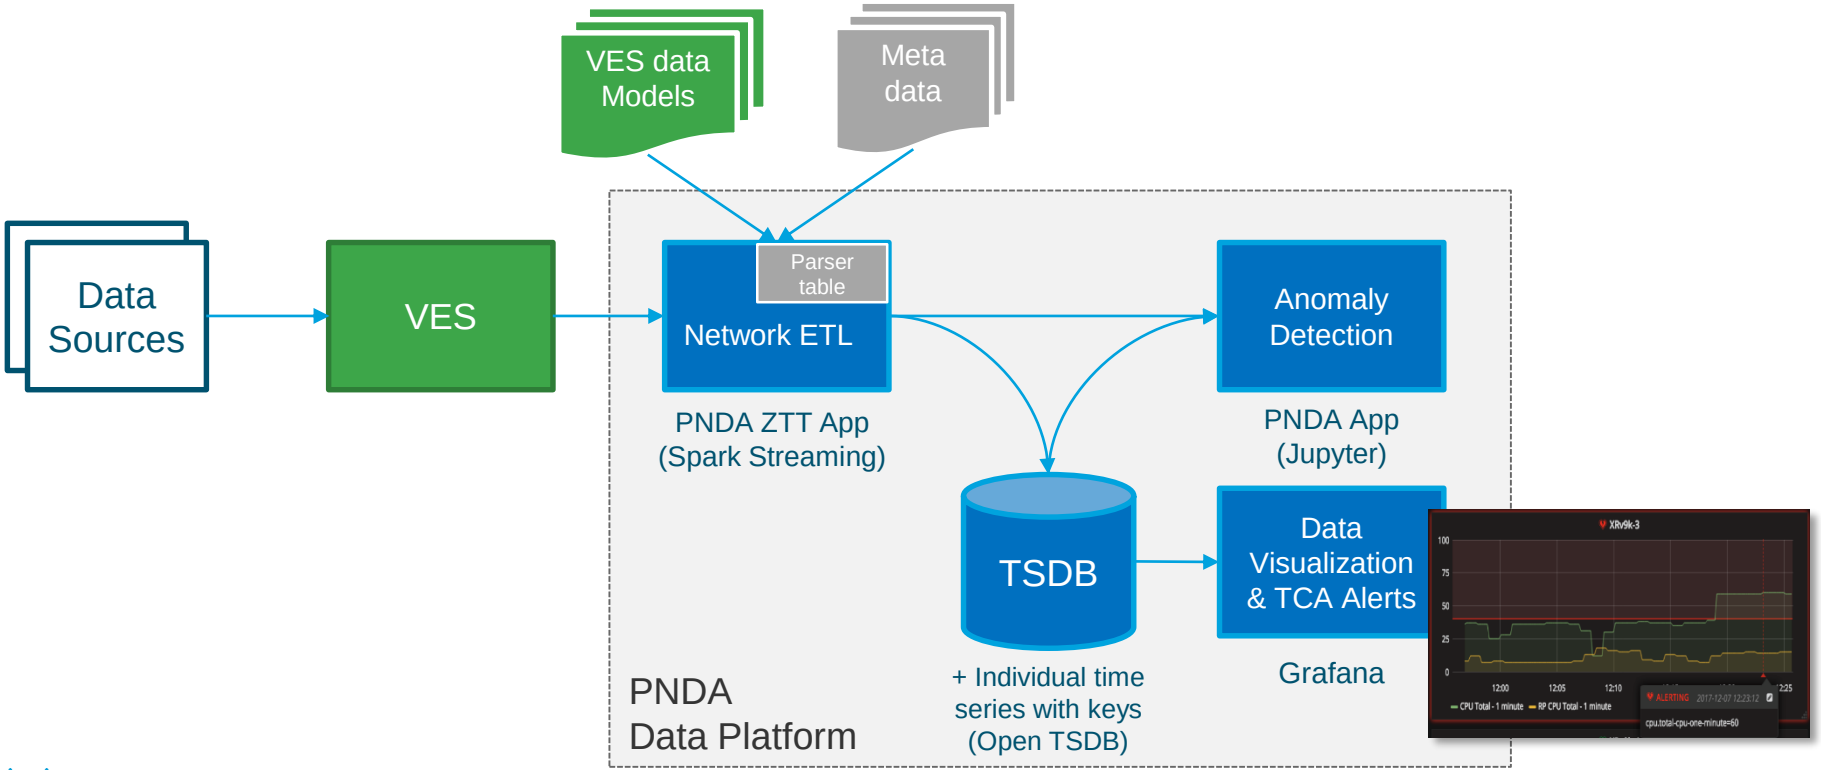

Increasing List of Use Cases – Supported by PNDA

- Analytics Based Service assurance
- ML-based Security Analytics with Apache SPOT on PNDA
- Path Anomaly detection in PNDA using in-band OAM
- Openstack Analytics with PNDA and Calipso
- Smart Transport – Connected Car Cloud Analytics with Machine Learning using PNDA
- BGP analytics with SNAS.io and PNDA.io
- ETSI NFV and Big Data Analytics with PNDA
- PNDA and Paris IOT Smart Cities Pilot
- Cable Plant Anomaly Detection with PNDA

More Information:

github.com/CiscoDevNet/iOAM

pnda.io

Thank You