Network Traffic Visibility and Anomaly Detection

@Scale: October 27th, 2016

Dan Ellis

kentik

kentik | Introduction

- Network traffic visibility?

kentik | Introduction

- Network traffic visibility?
  - What data is available on your network
  - What can you do with this data
  - Tools available

kentik | Introduction

- Network traffic visibility?
  - What data is available on your network
  - What can you do with this data
  - Tools available
- 20+ years running blind (ISP's, CDN's, enterprise)
- Who is Kentik

Goal of this talk: Make your life easier

kentik | Traffic Visibility Problem

- Data networks can be compared to FedEx
- Imagine FedEx without package tracking
- Majority of data networks operate in this vacuum of visibility
- Hard to believe? Problem is massive data scale, lack of tools, little network + systems collaboration

kentik | Not Helping...

- Interface Volume (Mb/s, pps)?

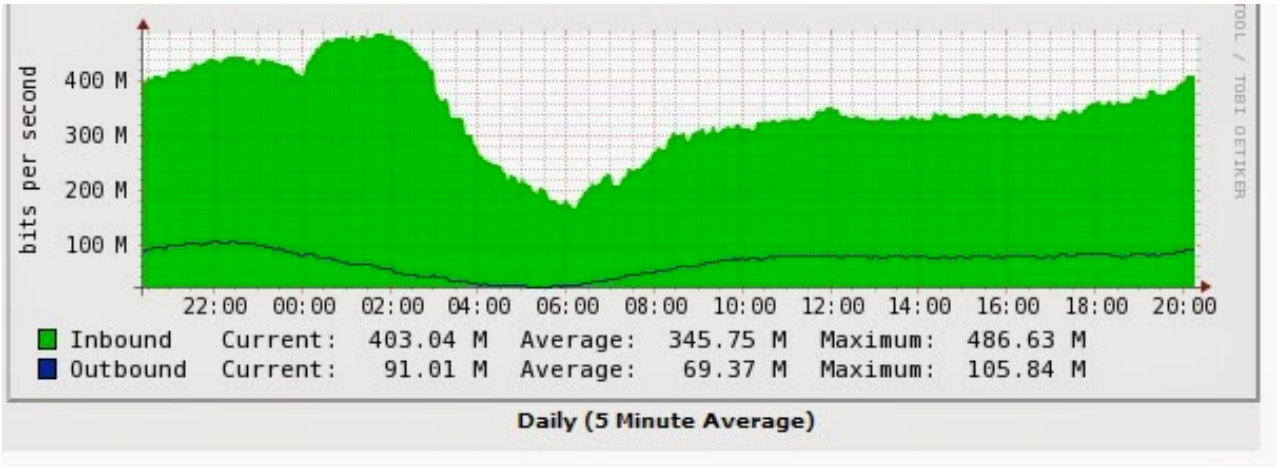

kentik | A Poll

- Interface Volume (Mb/s, pps)?
- `Src/Dst` IP+Port, ASN, BGP Path?

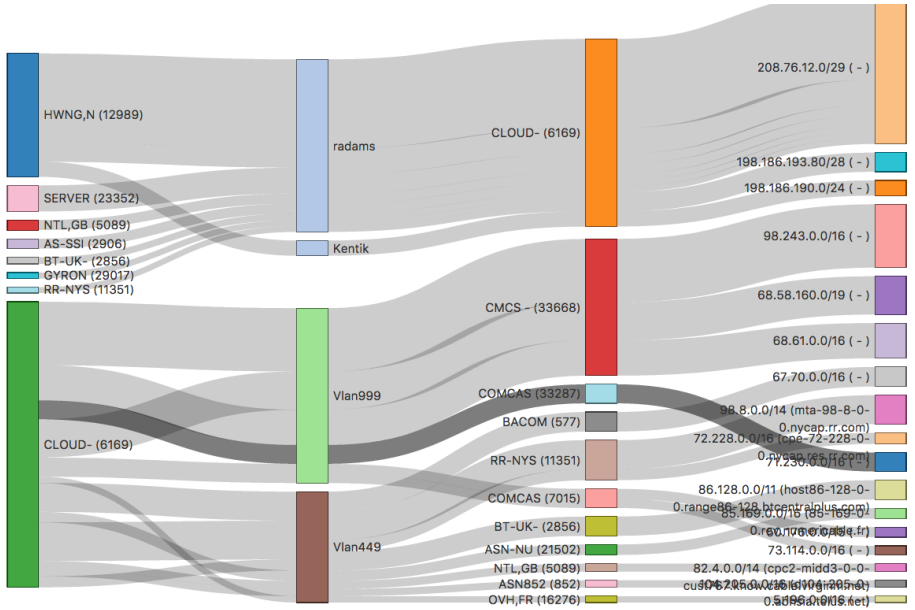

key	Avg Mb/sec	95th Percentile	Max Mb/sec	Last Datapoint
Total	148.09	215.11	248.82	119.15
CMCS - Comcast Cable Communications, LLC,US (33668) ---- 98.243.0.0/16 (-)	6.88 (4.6%)	70.19	89.83	0.03
CMCS - Comcast Cable Communications, LLC,US (33668) ---- 68.58.160.0/19 (-)	9.21 (6.2%)	45.71	54.15	0.04
CMCS - Comcast Cable Communications, LLC,US (33668) ---- 68.61.0.0/16 (-)	36.12 (24.4%)	48.31	49.98	43.49
RR-NYSREGION-ASN-01 - Time Warner Cable Internet LLC,US (11351) ---- 98.8.0.0/14 (mta-98-8-0-0.nycap.rr.com)	3.53 (2.4%)	29.45	41.07	0.01
BACOM - Bell Canada,CA (577) ---- 67.70.0.0/16 (-)	7.60 (5.1%)	25.89	27.70	9.63
COMCAST-33287 - Comcast Cable Communications, LLC,US (33287) ---- 71.230.0.0/16 (-)	6.33 (4.3%)	16.62	26.94	11.53

kentik | A Poll

- Interface Volume (Mb/s, pps)?
- `Src/Dst` IP+Port, ASN, BGP Path?
- IP, Port, ASN or Path Thresholds?

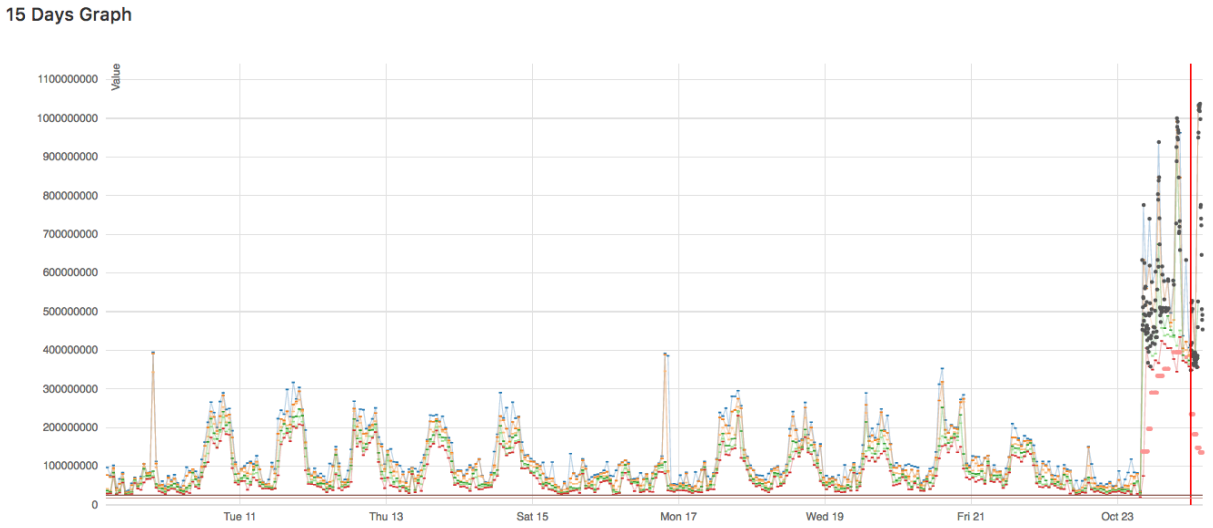

kentik | Is this really a problem?

