

Telemetry: Industry Status, Challenges, and IETF Opportunities

Frank Brockners, Benoit Claise

IETF 104 TechTalk, March 28, 2019

Monitoring & Tracing...


An Unlikely Target for Radical Innovation?

Let's start with network monitoring

Business

Application

Network




Too Slow

Incomplete

Device-Specific

Hard to Operationalize

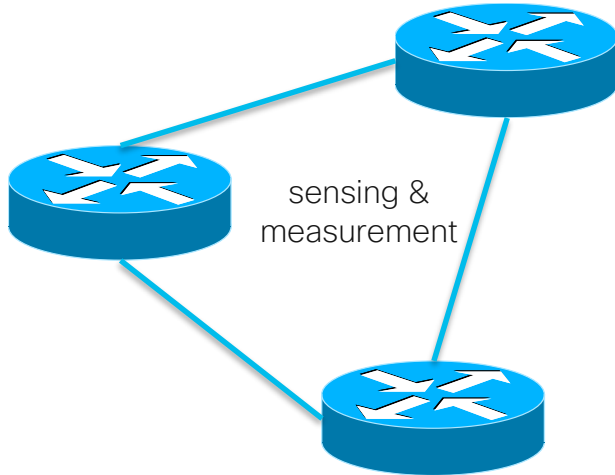


Really Hard to Operationalize

However, Automation is required

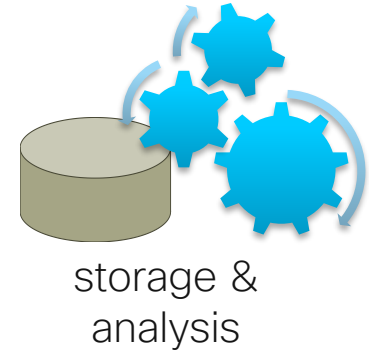
Free the Data!

Where Data Is Created



As Much Data
As Fast
As Relevant
As Easy
As Possible

Where Data Is Useful



So, What is Telemetry?

“**Telemetry** is an automated communications process by which measurements and other data are collected at remote or inaccessible points and transmitted to receiving equipment for monitoring.”*

- A buzzword in the industry, but nothing new
- Collection process of any useful operational data
- What’s (quite) new: Model-driven Telemetry

* <http://adsabs.harvard.edu/abs/1987STIN...8913455>;
<https://en.wikipedia.org/wiki/Telemetry>

Three Enablers for Network Telemetry




Don't Pull: Push!

Performance



Analytics-Ready Data

Tool-Chains



Data-Model Driven

Automation

Data-Model Driven

```
router>show ip interface  
  
Ethernet 0 is up, line protocol is up  
  
IP address is 10.210.93.51 /16  
.....
```

```
leaf if-index {  
  if-feature if-mib;  
  type int32 {  
    range "1..2147483647";  
  }  
  config false;  
m 5.5. NAS-Port  
d Description
```

ifIndex OB.
SYNTAX
MAX-AC
STATUS
DESCRIP

%PM-4-ER

10 ingressInte

port, port_len

This Attribute indicates the physical port number of the NAS which is authenticating the user. It is only used in Access-Request packets. Note that this is using "port" in its sense of a physical connection on the NAS, not in the sense of a TCP or UDP port number. Either NAS-Port or NAS-Port-Type (61) or both SHOULD be present in an Access-Request packet, if the NAS differentiates among its ports.

value
t
nay

The US-ASCII name taking place, and its length in bytes. The value of this field is client specific. (For example, Cisco uses "tty10" to denote the tenth tty line and "Async10" to denote the tenth async interface). The port_len indicates the length of the port field, in bytes.

14 egressInte

g sent. The value
2863]. Note that
the interfaces may

be renumbered every time the device's management system is re-initialized, as specified in [\[RFC2863\]](#).

What's the Solution?

An expensive mediation function

Or the same data model language

What is important is the semantic!

Analytics-Ready Data

- Next to the semantic definition, the context information
- **Data manifest:** all the necessary metadata required to interpret the telemetry information.
For example, how and when the data were measured.
- **Platform manifest:** all the necessary platform metadata required to interpret the telemetry information
For example, the equivalent of show version on a router

Data Models Define *What* You Want to Stream

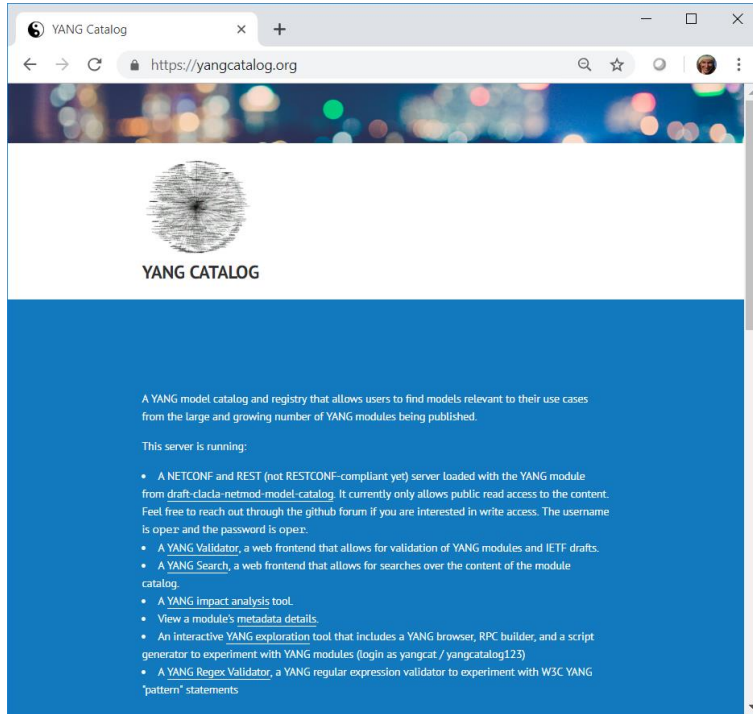
Tools Determine *How* You Consume The Data

There Are A Lot Data Model Tools Out There

OpenSource validators

```
$ pyang -f tree ietf-interfaces.yang --tree-path=interfaces-state/interface/statistics
module: ietf-interfaces
  x--ro interfaces-state
    x--ro interface* [name]
      x--ro statistics
        x--ro discontinuity-time      yang:date-and-time
        x--ro in-octets?              yang:counter64
        x--ro in-unicast-pkts?       yang:counter64
        x--ro in-broadcast-pkts?     yang:counter64
        x--ro in-multicast-pkts?     yang:counter64
        x--ro in-discards?           yang:counter32
        x--ro in-errors?             yang:counter32
        x--ro in-unknown-protos?     yang:counter32
        x--ro out-octets?            yang:counter64
        x--ro out-unicast-pkts?      yang:counter64
        x--ro out-broadcast-pkts?    yang:counter64
        x--ro out-multicast-pkts?    yang:counter64
        x--ro out-discards?          yang:counter32
        x--ro out-errors?            yang:counter32
```

There Are A Lot Of Data Models Out There... yangcatalog.org



The screenshot shows the homepage of the YANG Catalog. At the top, there is a navigation bar with the site name and a search bar. Below the navigation bar is a large image of a globe with the text "YANG CATALOG" underneath. The main content area is a blue box with white text. It starts with a paragraph: "A YANG model catalog and registry that allows users to find models relevant to their use cases from the large and growing number of YANG modules being published." This is followed by a section titled "This server is running:" which lists several tools and services: a NETCONF and REST server, a YANG Validator, a YANG Search tool, a YANG impact analysis tool, a module metadata viewer, an interactive YANG exploration tool, and a YANG Regex Validator.

YANG Catalog

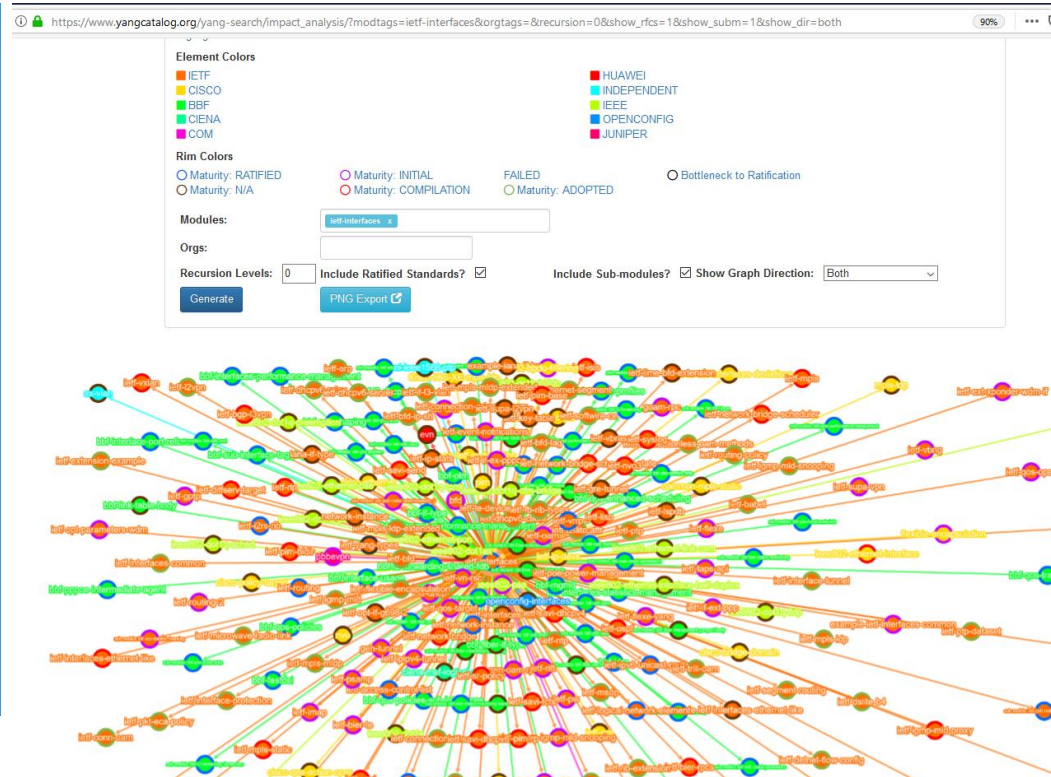
YANG CATALOG

A YANG model catalog and registry that allows users to find models relevant to their use cases from the large and growing number of YANG modules being published.

