Chicago Network Operators Group

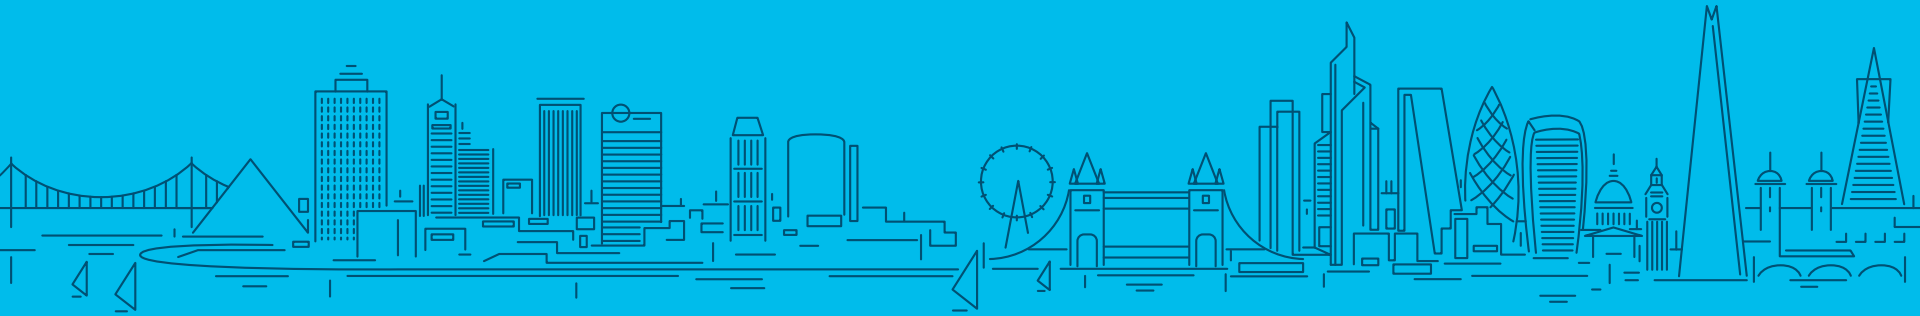

Streaming Telemetry Under the Hood:
Something to Think About

Viktor Osipchuk (vosipchu@cisco.com)

05-10-2018

Do They Select Cars By Color?

Agenda

1. Brief Telemetry Overview
2. Closer Look at Router
3. Closer Look at the Link Between
4. Closer Look at Collector
5. 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

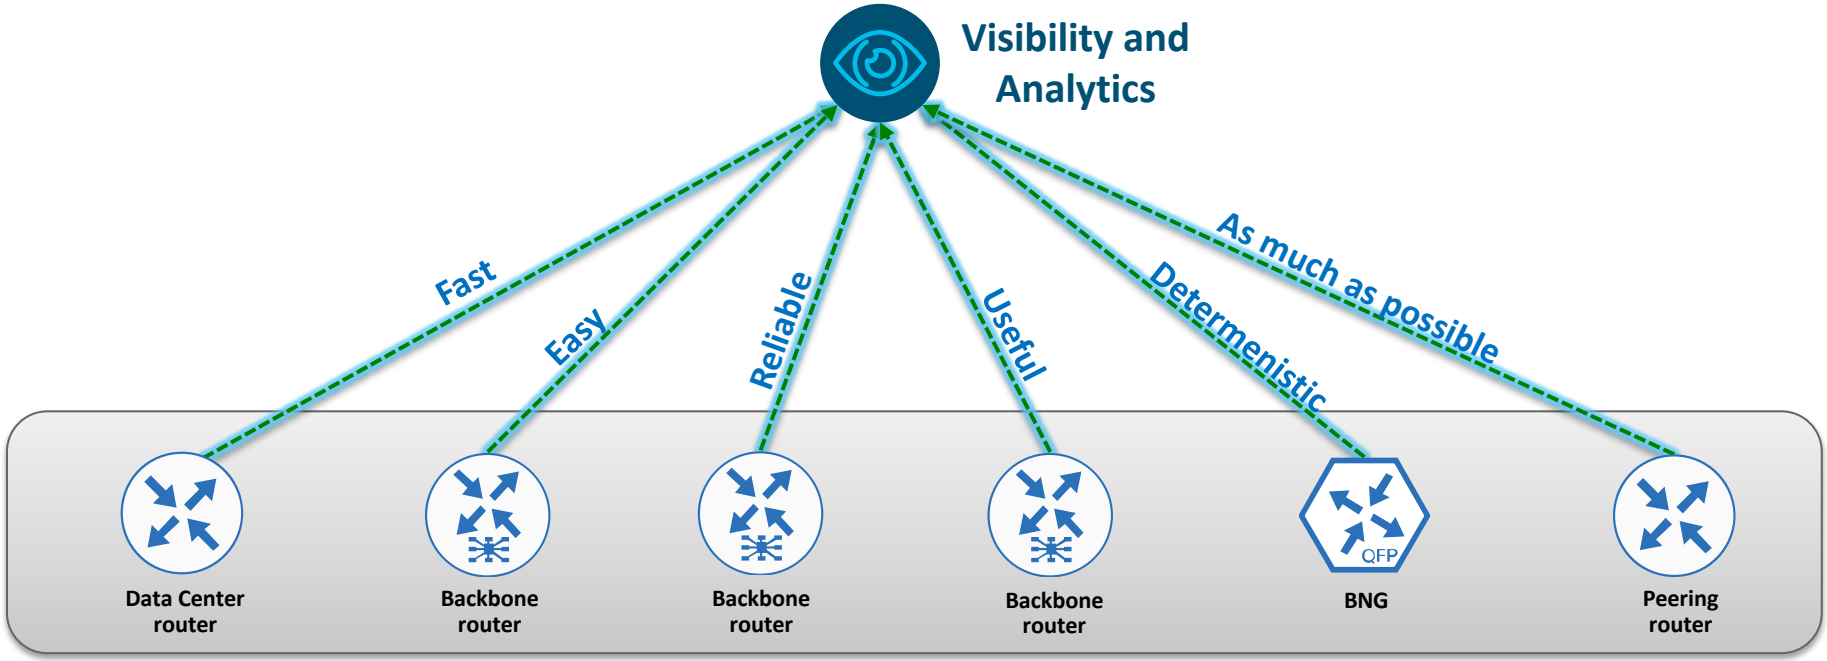

Agenda

1. Brief Telemetry Overview
2. **Closer Look at Router**
3. Closer Look at the Link Between
4. Closer Look at Collector
5. Final Thoughts

How Do You See Telemetry In a Router?

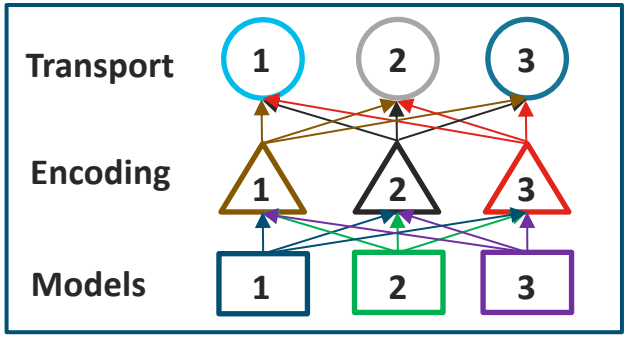

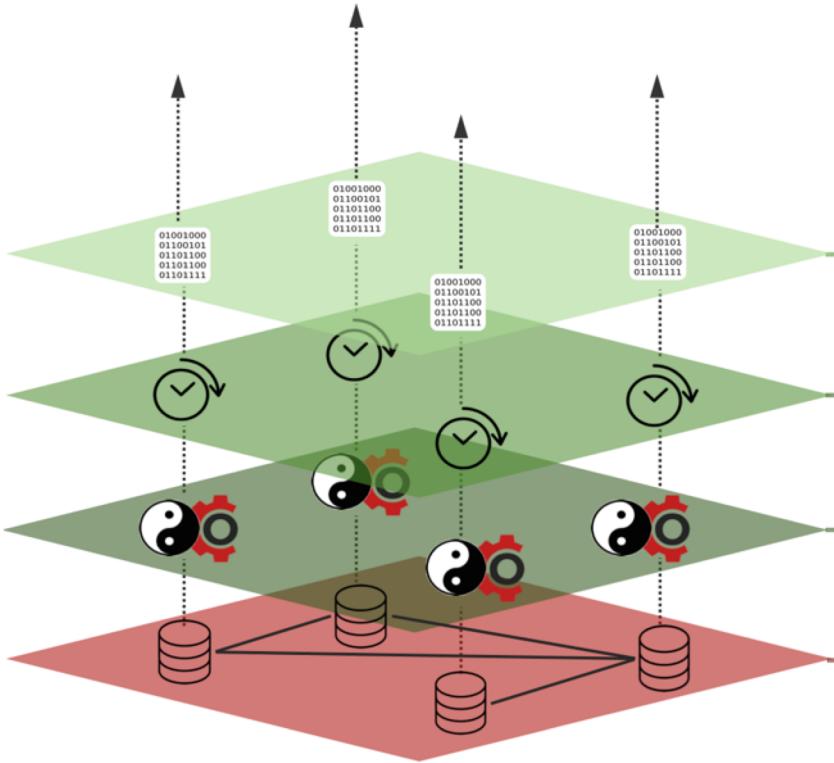

Exporter layer
Encoding and transportation for the models

Producer layer
Time intervals definitions for the models

Data model layer
Raw data mapped to a model
(YANG native, OpenConfig, etc)

