

# Streaming Telemetry: Considerations & Challenges

Mike Korshunov, TME @ Cisco  
[mkorshun@cisco.com](mailto:mkorshun@cisco.com)

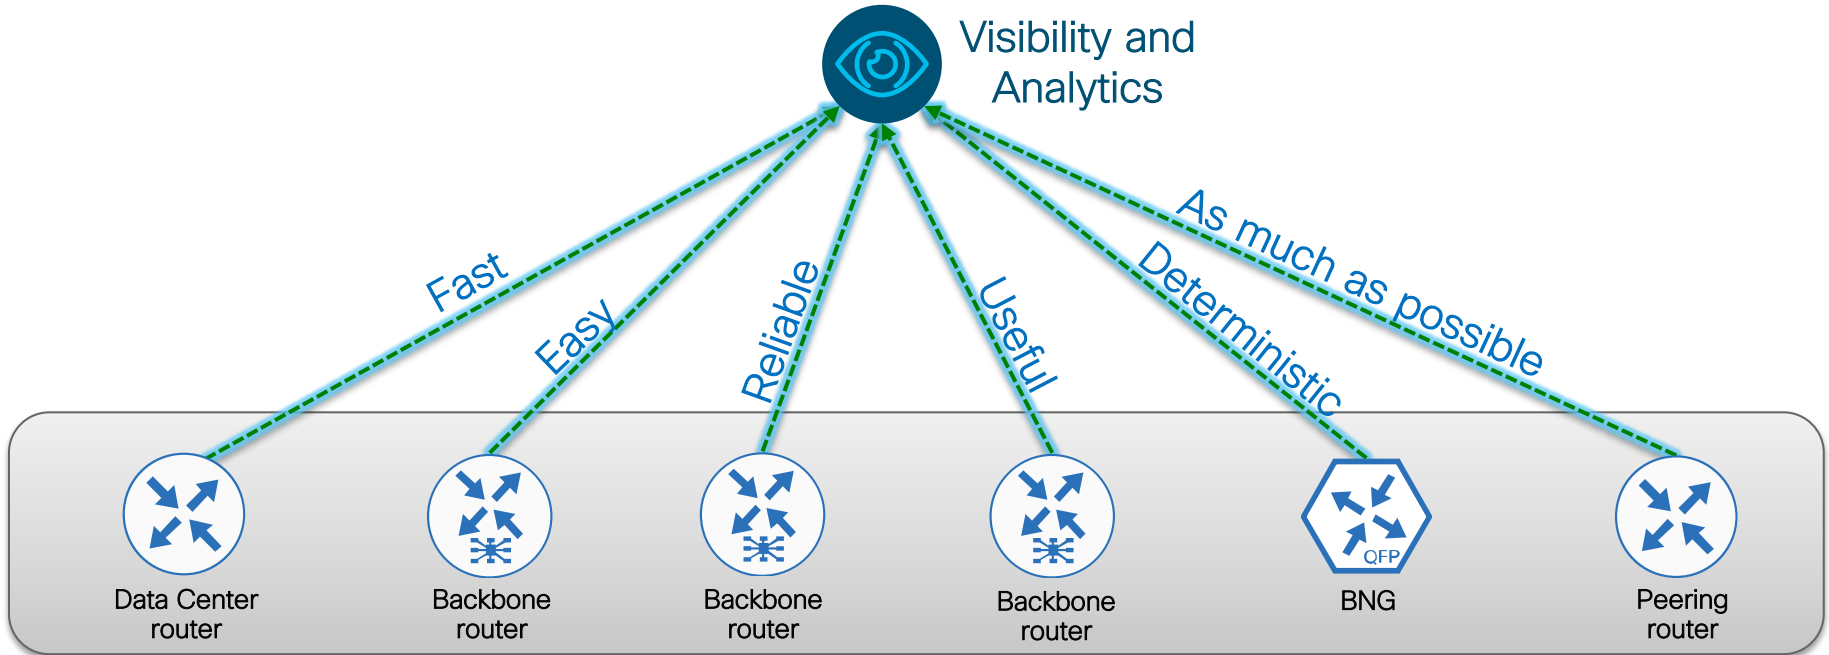
# Agenda

- 1 Brief Telemetry Overview
- 2 Closer Look at Telemetry Components
- 3 The Progress So Far
- 4 Final Thoughts

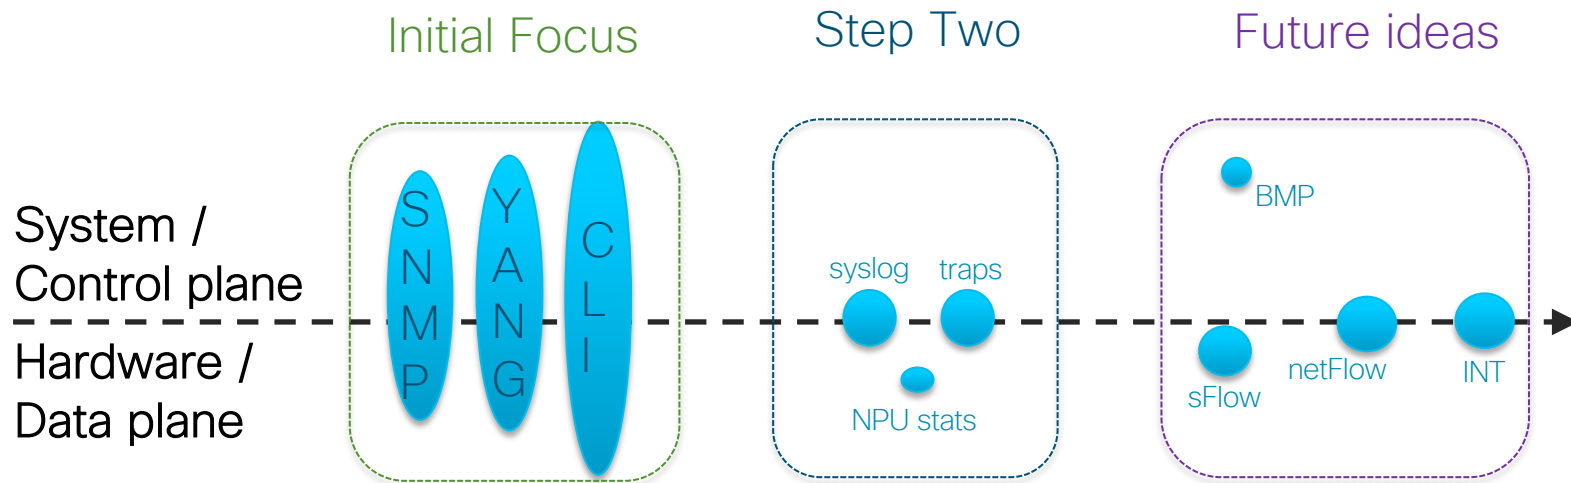
# ~~Scream~~ Stream If You Wanna Go Faster”

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

<https://en.wikipedia.org/wiki/Telemetry>



# Telemetry Evolution





# Agenda

- 1 Brief Telemetry Overview
- 2 Closer Look at Telemetry Components
- 3 The Progress So Far
- 4 Final Thoughts

# Two Approaches for Models



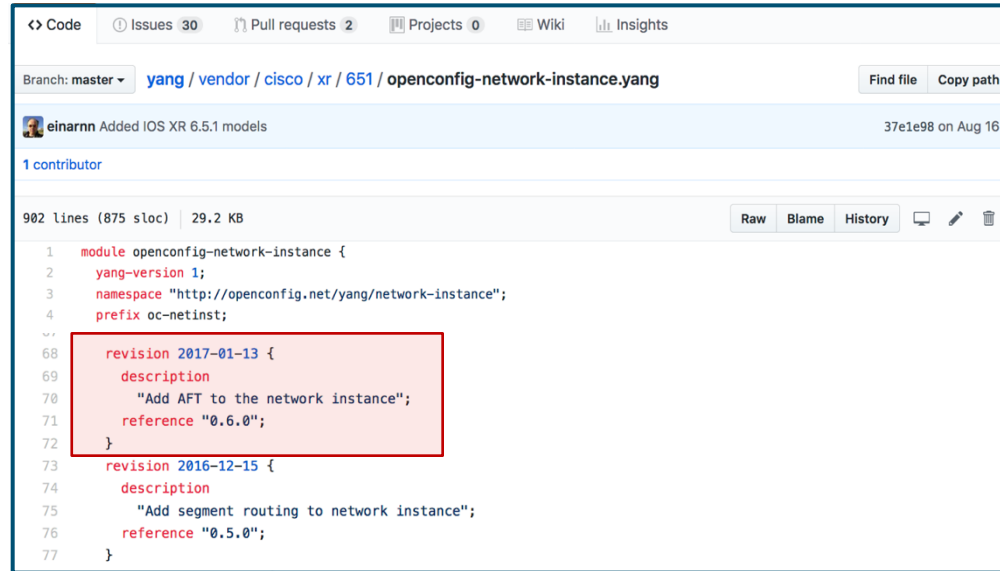
Native (Proprietary)  
Models



OpenConfig  
Models


# Pay Attention to Details

## Cisco OC-NI Model



<> Code ⓘ Issues 30 📄 Pull requests 2 📁 Projects 0 📖 Wiki 📊 Insights

Branch: master ▾ **yang / vendor / cisco / xr / 651 / openconfig-network-instance.yang** Find file Copy path