This server is running:

- A NETCONF and REST (not RESTCONF-compliant yet) server loaded with the YANG module from draft:clacla-netmod-model-catalog. It currently only allows public read access to the content. Feel free to reach out through the github forum if you are interested in write access. The username is open and the password is open.
- A YANG Validator, a web frontend that allows for validation of YANG modules and IETF drafts.
- A YANG Search, a web frontend that allows for searches over the content of the module catalog.
- A YANG impact analysis tool.
- View a module's metadata details.
- An interactive YANG exploration tool that includes a YANG browser, RPC builder, and a script generator to experiment with YANG modules (login as yangcat / yangcatalog123)
- A YANG Regex Validator, a YANG regular expression validator to experiment with W3C YANG "pattern" statements

<https://yangcatalog.org/>



The screenshot shows the search results page for the YANG Catalog. The URL is "https://www.yangcatalog.org/yang-search/impact_analysis/?modtags=ietf-interfaces&orgtags=&recursion=0&show_rfc=1&show_subm=1&show_dir=both". The page features a search filter panel on the left with various options for element colors, rim colors, modules, and recursion levels. Below the filter panel is a large, complex network graph. The graph consists of numerous nodes, each representing a YANG module, connected by lines representing dependencies. The nodes are color-coded according to the filter panel, and the connections are also color-coded. The graph is dense and shows a complex web of relationships between different YANG modules.

https://www.yangcatalog.org/yang-search/impact_analysis/?modtags=ietf-interfaces&orgtags=&recursion=0&show_rfc=1&show_subm=1&show_dir=both

Element Colors

- IETF
- CISCO
- BBF
- CIENA
- COM
- HUAWEI
- INDEPENDENT
- IEEE
- OPENCONFIG
- JUNIPER

Rim Colors

- Maturity: RATIFIED
- Maturity: N/A
- Maturity: INITIAL
- Maturity: COMPILATION
- Maturity: FAILED
- Maturity: ADOPTED
- Bottleneck to Ratification

Modules:

Orgs:

Recursion Levels: Include Ratified Standards? Include Sub-modules? Show Graph Direction:

Generate

There Are Some Telemetry Tools Out There

The screenshot displays the 'Advanced Netconf Explorer' web interface. The browser address bar shows 'Netconf: 10.49.234.115' and the page title is 'Advanced Netconf Explorer'. The interface is divided into several sections:

- Navigation:** 'Start' and 'NETCONF console' buttons.
- YANG Models:** A dropdown menu for selecting models, with 'View' and 'Download all' buttons.
- Telemetry Tools:** A dropdown menu for selecting tools, with 'Edit group' and 'Live data' buttons.
- GNMI Subscribe Interval:** A dropdown menu for selecting intervals, with 'GNMI Port' and 'GNMI Subscribe' buttons.
- Capabilities:** A list of YANG model URLs, including:
 - `http://cisco.com/calvados/Cisco-IOS-XR-sysadmin-issu?module=Cisco-IOS-XR-sysadmin-issu&revision=2018-07-05`
 - `http://cisco.com/calvados/Cisco-IOS-XR-sysadmin-time-of-day-timezone?module=Cisco-IOS-XR-sysadmin-time-of-day-timezone`
 - `http://cisco.com/ns/yang/Cisco-IOS-XR-Subscriber-infra-subdb-oper?module=Cisco-IOS-XR-Subscriber-infra-subdb-oper&revision=2018-09-04`
 - `http://cisco.com/ns/yang/Cisco-IOS-XR-aaa-aaacore-cfg?module=Cisco-IOS-XR-aaa-aaacore-cfg&revision=2018-09-04`
- Failed YANG models:** A list of failed models, with one highlighted: `talif-common@2017-08-23.yang` with the error message: `tries to imports missing model: talif-meta-extensions, talif-cli-extensions`.
- Search Results:** A search bar with 'wdsysmon' entered, and search buttons 'Apply', 'Clear', and 'Show Data'. The results show a tree structure for 'system-monitoring' with sub-nodes for 'cpu-utilization', 'process-cpu', and 'node-name'.

Netconf: 10.49.234.115

Not secure | timon:9269

Advanced Netconf Explorer

Device: 10.49.234.115 (514 YANG models)

Start NETCONF console

Search models Search nodes Apply Clear Show Data

YANG Models

Telemetry Tools

mdt-realtime

GNMI Subscribe Interval GNMI Port

change 57777

Capabilities

Failed YANG models

tallf-common@2017-08-23.yang tries to imports missing

Live Telemetry data: mdt-realtime

Measurement

[08:30:50] Cisco-IOS-XR-wdsysmon-fd-oper:system-monitoring/cpu-utilization

Telemetry Data (JSON format)

```
{
  "node_id_str": "IOSXRv9k-1",
  "subscription_id_str": "anx-1542356893720",
  "encoding_path": "Cisco-IOS-XR-wdsysmon-fd-oper:system-monitoring/cpu-utilization",
  "collection_id": 3850936,
  "collection_start_time": 1.542357050776E12,
  "msg_timestamp": 1.542357050963E12,
  "data_json": [
    {
      "timestamp": 1.542357050939E12,
      "keys": [
        {
          "node-name": "0/RP0/CPU0"
        }
      ],
      "content": {
        "total-cpu-one-minute": 3,
        "total-cpu-five-minute": 3,
        "total-cpu-fifteen-minute": 4,
        "process-cpu": [
          {
            "process-name": "Init",
            "process-id": 1,
            "process-cpu-one-minute": 0,
            "process-cpu-five-minute": 0,

```

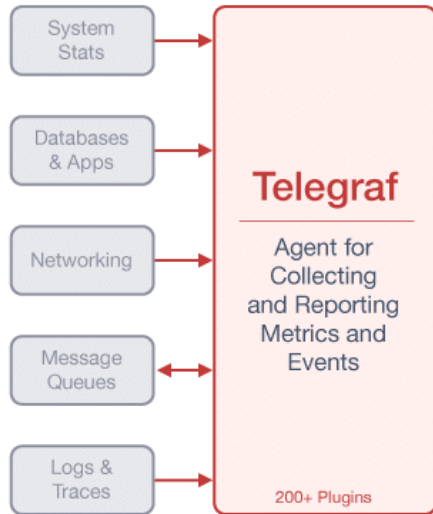
bfd (http://cisco.com/ns/yang/Cisco-IOS-XR-ip-bfd-cfg)

bfd (http://cisco.com/ns/yang/Cisco-IOS-XR-ip-bfd-oper)

bgp (http://cisco.com/ns/yang/Cisco-IOS-XR-ipv4-bgp-cfg)

bgp (http://cisco.com/ns/yang/Cisco-IOS-XR-ipv4-bgp-oper)

There Are Some Telemetry Tools Out There Telegraf Example



The image displays two overlapping screenshots. The background screenshot shows the GitHub repository for the 'telegraf-plugin' on GitHub, specifically the 'cisco_telegram.yml' file. The foreground screenshot shows the 'Cisco Nexus Telemetry - Device View' interface. This interface is a dashboard with multiple panels displaying various network metrics and configurations. The panels include:

- CPU Usage:** A line graph showing CPU usage over time, with a gauge indicating 12.9%.
- Memory Usage:** A gauge showing memory usage.
- Temperature Sensors:** A gauge showing temperature.
- Ingress Traffic - Packet Type:** A line graph showing traffic volume.
- Ingress Traffic - Bytes:** A line graph showing traffic volume.
- EGP Neighbors:** A table showing neighbor status.
- IPv4 Pools:** A table showing pool status.
- VLAN:** A gauge showing 1,048 VLANs.
- MAC Addresses:** A gauge showing 48,683 MAC addresses.
- VNI:** A gauge showing 1,092 VNIs.
- OSF Adjacencies:** A gauge showing 3 adjacencies.
- MAC Addresses Table:**

MAC Address	VLAN	Port
1	vlan-200	po200
1	vlan-123	po4095
1	vlan-300	po4095
- BGP L2VPN EVPN Routes:** A table showing route profiles.
- BGP L2VPN EVPN Routes:** A gauge showing 8 routes.
- VNI Table:**

VNI	Type	Status	VLAN	Mode	Multicast Group
10200	L2	Up	vlan-200	CP	233.1.0.0/22
102000	L3	Up	vlan-2000	CP	0.0.0.0/22
- VNI Traffic Statistics:** A line graph showing traffic volume.