Data store layer
Native (raw) data inside a router's database

Check Your Models

Native

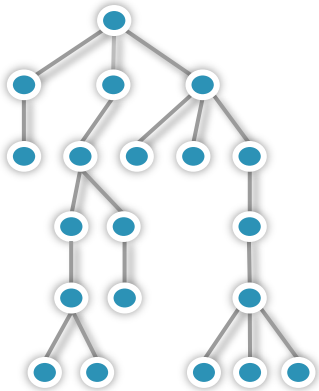

OpenConfig

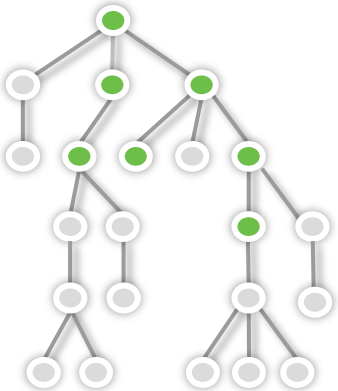

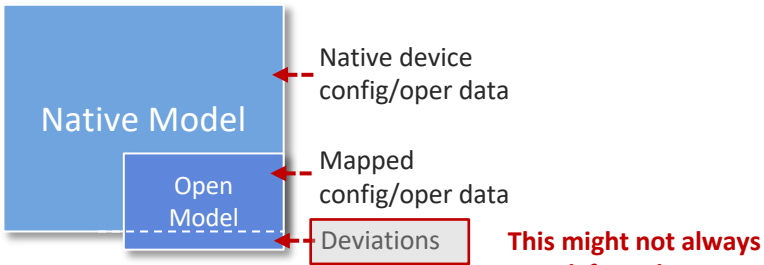

This might not always work for Telemetry

YangModels / yang

Code | Issues 26 | Pull requests 1 | Projects 0 | Wiki | In

Branch: master ▾ yang / vendor / cisco / xr / 632 /

einarnn Typos

..

cisco-xr-openconfig-lacp-deviations.yang	Initial IOS-XR 6.3.2 model commit
cisco-xr-openconfig-lldp-deviations.yang	Initial IOS-XR 6.3.2 model commit
cisco-xr-openconfig-local-routing-deviations.yang	Initial IOS-XR 6.3.2 model commit
cisco-xr-openconfig-mpls-deviations.yang	Initial IOS-XR 6.3.2 model commit
cisco-xr-openconfig-network-instance-deviations.y...	Initial IOS-XR 6.3.2 model commit
cisco-xr-openconfig-optical-amplifier-deviations.ya...	Initial IOS-XR 6.3.2 model commit
cisco-xr-openconfig-platform-deviations.yang	Initial IOS-XR 6.3.2 model commit
cisco-xr-openconfig-platform-transceiver-deviatio...	Initial IOS-XR 6.3.2 model commit
cisco-xr-openconfig-rib-bgp-deviations.yang	Initial IOS-XR 6.3.2 model commit

aristanetworks / yang

Code | Issues 0 | Pull requests 0 | Wiki | Insights

Branch: master ▾ yang / release / openconfig / models / interfaces /

aaronbee EOS 4.20.2F

..

arista-intf-augments.yang

arista-intf-deviations.yang

What Does Sample Interval Really Mean?

- Works fine for small collections
- You should never see missed collections
- Hard to automate, no consistent behavior

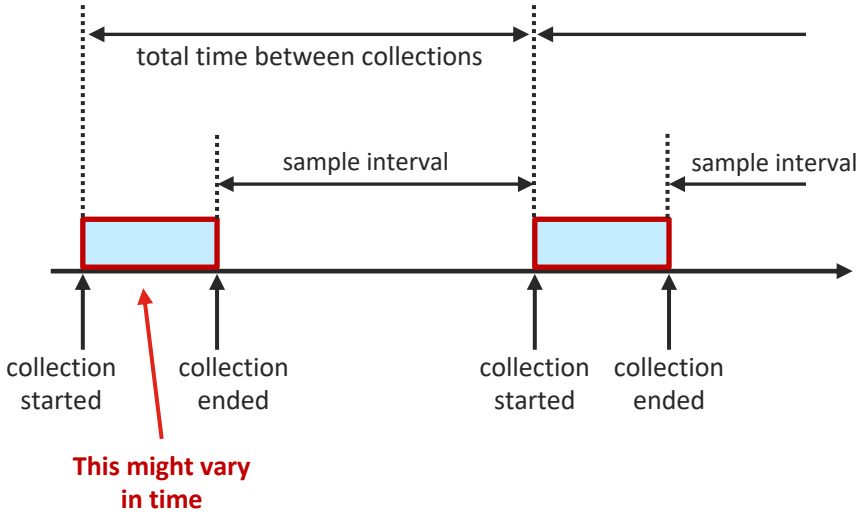

- Consistent behavior, easy to automate
- Sample interval must be more than max collection time
- You might see missed collections

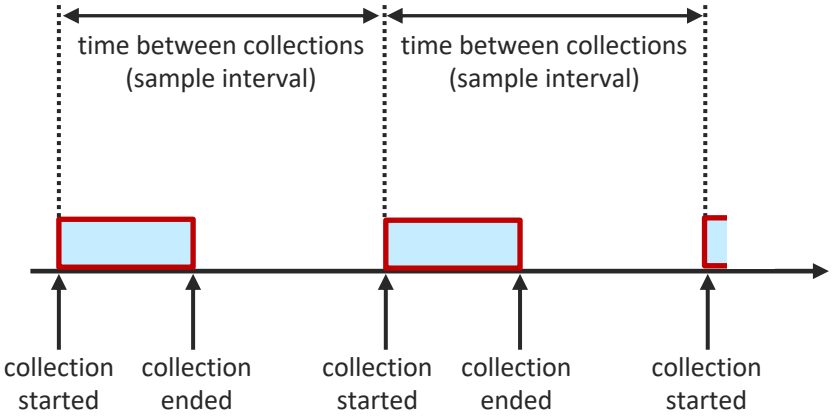

Missed Collections. What To Do?

What do you expect from your router?

- Send all 1s / all 0s?
- Send previous collection values?
- Leave things as is and don't send anything?

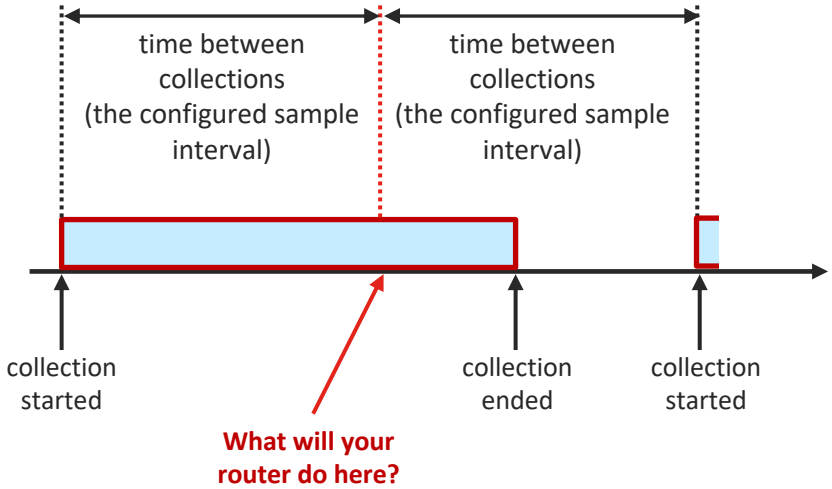

Timestamp	CPU	(Epoch converted)
1522936669977000000	9	(April 5, 2018 1:57:49.977 PM)
1522936674977000000	9	(April 5, 2018 1:57:54.977 PM)
1522936679979000000	9	(April 5, 2018 1:57:59.979 PM)
1522936684976000000	9	(April 5, 2018 1:58:04.976 PM)
1522936689978000000	9	(April 5, 2018 1:58:09.978 PM)
1522936694973000000	9	(April 5, 2018 1:58:14.973 PM)
1522936699983000000	9	(April 5, 2018 1:58:19.983 PM)
1522936704976000000	9	(April 5, 2018 1:58:24.976 PM)
1522936709976000000	9	(April 5, 2018 1:58:29.976 PM)
1522936714976000000	9	(April 5, 2018 1:58:34.976 PM)
1522936719975000000	9	(April 5, 2018 1:58:39.975 PM)
1522936724975000000	9	(April 5, 2018 1:58:44.975 PM)
1522936729982000000	9	(April 5, 2018 1:58:49.982 PM)
1522936734976000000	8	(April 5, 2018 1:58:54.976 PM)
1522936739978000000	8	(April 5, 2018 1:58:59.978 PM)
1522936749977000000	8	(April 5, 2018 1:59:09.977 PM)
1522936754975000000	8	(April 5, 2018 1:59:14.975 PM)

Do you want to fill the gap or not?

How Do You Want Your Data Out?