 einarinn Added IOS XR 6.5.1 models 37e1e98 on Aug 16

1 contributor

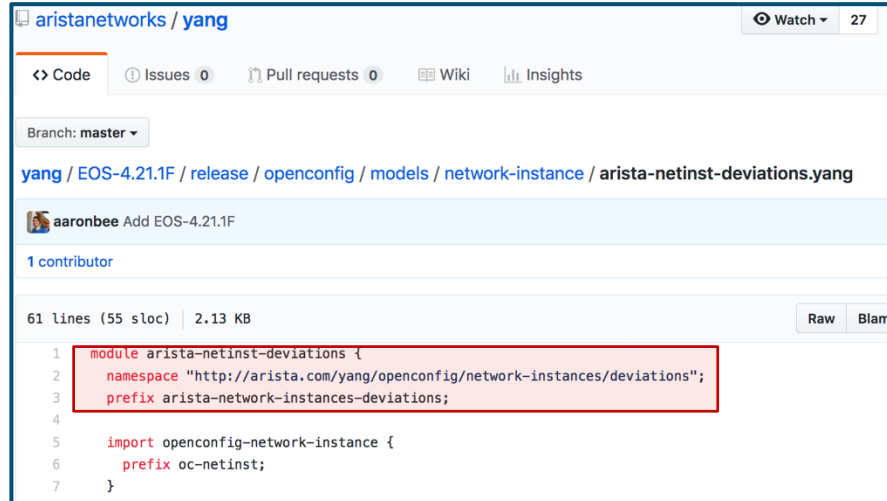
902 Lines (875 sloc) 29.2 KB Raw Blame History

```
1 module openconfig-network-instance {
2   yang-version 1;
3   namespace "http://openconfig.net/yang/network-instance";
4   prefix oc-netinst;
5
6   ...
68   revision 2017-01-13 {
69     description
70       "Add AFT to the network instance";
71     reference "0.6.0";
72   }
73   revision 2016-12-15 {
74     description
75       "Add segment routing to network instance";
76     reference "0.5.0";
77   }
```

<https://github.com/YangModels/yang/blob/master/vendor/cisco/xr/651/openconfig-network-instance.yang>

# Check Deviations For Not Supported Leafs

## Arista OC-NI Model Deviation



aristanetworks / yang

Watch 27

Code Issues 0 Pull requests 0 Wiki Insights

Branch: master

yang / EOS-4.21.1F / release / openconfig / models / network-instance / arista-netinst-deviations.yang

aaronbee Add EOS-4.21.1F

1 contributor

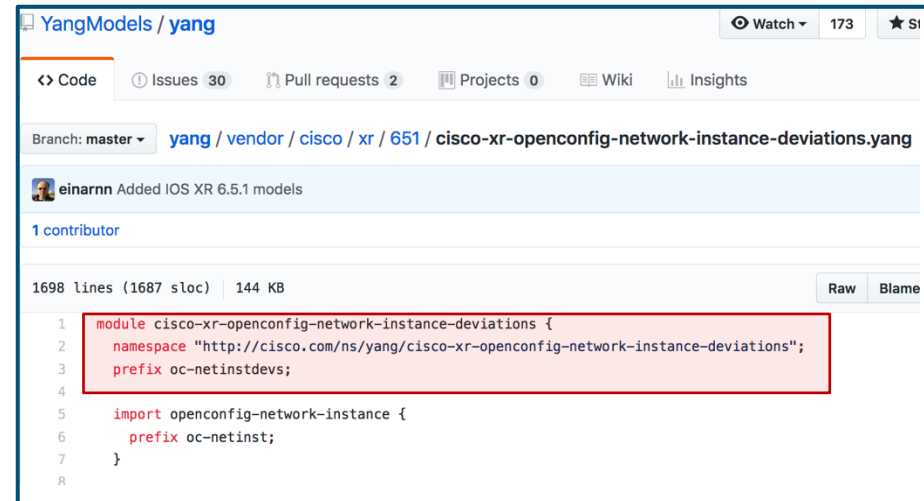
61 lines (55 sloc) | 2.13 KB

Raw Blame

```
1 module arista-netinst-deviations {
2   namespace "http://arista.com/yang/openconfig/network-instances/deviations";
3   prefix arista-network-instances-deviations;
4
5   import openconfig-network-instance {
6     prefix oc-netinst;
7   }
```

<https://github.com/aristanetworks/yang/blob/master/EOS-4.21.1F/release/openconfig/models/network-instance/arista-netinst-deviations.yang>

## Cisco OC-NI Model Deviation



YangModels / yang

Watch 173

Code Issues 30 Pull requests 2 Projects 0 Wiki Insights

Branch: master yang / vendor / cisco / xr / 651 / cisco-xr-openconfig-network-instance-deviations.yang

einarann Added IOS XR 6.5.1 models

1 contributor

1698 lines (1687 sloc) | 144 KB

Raw Blame

```
1 module cisco-xr-openconfig-network-instance-deviations {
2   namespace "http://cisco.com/ns/yang/cisco-xr-openconfig-network-instance-deviations";
3   prefix oc-netinstdevs;
4
5   import openconfig-network-instance {
6     prefix oc-netinst;
7   }
```

<https://github.com/YangModels/yang/blob/master/vendor/cisco/xr/651/cisco-xr-openconfig-network-instance-deviations.yang>

# How to Select the Protocols



source: [www.kisspng.com](http://www.kisspng.com)

**GRPC**



source: [www.novatoys.ru](http://www.novatoys.ru)

**TCP**



source: [www.adventuremotorcycle.com](http://www.adventuremotorcycle.com)

**UDP**

# TCP And UDP Are Simple

No.	Time	Source	Destination	Protocol	Leng	Info
1	0.000000000	10.30.110.41	10.30.110.40	TCP	62	18577 → 5432 [SYN] Seq=0 Win=16384 Len=0 MSS=1460 ...
2	0.0000000863	10.30.110.40	10.30.110.41	TCP	62	5432 → 18577 [SYN, ACK] Seq=0 Ack=1 Win=29200 Len=...
3	0.000696242	10.30.110.41	10.30.110.40	TCP	60	18577 → 5432 [ACK] Seq=1 Ack=1 Win=16384 Len=0
4	0.010749238	10.30.110.41	10.30.110.40	TCP	239	18577 → 5432 [PSH, ACK] Seq=1 Ack=1 Win=16384 Len=...
5	0.010788711	10.30.110.40	10.30.110.41	TCP	54	5432 → 18577 [ACK] Seq=1 Ack=186 Win=30336 Len=0
6	9.072228347	10.30.110.41	10.30.110.40	TCP	60	18577 → 5432 [FIN, ACK] Seq=186 Ack=1 Win=16384 Le...
7	9.072517389	10.30.110.40	10.30.110.41	TCP	54	5432 → 18577 [FIN, ACK] Seq=1 Ack=187 Win=30336 Le...
8	9.073190068	10.30.110.41	10.30.110.40	TCP	60	18577 → 5432 [ACK] Seq=187 Ack=2 Win=16384 Len=0

0030	40 00 b7 e4 00 00	00 01 00 01 00 01 00 00 00 00	@.....
0040	00 ad 0a 0b 4e 43	53 35 30 31 5f 74 6f 70 1a	...NCS5 501_top.
0050	04 74 65 73 74 32 2e 43	69 73 63 6f 2d 49 4f 53	.test2.Cisco IOS
0060	2d 58 52 2d 73 68 65 6c	6c 75 74 69 6c 2d 6f 70	-XR-shell lutil-op
0070	65 72 3a 73 79 73 74 65	6d 2d 74 69 6d 65 2f 75	er:system-time/u
0080	70 74 69 6d 65 3a 0a 32	30 31 35 2d 30 31 2d 30	ptime: 2 015-01-0
0090	37 40 b5 a8 cd 09 48 ff	83 80 cb a9 2c 50 ff 83	7@....H. ....,P..

No.	Time	Source	Destination	Protocol	Leng	Info
1	0.000000000	10.30.110.41	10.30.110.40	UDP	227	26539 → 18765 Len=165

0020	6e 28 67 ab 49 4d 00 c1	1a 66 00 01 00 01 00 01	m(g.1M...f.....
0030	00 00 00 00 00 ad	0a 0b 4e 43 53 35 35 30 31 5f	..... NCS5501
0040	74 6f 70 1a 04 74 65 73	74 32 2e 43 69 73 63 6f	top..tes t2.Cisco
0050	2d 49 4f 53 2d 58 52 2d	73 68 65 6c 6c 75 74 69	-IOS-XR- shelluti
0060	6c 2d 6f 70 65 72 3a 73	79 73 74 65 6d 2d 74 69	l-oper:system-ti
0070	6d 65 2f 75 70 74 69 6d	65 3a 0a 32 30 31 35 2d	me/uptime: .2015-
0080	30 31 2d 30 37 40 f8 ea	cd 09 48 d8 f3 ac cb a9	01-07@...H.....
0090	2c 50 d8 f3 ac cb a9 2c	5a 42 08 df f3 ac cb a9	,P....., ZB.....