At the bottom of the foreground screenshot, a code editor shows the following configuration snippet:

```
## redial in case of failures after  
redial - "10s"
```

<https://github.com/ios-xr/telegraf-plugin>

There Are Some Telemetry Tools Out There The Industry Needs More

Cisco Telemetry Data Mapper

Input the paths you wish to check TDM for matches for, and it will report back direct matches, calculations, and those that do not yet have matches.

The below form will resize dynamically with the amount of data paths you input.

Paths

[Get Matches](#)

Matches

Submitted	Returned
ifAdminStatus	Cisco-IOS-XR-pf1-im-cmd-oper:interfaces/interface-xr/interface/state
	Cisco-IOS-XR-pf1-im-cmd-oper:interfaces/interfaces/interface
	ietf-interfaces:interfaces-state/interface/admin-status
	openconfig-interfaces:interfaces/interface/state/admin-status
	openconfig-interfaces:interfaces/interface/subinterfaces/subinterface/state/admin-status
ifIndex	Cisco-IOS-XR-snmp-agent-oper:snmp/if-indexes
	Cisco-IOS-XR-snmp-agent-oper:snmp/interface-indexes/interface-index/interface-index
	ietf-interfaces:interfaces-state/interface/if-index
ifOperStatus	openconfig-interfaces:interfaces/interface/subinterfaces/subinterface/state/index
	Cisco-IOS-XR-pf1-im-cmd-oper:interfaces/interface-xr/interface/state
	Cisco-IOS-XR-pf1-im-cmd-oper:interfaces/interfaces/interface
	ietf-interfaces:interfaces-state/interface/oper-status
	openconfig-interfaces:interfaces/interface/state/oper-status
	openconfig-interfaces:interfaces/interface/subinterfaces/subinterface/state/oper-status

IETF Specifications

Tools

(OSS) Reference Code

Industry Coordination

WHAT?

(Model)



YANG



YANG
Tools



YANG
Tools



HOW?

(Associated Operations)


No telemetry
spec. ... yet



Yes, but ...

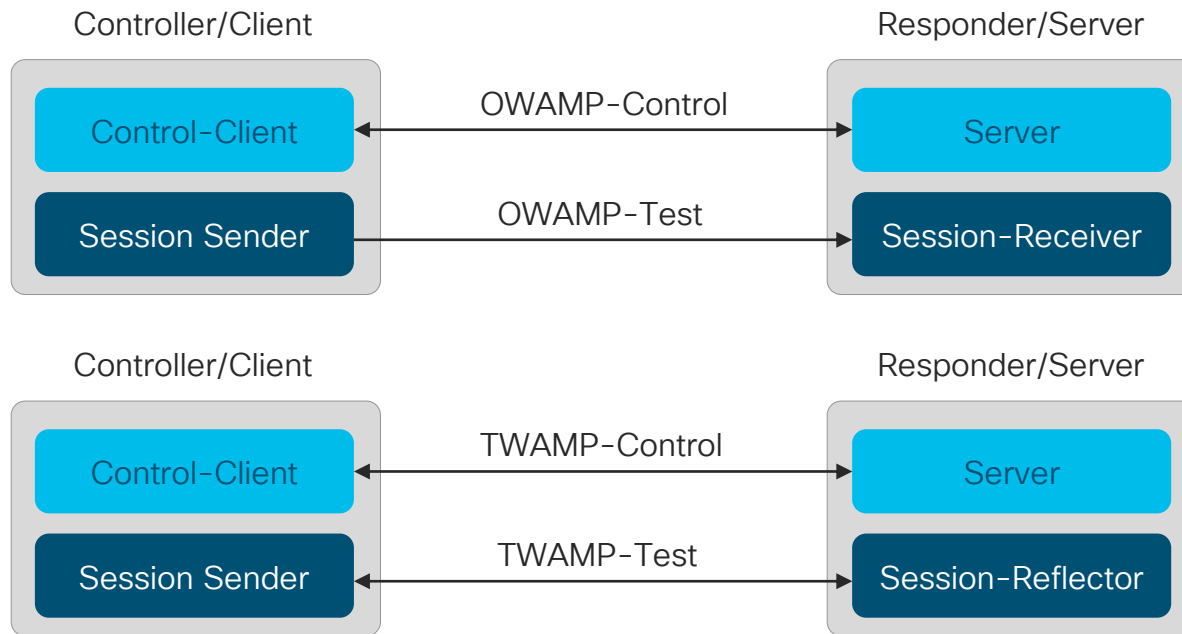
Active OAM

- Ping
- Traceroute
- TWAMP
- OWAMP
- ...



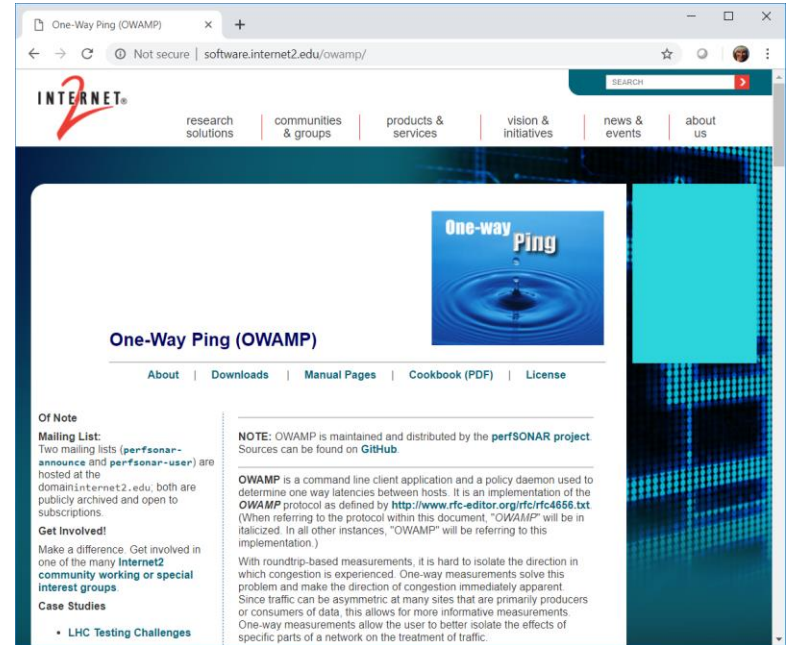
Probe
(ping, traceroute)


Active OAM Examples: OWAMP (RFC 4656) & TWAMP (RFC 5357)




Tools...

- OWAMP sample OSS code:
<https://github.com/perfsonar/owamp>
- TWAMP sample OSS code:
<https://github.com/emirica/twamp-protocol.git>
- And there are many more active OAM tools, incl. [Paris Traceroute](#), etc..






Observe
(SNMP, Syslog,
Streaming Telemetry..)




Probe
(ping, traceroute,
OWAMP, TWAMP,..)

But what
about my
live user
traffic?



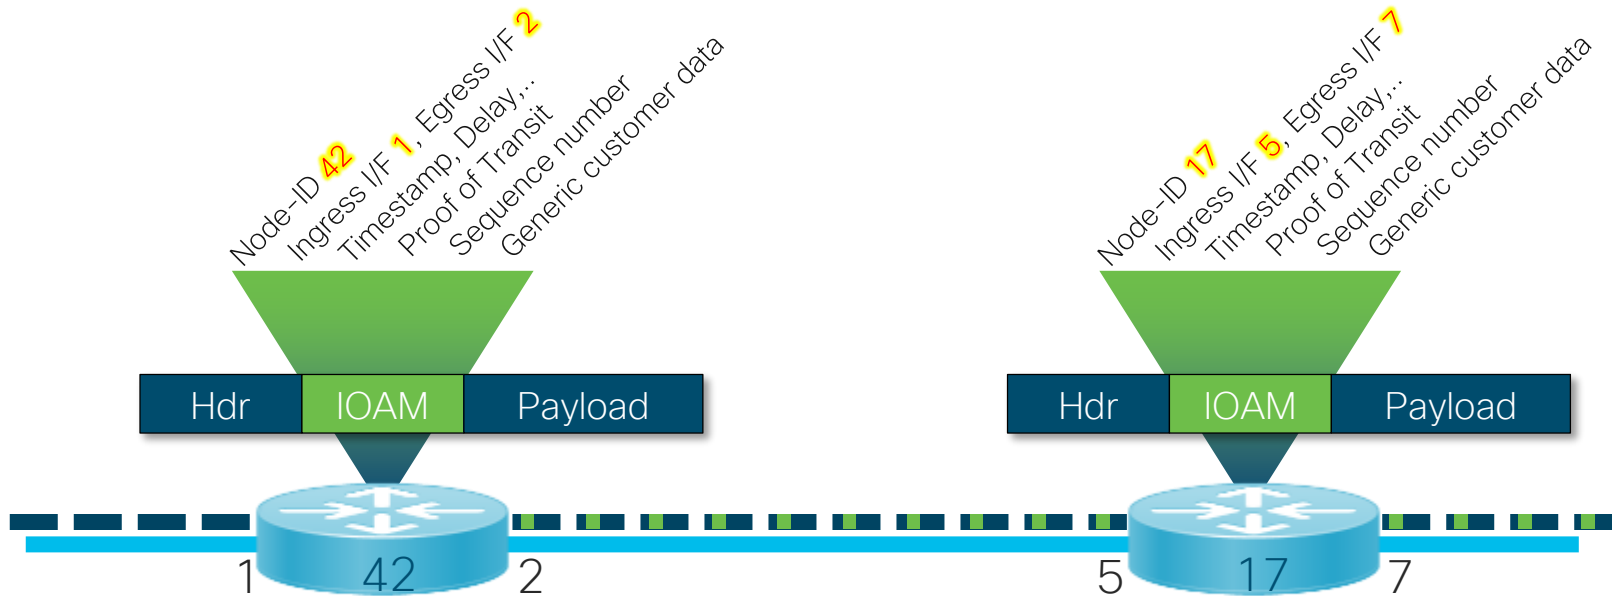
Observe
(SNMP, Syslog
Streaming Telemetry,..)