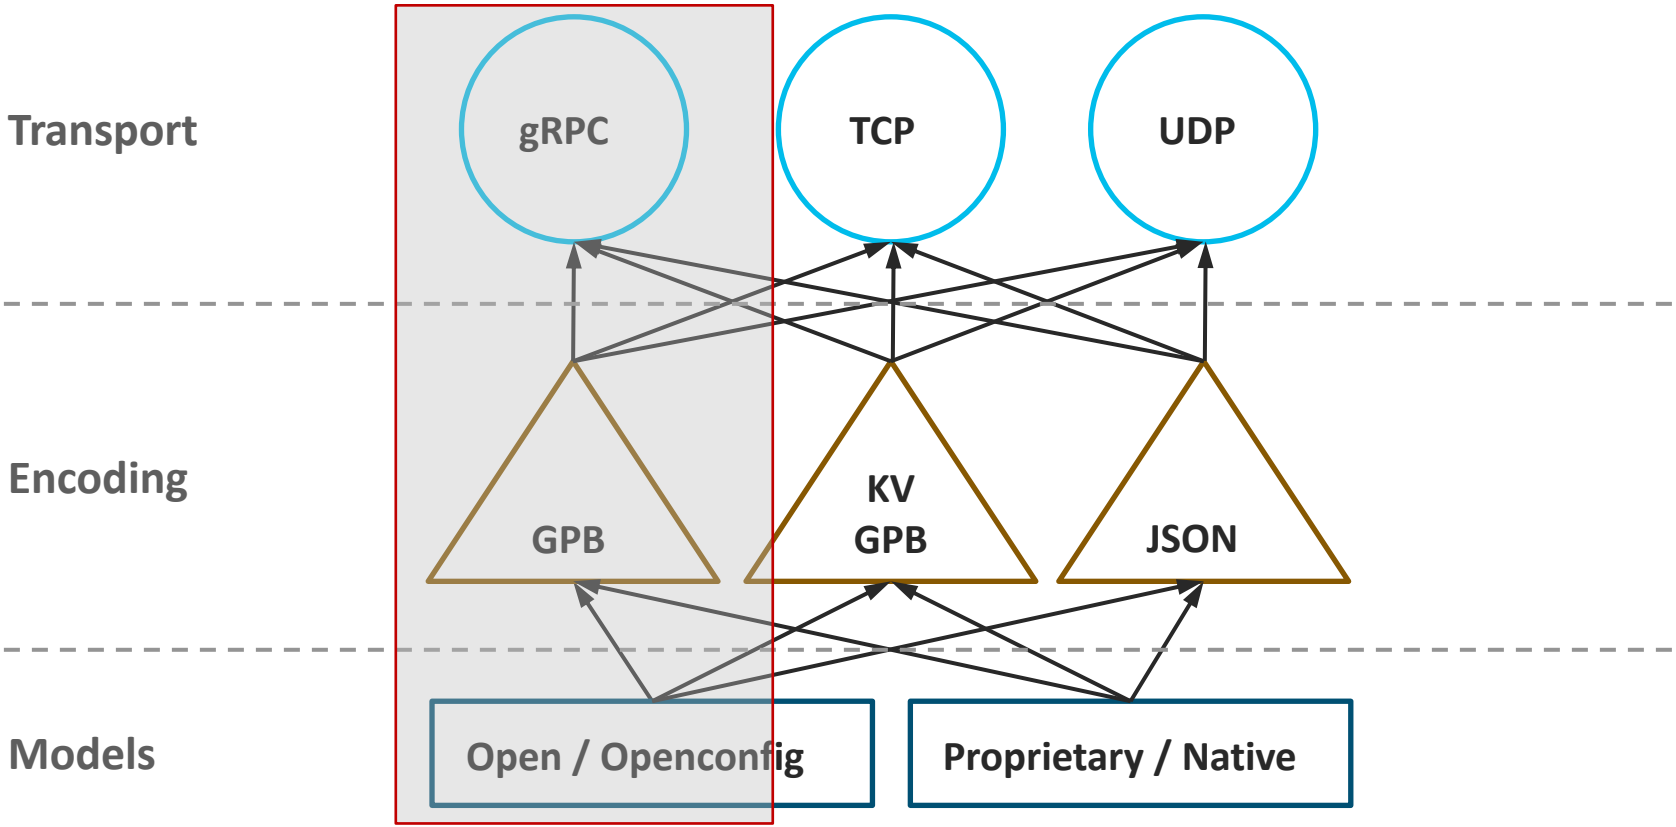

Is It Enough To State gRPC/GPB Support?

Define your Telemetry message

```
syntax = "proto3";
option go_package = "telemetry_bis";

/* Common Telemetry message */   // this is common for both
message Telemetry {
  oneof node_id {
    string node_id_str = 1;
}
  oneof subscription {
    string   subscription_id_str    = 3;
  }
  string   encoding_path          = 6;
  uint64   collection_id          = 8;
  uint64   collection_start_time    = 9;
  uint64   msg_timestamp          = 10;
  repeated TelemetryField data_gpbkv = 11;
  TelemetryGPBTable data_gpb       = 12;
  uint64   collection_end_time     = 13;
}
```

https://github.com/cisco/bigmuddy-network-telemetry-proto/blob/master/staging/telemetry.proto

Define your gRPC calls

```
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) {};
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

Is It Enough To State gRPC/GPB Support?

Juniper's OC .proto

```
package telemetry;

// Interface exported by Agent
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)

    // Terminates and removes an exisiting telemetry subscription
    rpc cancelTelemetrySubscription(CancelSubscriptionRequest)

    // 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)

    // Get Telemetry Agent Operational States
    rpc getTelemetryOperationalState(GetOperationalStateRequest)

    // Return the set of data encodings supported by the device for
    // telemetry data
    rpc getDataEncodings(DataEncodingRequest)
}
```

https://github.com/nileshsimaria/jtimon/blob/master/telemetry/telemetry.proto

Juniper's UDP .proto

```
import "pbj.proto";
import "google/protobuf/descriptor.proto";

extend google.protobuf.FieldOptions {
    optional TelemetryFieldOptions telemetry_options = 1024;
}

message TelemetryFieldOptions {
    optional bool is_key             = 1;
    optional bool is_timestamp       = 2;
    optional bool is_counter         = 3;
    optional bool is_gauge           = 4;
}

message TelemetryStream {
    // router name or export IP address
    required string system_id        = 1 [(telemetry_options).is_key = true,
                                          (pbj_field_option).type = FT_POINTER];

    // line card / RE (slot number)
    optional uint32 component_id     = 2 [(telemetry_options).is_key = true];

    // PFE (if applicable)
    optional uint32 sub_component_id = 3 [(telemetry_options).is_key = true];

    // configured sensor name
    optional string sensor_name      = 4 [(telemetry_options).is_key = true,
                                          (pbj_field_option).type = FT_POINTER];
```

https://github.com/nileshsimaria/jmonudp/blob/master/protos/telemetry_top/telemetry_top.proto

GNMI Should Be The Answer. Right?

Defines the gRPC call

```
service gNMI {
  // Capabilities allows the client to retrieve the set of capabilities that
  // is supported by the target. This allows the target to validate the
  // service version that is implemented and retrieve the set of models that
  // the target supports. The models can then be specified in subsequent RPCs
  // to restrict the set of data that is utilized.
  // Reference: gNMI Specification Section 3.2
  rpc Capabilities(CapabilityRequest) returns (CapabilityResponse);
  // Retrieve a snapshot of data from the target. A Get RPC requests that the
  // target snapshots a subset of the data tree as specified by the paths
  // included in the message and serializes this to be returned to the
  // client using the specified encoding.
  // Reference: gNMI Specification Section 3.3
  rpc Get(GetRequest) returns (GetResponse);
  // Set allows the client to modify the state of data on the target. The
  // paths to modified along with the new values that the client wishes
  // to set the value to.
  // Reference: gNMI Specification Section 3.4
  rpc Set(SetRequest) returns (SetResponse);
  // Subscribe allows a client to request the target to send it values
  // of particular paths within the data tree. These values may be streamed
  // at a particular cadence (STREAM), sent one off on a long-lived channel
  // (POLL), or sent as a one-off retrieval (ONCE).
  // Reference: gNMI Specification Section 3.5
  rpc Subscribe(stream SubscribeRequest) returns (stream SubscribeResponse);
}
```

https://github.com/openconfig/gnmi/blob/master/proto/gnmi/gnmi.proto

What to select here?

Defines the message

```
message Update {
  Path path = 1;                        // The path (key) for the update.
  Value value = 2 [deprecated=true];    // The value (value) for the update.
  TypedValue val = 3;                   // The explicitly typed update value.
  uint32 duplicates = 4;                // Number of coalesced duplicates.
}