Good to know  
if there is an  
additional  
header inside

# gRPC Comes With an Overhead...

No.	Time	Source	Destination	Protocol	Length	Info
1	0.000000000	10.30.110.41	10.30.110.40	TCP	74	64652 → 57500 [SYN] Seq=0 Win=...
2	0.000093450	10.30.110.40	10.30.110.41	TCP	74	57500 → 64652 [SYN, ACK] Seq=0...
3	0.000514962	10.30.110.41	10.30.110.40	TCP	66	64652 → 57500 [ACK] Seq=1 Ack=...
4	0.000603066	10.30.110.41	10.30.110.40	HTTP2	90	Magic
5	0.000628645	10.30.110.40	10.30.110.41	TCP	66	57500 → 64652 [ACK] Seq=1 Ack=...
6	0.000639815	10.30.110.41	10.30.110.40	HTTP2	75	SETTINGS
7	0.000652105	10.30.110.40	10.30.110.41	TCP	66	57500 → 64652 [ACK] Seq=1 Ack=...
8	0.000660746	10.30.110.41	10.30.110.40	HTTP2	79	WINDOW_UPDATE
9	0.000670352	10.30.110.40	10.30.110.41	TCP	66	57500 → 64652 [ACK] Seq=1 Ack=...
10	0.000686445	10.30.110.41	10.30.110.40	HTTP2	152	HEADERS
11	0.000694705	10.30.110.40	10.30.110.41	TCP	66	57500 → 64652 [ACK] Seq=1 Ack=...
12	0.000810032	10.30.110.40	10.30.110.41	HTTP2	75	SETTINGS
13	0.000839324	10.30.110.40	10.30.110.41	HTTP2	79	WINDOW_UPDATE
14	0.001087631	10.30.110.41	10.30.110.40	TCP	66	64652 → 57500 [ACK] Seq=133 Ac...
15	0.001118306	10.30.110.40	10.30.110.41	HTTP2	75	SETTINGS
16	0.001134535	10.30.110.41	10.30.110.40	TCP	66	64652 → 57500 [ACK] Seq=133 Ac...
17	0.001409008	10.30.110.41	10.30.110.40	HTTP2	75	SETTINGS
18	0.011716945	10.30.110.41	10.30.110.40	HTTP2	259	DATA
19	0.011798380	10.30.110.40	10.30.110.41	TCP	66	57500 → 64652 [ACK] Seq=32 Ack=...
20	24.000822908	10.30.110.41	10.30.110.40	TCP	66	64652 → 57500 [FIN, ACK] Seq=3...
21	24.001077163	10.30.110.40	10.30.110.41	TCP	66	57500 → 64652 [FIN, ACK] Seq=3...
22	24.001617408	10.30.110.41	10.30.110.40	TCP	66	64652 → 57500 [ACK] Seq=336 Ac...

0040	62 41 00 00 b8 00 00 00 00 00 01 00 00 00 00 b3	bA..... .....
0050	08 ff 05 12 ad 01 0a 0b 4e 43 53 35 35 30 31 5f	..... NCS5501
0060	74 6f 70 1a 04 74 65 73 74 32 2e 43 69 73 63 6f	top..tes t2.Cisco
0070	2d 49 4f 53 2d 58 52 2d 73 68 65 6c 6c 75 74 69	-IOS-XR- shelluti
0080	6c 2d 6f 70 65 72 3a 73 79 73 74 65 6d 2d 74 69	l-oper:s system-ti
0090	6d 65 2f 75 70 74 69 6d 65 3a 0a 32 30 31 35 2d	me/uptime: .2015-

Magic number to start HTTP2 phase

Settings from the router

Window size from the router

HTTP2 details

Window size/settings from the collector

Settings confirmations

Negotiation happens before the data is streamed



## ...But Brings Some Good Benefits

No.	Time	Source	Destination	Protocol	Length	Info
4...	16.020449012	10.30.110.40	10.30.110.41	HTTP2	79	WINDOW_UPDATE
4...	16.020453129	10.30.110.41	10.30.110.40	HTTP2	16459	DATA
4...	16.020465328	10.30.110.40	10.30.110.41	HTTP2	79	WINDOW_UPDATE
4...	16.020469649	10.30.110.41	10.30.110.40	HTTP2	16459	DATA
4...	16.020476567	10.30.110.41	10.30.110.40	TCP	1514	64969 → 57500 [ACK] Seq=768790 Ack=534 W...
4...	16.020620532	10.30.110.40	10.30.110.41	TCP	66	57500 → 64969 [ACK] Seq=534 Ack=770238 W...
4...	16.021387959	10.30.110.41	10.30.110.40	TCP	66	64945 → 57500 [ACK] Seq=1179628 Ack=612 ...
4...	16.021406127	10.30.110.41	10.30.110.40	TCP	4410	64969 → 57500 [ACK] Seq=770238 Ack=534 W...
4...	16.021410692	10.30.110.41	10.30.110.40	HTTP2	16459	DATA
4...	16.021412266	10.30.110.40	10.30.110.41	TCP	66	57500 → 64969 [ACK] Seq=534 Ack=774582 W...
4...	16.021419367	10.30.110.41	10.30.110.40	HTTP2	10293	DATA
4...	16.021421080	10.30.110.41	10.30.110.40	HTTP2	527	DATA
4...	16.021424808	10.30.110.41	10.30.110.40	HTTP2	237	DATA
4...	16.021430133	10.30.110.41	10.30.110.40	HTTP2	679	DATA
4...	16.021472771	10.30.110.40	10.30.110.41	HTTP2	79	WINDOW_UPDATE
4...	16.021480055	10.30.110.41	10.30.110.40	HTTP2	1048	DATA
4...	16.021509967	10.30.110.41	10.30.110.40	HTTP2	676	DATA
4...	16.021513831	10.30.110.40	10.30.110.41	TCP	66	57500 → 64945 [ACK] Seq=612 Ack=1198248 ...
4...	16.021583417	10.30.110.41	10.30.110.40	HTTP2	670	DATA
4...	16.021604650	10.30.110.40	10.30.110.41	TCP	66	57500 → 64945 [ACK] Seq=612 Ack=1199462 ...
4...	16.021708748	10.30.110.40	10.30.110.41	HTTP2	79	WINDOW_UPDATE
4...	16.021836970	10.30.110.41	10.30.110.40	TCP	1514	64969 → 57500 [ACK] Seq=784809 Ack=547 W...
4...	16.021871368	10.30.110.41	10.30.110.40	HTTP2	1223	DATA

Speed  
Control  
(from the  
collector side)

► Flags: 0x00

0... .. = Reserved: 0x0

.000 0000 0000 0000 0000 0000 0001 = Stream Identifier: 1

0... .. = Reserved: 0x0

.000 0000 0000 0000 0110 1001 1110 0101 = Window Size Increment: 27109



# Many Asked About Security

## gRPC Dial-in (NO-TLS)

### Password exchange

No.	Time	Source	Destination	Protocol	Length	Info
10	0.001450562	10.30.110.41	10.30.110.40	HTTP2	79	WINDOW_UPDATE
11	0.001457001	10.30.110.41	10.30.110.40	TCP	66	43738 → 57500 [ACK] Seq=47 Ack=23 Win=29..
12	0.001478540	10.30.110.41	10.30.110.40	HTTP2	75	SETTINGS
13	0.001485483	10.30.110.40	10.30.110.41	TCP	66	43738 → 57500 [ACK] Seq=47 Ack=32 Win=29..
14	0.001518612	10.30.110.40	10.30.110.41	HTTP2	75	SETTINGS
15	0.001729069	10.30.110.40	10.30.110.41	HTTP2	228	HEADERS, DATA
16	0.00172411	10.30.110.40	10.30.110.41	HTTP2	75	DATA
17	0.007733921	10.30.110.41	10.30.110.40	TCP	66	57500 → 43738 [ACK] Seq=32 Ack=227 Win=4..
18	0.019506700	10.30.110.41	10.30.110.40	HTTP2	281	HEADERS, DATA
19	0.058994693	10.30.110.40	10.30.110.41	TCP	66	43738 → 57500 [ACK] Seq=227 Ack=247 Win=...

```
Header Block Fragment: 838644ae326add9b7079496a41a3a0ba0750e61c66a0c9...
[Header Length: 267]
[Header Count: 9]
  ▶ Header: :method: POST
  ▶ Header: :scheme: http
  ▶ Header: :path: /IOSXRExtensibleManagabilityService.gRPCConfigOper/CreateSubs
  ▶ Header: :authority: 10.30.110.41
  ▶ Header: content-type: application/grpc
  ▶ Header: user-agent: grpc-go/1.0
  ▶ Header: :status:
  ▶ Header: password: cisco
  ▶ Header: username: cisco
  ▶ Padding: "SSSSSS"
  ▶ Stream: DATA, Stream ID: 1, Length 23
  Length: 23
```

### Message content

```
&=LMed.....
.NCS5501_top..test2.Cisco-IOS-XR-shellutil-oper:system-time/uptime:
2015-01-07@...
H.....P.....ZB.....,Z...keysz1..contentz...hostname*.NCS5501_topz
..uptime8.....h.....
```