Probe
(ping, traceroute,
OWAMP, TWAMP,..)



In-situ OAM
(per packet telemetry)



IOAM data fields are defined in a protocol
independent way

draft-ietf-ippm-ioam-data-05

In-situ OAM Data Fields Overview

- Per node scope

- Hop-by-Hop information processing
 - Hop Limit
 - Node_ID (long/short)
 - Ingress Interface ID (long/short)
 - Egress Interface ID (long/short)
 - Timestamp
 - Wall clock (seconds, nanoseconds)
 - Transit delay
 - Queue length
 - Opaque data
 - Namespace specific data (long/short)

Two transport options:

- Pre-allocated array (SW friendly)
- Incrementally grown array (HW friendly)

- Set of nodes scope

- Hop-by-Hop information processing
 - Service Chain Validation (Random, Cumulative)

- Edge to Edge scope

- Edge-to-Edge information processing
 - Sequence Number
 - Timestamp

IOAM data fields use a dedicated namespace. IOAM data fields are layer independent and can be filled by any node capable of filling-in IOAM data fields.

IOAM data fields can be carried in various protocols

NSH - draft-ietf-sfc-ioam-nsh-01

IPv6 - draft-ioametal-ippm-6man-ioam-ipv6-options-01, draft-ioametal-ippm-6man-ioam-ipv6-deployment-00

IPv4 - draft-gafni-ippm-ioam-ipv4-options-00

Ethertype based protocol encap - draft-weis-ippm-ioam-eth-01

Geneve - draft-brockners-ippm-ioam-geneve-02

VXLAN-GPE - draft-brockners-ippm-ioam-vxlan-gpe-01

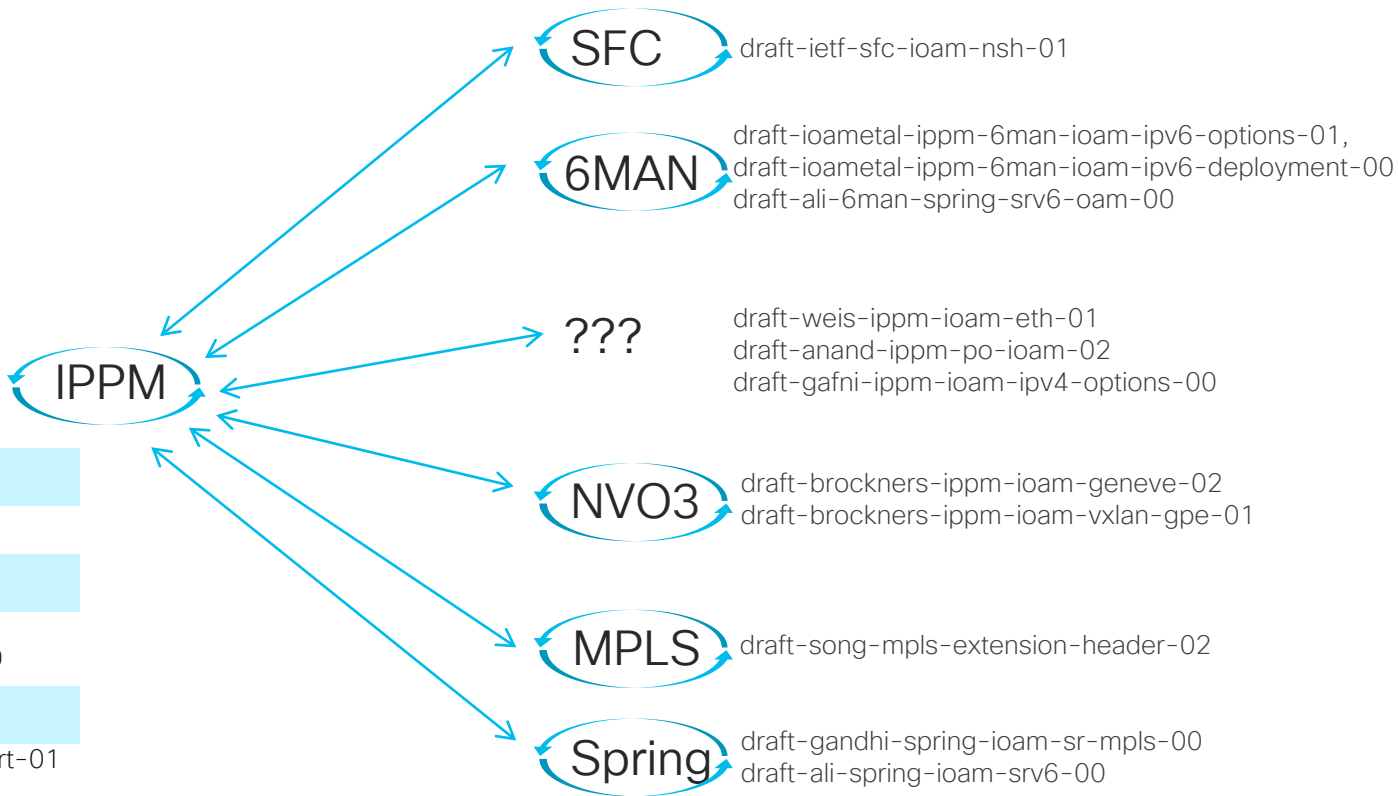
SRv6 - draft-ali-6man-spring-srv6-oam-00, draft-gandhi-spring-ioam-sr-mpls-00, draft-ali-spring-ioam-srv6-00

Optical - draft-anand-ippm-po-ioam-02

MPLS - draft-song-mpls-extension-header-02

...

Encapsulations



Data Formats

draft-ietf-ippm-ioam-data-05

YANG/Operations

draft-zhou-ippm-ioam-yang-03

draft-mizrahi-ippm-ioam-profile-00

Data Export

draft-spiegel-ippm-ioam-rawexport-01

IOAM Adoption:

Standards  OSS  Tool-Chains

IOAM OSS Implementation – Seeding Adoption...



- Dataplane:
 - Open-Source:
 - FD.io/VPP (see fd.io)
 - Linux Kernel ([PoC for 4.12](#))
- Equipment vendors active in IOAM
 - Cisco, Broadcom, Barefoot, Mellanox, Huawei, ZTE, ...
- Controller:
 - OpenDaylight (Carbon release)

Next Step: Evolve tools / tool chains for IOAM YANG Data Models for Operations and Data-Plane

Operational
Template/Model

YANG model for the
configuration of a
deployment



Data-Plane
Template/Model

YANG model for the
data-plane/on-wire
definition

Early seed work: [Draft-zhou-ippm-ioam-yang](#), [draft-mizrahi-ippm-ioam-profile](#)

Network Tracing: Observation:

OSS/Tool-Chains ➡ Adopted Standards

WHAT?

(Model)

HOW?

(Associated Operations)

IETF Specifications

(✓)

Tools

(OSS) Reference Code

✓

✓

Industry Coordination

How to rope
application visibility in?

Business

Application

Network

Application Level Tracing Asks Resemble Network Level Tracing Asks...

*“Which requests are slow?
Why?”*

“DB lookup was slow, but why?”