// TypedValue is used to encode a value being sent between the client and
// target (originated by either entity).
message TypedValue {
  // One of the fields within the val oneof is populated with the value
  // of the update. The type of the value being included in the Update
  // determines which field should be populated. In the case that the
  // encoding is a particular form of the base protobuf type, a specific
  // field is used to store the value (e.g., json_val).
  oneof value {
    string string_val = 1;              // String value.
    int64 int_val = 2;                  // Integer value.
    uint64 uint_val = 3;                // Unsigned integer value.
    bool bool_val = 4;                  // Bool value.
    bytes bytes_val = 5;                // Arbitrary byte sequence value.
    float float_val = 6;                // Floating point value.
    Decimal64 decimal_val = 7;          // Decimal64 encoded value.
    ScalarArray leaflist_val = 8;       // Mixed type scalar array value.
    google.protobuf.Any any_val = 9;    // protobuf.Any encoded bytes.
    bytes json_val = 10;                // JSON-encoded text.
    bytes json_ietf_val = 11;           // JSON-encoded text per RFC7951.
    string ascii_val = 12;              // Arbitrary ASCII text.
```

https://github.com/openconfig/gnmi/blob/master/proto/gnmi/gnmi.proto#L90-L119

Agenda

1. Brief Telemetry Overview
2. Closer Look at Router
3. **Closer Look at the Link Between**
4. Closer Look at Collector
5. Final Thoughts

How Can a Router Send Its Data Out?

gRPC

TCP

UDP

gRPC Comes With an Overhead…

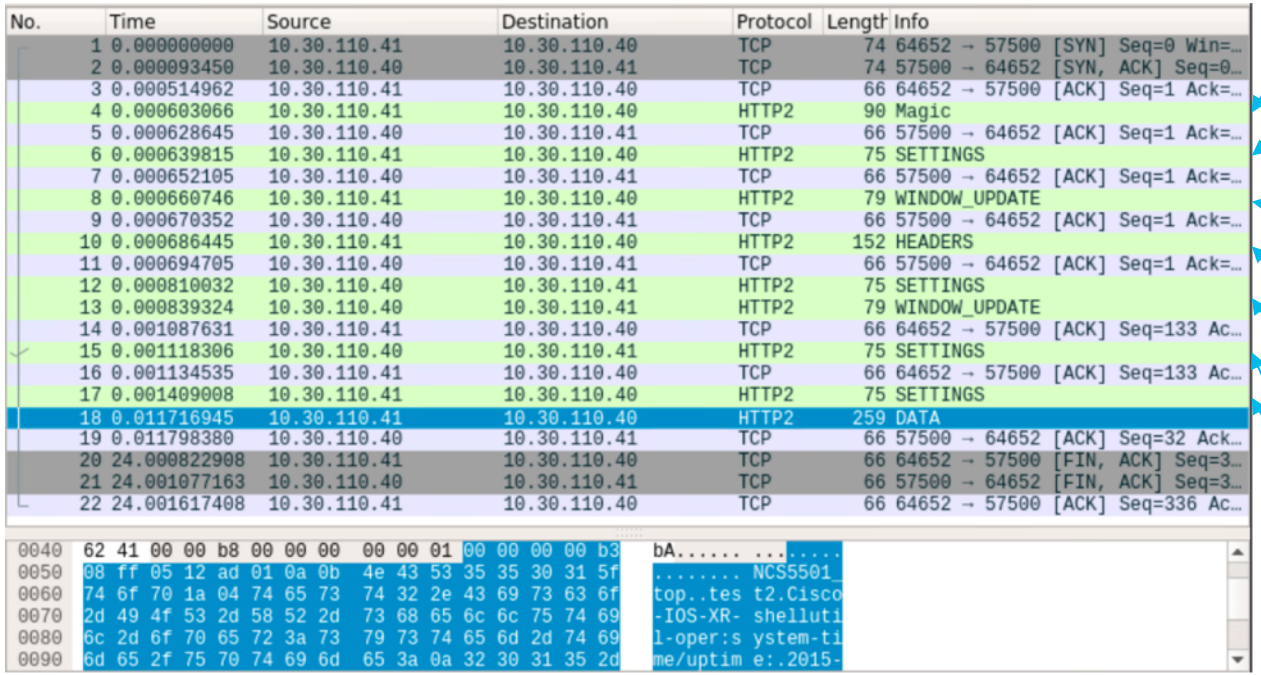

No.	Time	Source	Destination	Protocol	Length	Info
1	0.000000000	10.30.110.41	10.30.110.40	TCP	74	64652 → 57500 [SYN] Seq=0 Win=…
2	0.000093450	10.30.110.40	10.30.110.41	TCP	74	57500 → 64652 [SYN, ACK] Seq=0…
3	0.000514962	10.30.110.41	10.30.110.40	TCP	66	64652 → 57500 [ACK] Seq=1 Ack=…
4	0.000603066	10.30.110.41	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 ystem-ti
0090  6d 65 2f 75 70 74 69 6d  65 3a 0a 32 30 31 35 2d   me/uptim e:.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

```
▸ Flags: 0x00
  0... .... .... .... .... .... .... .... = Reserved: 0x0
  .000 0000 0000 0000 0000 0000 0000 0001 = Stream Identifier: 1
  0... .... .... .... .... .... .... .... = Reserved: 0x0
  .000 0000 0000 0000 0110 1001 1110 0101 = Window Size Increment: 27109
```

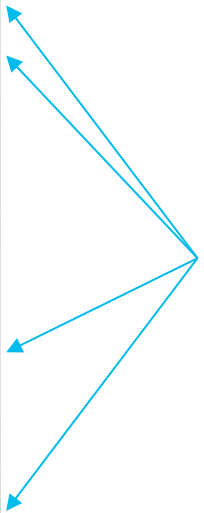

Speed Control (from the collector side)

TCP And UDP Are Simple

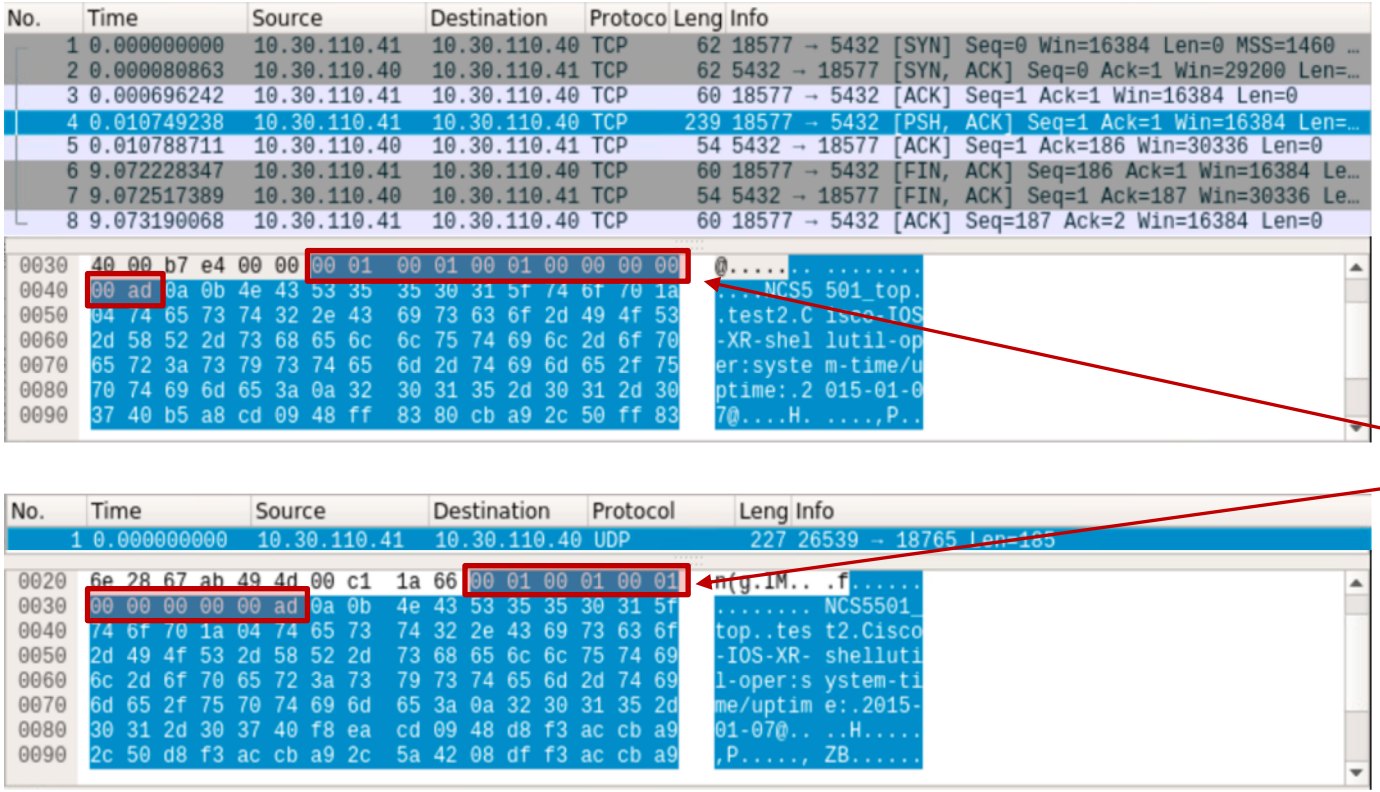

Good to know if there is any additional header inside

UDP Is Fast, But…

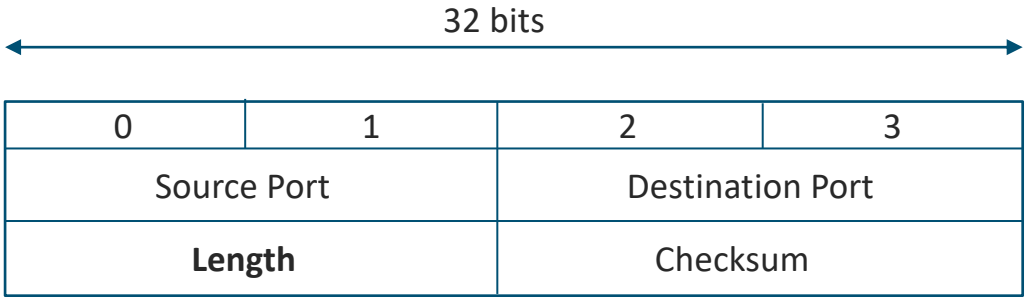

Max datagram length is:
2^{16} – 20 (IPH) - 8 (UDPH) = **65,507** bytes

Sep 19 01:35:01.452 m2m/mdt/backend-timer 0/RP0/CPU0 t15234 45924 [mdtbk_bte_encode_cb]: sub_id 5, /oper/optics/if/*/optics_info, **len 77580**: **mdt_send_encoded_data returned error**

To Encrypt or Not To Encrypt?

gRPC Dial-in (NO-TLS)

Password exchange

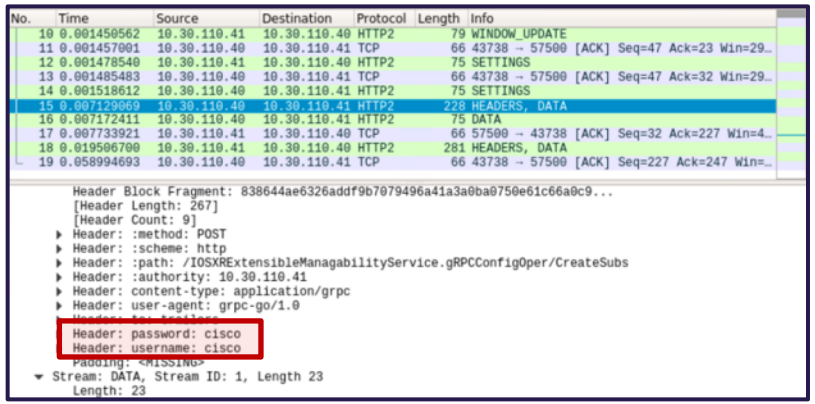

Message content

gRPC Dial-in (TLS)

Password exchange

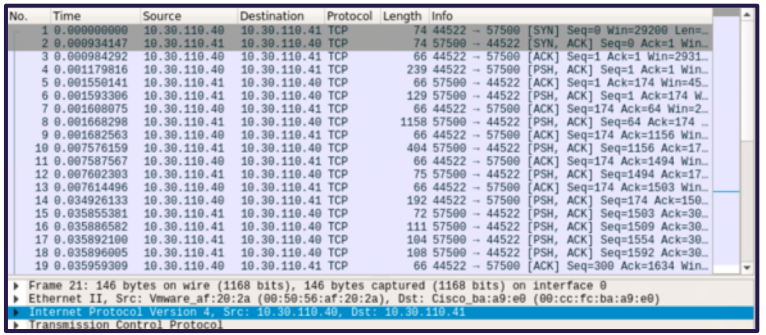

Message content

Which Encoding To Use?

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 67  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 d0 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 8d 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 e6 03  b8 03 04 c0 03 fd 8b e6  .9...... ........
03 c8 03 05 d0 03 00 d8  03 00 e0 03 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 d0 04 00 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 67  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 d0 09  48 fe b9 af da a9 2c 50  -09@.... H.....,P
fe b9 af da a9 2c 5a dc  07 08 82 ba af da a9 2c  .....,Z. .......,
7a 2c 12 04 6b 65 79 73  7a 24 12 0e 69 6e 74 65  z,..keys z$..inte
72 66 61 63 65 2d 6e 61  6d 65 2a 12 48 75 6e 64  rface-na me*.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 40 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 40 dc aa  c2 eb b4 01 7a 15 12 0c  eived@.. ....z...
70 61 63 6b 65 74 73 2d  73 65 6e 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  39 7a 21 12 1a 6d 75 6c  t@...... 9z!..mul
74 69 63 61 73 74 2d 70  61 63 6b 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  broadcas t-packet
73 2d 72 65 63 65 69 76  65 64 40 04 7a 1d 12 16  s-receiv ed@.z...
6d 75 6c 74 69 63 61 73  74 2d 70 61 63 6b 65 74  multicas t-packet
73 2d 73 65 6e 74 40 d3  8f e6 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
00 7a 0f 12 0b 69 6e 70  75 74 2d 64 72 6f 70 73  .z...inp ut-drops
38 00 7a 15 12 11 69 6e  70 75 74 2d 71 75 65 75  8.z...in put-queu
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-re ceived8.
7a 1e 12 1a 74 68 72 6f  74 74 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 ceived8.
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  !unknown -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  12 0a 63 72 63 2d 65 72  rors8.z. ..crc-er
72 6f 72 73 38 00 7a 12  12 0e 69 6e 70 75 74 2d  rors8.z. ..input-
6f 76 65 72 72 75 6e 73  38 00 7a 1b 12 17 66 72  overruns 8.z...fr
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 00 01 00 00 00 05 33  ....8... .......3
08 f6 06 12 ad 0a 7b 22  6e 6f 64 65 5f 69 64 5f  ......{" node_id_
73 74 72 22 3a 22 4e 43  53 35 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  67 5f 70 61 74 68 22 3a  "encodin 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 67 65 6e  rface/la test/gen
65 72 69 63 2d 63 6f 75  6e 74 65 72 73 22 2c 22  eric-cou 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 67 5f 74 69 6d 65 73  74 61 6d 70 22 3a 31 35  sg_times tamp":15
32 33 30 32 34 33 32 35  32 38 39 2c 22 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 30 2f 31 2f 30  ndredGig E0/0/1/0
22 7d 2c 22 63 6f 6e 74  65 6e 74 22 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 79 74 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 2d 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  25428671 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- received
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  ackets-s ent":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  22 69 6e 70 75 74 2d 64  rops":0, "input-d
72 6f 70 73 22 3a 30 2c  22 69 6e 70 75 74 2d 71  rops":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 65  unt-pack ets-rece
69 76 65 64 22 3a 30 2c  22 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":
```