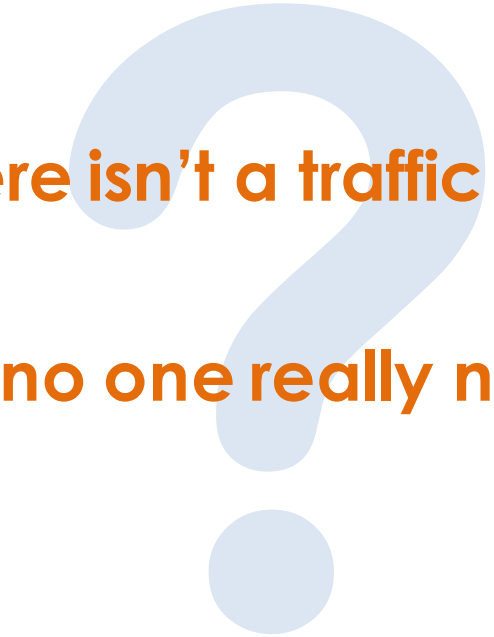

Maybe there isn't a traffic visibility problem

Maybe no one really needs this data

kentik | Complaints of high latency… BGP Path to Comcast

Top Dest BGP AS_Path by 95th Percentile Bits/s

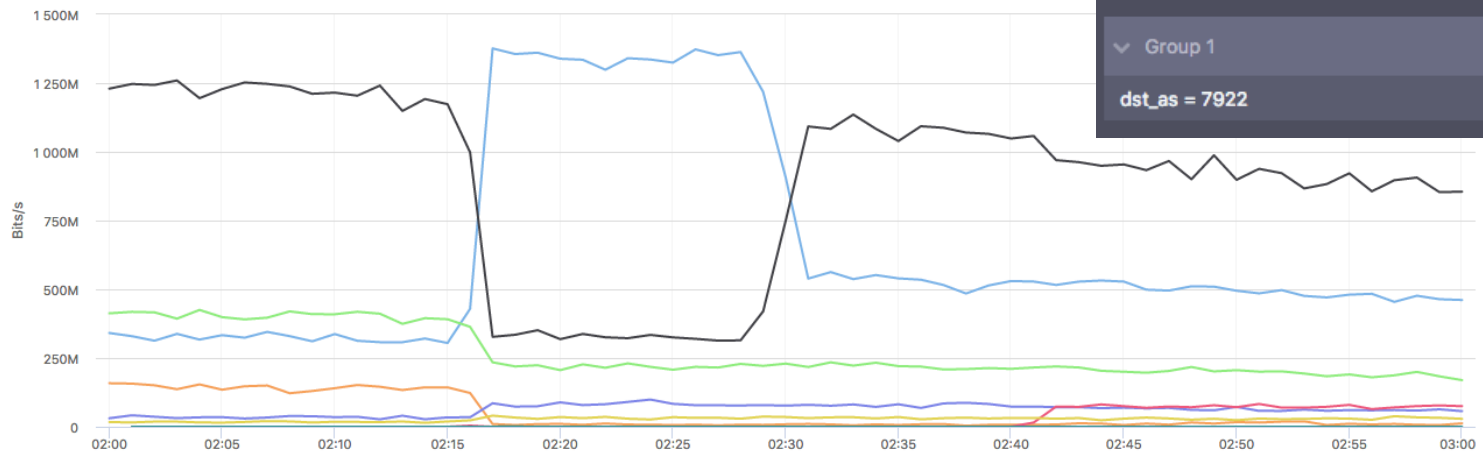

Left +Y Axis

key	Avg Mb/sec	95th Percentile	Max Mb/sec	Last Datapoint
13789 1299 7922	651.32	1,358.16	1,376.25	461.08
13789 701 7922	914.94	1,245.05	1,259.58	855.10
13789 6461 7922	268.25	417.08	424.98	169.76
13789 209 7922	47.26	151.36	158.45	12.25
13789 7018 7922	63.32	86.45	99.01	56.83
13789 7922	24.11	78.24	82.92	75.02
13789 174 7922	27.99	35.74	40.44	28.90
174 7922	0.02	0.02	0.82	0.01

kentik | Dyn attack last week – ISP recursive inbound

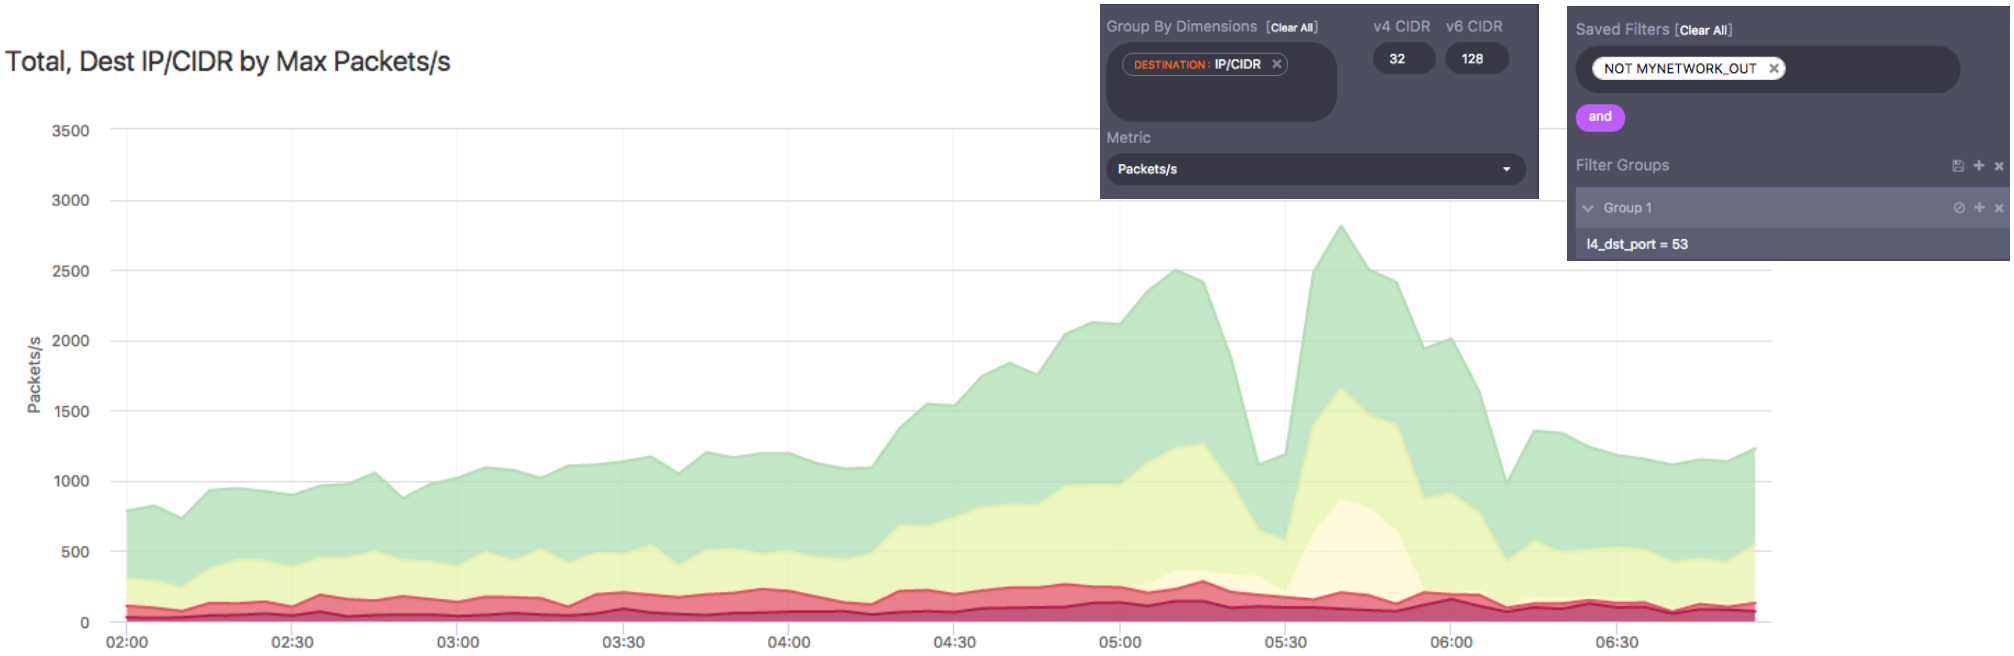

Left +Y Axis

name	Avg pps	95th Percentile	Max pps	Last Datapoint
Total ---- 204.186.0.203/32 (dns3.ptd.net)	765	1,157	1,270	683
Total ---- 204.186.0.180/32 (dns.pal.ptd.net)	434	792	901	403
Total ---- 75.97.132.95/32 (75.97.132.95.res-cmts.sewb.ptd.net)	52	493	655	14