*“We know which operations
were slow, but can’t correlate
them to a particular request”*

*“How to aggregate application
performance data?”*

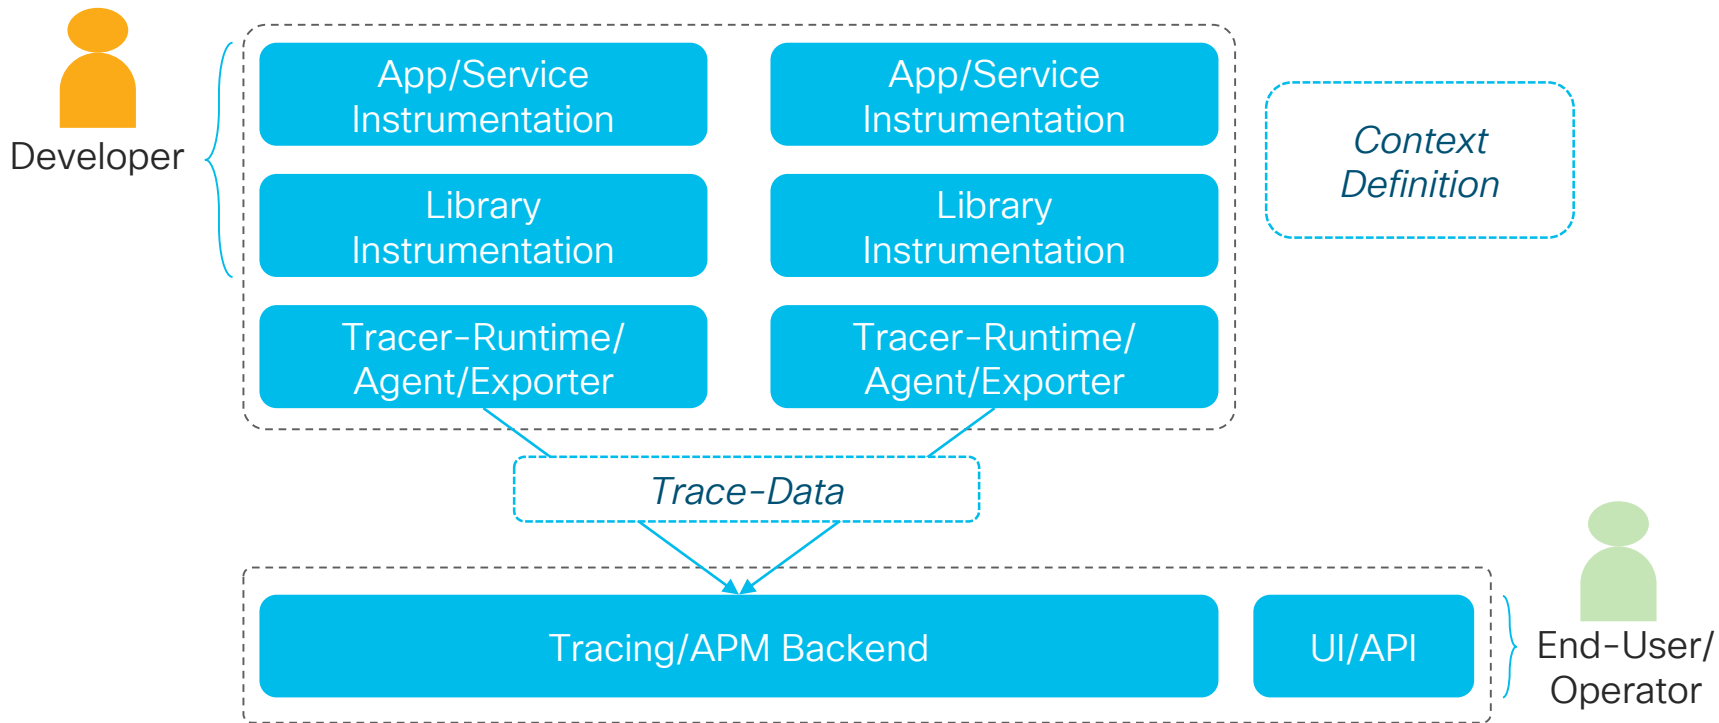
*“Single approach for different
applications/deployments?”*

Background: Application Tracing / Performance Management

- Enterprises use a wide range of technologies across different teams
 - Java, .NET, etc. – of different age
 - Message brokers like Tibco, RabbitMQ, ZeroMQ, ...
- Tools have so far driven Tracing/APM API and schema definitions
 - Many different implementations/tools for tracing / end-to-end transaction monitoring
 - Manual instrumentation with language/tool specific API

Tracing Solutions/ Backends; Examples:

- AppDynamics
- AWS X-Ray
- Azure Monitor
- Datadog
- Elasticsearch
- Honeycomb
- Instana
- Jaeger
- Prometheus
- SignalFX
- Stackdriver
- Zipkin
-



Can There Be Standards?

- Standard tracing API with associated language bindings?
 - Developer writes the instrumentation using that API
- Open interchange of trace data between different
 - Applications/Services?
 - Tracing systems?
 - Network Layers?

OpenTracing

“OpenTracing allows developers to add instrumentation to their application code using APIs that do not lock them into any one particular product or vendor.”

- Vendor-neutral API specification
- frameworks and libraries that have implemented the specification
- documentation for the project.

CNCF Project

OpenCensus

“OpenCensus currently provides libraries for a number of languages that allow you to capture, manipulate, and export metrics and distributed traces to the backend(s) of your choice.”

- Vendor-agnostic single distro of libraries to provide:
- metrics collection
 - tracing for services
 - pluggable exporters to ship data to external systems