## gRPC Dial-in (TLS)

### Password exchange

No.	Time	Source	Destination	Protocol	Length	Info
1	0.000000000	10.30.110.40	10.30.110.41	TCP	74	44522 → 57500 [SYN] Seq=0 Win=29200 Len=...
2	0.000934147	10.30.110.41	10.30.110.40	TCP	74	57500 → 44522 [SYN, ACK] Seq=0 Ack=1 Win=...
3	0.000994292	10.30.110.40	10.30.110.41	TCP	66	44522 → 57500 [ACK] Seq=1 Ack=1 Win=2931..
4	0.001179816	10.30.110.40	10.30.110.41	TCP	239	44522 → 57500 [PSH, ACK] Seq=1 Ack=1 Win=...
5	0.001550141	10.30.110.41	10.30.110.40	TCP	66	57500 → 44522 [ACK] Seq=1 Ack=174 Win=45..
6	0.001593306	10.30.110.41	10.30.110.40	TCP	129	57500 → 44522 [PSH, ACK] Seq=1 Ack=174 W..
7	0.001608075	10.30.110.40	10.30.110.41	TCP	66	44522 → 57500 [ACK] Seq=174 Ack=64 Win=2..
8	0.001668298	10.30.110.41	10.30.110.40	TCP	1158	57500 → 44522 [PSH, ACK] Seq=64 Ack=174 ..
9	0.001682563	10.30.110.40	10.30.110.41	TCP	66	44522 → 57500 [ACK] Seq=174 Ack=1156 Win=...
10	0.007576159	10.30.110.41	10.30.110.40	TCP	404	57500 → 44522 [PSH, ACK] Seq=1156 Ack=17..
11	0.007587567	10.30.110.40	10.30.110.41	TCP	66	44522 → 57500 [ACK] Seq=174 Ack=1494 Win=...
12	0.007602303	10.30.110.41	10.30.110.40	TCP	75	57500 → 44522 [PSH, ACK] Seq=1494 Ack=17..
13	0.007614496	10.30.110.40	10.30.110.41	TCP	66	44522 → 57500 [ACK] Seq=174 Ack=1503 Win=...
14	0.034926133	10.30.110.40	10.30.110.41	TCP	192	44522 → 57500 [PSH, ACK] Seq=174 Ack=150..
15	0.035855381	10.30.110.41	10.30.110.40	TCP	72	57500 → 44522 [PSH, ACK] Seq=1503 Ack=30..
16	0.035806502	10.30.110.41	10.30.110.40	TCP	111	57500 → 44522 [PSH, ACK] Seq=1509 Ack=30..
17	0.035892100	10.30.110.41	10.30.110.40	TCP	104	57500 → 44522 [PSH, ACK] Seq=1554 Ack=30..
18	0.035896005	10.30.110.41	10.30.110.40	TCP	108	57500 → 44522 [PSH, ACK] Seq=1592 Ack=30..
19	0.035959309	10.30.110.40	10.30.110.41	TCP	66	44522 → 57500 [ACK] Seq=300 Ack=1634 Win=...

```
▶ Frame 21: 146 bytes on wire (1168 bits), 146 bytes captured (1168 bits) on interface 0
▶ Ethernet II, Src: Vmware af:20:2a (00:50:56:af:20:2a), Dst: Cisco ba:a9:e0 (00:cc:fc:ba:a9:e0)
▶ Internet Protocol Version 4, Src: 10.30.110.40, Dst: 10.30.110.41
▶ Transmission Control Protocol
```

### Message content

```
3.R.e!..UY..=1.....Z.5J.....OM4.....A."v.60...|. @...E...Z...4...L.G...h.7...K.
[...Rz...'k.F...u...v-v-.(.4P).Z+
0TR.../.BPbpT7.Yk8.....F..BA..e.....#.d..i.....M..4...A..3...m.....QM..P....>..k.
0.T'.....aG[...K.../...F.CY'...../.....M...I)...;...y..I...
{...!...I...+...*...y"
[...[...%.....Q...S.J...G..b..z.Pc;.....0.....Wf+a..)f).....]...91.
(..).....R..d.....f.....S...%.....3'..Z...$.Q...%.....]..z..oA
S.....q...z.|U...Y..f..BY...c.....!.....h.q..4
0'.....!S...Y...<...#.a.....W..b.....e...b.....iH..[...[@!gl..d..
yD.....P.X...PC...].1 I..S..c].1..L4.....vx..0.S...W.Oa.....b.....b.....I'b.....f...:6...
d.4...p..ll.....f.....f.....0...a.....W...o...a]WE.....S]...~%.../...*...F.
5.b..S'S'.....L.....R.....R.....R.....a]WE.....S]...~%.../...*...F.
9.h.....u."K..j7K.....A%..h.T.x..|..XF...k..T.....SB.&t.w...AS..0..
.....0...+DJv.9|...9..R..j.Z..k"...WmB...'.xL=Y.P.<.(Za.
...>...='f..
```

# Is It Enough To State gRPC Support?

## Cisco gRPC call proto

```
service gRPCConfigOper {  
  // Configuration related commands  
  rpc GetConfig(ConfigGetArgs) returns(stream ConfigGetReply) {};  
  rpc MergeConfig(ConfigArgs) returns(ConfigReply) {};  
  rpc DeleteConfig(ConfigArgs) returns(ConfigReply) {};  
  rpc ReplaceConfig(ConfigArgs) returns(ConfigReply) {};  
  rpc CliConfig(CliConfigArgs) returns(CliConfigReply) {};  
  rpc CommitReplace(CommitReplaceArgs)  
    returns (CommitReplaceReply) {};  
  // Do we need implicit or explicit commit  
  rpc CommitConfig(CommitArgs) returns(CommitReply) {};  
  rpc ConfigDiscardChanges(DiscardChangesArgs)  
    returns(DiscardChangesReply) {};  
  // Get only returns oper data  
  rpc GetOper(GetOperArgs) returns(stream GetOperReply) {};  
  // Get Telemetry Data  
  rpc CreateSubs(CreateSubsArgs) returns(stream CreateSubsReply)  
{};  
}
```

[https://github.com/cisco/bigmuddy-network-telemetry-  
proto/blob/master/staging/mdt\\_grpc\\_dialin/mdt\\_grpc\\_dialin.proto](https://github.com/cisco/bigmuddy-network-telemetry-<br/>proto/blob/master/staging/mdt_grpc_dialin/mdt_grpc_dialin.proto)

## Juniper gRPC call proto

```
service OpenConfigTelemetry {  
  // Request an inline subscription for data at the specified path.  
  // The device should send telemetry data back on the same  
  // connection as the subscription request.  
  rpc telemetrySubscribe(SubscriptionRequest)  
    returns (stream OpenConfigData) {}  
  // Terminates and removes an existing telemetry subscription  
  rpc cancelTelemetrySubscription(CancelSubscriptionRequest)  
    returns (CancelSubscriptionReply) {}  
  // Get the list of current telemetry subscriptions from the  
  // target. This command returns a list of existing subscriptions  
  // not including those that are established via configuration.  
  rpc getTelemetrySubscriptions(GetSubscriptionsRequest)  
    returns (GetSubscriptionsReply) {}  
  // Get Telemetry Agent Operational States  
  rpc getTelemetryOperationalState(GetOperationalStateRequest)  
    returns (GetOperationalStateReply) {}  
  // Return the set of data encodings supported by the device for telemetry  
  rpc getDataEncodings(DataEncodingRequest)  
    returns (DataEncodingReply) {}  
}
```

[https://github.com/Juniper/itimmon/blob/master/t  
elemetry/telemetry.proto](https://github.com/Juniper/itimmon/blob/master/t<br/>elemetry/telemetry.proto)

# Which Encoding To Use?

Encoding	Description	Wire Efficiency	Other Considerations
GPB-Compact	Everything binary (except values that are strings)	High	Proto file per model. Extra Ops complexity.
GPB-KV	String keys and binary values (except values that are strings)	Medium Low	Single .proto file for decoding header.
JSON	Everything strings: keys and values	Low	Friendly. Human readable, easy for humans and code to parse

# In Numbers?

## GPB

Message length: 330 bytes

```
08 f4 06 12 ca 02 0a 0b 4e 43 53 35 35 30 31 5f ..... NCS5501_
74 6f 70 1a 04 74 65 73 74 32 5c 43 69 73 63 6f top..tes t2/Cisco
2d 49 4f 53 2d 58 52 2d 69 6e 66 72 61 2d 73 74 -IOS-XR- infra-st
61 74 73 64 2d 6f 70 65 72 3a 69 6e 66 72 61 2d atsd-ope r:infra-
73 74 61 74 69 73 74 69 63 73 2f 69 6e 74 65 72 statisti cs/inter
66 61 63 65 73 2f 69 6e 74 65 72 66 61 63 65 2f faces/in terface/
6c 61 74 65 73 74 2f 6f 65 6e 65 72 69 63 2d 63 latest/g eneric-c
6f 75 6e 74 65 72 73 3a 0a 32 30 31 35 2d 31 31 ounters: .2015-11
2d 30 39 40 b8 88 0d 09 48 b8 9a 9b da a9 2c 50 -09@.... H.....P
b8 9a 9b da a9 2c 68 be 9a 9b da a9 2c 62 b0 01 .....h. ....b..
0a ad 01 08 bd 9a 9b da a9 2c 52 14 0a 12 48 75 ..... ,R...Hu
6e 64 72 65 64 47 69 67 45 30 2f 30 2f 31 2f 30 ndredGig E0/0/1/0
5a 80 01 90 03 f1 8f df 17 98 03 95 e5 a8 eb b4 Z.....
01 a0 03 f2 a0 c0 b9 b0 07 a8 03 8f f0 ed e8 fc .....
e1 39 b0 03 a7 94 e5 03 b8 03 04 c0 03 df 8b e6 .....
03 c0 03 05 d0 03 00 d8 03 00 03 00 00 e8 03 00 .....
f0 03 00 f8 03 00 80 04 00 88 04 00 90 04 00 98 .....
04 00 a0 04 00 a8 04 00 b0 04 00 b8 04 00 c0 04 .....
00 c8 04 00 04 00 80 d8 04 00 e0 04 00 e8 04 00 .....
f0 04 00 f8 04 05 80 05 00 88 05 ab fd 9d d6 05 .....
90 05 00 98 05 f2 ad c6 d3 05 a0 05 00 a8 05 00 .....
```