Design Your Transport Network Properly

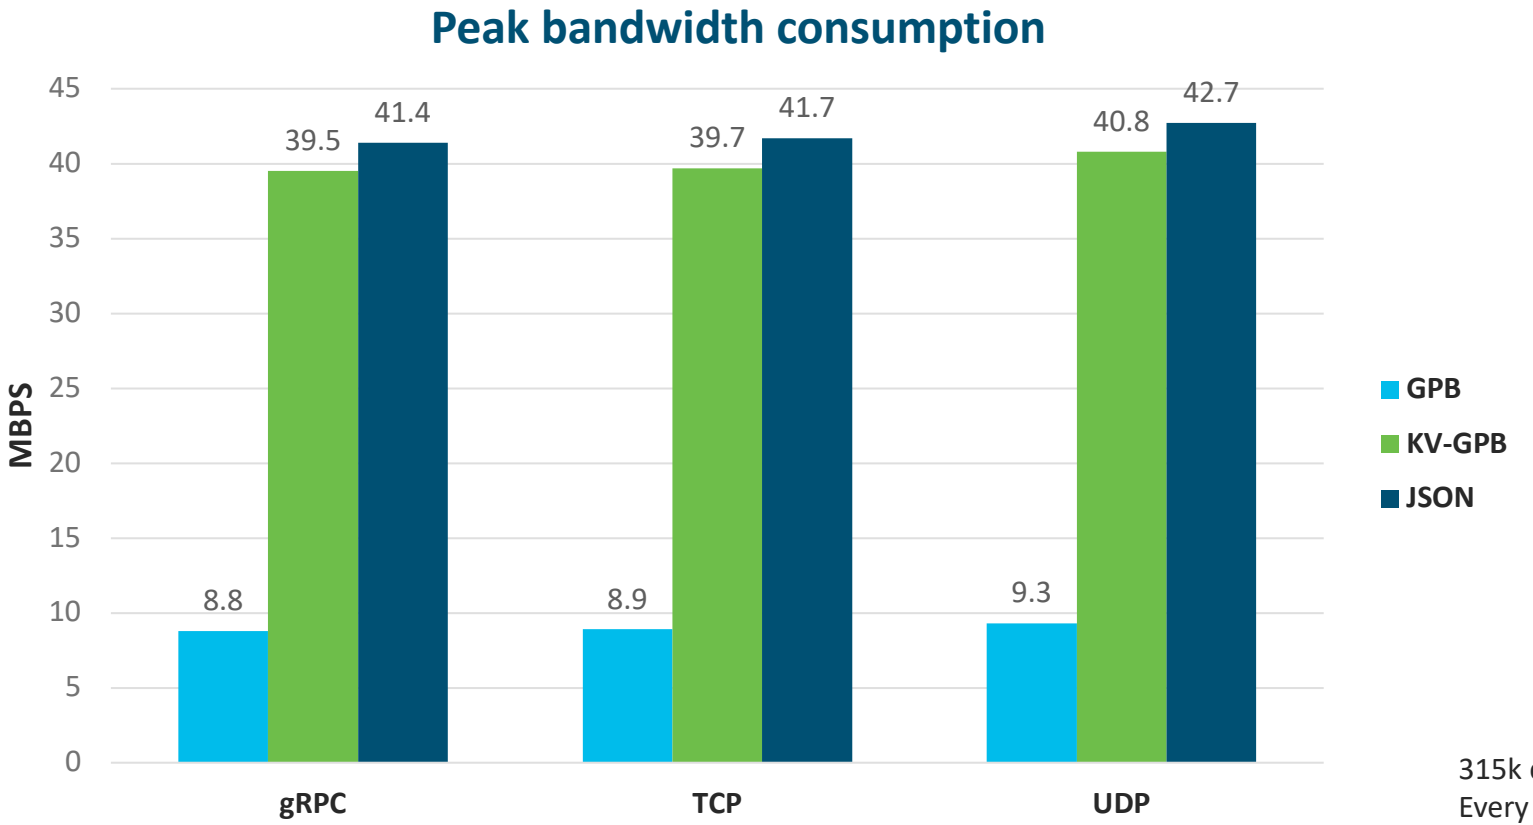

315k counters
Every 5 seconds

Agenda

1. Brief Telemetry Overview
2. Closer Look at Router
3. Closer Look at the Link Between
4. **Closer Look at Collector**
5. Final Thoughts

Is Your Collector Fast Enough?

Decoded messages volume

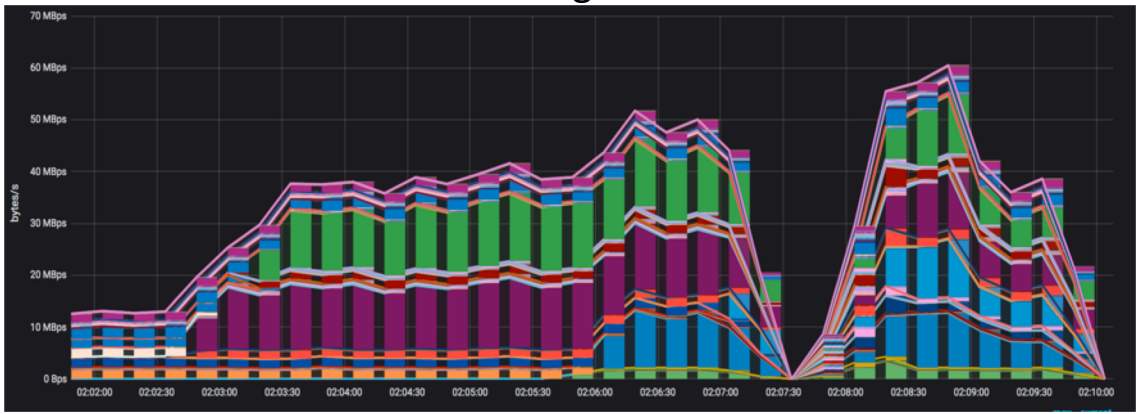


```
RP/0/RP0/CPU0:NCS5501_bottom#sh telemetry model destination internal
Wed Apr 11 02:07:19.543 PDT
  dest Id          Sub              IP                Port  Trans     State  Chan  Buff  Full
-------------------------------------------------------------------------------------------
des1             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
```

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

Is Your Hard Drive Write Speed Fast Enough?