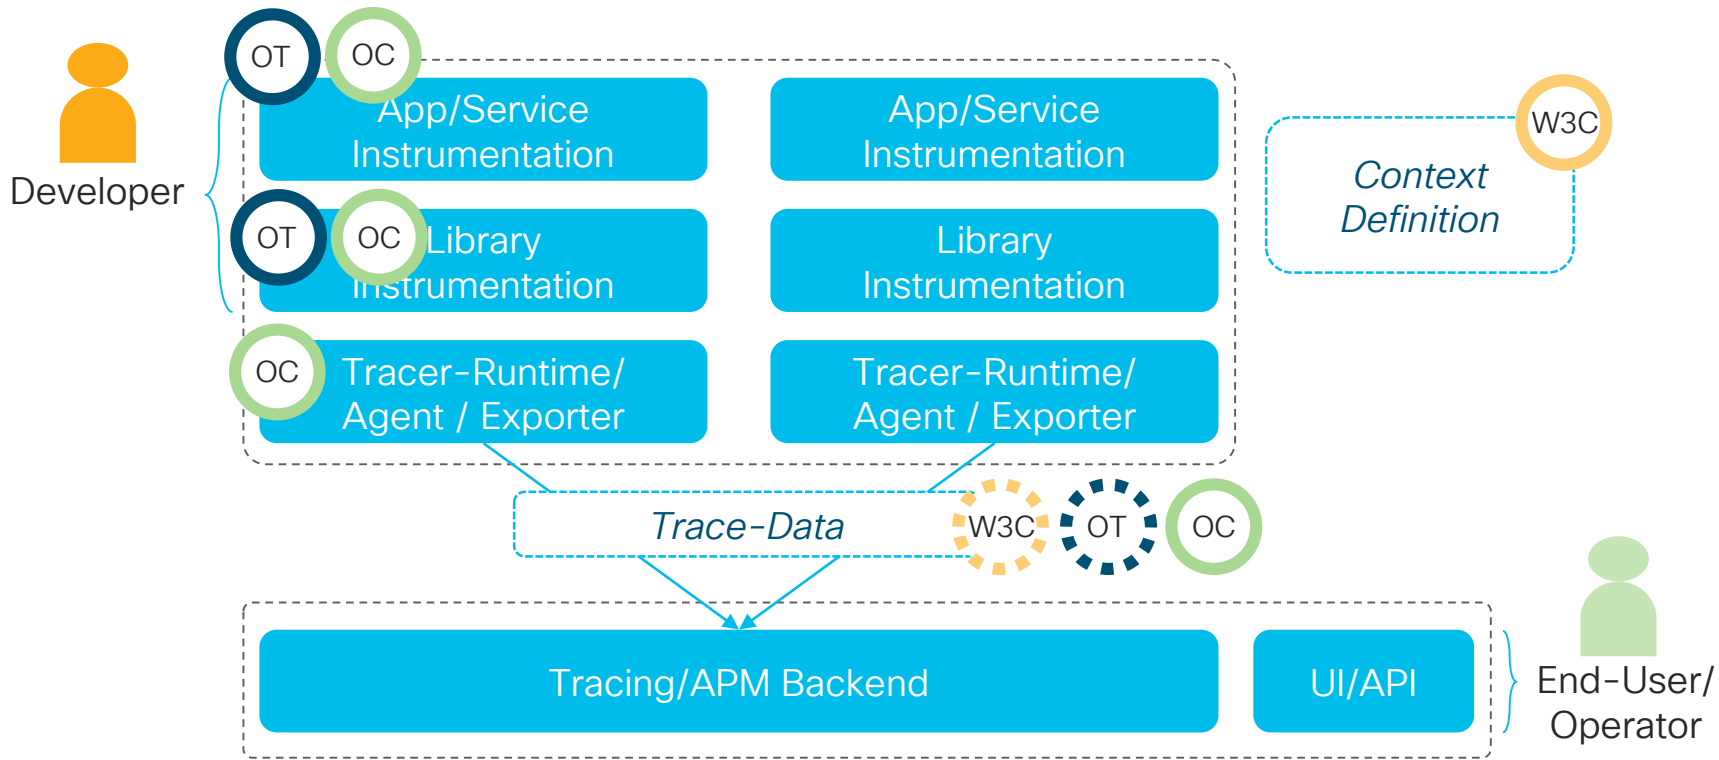
CNCF Project

W3C Distributed Tracing WG

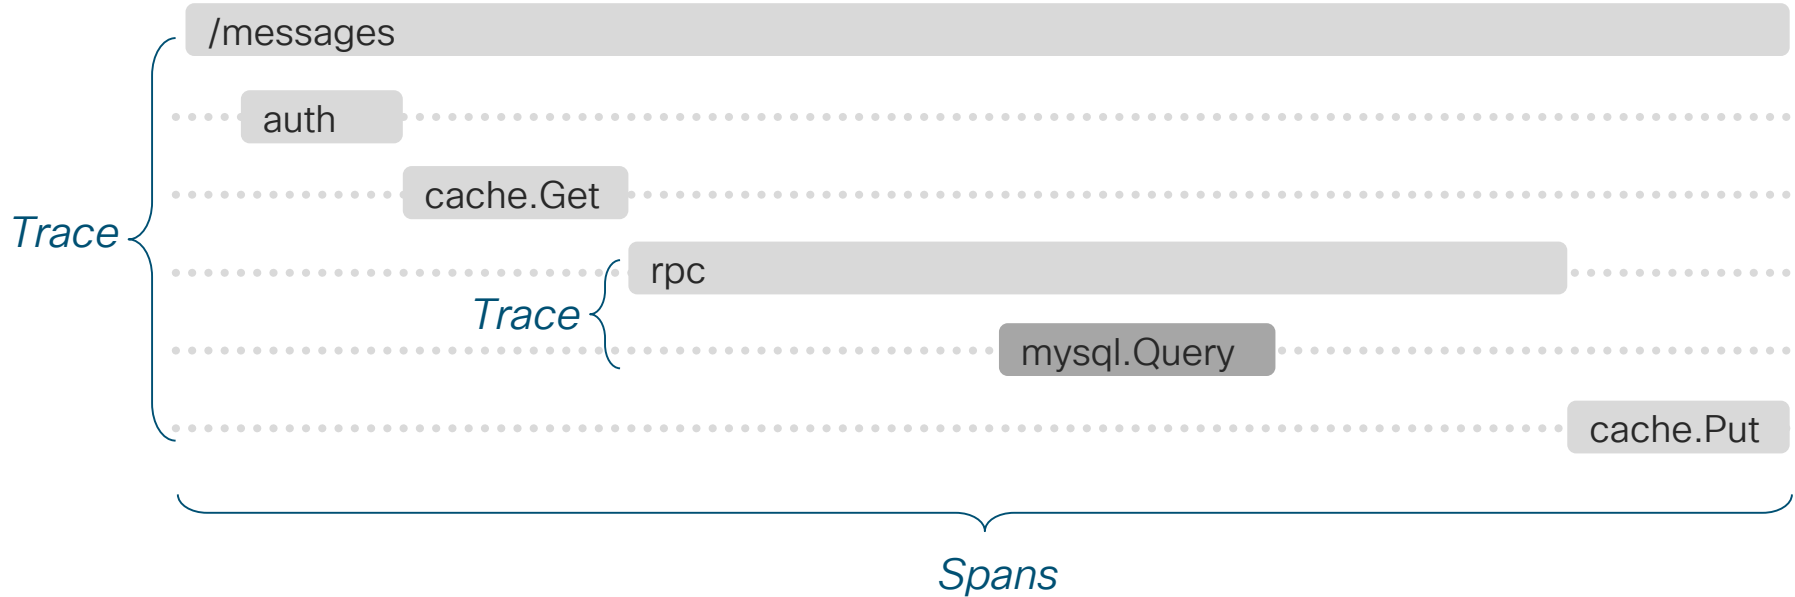
“The mission of the Distributed Tracing WG is to define standards for interoperability between tracing tools.”

Specifications:

- HTTP headers for use in transporting trace data (traceparent, tracestate)
- Rules for constructing and mutating these headers