## KV-GPB

Message length: 1142 bytes

```
08 f5 06 12 f6 08 0a 0b 4e 43 53 35 35 30 31 5f ..... NCS5501_
74 6f 70 1a 04 74 65 73 74 32 5c 43 69 73 63 6f top..tes t2/Cisco
2d 49 4f 53 2d 58 52 2d 69 6e 66 72 61 2d 73 74 -IOS-XR- infra-st
61 74 73 64 2d 6f 70 65 72 3a 69 6e 66 72 61 2d atsd-ope r:infra-
73 74 61 74 69 73 74 69 63 73 2f 69 6e 74 65 72 statisti cs/inter
66 61 63 65 73 2f 69 6e 74 65 72 66 61 63 65 2f faces/in terface/
6c 61 74 65 73 74 2f 6f 65 6e 65 72 69 63 2d 63 latest/g eneric-c
6f 75 6e 74 65 72 73 3a 0a 32 30 31 35 2d 31 31 ounters: .2015-11
2d 30 39 40 c6 88 0d 09 48 fe b9 af da a9 2c 50 -09@.... H.....P
fe b9 af da a9 2c 5a cd 07 08 82 ba af da a9 2c .....Z. ....
7a 2c 12 04 6a 65 79 73 7a 24 12 0e 69 6e 74 65 z...keys zS...inte
72 66 61 63 65 2d 6e 61 6d 65 2a 12 48 75 6e 64 rf-face-na mes:Hund
72 65 64 47 69 67 45 30 2f 30 2f 31 2f 30 7a a4 redGigE0 /0/1/0z..
07 12 07 63 6f 6e 74 65 6e 74 7a 17 12 10 70 61 .....conte ntz...pa
63 6b 65 74 73 2d 72 65 63 65 69 76 65 64 a0 e0 ckets-re ceived@.
93 df 17 7a 17 12 0e 62 79 74 65 73 2d 72 65 63 ...z..b ytes-rec
65 69 76 65 64 a0 dc aa c2 eb 64 01 7a 15 12 0c eived@. ....z...
70 61 63 6b 65 74 73 2d 73 65 64 74 40 97 a8 c7 packets- sent@.
e7 b0 07 7a 14 12 0a 62 79 74 65 73 2d 73 65 6e ...z..b ytes-sen
74 40 c3 e2 ee b0 e5 e4 7a 7a 21 12 1a 6d 75 6c t@..... 9z!..mul
74 69 63 61 73 74 6d 70 61 63 65 74 73 2d 72 ticast-p ackets-r
65 63 65 69 76 65 64 40 fd 97 e6 03 7a 1e 12 1a eceived@. ....z...
62 72 6f 61 64 63 61 73 74 2d 70 61 63 6b 65 74 broadcast-p acket
73 2d 72 65 63 65 69 76 65 64 a0 84 7a 1d 12 16 s-receiv ed@z...
6d 75 6c 74 69 63 61 73 74 2d 70 61 63 6b 65 74 multica s-t-packe
73 2d 73 65 6e 74 40 d3 8f e5 03 7a 1a 12 16 62 s-sent@. ....z..b
72 6f 61 64 63 61 73 74 2d 70 61 63 6b 65 74 73 roadcast -packets
2d 73 65 6e 74 40 05 7a 10 12 0c 6f 75 74 70 75 -sent@z ...outpu
74 2d 64 72 6f 70 73 38 00 7a 16 12 12 6f 75 74 t-drops8 .z...out
70 75 74 2d 71 75 65 75 65 2d 64 72 6f 70 73 38 put-queu e-drops8
70 0a 8f 12 0b 69 6e 70 75 74 2d 64 72 6f 70 73 38 ..z...inp ut-drops
38 00 7a 15 12 11 69 6e 70 75 74 2d 71 75 65 75 8.z...in put-que
65 2d 64 72 6f 70 73 38 00 7a 19 12 15 72 75 6e e-drops8 .z...run
74 2d 70 61 63 6b 65 74 73 2d 72 65 63 65 69 76 t-packet s-receiv
65 64 38 00 7a 1a 12 16 67 69 61 6e 74 2d 70 61 ed8.z... giant-pa
63 6b 65 74 73 2d 72 65 63 65 69 76 65 64 38 00 ckets-receiv ed8.
7a 1e 12 1a 74 68 72 6f 74 64 6c 65 64 2d 70 61 z...thro ttled-pa
63 6b 65 74 73 2d 72 65 63 65 69 76 65 64 38 00 ckets-re ceiv ed8.
7a 1b 12 17 70 61 72 69 74 79 2d 70 61 63 6b 65 z...pari ty-packe
74 73 2d 72 65 63 65 69 76 65 64 38 00 7a 25 12 ts-recei ved8.z%.
21 75 6e 6b 6e 6f 77 6e 2d 70 72 6f 74 6f 63 6f lunknow n-protoco
6c 2d 70 61 63 6b 65 74 73 2d 72 65 63 65 69 76 l-packet s-receiv
65 64 38 00 7a 10 12 0c 69 6e 70 75 74 2d 65 72 ed8.z... input-er
72 6f 72 73 38 00 7a 0e 10 0a 63 72 63 2d 65 72 rors8.z. ...crc-er
72 6f 72 73 38 00 7a 12 10 0e 69 6e 70 75 74 2d rors8.z. ...input-
6f 76 65 72 72 75 6e 73 38 00 7a 1b 12 0e 69 72 over-ans 8.z...
61 6d 69 6e 67 2d 65 72 72 6f 72 73 2d 72 65 63 aming-er rors-rec
65 69 76 65 64 38 00 7a 19 12 15 69 6e 70 75 74 eived8.z ...input
2d 69 67 6e 6f 72 65 64 2d 70 61 63 6b 65 74 73 -ignored -packets
38 00 7a 10 12 0c 69 6e 70 75 74 2d 61 62 6f 72 8.z...in put-abor
```

## JSON

Message length: 1325 bytes


```
88 af 00 05 38 00 00 00 00 00 01 00 00 00 05 33 ....8.....3
08 f6 06 12 ad 0a 7b 2d 6e 6f 64 65 5f 69 64 5f .....{' node_id
73 74 72 22 3a 22 4e 43 53 35 30 31 5f 74 6f str":"NC S5501.to
70 22 2c 22 73 75 62 73 63 72 69 70 74 69 6f 6e p","subs cription
5f 69 64 5f 73 74 72 22 3a 22 74 65 73 74 22 2c _id_str":"test",
22 65 6e 63 6f 64 69 6e 6f 5f 70 61 74 68 22 3a "encoding _g_path":
22 43 69 73 63 6f 2d 49 4f 53 2d 58 52 2d 69 6e "Cisco-I OS-XR-in
66 72 61 2d 73 74 61 74 73 64 2d 6f 70 65 72 3a fra-stat sd-oper:
69 6e 66 72 61 2d 73 74 61 74 69 73 74 69 63 73 infra-st atistics
2f 69 6e 74 65 72 66 61 63 65 73 2f 69 6e 74 65 /interfa ces/inte
72 66 61 63 65 2f 6c 61 74 65 73 74 2f 6f 65 6e rf-face/la test/gen
65 72 69 63 2d 63 6f 75 6e 74 65 72 73 22 2c 22 eric-co nters","
63 6f 6c 6c 65 63 74 69 6f 6e 5f 69 64 22 3a 32 collecti on_id":2
30 31 38 36 31 39 32 2c 22 63 6f 6c 6c 65 63 74 0186192,"collect
69 6f 6e 5f 73 74 61 72 74 5f 74 69 6d 65 22 3a ion_star t_time":
31 35 32 33 30 32 34 33 32 35 32 38 33 2c 22 6d 15230243 25283,"m
73 63 5f 74 69 6d 65 73 74 61 69 70 22 3a 31 35 sg_times tamp":15
32 33 30 32 34 33 32 35 32 38 32 2d 64 61 74 23024325 289,"dat
61 5f 6a 73 6f 6e 22 3a 5b 7b 22 74 69 6d 65 73 a_json": {{"times
74 61 6d 70 22 3a 31 35 32 33 30 32 34 33 32 35 tamp":15 23024325
32 38 38 2c 22 6b 65 79 73 22 3a 7b 22 69 6e 74 288,"key s":{"int
65 72 66 61 63 65 2d 6e 61 6d 65 22 3a 22 48 75 erface-n ame":"Hu
6e 64 72 65 64 47 69 67 45 30 2f 31 2f 30 0 ndredGig E0/0/1/0
22 7d 2c 22 63 6f 6e 74 65 6e 74 2c 3a 7b 22 70 "},"cont ent":{"p
61 63 6b 65 74 73 2d 72 65 63 65 69 76 65 64 22 ackets-r eceived":
3a 34 39 37 39 34 31 38 36 2c 22 62 74 73 65 73 :4979418 6,"bytes
2d 72 65 63 65 69 76 65 64 22 3a 34 38 35 34 34 -receive d":48544
34 32 32 38 30 32 2c 22 70 61 63 6b 65 74 73 2d 422802,"packets-
73 65 6e 74 22 3a 32 35 33 37 35 34 38 33 35 33 sent":25 37548353
38 30 2c 22 62 79 74 65 73 2f 73 65 6e 74 22 3a 80,"byte s-sent":
32 35 34 32 38 36 37 31 32 33 39 36 39 32 32 2c 254238671 2396922,
22 6d 75 6c 74 69 63 61 73 74 2d 70 61 63 6b 65 "multica st-packe
74 73 2d 72 65 63 65 69 76 65 64 22 3a 37 39 36 ts-recei ved":796
36 33 34 33 2c 22 62 72 6f 61 64 63 61 73 74 2d 6343,"br oadcast-
70 61 63 6b 65 74 73 2d 72 65 63 65 69 76 65 64 packets- receiv
22 3a 34 2c 22 6d 75 6c 74 69 63 61 73 74 2d 70 "4","mul ticast-p
61 63 6b 65 74 73 2d 73 65 6e 74 22 3a 37 39 36 packets- sent":796
35 32 37 36 2c 22 62 72 6f 61 64 63 61 73 74 2d 5276,"br oadcast-
70 61 63 6b 65 74 73 2d 73 65 6e 74 22 3a 35 2c packets- sent":5,
22 6f 75 74 70 75 74 2d 64 72 6f 70 73 22 3a 30 "output- drops":0
2c 22 6f 75 74 70 75 74 2d 71 75 65 75 65 2d 64 "output -queue-d
72 6f 70 73 22 3a 30 2c 2d 69 6e 70 75 74 2d 64 rops":0, "input-d
72 6f 70 73 22 3a 30 2c 2d 69 6e 70 75 74 2d 71 rop":0, "input-q
75 65 75 65 2d 64 72 6f 70 73 22 3a 30 2c 22 72 ueue-dro ps":0,"r
75 6e 74 2d 70 61 63 6b 65 74 73 2d 72 65 63 unt-pack ets-rece
69 76 65 64 22 3a 30 2c 2d 67 69 61 6e 74 2d 70 ived":0, "giant-p
61 63 6b 65 74 73 2d 72 65 63 65 69 76 65 64 22 ackets-r eceived"
3a 30 2c 22 74 68 72 6f 74 74 6c 65 64 2d 70 61 :0,"thro ttled-pa
63 6b 65 74 73 2d 72 65 63 65 69 76 65 64 22 3a ckets-re ceived":
```