HDD-based server (SAS)

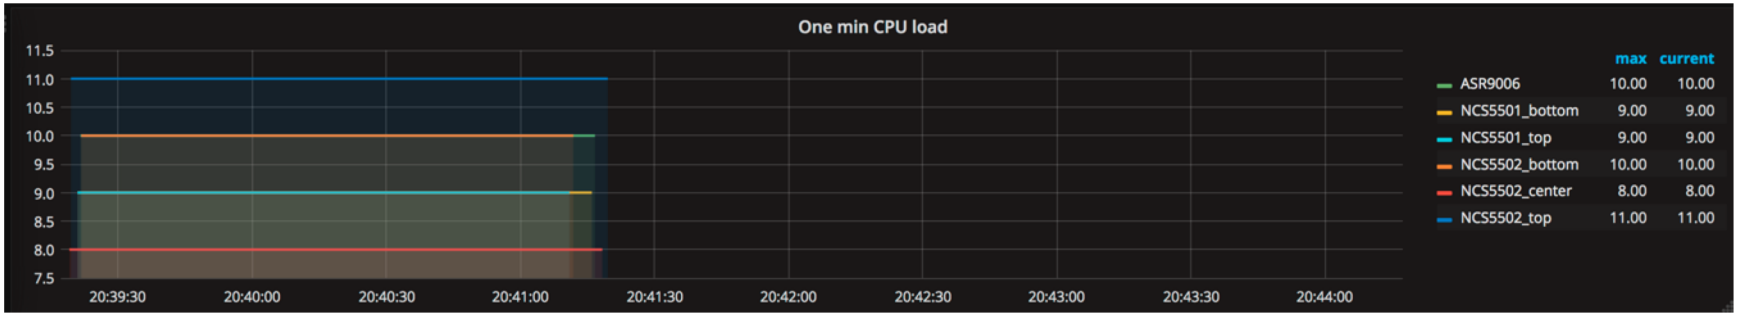

SSD-based server (SAS)

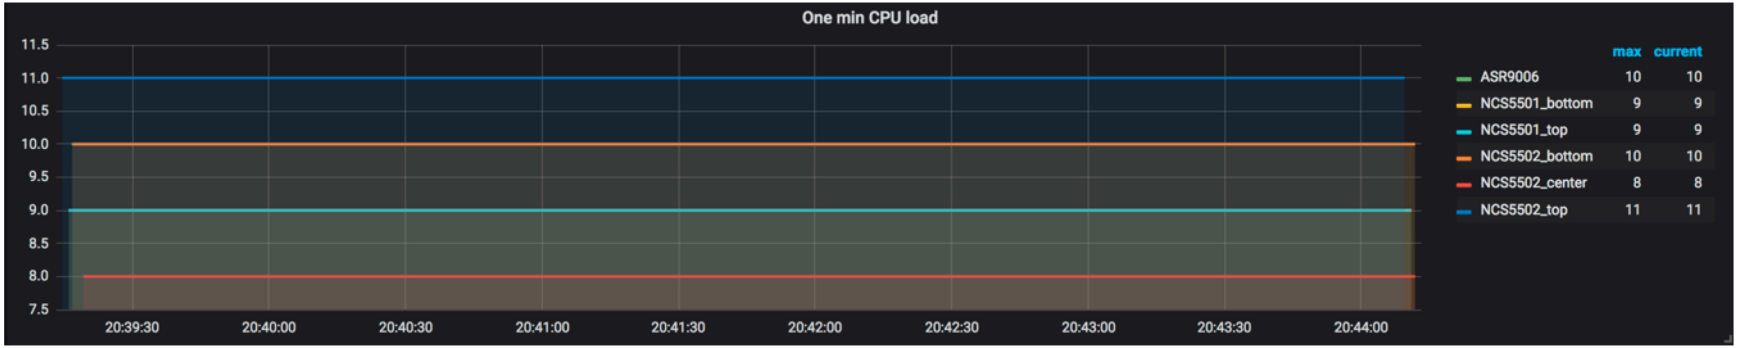

What To Think About Selecting a TSDB

Metric	InfluxDB	Prometheus	Elasticsearch	OpenTSDB
Website	https://influxdata.com/	https://prometheus.io/	https://www.elastic.co/products/elasticsearch	http://opentsdb.net
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	nano second	milli second	milli second	milli second
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/y8ofbjyy
InfluxDB vs Cassandra: https://tinyurl.com/y83vv9ys
DB ranking: https://tinyurl.com/ya8rrrjp
InfluxDB vs Elasticsearch: https://tinyurl.com/y7yxjf6v

Everything is working, where are my stats?

```
RP/0/RP0/CPU0:ios-xr# sh clock
Sun Apr  1 20:56:15.074 PDT
20:56:15.167 PDT Sun Apr 1 2018

cisco@ubuntu51-1:~$ date
Sun Apr  1 23:13:11 PDT 2018
```

```
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
```

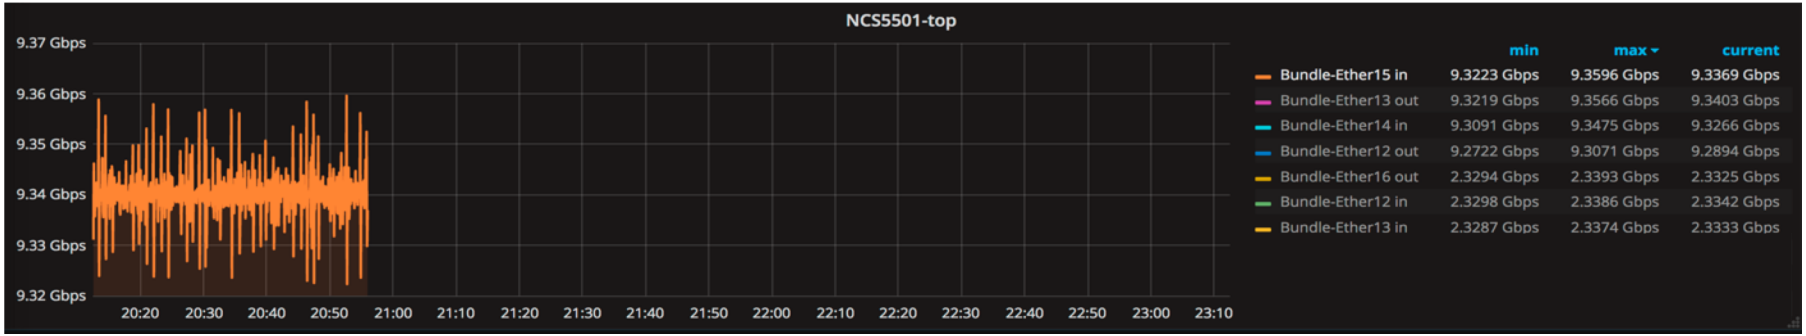

Agenda

1. Brief Telemetry Overview
2. Closer Look at Router
3. Closer Look at the Link Between
4. Closer Look at Collector
5. **Final Thoughts**