kentik | Dyn attack last week – Traffic / source_ip

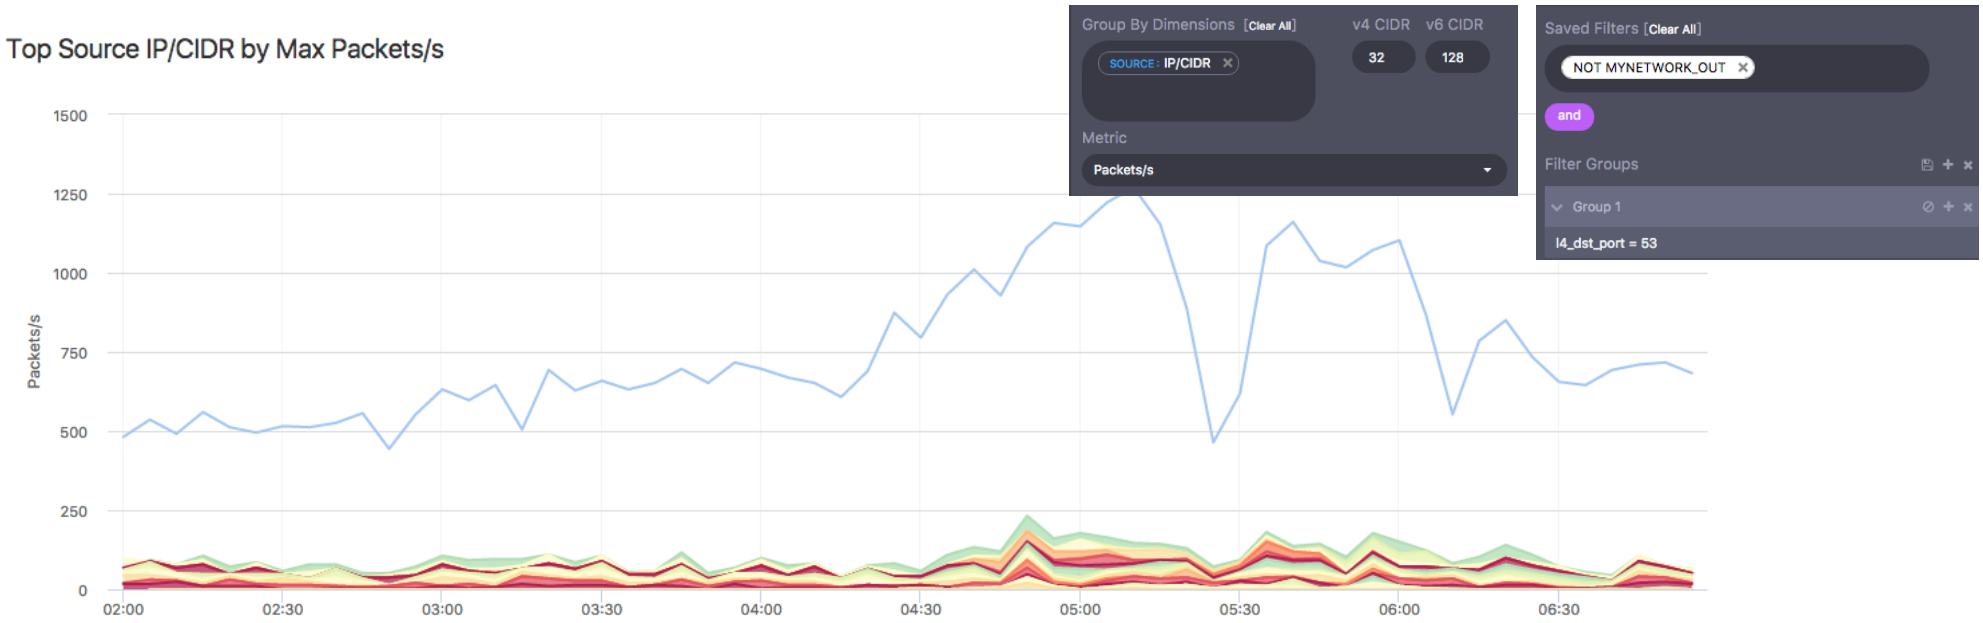

Left +Y Axis

key	Avg pps	95th Percentile	Max pps	Last Datapoint
Total	765	1,157	1,270	683
75.97.108.23/32 (75.97.108.23.res-cmts.t132.ptd.net)	15 (1.9%)	34	48	7
24.102.242.141/32 (24.102.242.141.res-cmts.t132.ptd.net)	2 (0.3%)	24	38	38
24.102.242.148/32 (24.102.242.148.res-cmts.t132.ptd.net)	9 (1.2%)	20	38	14

kentik | **Dyn attack last week** – ISP recursive outbound

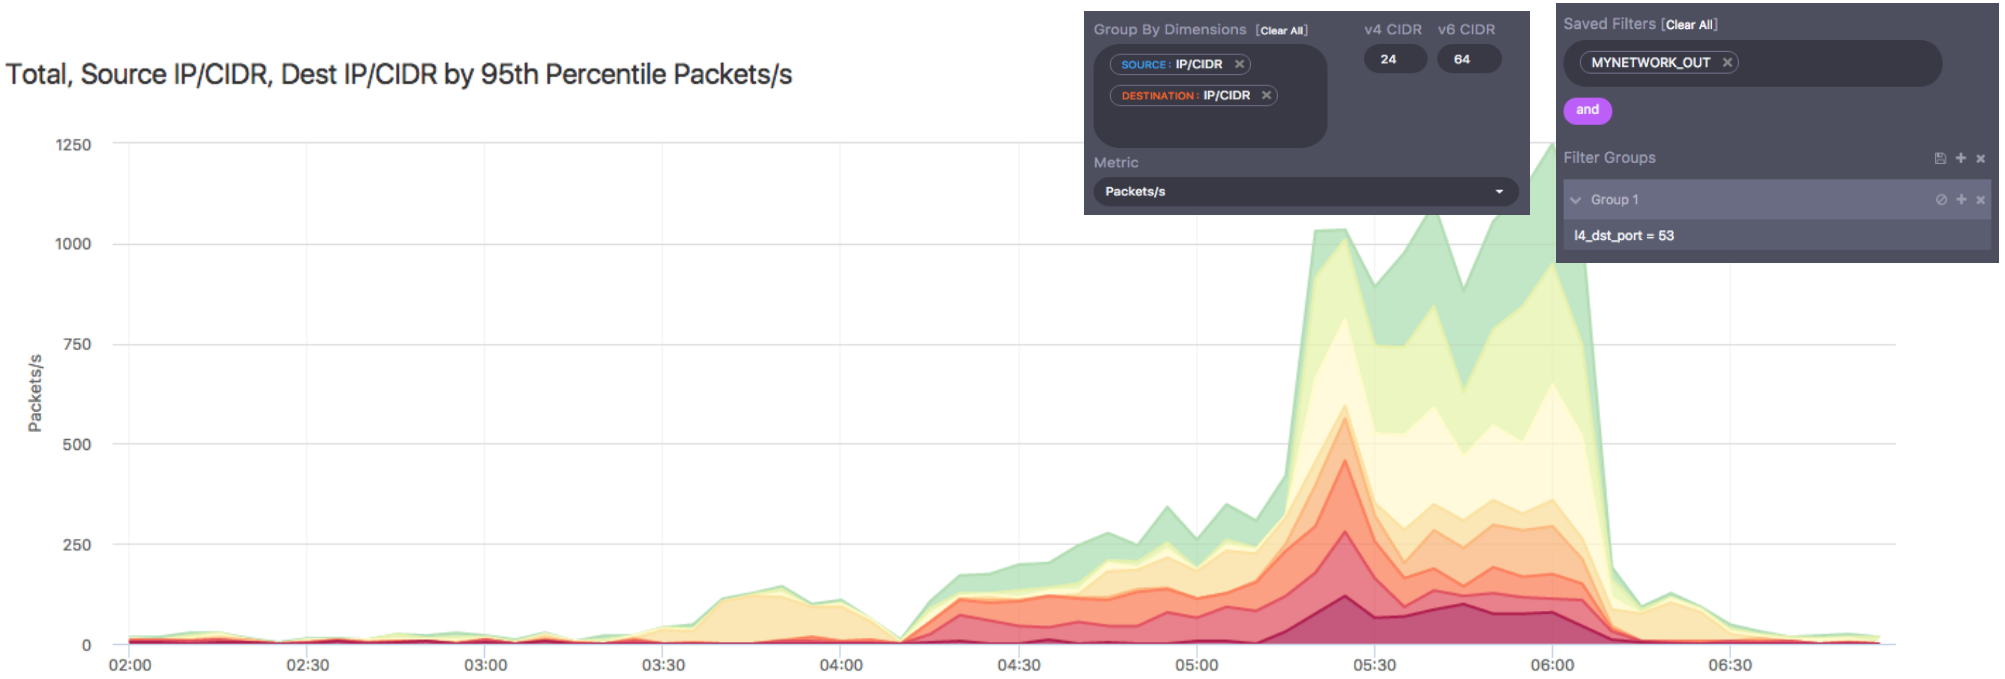

Left +Y Axis

name	Avg pps	95th Percentile	Max pps	Last Datapoint
Total ---- 207.44.124.0/24 (-) ---- 204.13.250.0/24 (ns2.p00.dynect.net)	55	270	300	7
Total ---- 2606:9400:0:e::/64 (-) ---- 2001:500:90:1::/64 (ns1.p00.dynect.net)	47	253	341	17
Total ---- 2606:9400:0:e::/64 (-) ---- 2001:500:94:1::/64 (ns3.p00.dynect.net)	42	239	290	7

kentik |

Use cases of traffic visibility

- Network Planning
- Peering Analytics and Abuse
- Congestion detection
- Is it the network?
- Where on the network?
- Proactive alerting
- Distributed DDoS Detection
- What Changed Post Deploy?
- Security and Breach Detection
- Cost Analytics
- Revenue Identification (New + Risk)
- Enabling Internal Groups

kentik | Tenets

- Infinite granularity storage for months
- Drillable visibility, network specific UI
- Real-time and fast (< 10s queries)
- Anomaly detection + actions
- Open / API
- Scale