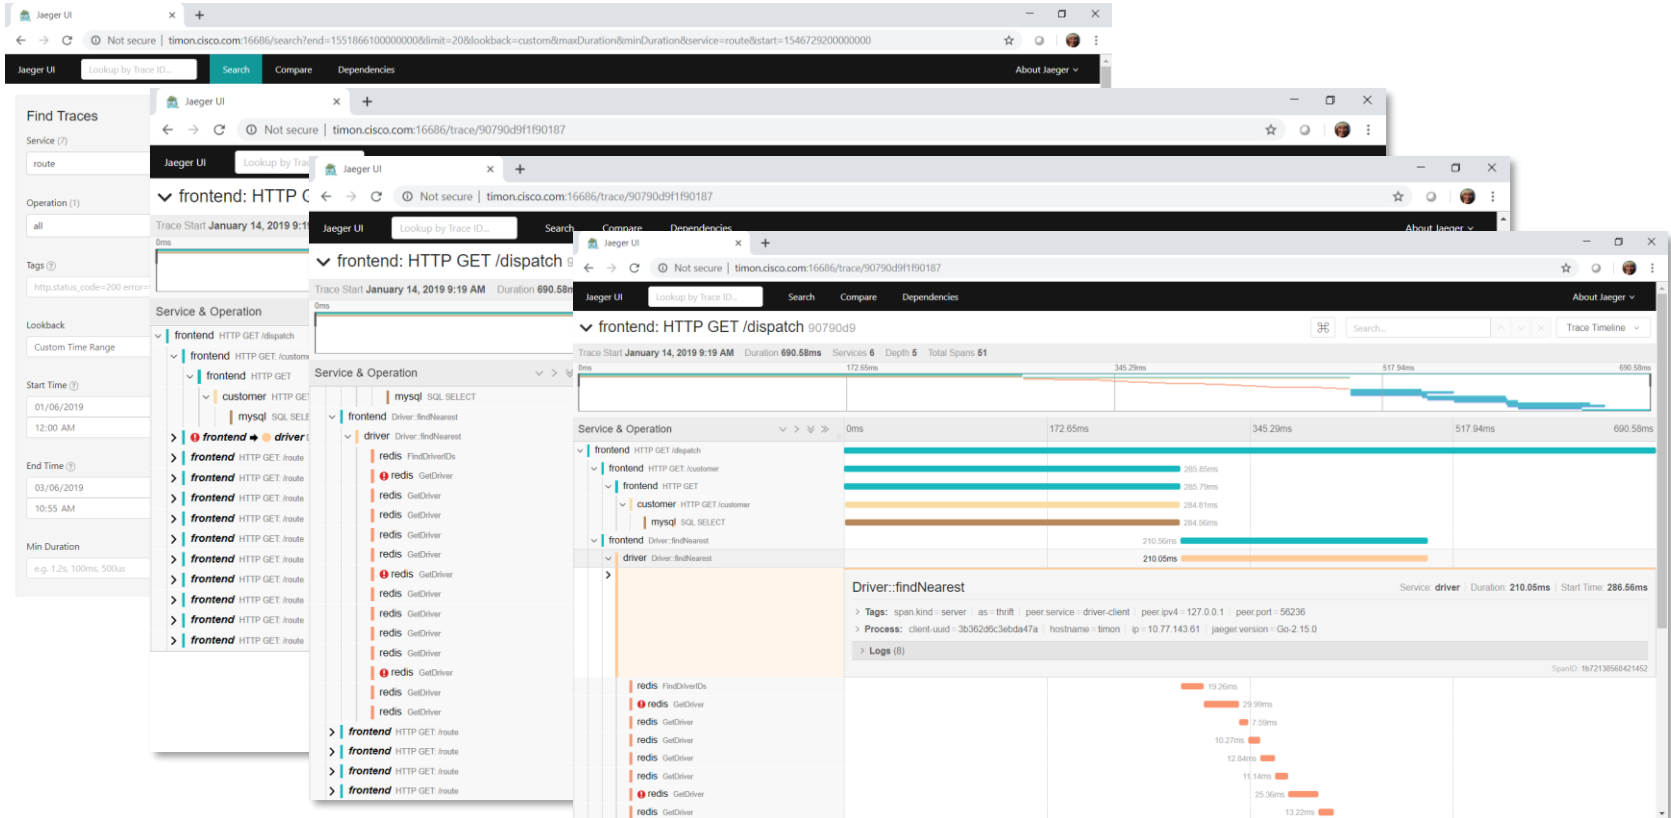


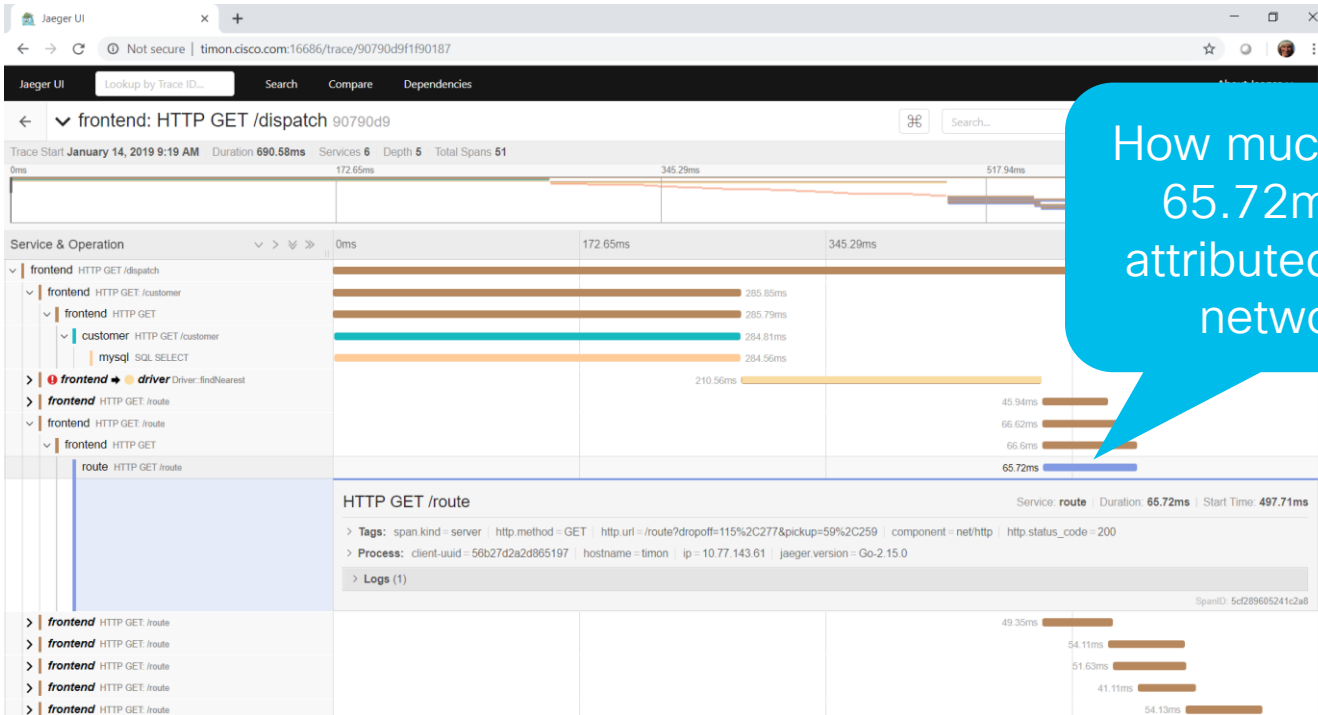
All 3 Groups Share A Similar Tracing Model



Pass around a unique ID (Trace ID): From service entry point to the end of execution; Build a DAG of related operations; Contextualize metadata

Tracing Tool Example: Jaeger

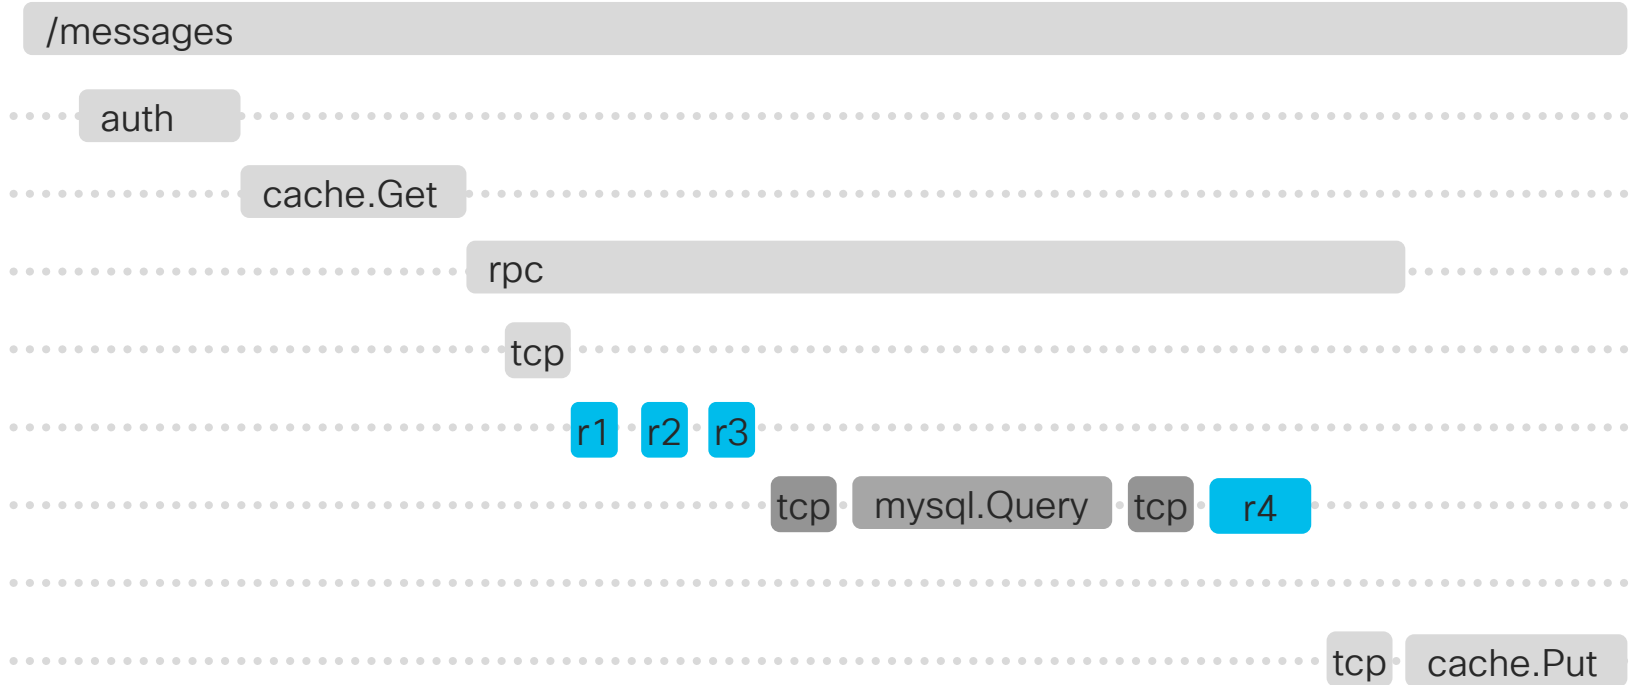




How much of the 65.72ms are attributed to the network?

Network Telemetry  Application Telemetry ?