# Agenda

- 1 Brief Telemetry Overview
- 2 Closer Look at Telemetry Components
- 3 The Progress So Far
- 4 Final Thoughts



# Design Your Transport Network Properly

- Peak bandwidth consumption



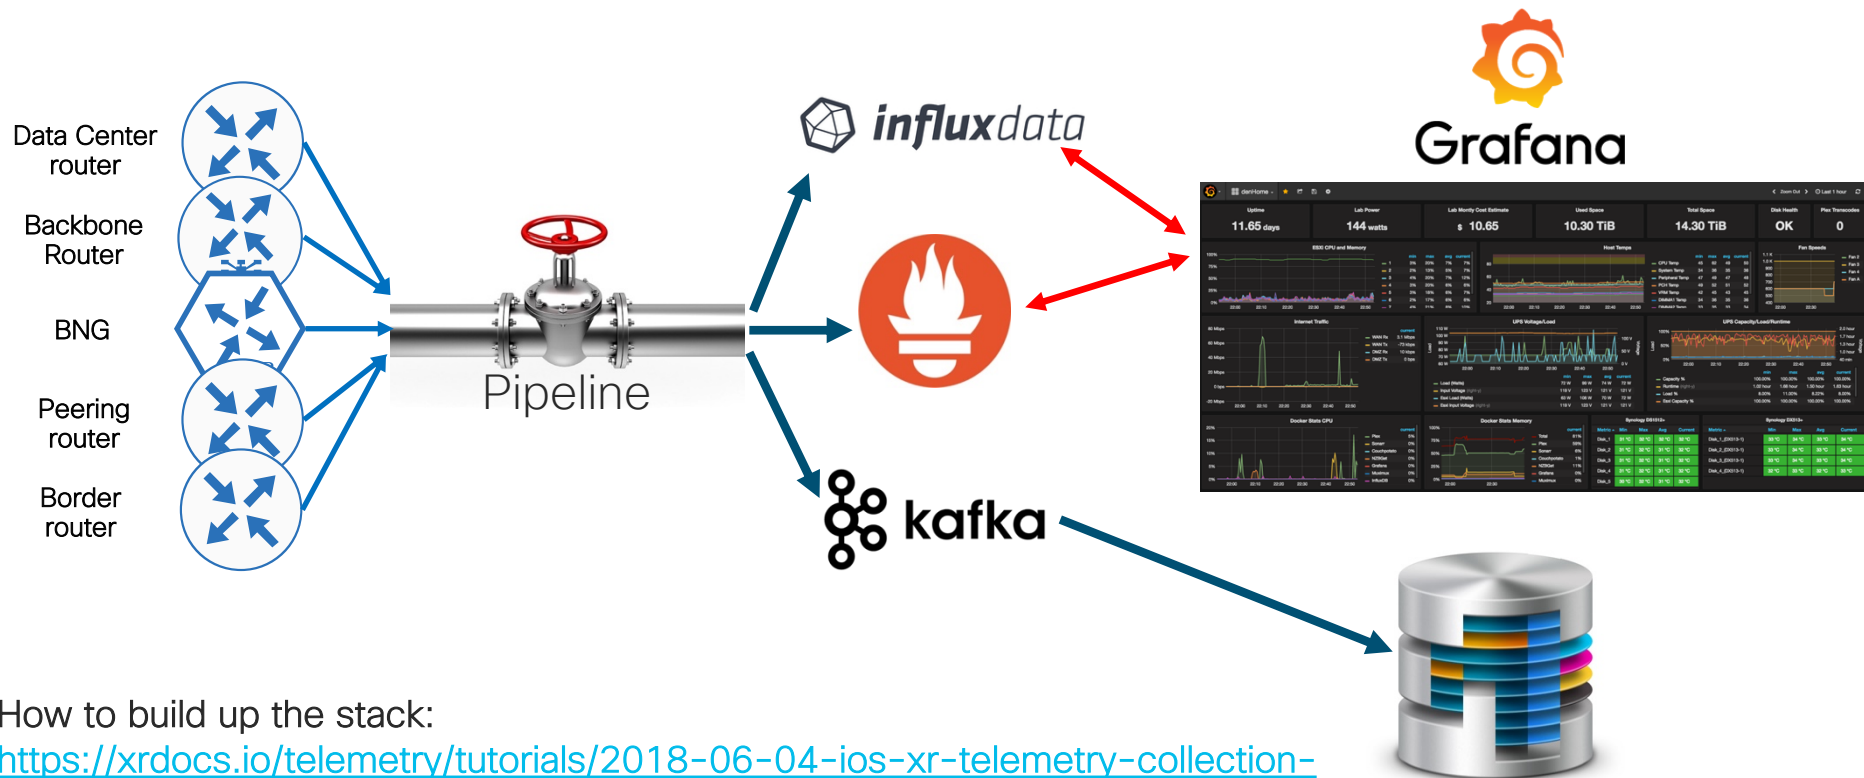
315k counters  
Every 5 seconds

# How Will Telemetry Fill Your Links?



# Start Exploring Telemetry Today

## Go With Open Source Tools




How to build up the stack:

<https://xrdocs.io/telemetry/tutorials/2018-06-04-ios-xr-telemetry-collection-stack-intro>



# Is Your Collector Fast Enough?

Decoded messages volume




```
RP/0/RP0/CPU0:NCSS501_bottom#sh telemetry model destination internal
Wed Apr 11 02:07:19.543 PDT
```

dest Id	Sub	IP	Port	Trans	State	Chan	Buff	Full
des1	subs21	10.30.110.55	57500	grpc	1	398	16	0
des1	subs14	10.30.110.55	57500	grpc	1	401	0	0
des1	subs40	10.30.110.55	57500	grpc	1	431	3	0
des1	subs39	10.30.110.55	57500	grpc	1	440	20	0
des1	subs53	10.30.110.55	57500	grpc	1	435	0	0
des1	subs29	10.30.110.55	57500	grpc	1	436	1	0
des1	subs60	10.30.110.55	57500	grpc	1	386	564	0
des1	subs12	10.30.110.55	57500	grpc	1	409	0	0
des1	subs37	10.30.110.55	57500	grpc	1	415	0	0
des1	subs35	10.30.110.55	57500	grpc	1	421	1	0
des1	subs50	10.30.110.55	57500	grpc	1	433	0	0
des1	subs43	10.30.110.55	57500	grpc	1	426	0	0
des1	subs46	10.30.110.55	57500	grpc	1	429	5	0
des1	subs7	10.30.110.55	57500	grpc	1	390	0	0
des1	subs20	10.30.110.55	57500	grpc	1	407	364	0
des1	subs2	10.30.110.55	57500	grpc	1	399	0	0
des1	subs15	10.30.110.55	57500	grpc	1	403	0	0
des1	subs22	10.30.110.55	57500	grpc	1	410	1733	0
des1	subs34	10.30.110.55	57500	grpc	1	441	21	0
des1	subs44	10.30.110.55	57500	grpc	1	419	0	0
des1	subs33	10.30.110.55	57500	grpc	1	420	0	0
des1	subs32	10.30.110.55	57500	grpc	1	439	21	0
des1	subs5	10.30.110.55	57500	grpc	1	389	0	0
des1	subs17	10.30.110.55	57500	grpc	1	406	40	0
des1	subs0	10.30.110.55	57500	grpc	1	391	0	0
des1	subs11	10.30.110.55	57500	grpc	1	394	0	0
des1	subs47	10.30.110.55	57500	grpc	1	424	0	0
des1	subs27	10.30.110.55	57500	grpc	1	428	0	0
des1	subs6	10.30.110.55	57500	grpc	1	404	185	0
des1	subs30	10.30.110.55	57500	grpc	1	408	0	0
des1	subs48	10.30.110.55	57500	grpc	1	423	0	0
Total:							2974	0


Make sure the collector has enough power to process your telemetry data.