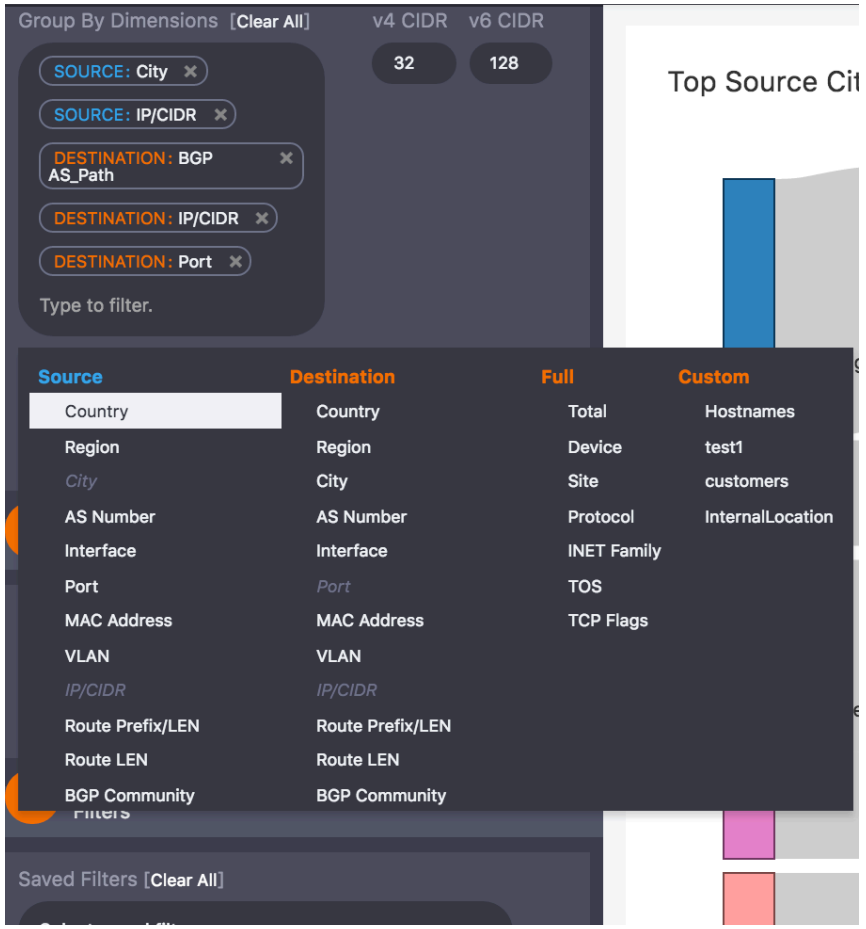

Now we know what we need,
how do we do it?

kentik | Where to find this data ?

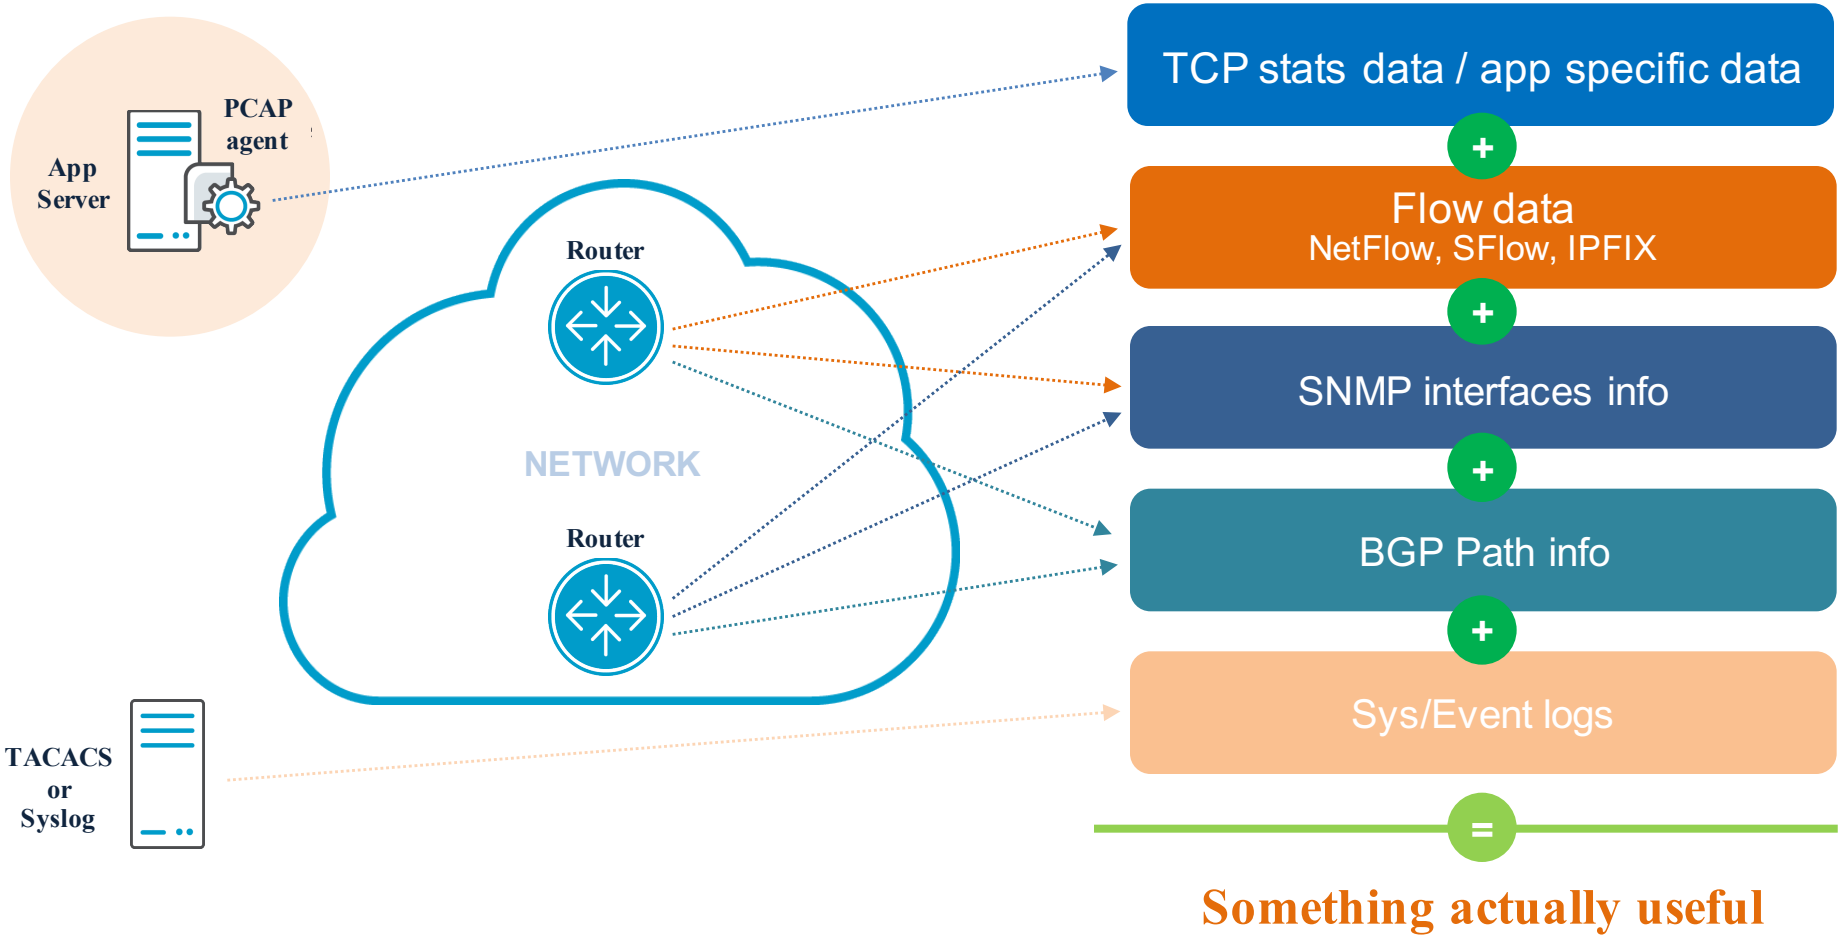

kentik | What kind of tools

- **Current Open Source:** pmacct, ntop, SiLK, cacti
- **Older Open Source:** cflowd, AS-PATH, RRDtool
- **Commercial software:** Arbor, Plixer, SevOne, Solarwinds, ManageEngine

- **DIY Big Data:** Kafka + ELK, Hadoop, druid, grafana, tsdb
- **On-Prem Big Data**: Cisco Tetration, Deepfield…
- **SaaS Big Data:** Kentik, Datadog, Appneta, Splunk

kentik | Many tools gets you almost there

kentik | Really though… (tools)

Open source (ish):

- Pmacct
- Nprobe / Ntop
- Elastic Search + Kibana (ELK)

Commercial:

- Arbor
- Kentik

kentik | Three layer approach

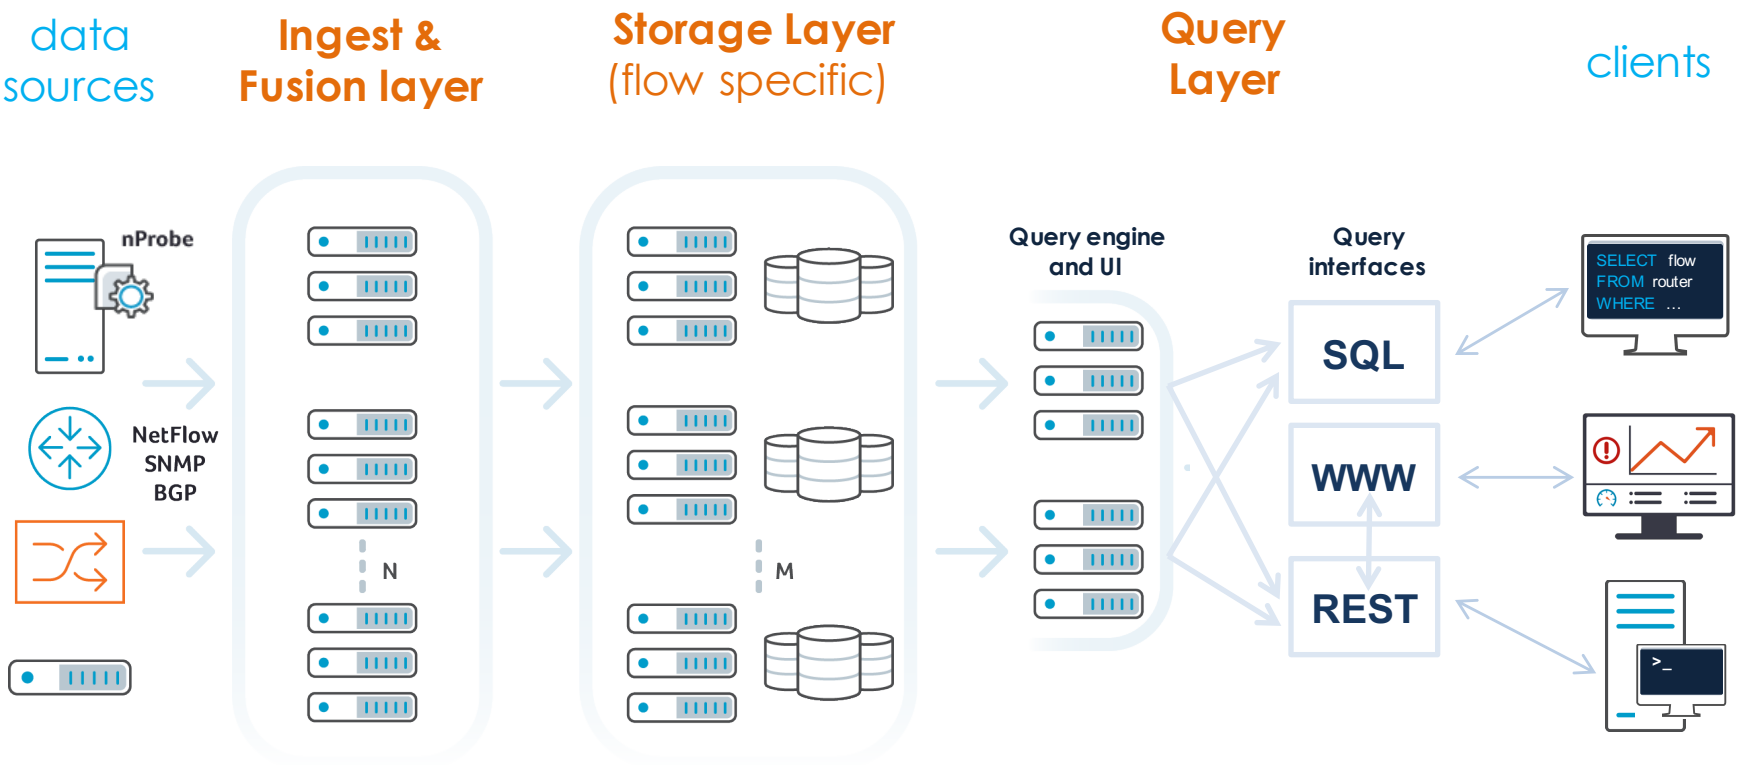

Each layer has separate and different scaling characteristics

kentik

Seriously?

How much data

- **Small network (< 10Gb/s traf.)** 10k flows/sec (+rows/sec)
- **Large network (1 Tb/s traf.)** 500k flows/sec
- **Querying over 30+ days** @ 200k fps (518 B rows, 207 TB) in < 10s

Data fusion

is a key enabler to useful data

kentik | DATA FUSION

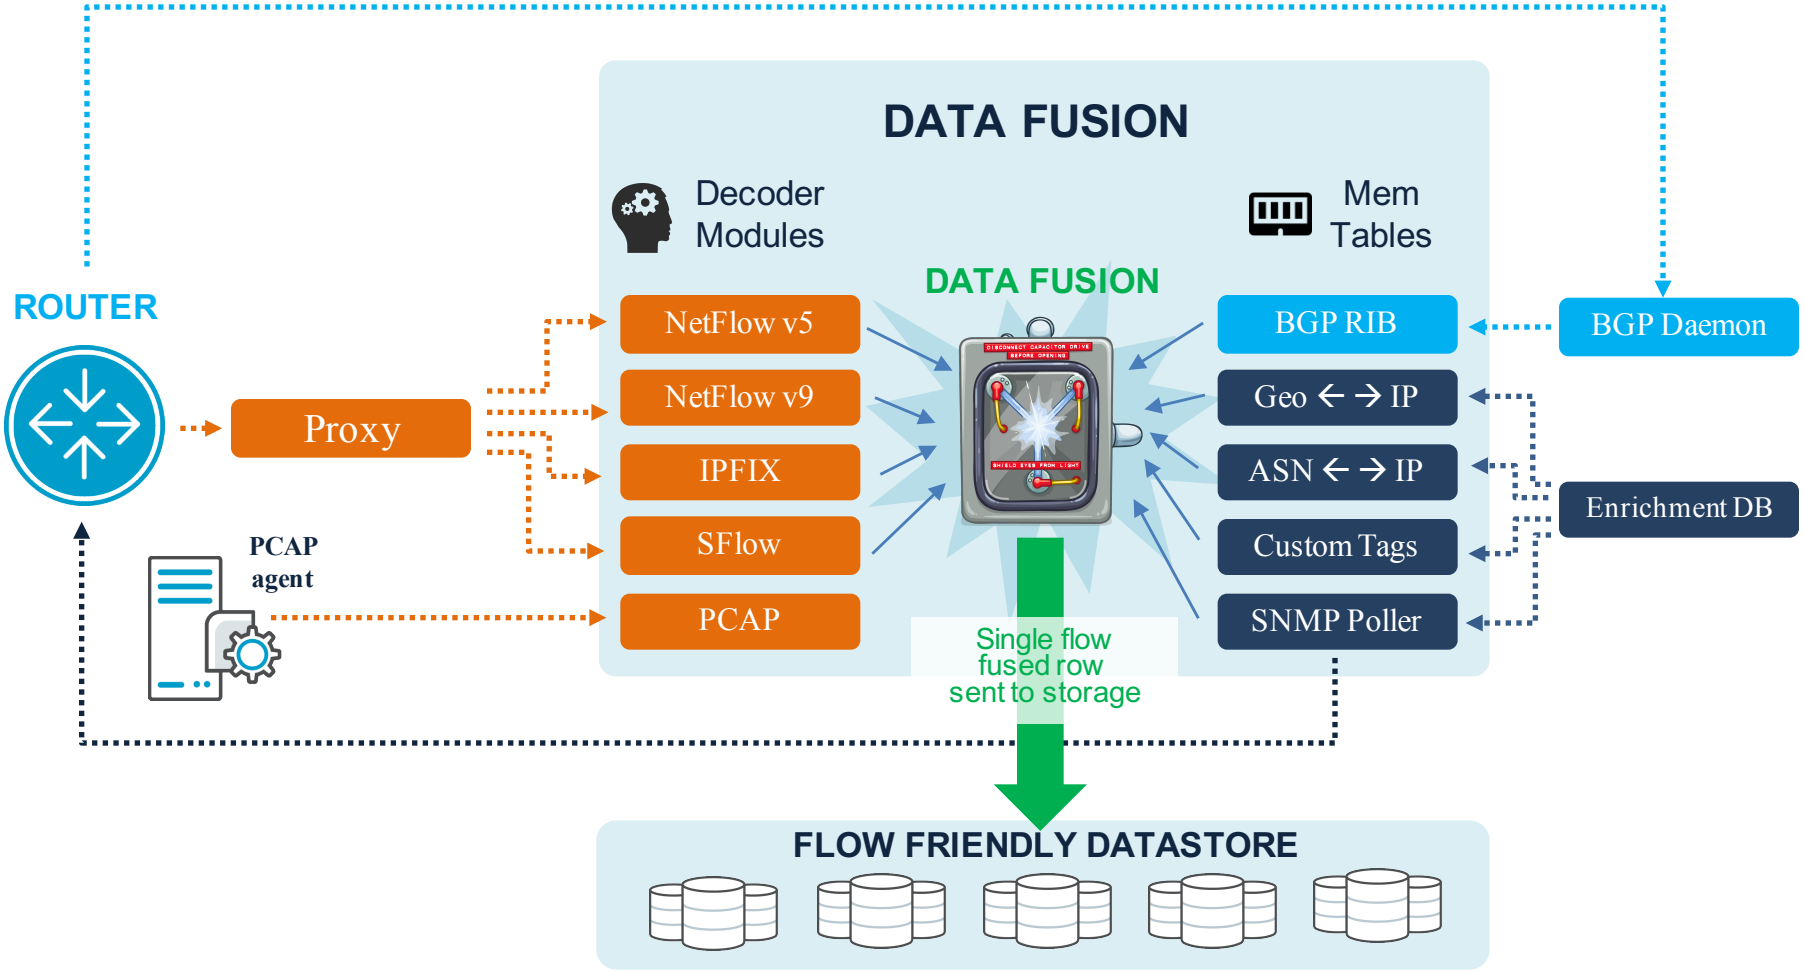

kentik | DATA FUSION

Serious. Flux Capacitor. Is. Serious.