Include L2/L3/L4 Traces as Child-Traces

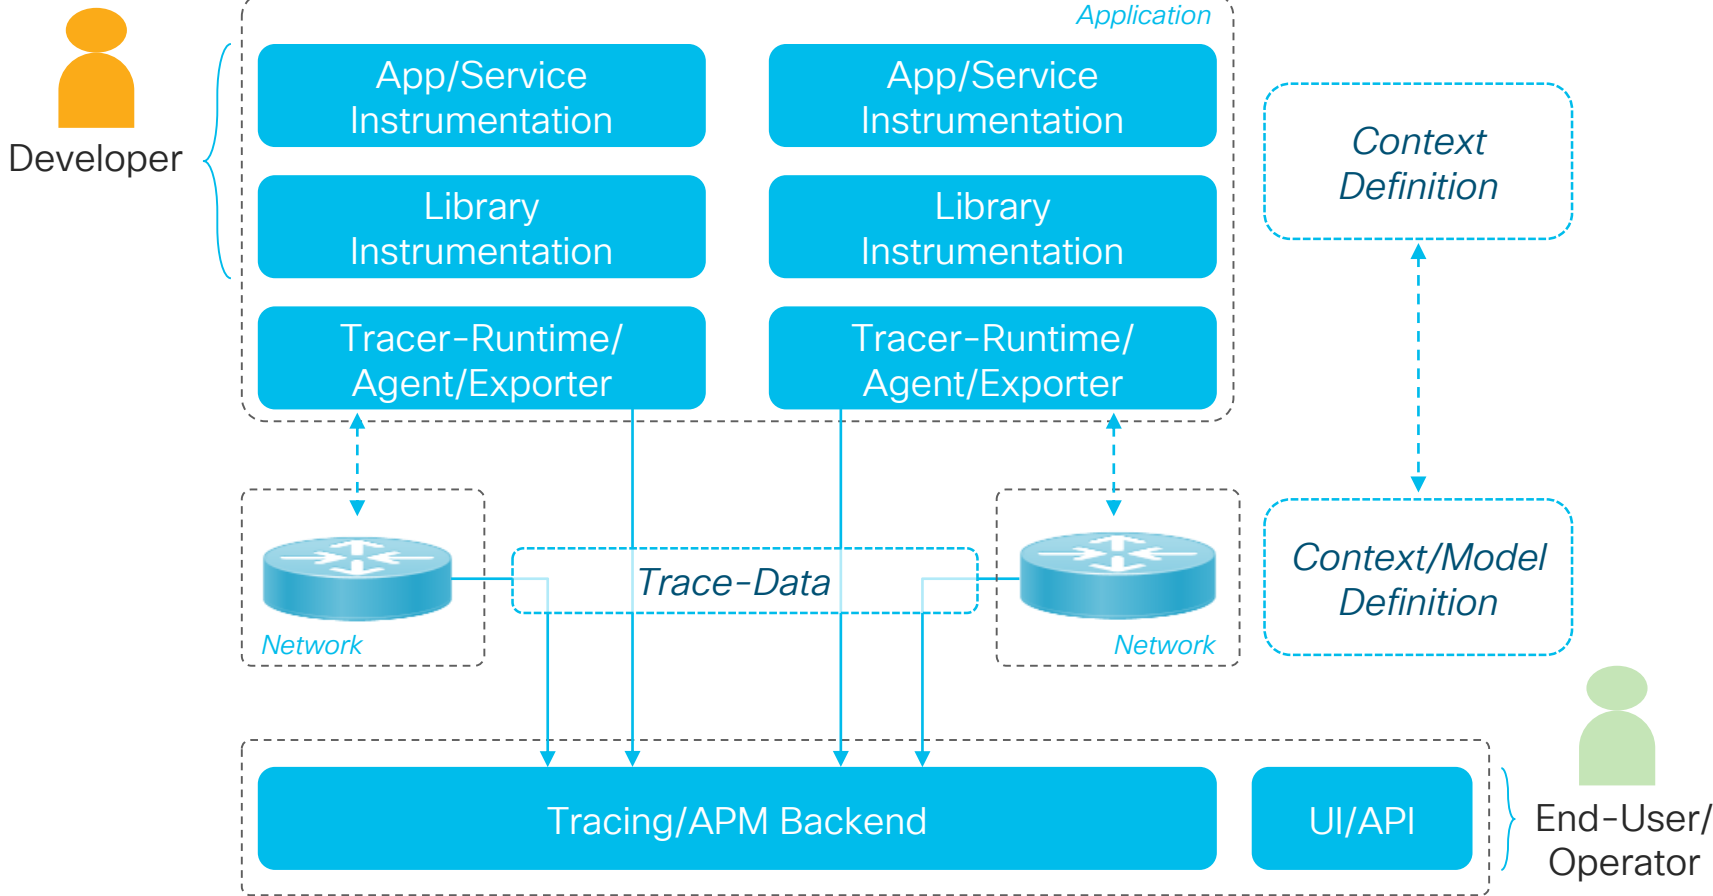


Questions

- Association and propagation of TraceIDs across layers?
 - Application layer tracing operates L7-L5 (OC, OT, ..)
 - Transport layer tracing – L4 ([Robin Marx, TSV-IETF104](#))
 - Network layer tracing operates at L1-3 (IOAM, ..)
- Standard Interchange Format?
 - E.g. OpenTracing [Issue #64](#)
- Metrics correlation?
 - Tags (in OpenCensus) – Key/Value pairs which are used to tag metrics; Tags can be propagated on the wire
- Security / Dealing with E2E encryption

[Network Service Mesh](#) – integrates Application and Network services in K8s, and elevates the above questions...

```
{
  "baggage": [],
  "duration": 11726,
  "logs": [
    {
      "fields": {},
      "message": "hello-world",
      "time": 0
    }
  ],
  "operationName": "some-span",
  "references": [],
  "spanId": 4866369132225493238,
  "start": 1551079498409000,
  "tags": {
    "sampler.type": "const",
    "sampler.param": "true"
  },
  "traceId": 4866369132225493238
}
```



Application Tracing / Transaction Monitoring: Observation:

OSS/Tool-Chains ➡ Adopted Standards

Specifications

Tools

(OSS) Reference Code

Industry Coordination

WHAT?

(API/Model)

✓ (W3C, OT, OC)

✓ (OC, ...)

✓ (OC, OT)

✓ (CNCF)

HOW?

(Associated Operations)

✓ (OC)

✓ (OC, ...)

✓ (OC, OT)

✓ (CNCF)

Another layer up: Business Telemetry


Business

Application

Network

Business-Data
Operational-Data
Configuration-Data

all go hand in hand,
across all layers



Picture source: <https://ayoqq.org/handshake-drawing.html>

Associating a Service KPI to Telemetry

Service Assurance for Intent-based Networking

- A compulsory step for intent-based networking is: closing the loop with telemetry for service assurance
 - Service assurance is difficult
 - Telemetry, with context information
- From streaming data to streaming relevant information
- Tie telemetry to the intent, based on the service KPI.
- Service-specific monitoring, troubleshooting, assurance, and prediction
- Treat the network as a holistic service-delivery system (physical and virtual)

So, What is Telemetry?

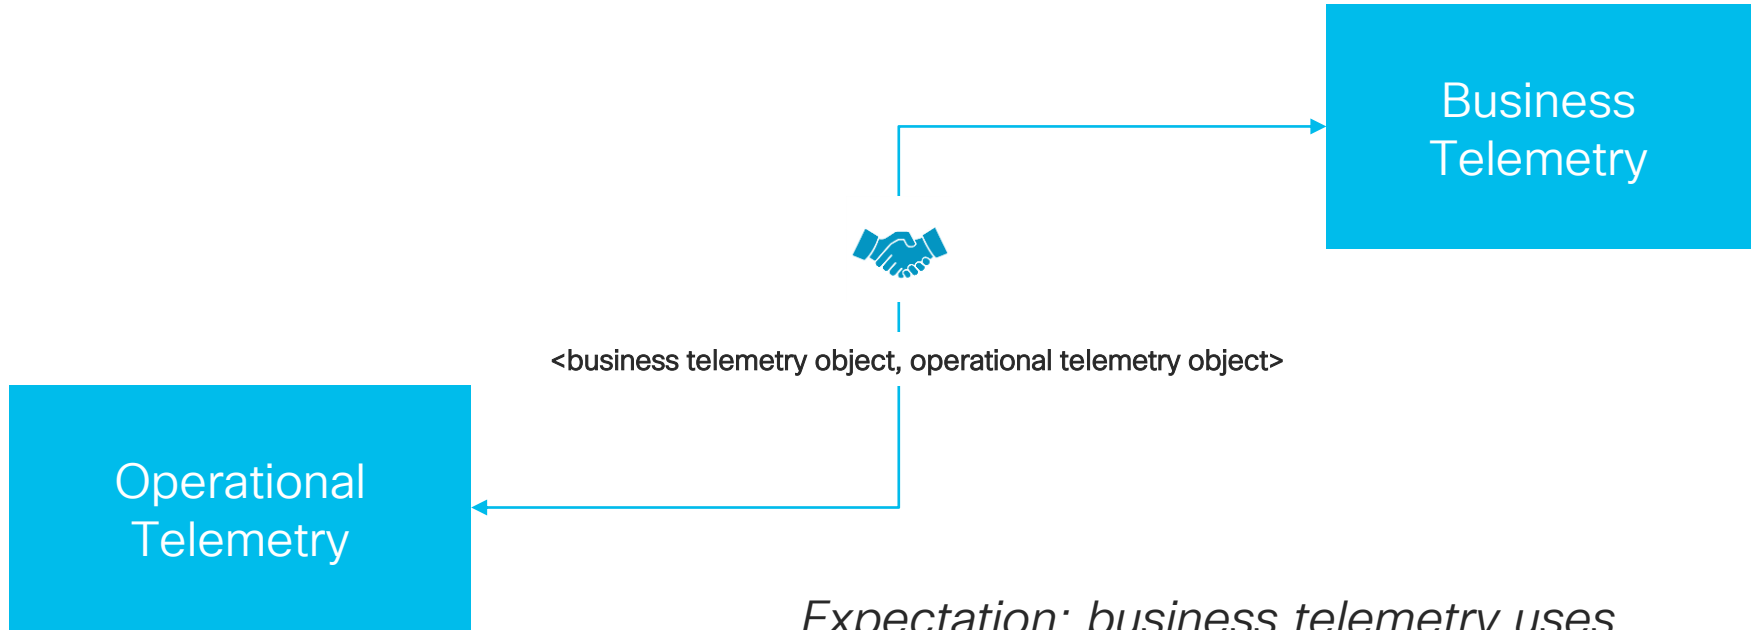
- **(Operational) telemetry:** by default, telemetry deals with operational information. Therefore, it could be flagged as operational telemetry, even if nobody uses that term in the industry. **Network and operations engineers speak of telemetry.**
- **Business telemetry** refers the use of a telemetry to stream information useful for business developments. **Business developers, (senior) vice presidents, and top execs speak of telemetry but they actually mean business telemetry.**

Business Telemetry

Report the business value of an asset

- Asset Identification – a unique product, feature, user, ...
- Associated entities / dependencies
- License – one time purchase, subscription
- State – (de-)activated
- Usage – Usage information for the entity, utilization, performance, ...

Business / Operational Telemetry Mapping



Expectation: business telemetry uses operational telemetry objects

Are we ready for a broader agenda?

What?

+

How?

+

Who?

Adopted Solutions

Specifications

+ Tooling

+ Code

+ Industry coordination

= ADOPTED SOLUTIONS

A specification-only focus has limited impact. You can even call it a mistake...

Thank You