# Is Your Hard Drive Write Speed Fast Enough?

HDD-based server (SAS)



SSD-based server (SAS)



More about hard drives, DRAM and CPU for MDT here:

<https://xrdocs.io/telemetry/tutorials/2018-07-10-is-your-infra-ready-for-telemetry/>


# Don't Forget To Set The Correct Time!

```
RP/0/RP0/CPU0:ios-xr# sh clock
Sun Apr 1 20:56:15.074 PDT
20:56:15.167 PDT Sun Apr 1 2019
```

```
cisco@ubuntu51-1:~$ date
Sun Apr 1 23:13:11 PDT 2019
```

```
RP/0/RP0/CPU0:ios-xr#sh tele m subscription if-stats
Sun Apr 1 20:50:17.883 PDT
Subscription: if-stats
-----
State: ACTIVE
DSCP/Qos marked value: Default
Sensor groups:
Id: if-stats
Sample Interval: 5000 ms
Sensor Path: Cisco-IOS-XR-infra-statsd-oper:infra-
statistics/interfaces/interface[interface-name='Bundle-Ether*']/latest/generic-counters
Sensor Path State: Resolved

Destination Groups:
Group Id: DGroup1
Destination IP: 10.30.110.40
Destination Port: 57500
Encoding: self-describing-gpb
Transport: grpc
State: Active
```



# What To Think About Selecting a TSDB

Metric	InfluxDB	Prometheus	Elasticsearch	OpenTSDB
Website	<a href="https://influxdata.com/">https://influxdata.com/</a>	<a href="https://prometheus.io/">https://prometheus.io/</a>	<a href="https://www.elastic.co/products/elasticsearch">https://www.elastic.co/products/elasticsearch</a>	<a href="http://opentsdb.net">http://opentsdb.net</a>
Category	Real-time Analytics	Monitoring System	Real-time Search	Real-time Analytics
Supported Measurements	metrics, events	metrics	metrics, events	metrics
High Availability (HA)	Double writing 2 servers	Double writing 2 servers	Clustering	Clustering
Underlying Technology	Golang	Golang	Java	Java, Hadoop
Storage Backend	Custom	Custom	Document	Hadoop (Columnar)
Supported Data Types	int64, float64, bool, and string	float64	string, int32, int64, float32, float64, bool, null	int64, float32, float64
Bytes per point after compression	2.2	1.3	22	12
Metric Precision	nanosecond	millisecond	millisecond	millisecond
Write Performance - Single Node	470k metrics / sec (custom HW)	800k metrics / sec	30k metrics / sec	32k metrics /sec (calculated)
Query Performance (1 host, 12hr by 1m)	3.78 ms (min), 8.17 (avg)	tbd	13.23 ms (min), 28.6 (avg)	tbd
Query Language	InfluxQL (SQL like)	PromQL	Query DSL	lookup only
Community Size	large	large	large	medium
Maturity	Stable	Stable	Stable	stable

Full table: <https://tinyurl.com/jsd4esy>

Good to read: <https://tinyurl.com/ybaw4ww6>

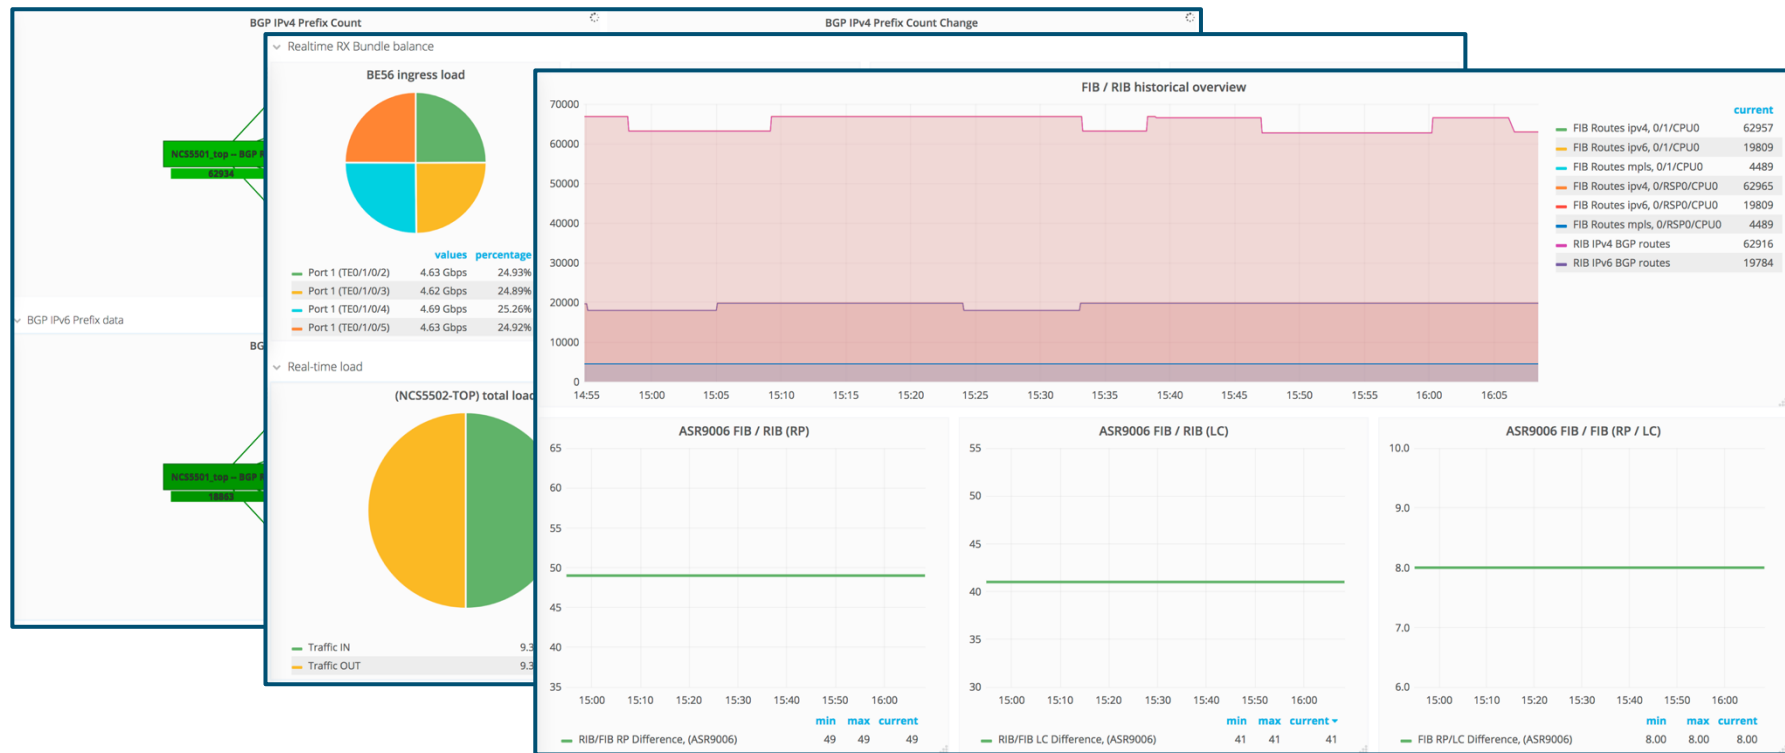
InfluxDB vs OpenTSDB: <https://tinyurl.com/y8ofbiyv>