Fusing should be:

near real-time
performed at ingest
data specific

kentik | Network planning: traffic by BGP hop

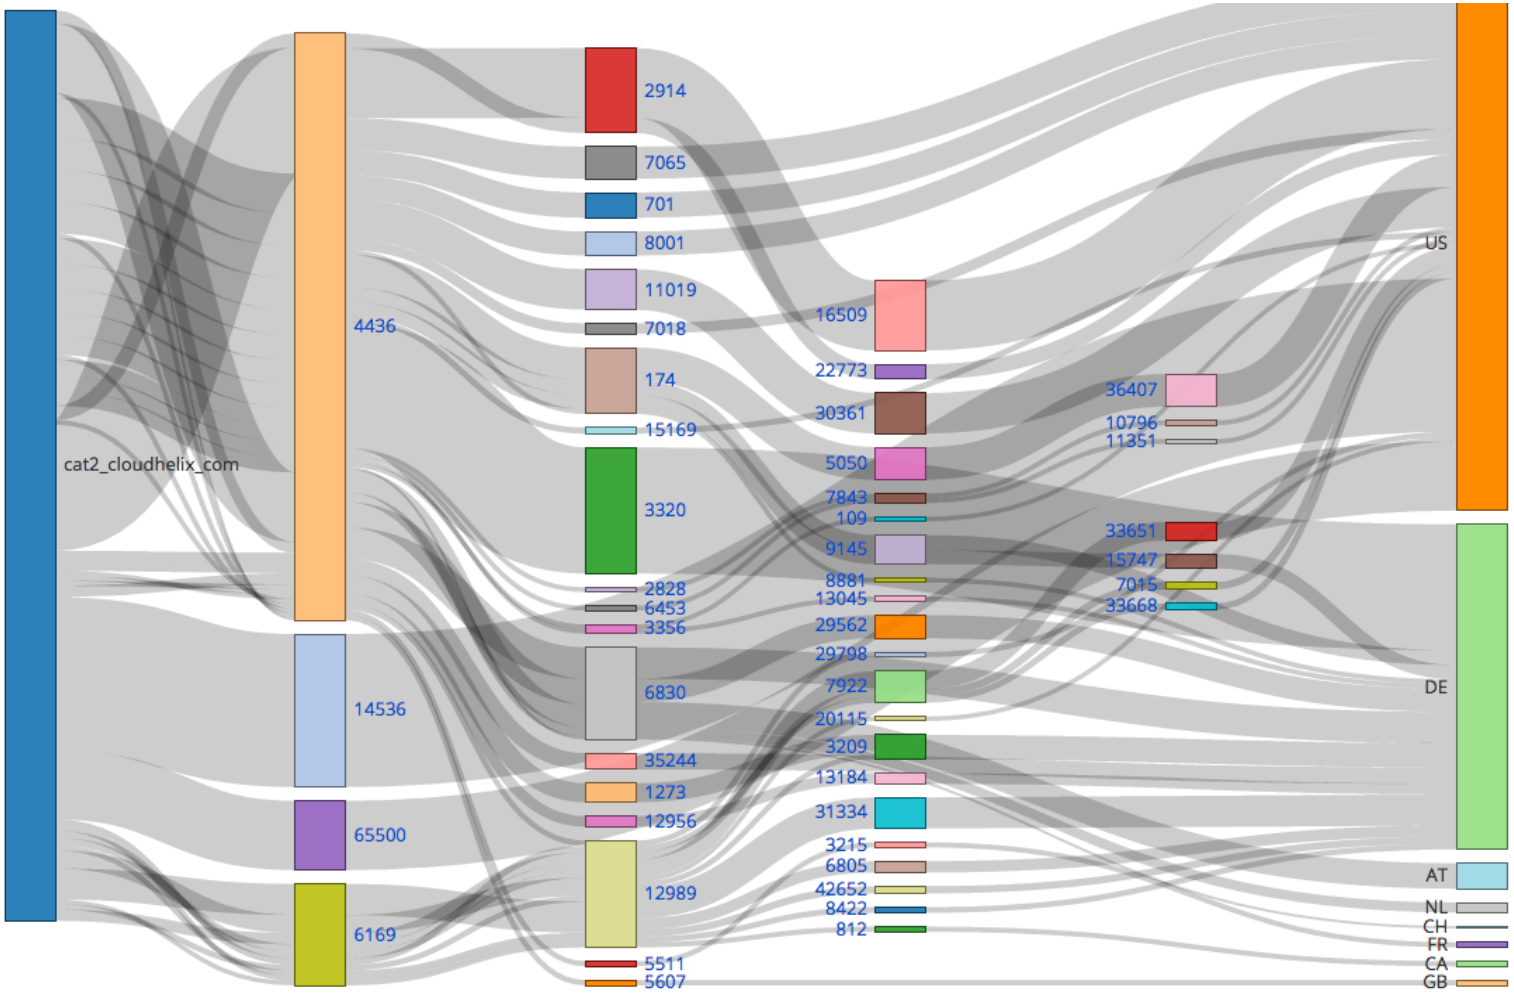

kentik | Network planning: collapsed path, exclude 1st

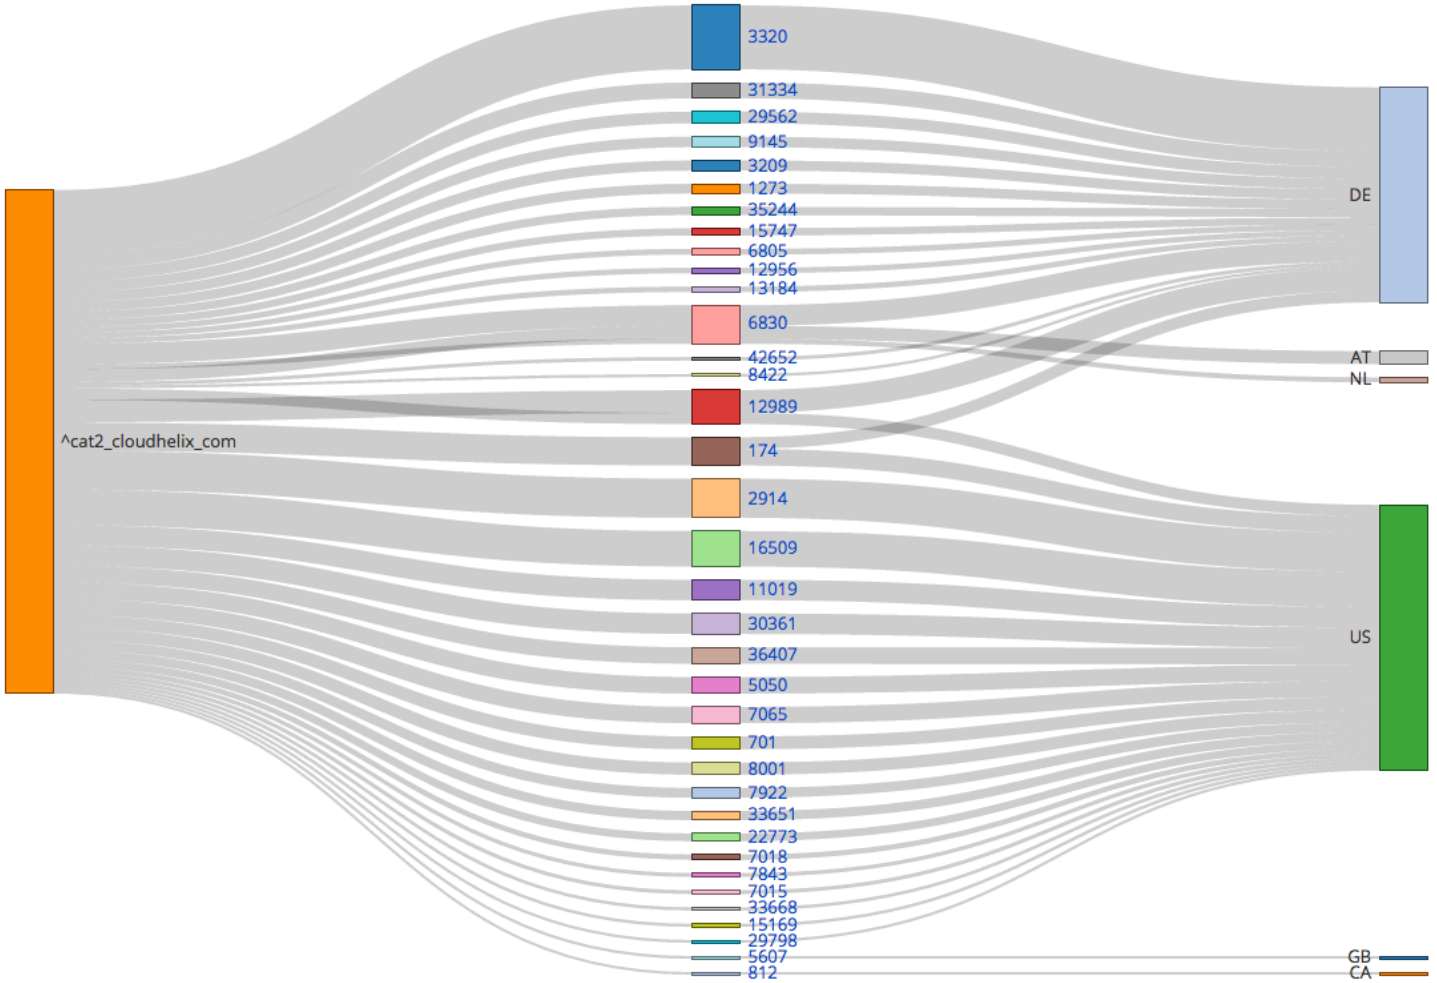

Looking at existing architectures out there

kentik | PMAcct-based **implementation**

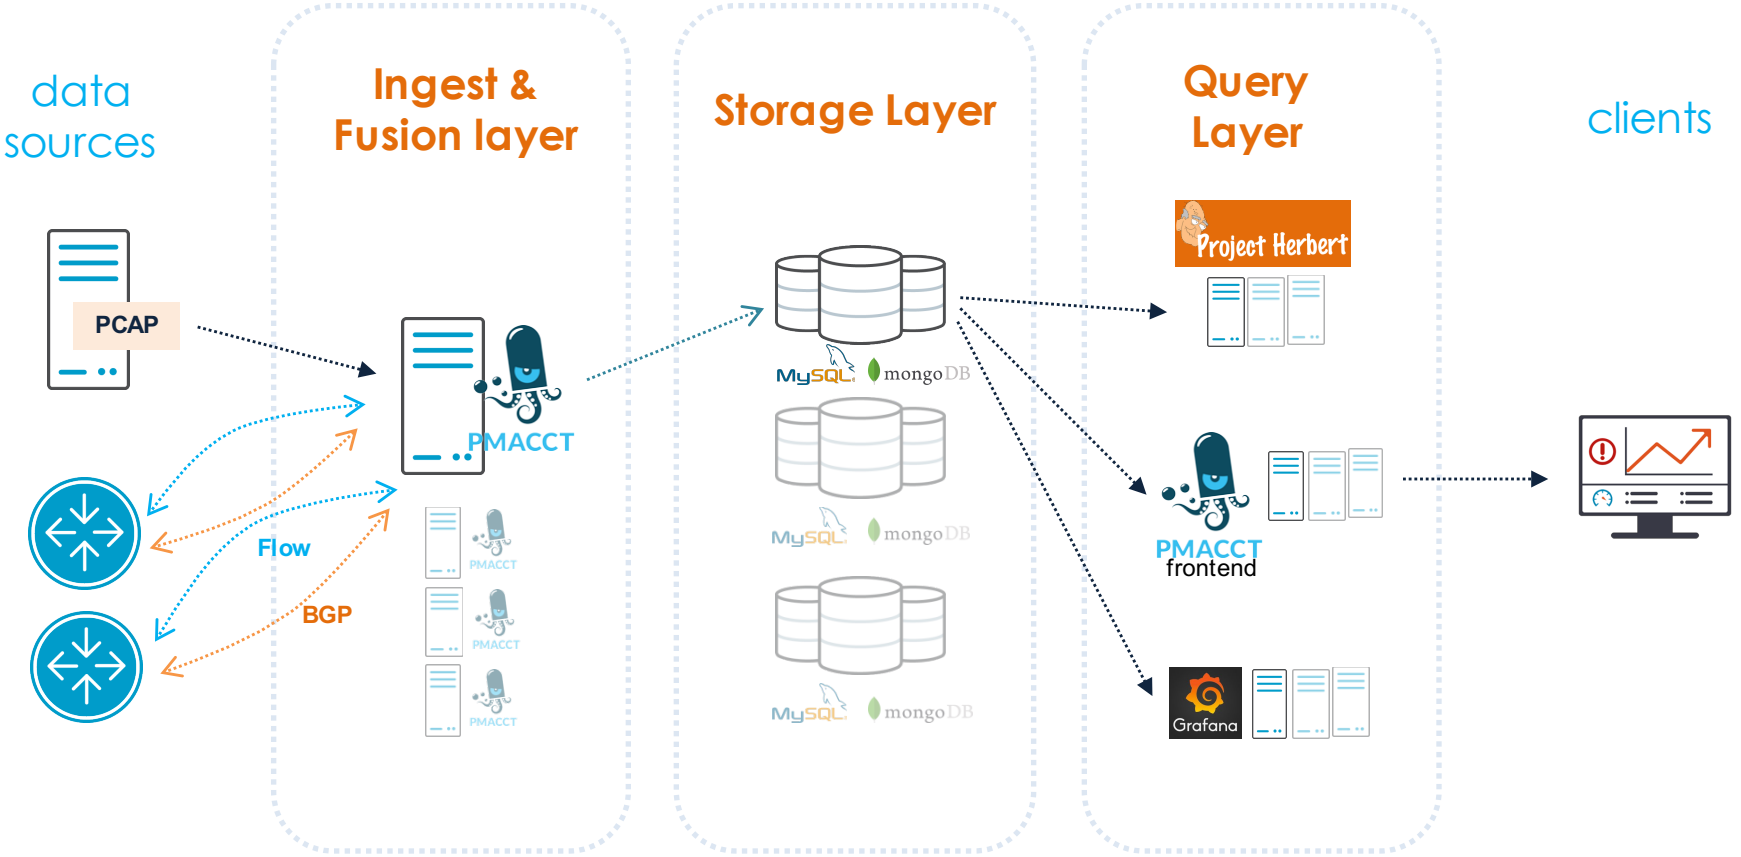

kentik | Nprobe + Ntop + ElasticSearch

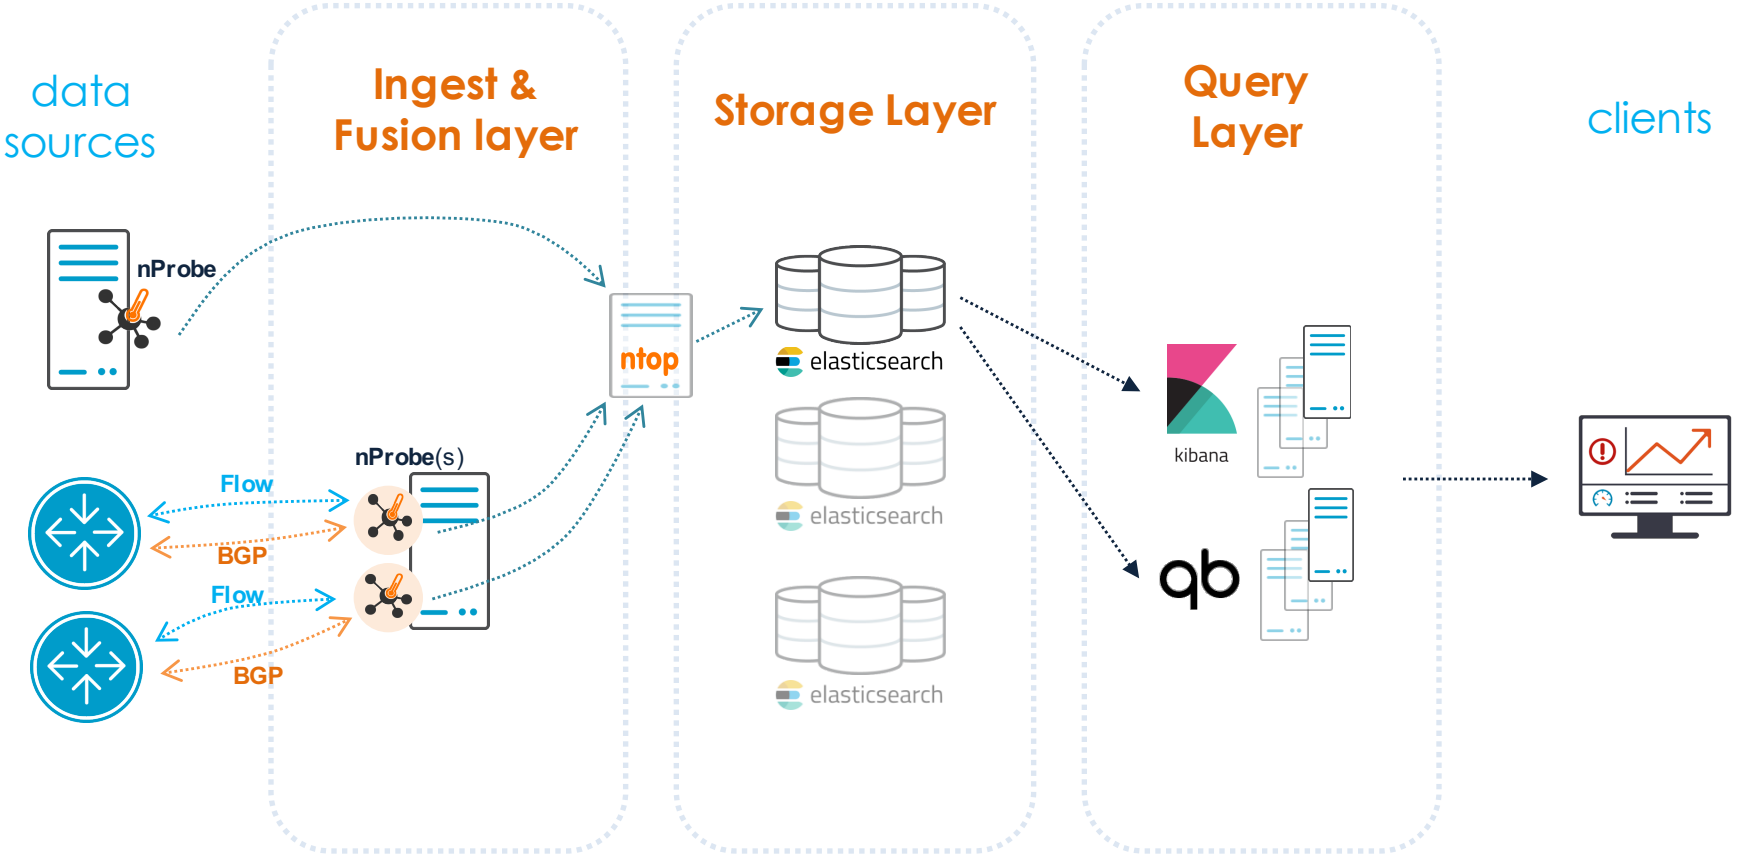

kentik | Kafka + Elastic Search

- Dropbox implementation of a (mostly) open-source NetFlow solution here: [Dropbox blog]

- Requires custom ingest, fusing, UI

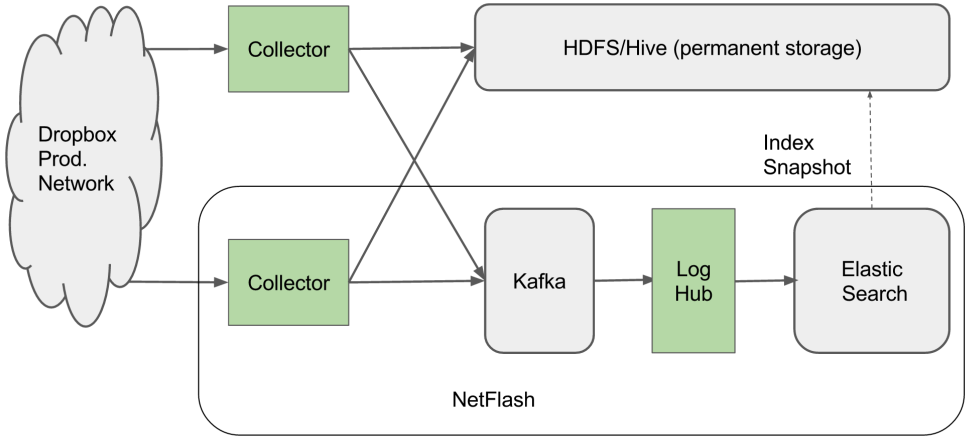

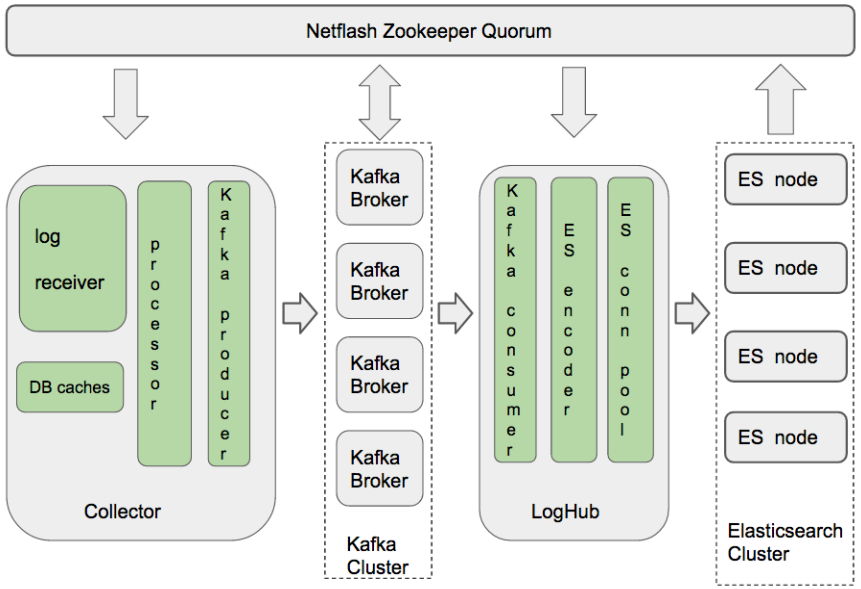

kentik | Caveats

- **Ingest:**
  Distributing and scaling (1xNProbe = 1xDevice)