Think About Your Overall Design

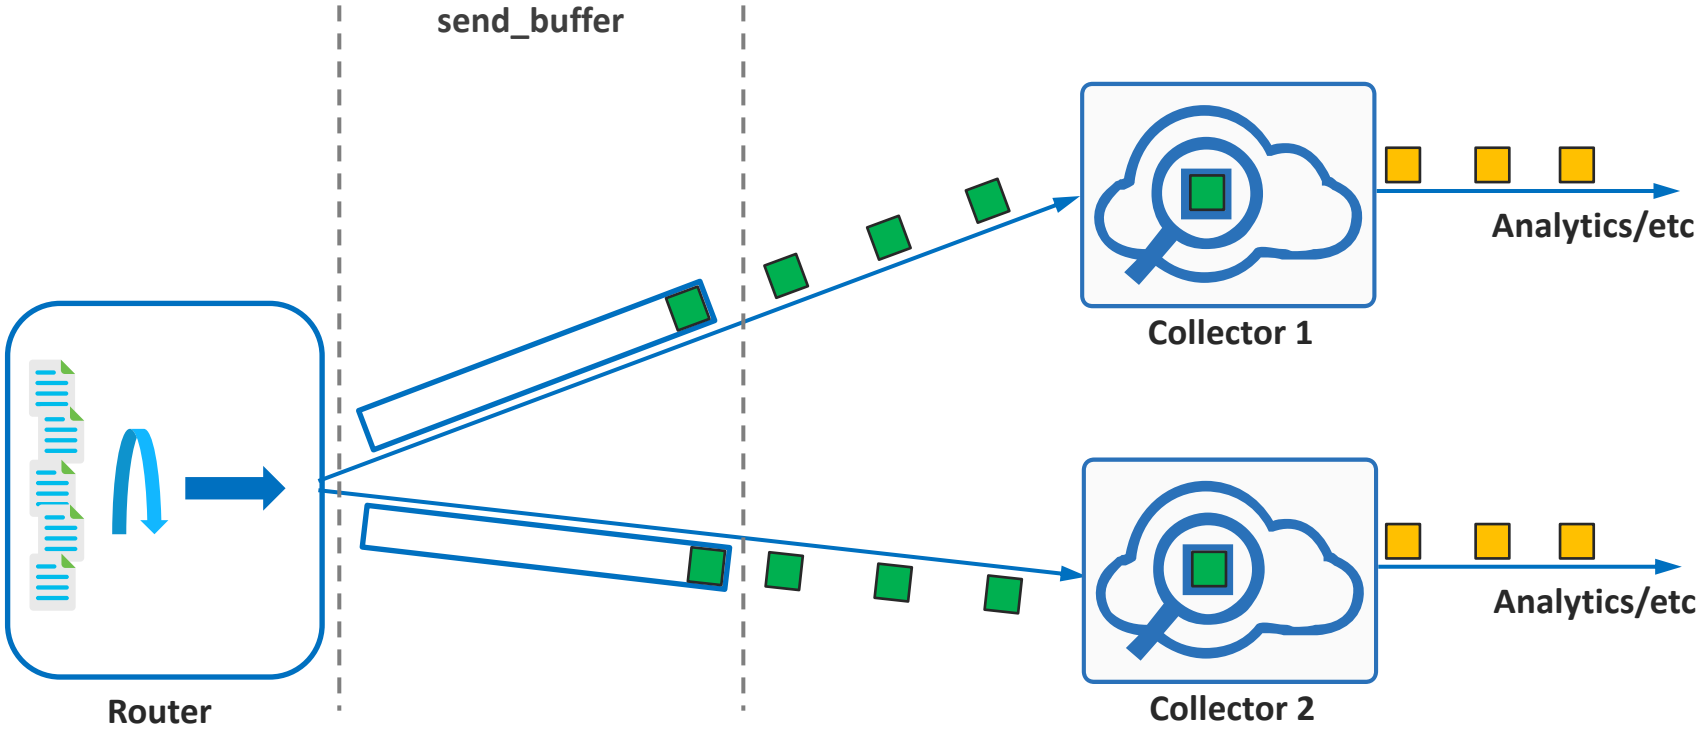

Think About Your Overall Design

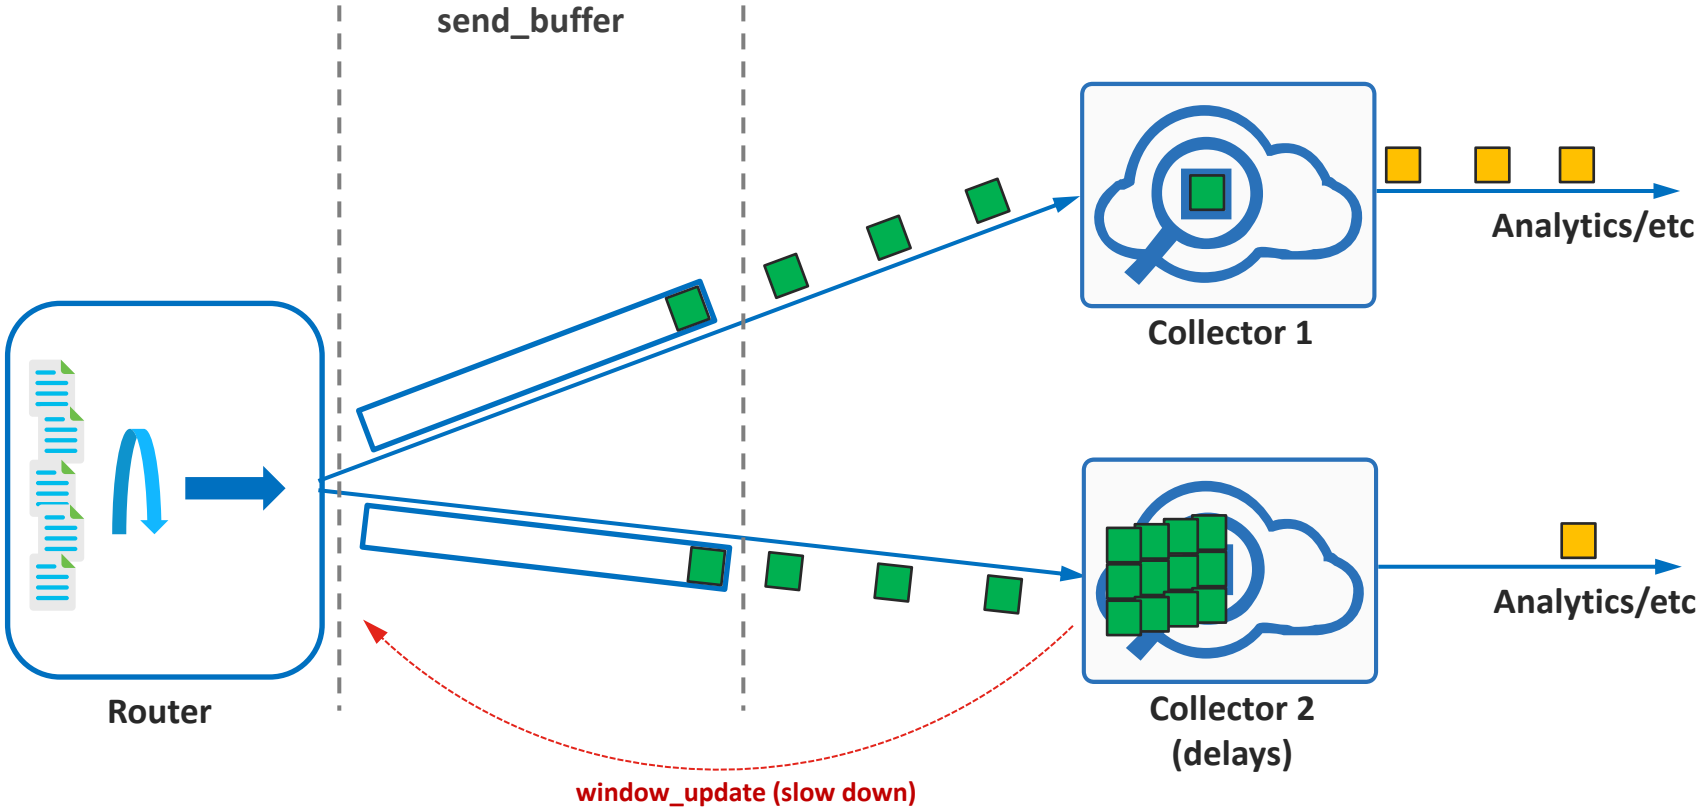

Think About Your Overall Design

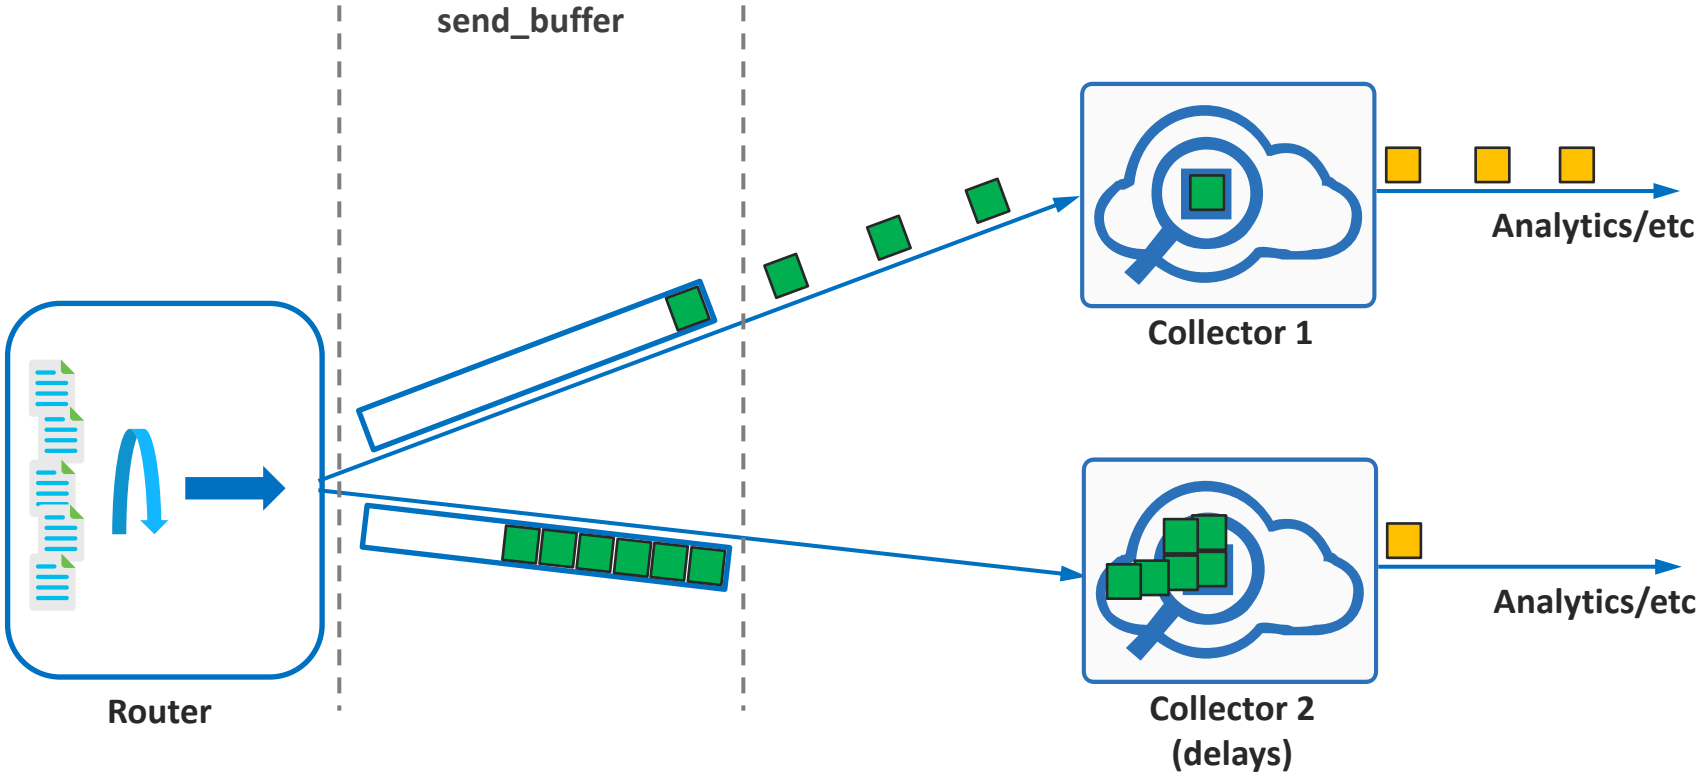