InfluxDB vs Cassandra: <https://tinyurl.com/y83vv9ys>

DB ranking: <https://tinyurl.com/ya8rrrip>

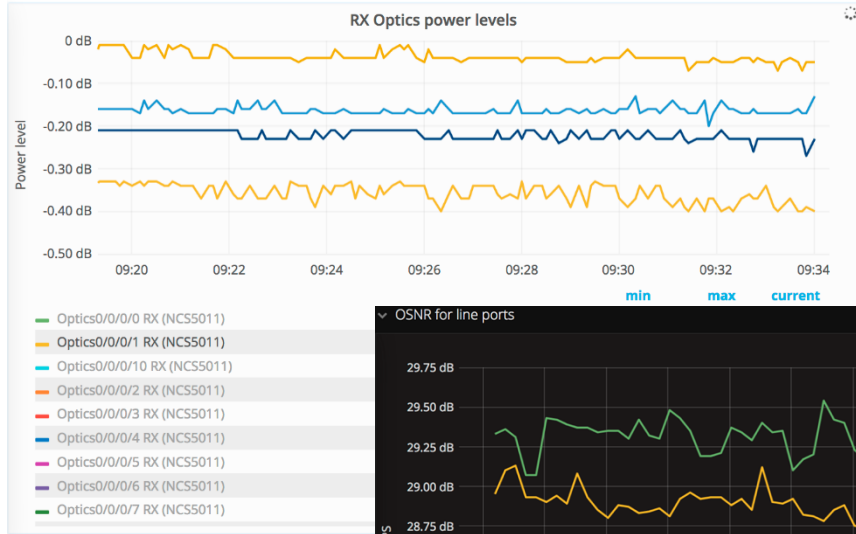
InfluxDB vs Elasticsearch: <https://tinyurl.com/y7yxjf6v>

# You Can See a Lot. In Real Time




RIB/FIB inconsistency check

# Telemetry For Optical Transceivers and Platforms

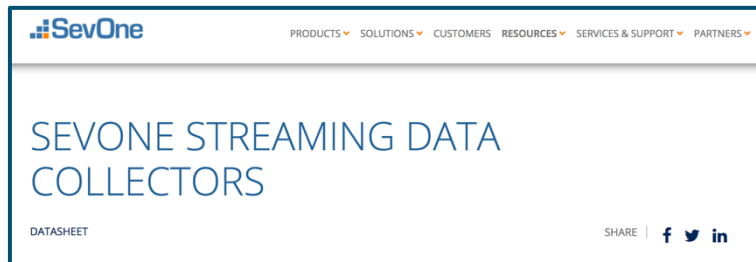


*sensor-path Cisco-IOS-XR-controller-optics-oper:optics-oper/optics-ports/optics-port/optics-info*



<https://xrdocs.io/telemetry/tutorials/2017-10-25-ncs1002-telemetry-deep-dive>

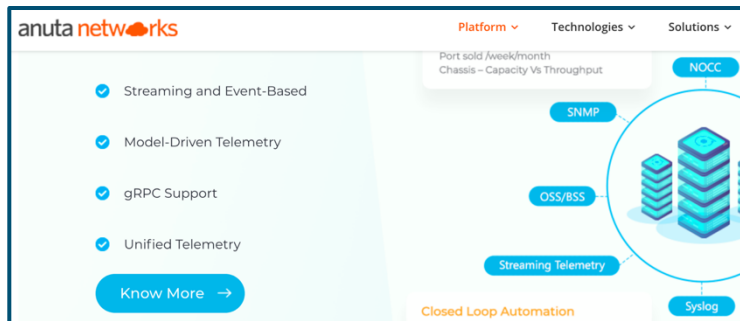
# Different Companies Are Starting To Be Involved



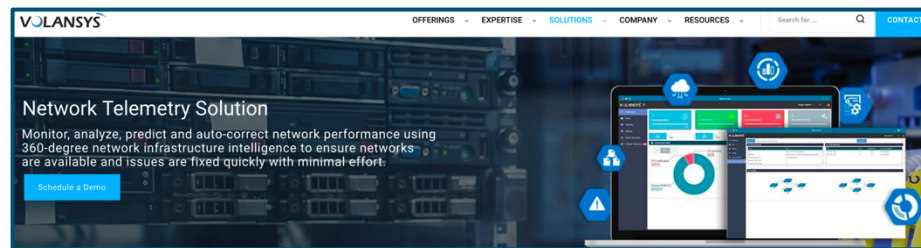
The screenshot shows the SevOne website header with navigation links: PRODUCTS, SOLUTIONS, CUSTOMERS, RESOURCES, SERVICES & SUPPORT, and PARTNERS. The main content area features the title 'SEVONE STREAMING DATA COLLECTORS' and a 'DATASHEET' link. Social media sharing icons for Facebook, Twitter, and LinkedIn are also present.



The screenshot shows the Packet Design website header with navigation links: SOLUTIONS, PRODUCTS, ABOUT, and RESOURCES. The main content area features the title 'Network Basics by Packet Design: What is Streaming Telemetry?'.



The screenshot shows the Anuta Networks website header with navigation links: Platform, Technologies, and Solutions. The main content area features a list of features: Streaming and Event-Based, Model-Driven Telemetry, gRPC Support, and Unified Telemetry. A diagram illustrates the network telemetry components: Port, SNMP, OSS/BSS, Streaming Telemetry, Syslog, and NOCC. The text 'Port sold /week/month Chassis - Capacity Vs Throughput' is also visible. A 'Know More' button is present at the bottom.



The screenshot shows the Volansys website header with navigation links: OFFERINGS, EXPERTISE, SOLUTIONS, COMPANY, RESOURCES, and CONTACT. The main content area features the title 'Network Telemetry Solution' and a description: 'Monitor, analyze, predict and auto-correct network performance using 360-degree network infrastructure intelligence to ensure networks are available and issues are fixed quickly with minimal effort.' A 'Schedule a Demo' button is present.

# gNMI part of the Solution

- Network management interface defined by OpenConfig (mostly lead by Google)
- Configuration management and streaming telemetry in a single protocol
- Data model independent
- Based on Google RPC framework and HTTP/2

The main goal for Telemetry is to provide a “standard” approach for encoding and transport protocols support across different vendors.



# gNMI Implementation in Cisco IOS XR Telemetry

- MDT is based on gNMI v0.4.0
- Introduced in release IOS XR 6.5.1
- The only IOS XR Configuration needed:

```
grpc  
port <10000 .. 57999>
```

- TLS is enabled by default. To disable TLS:

```
grpc  
no-tls
```

# Agenda

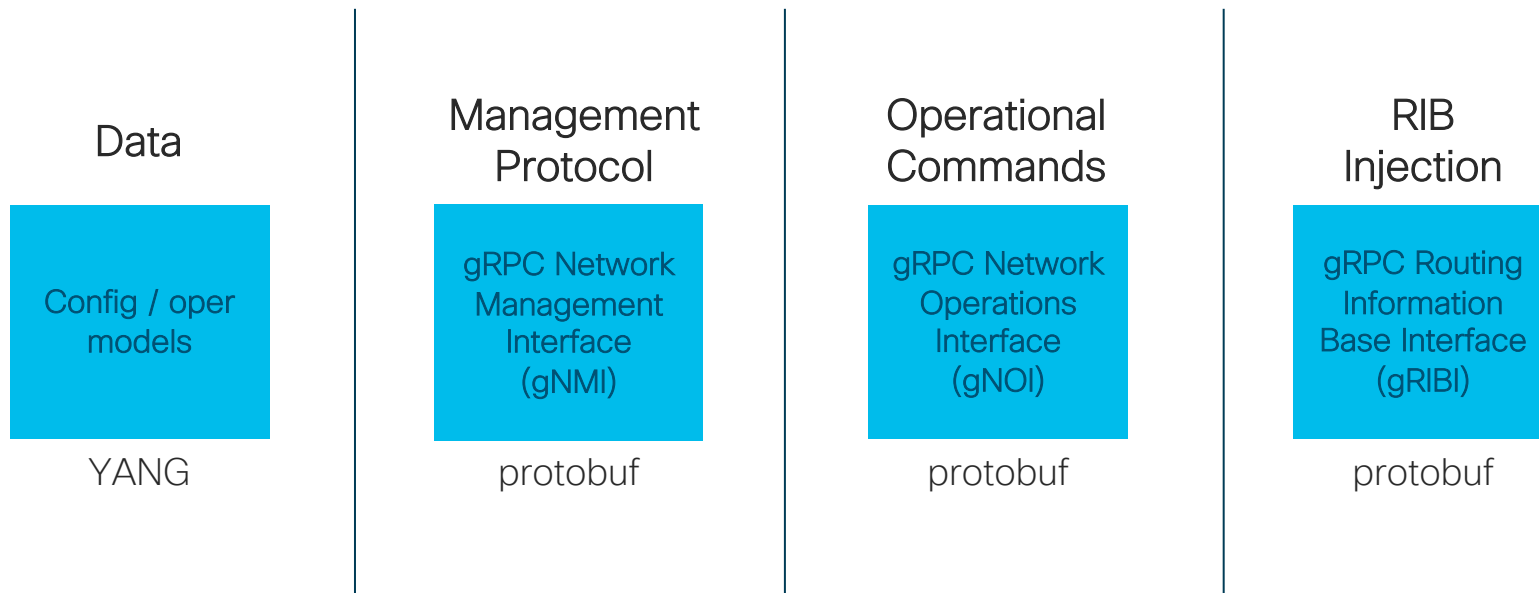
- 1 Brief Telemetry Overview
- 2 Closer Look at Router
- 3 The Progress So far
- 4 Final Thoughts

# Key Messages

- Streaming Telemetry is here for you. Start to receive benefits from it today!
- Select encoding and transport wisely. Good start for beginners with Key-Value GPB & gRPC.
- It's easy to explore. Scripts to bring up stack [available](#). Should take less than 15 mins to provision.
- Read for more materials: <https://xrdocs.io/telemetry/>



# Have a Full Picture in Your Head



Telemetry is just a piece of a puzzle

# Thank you!

Questions / Comments?

Drop me a note: [mkorshun@cisco.com](mailto:mkorshun@cisco.com)