Think About Your Overall Design

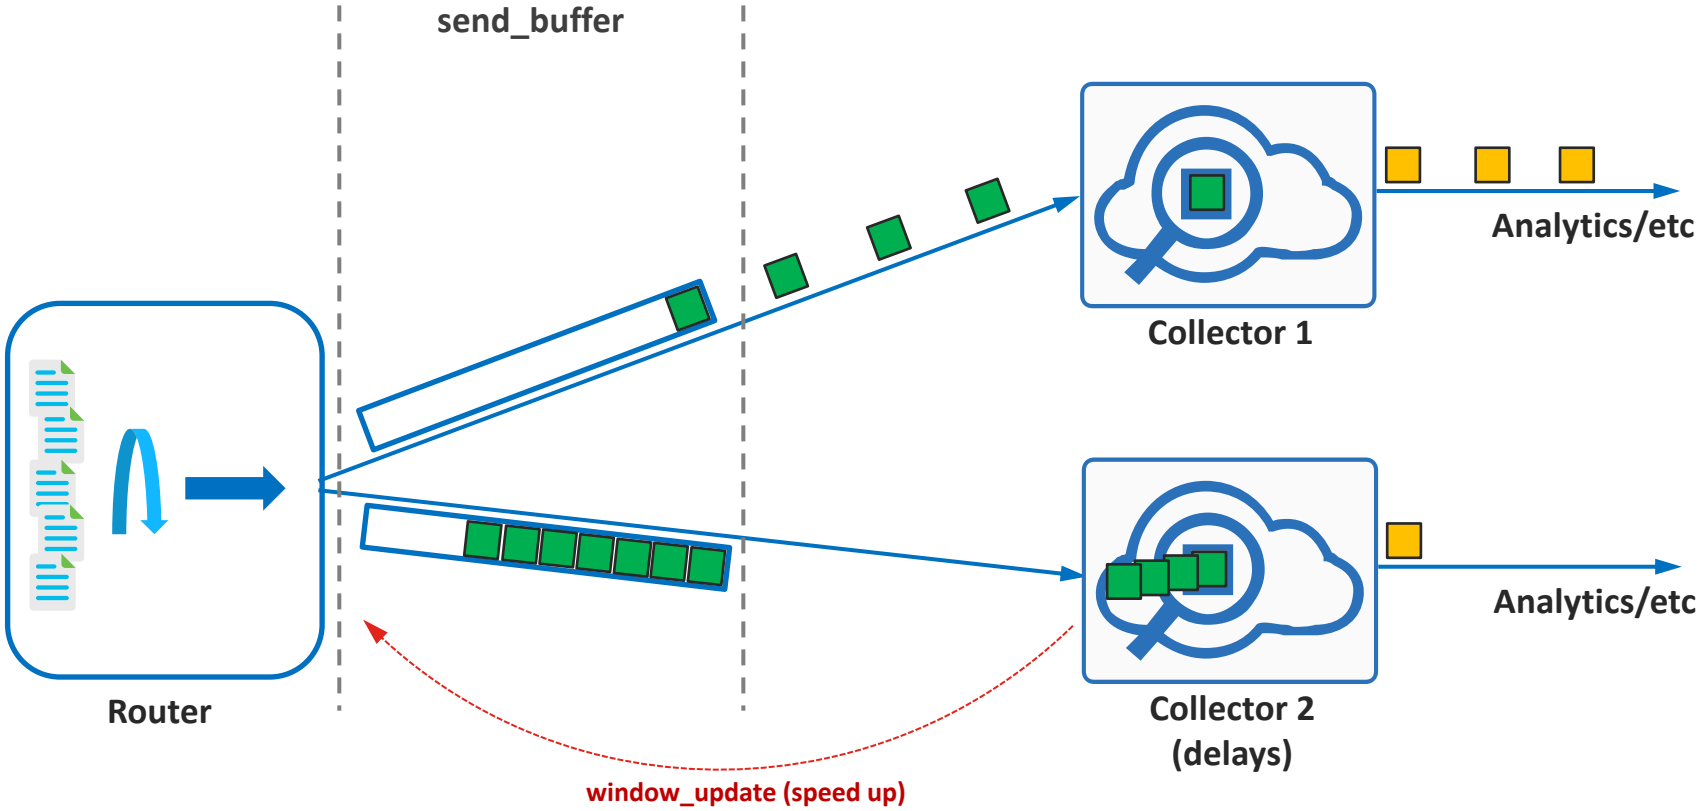

Think About Your Overall Design

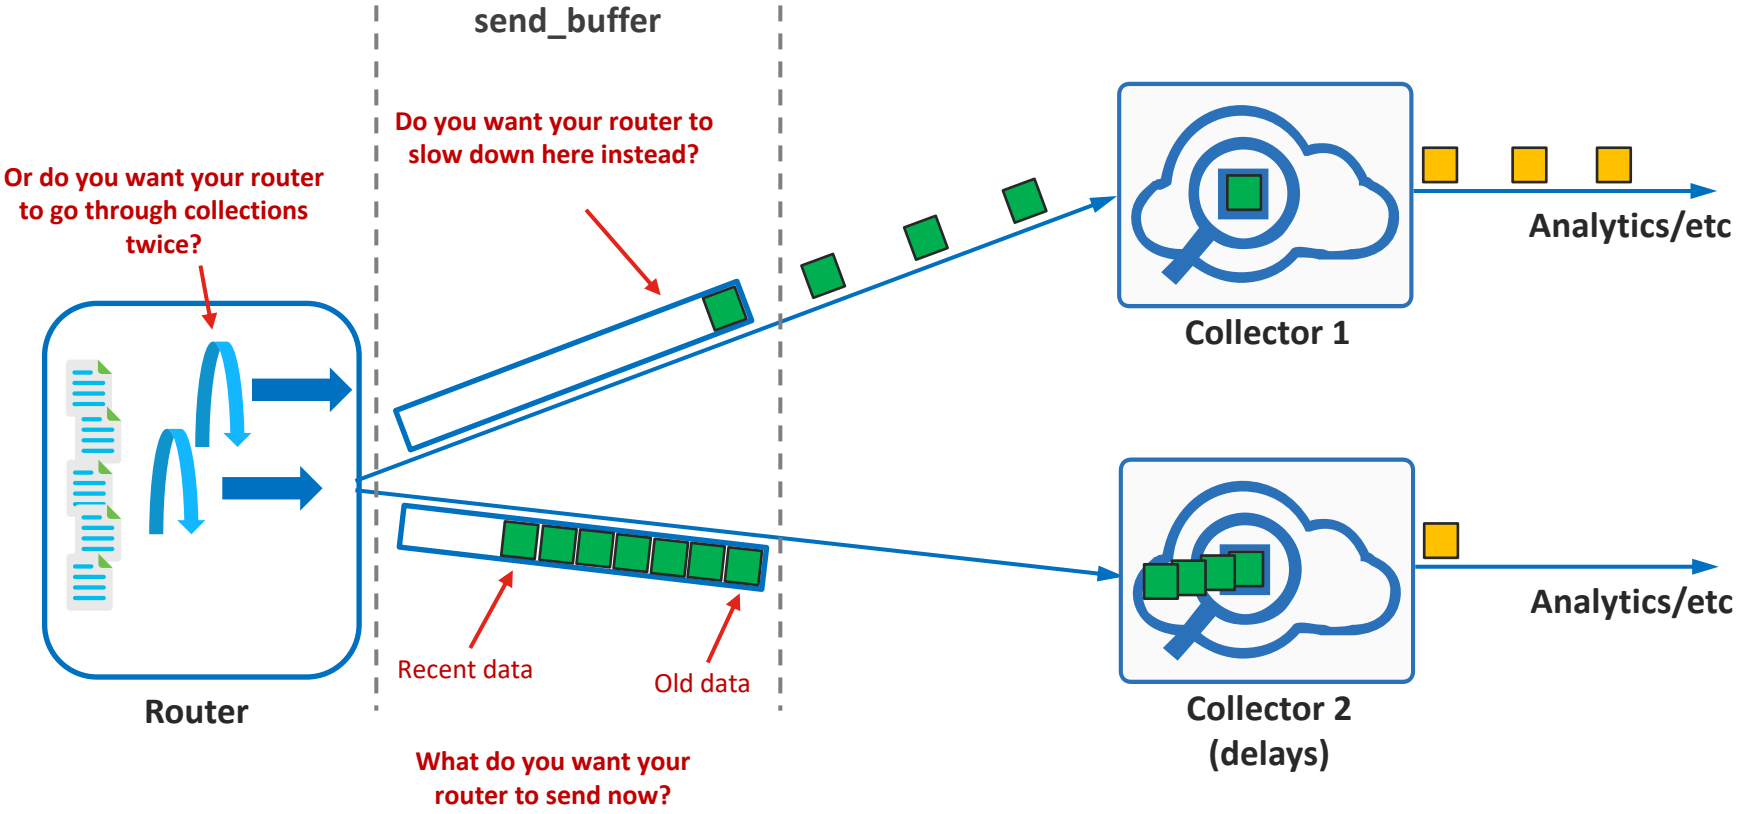

It Is Good to Know More

A snapshot from https://www.youtube.com/watch?v=ZMZJ3ZaEcIQ