  No SNMP (= no IF info available for fusion)
  Aggregation (no infinite granularity)
- **Data-store:**
  Challenging at scale when ES
  very hard for MySQL/MongoDB
- **Query frontends very generic:**
  Tailoring of meaningful dashboards difficult
  No anomaly detection

Commercial HW solutions (Arbor)
Appliance based
not truly distributed
pre-determined list of aggregated data (no infinite granularity)

And so…

kentik | Why isn't everybody already doing it?

Required areas of expertise

(because every presentation needs a Vin diagram)

Train all the other teams on the involved network protocols and their usage

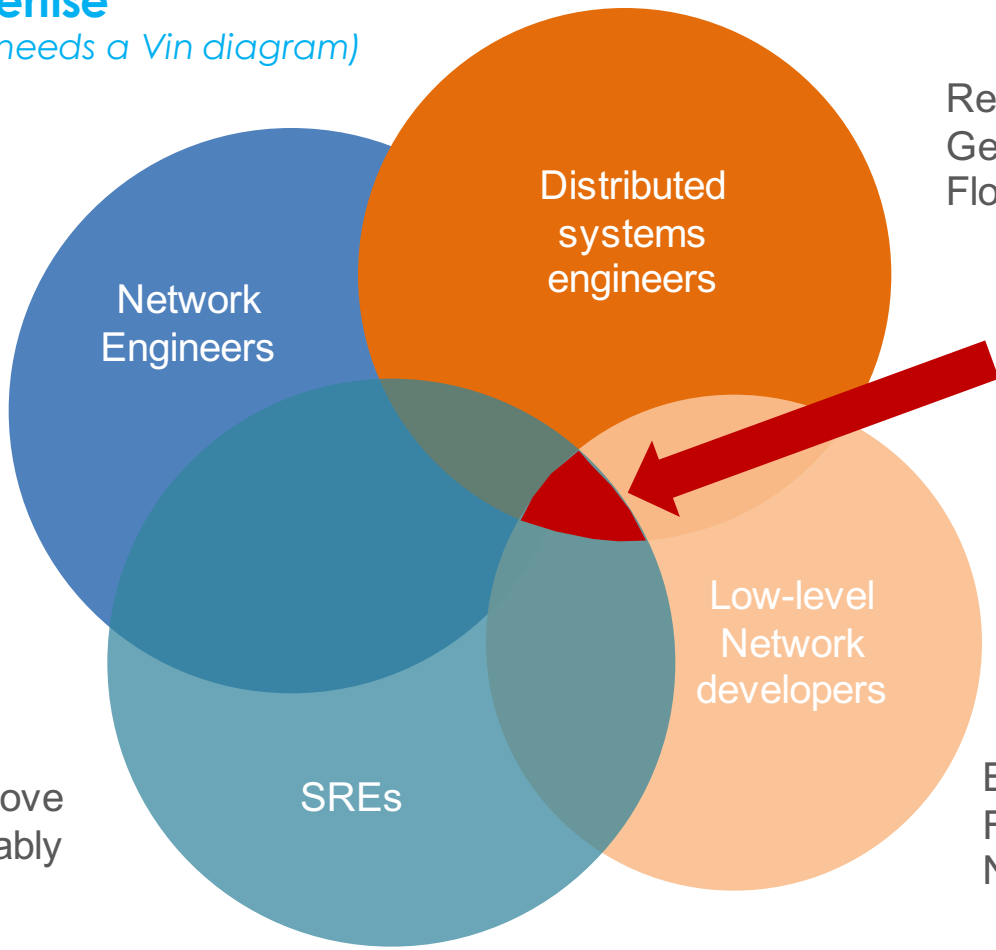

Resilience / Reliability
Geo-distributed ingest
Flow friendly data-store

Make all of the above work reliably

BGP Daemon
Flow inspection & conversion
Network protocols hacking

Looking beyond the basics

kentik | # Building on top of the platform

Once you have a platform, what's next?

- Augmented flow (retransmits, latency, URL, DNS)
- Anomaly detection
- Multi-hop exit determination
- BGP-path congestion detection

kentik | Building on top of the platform

Imagine if we could get performance data from the network:

- Q Depth
- Retransmits per flow
- TCP latency
- Application Latency

You can:

- Nprobe (ntop) collects Latency, Rxmits, URL, DNS -> IPFIX flow
  - Deploy on a host or a sensor

- Cisco, Juniper, Arista working to expose Q Depth into flow

kentik | Retransmits enhanced flow: rexmits / interface

Top Dest Interface by Average % Retransmits (Beta)

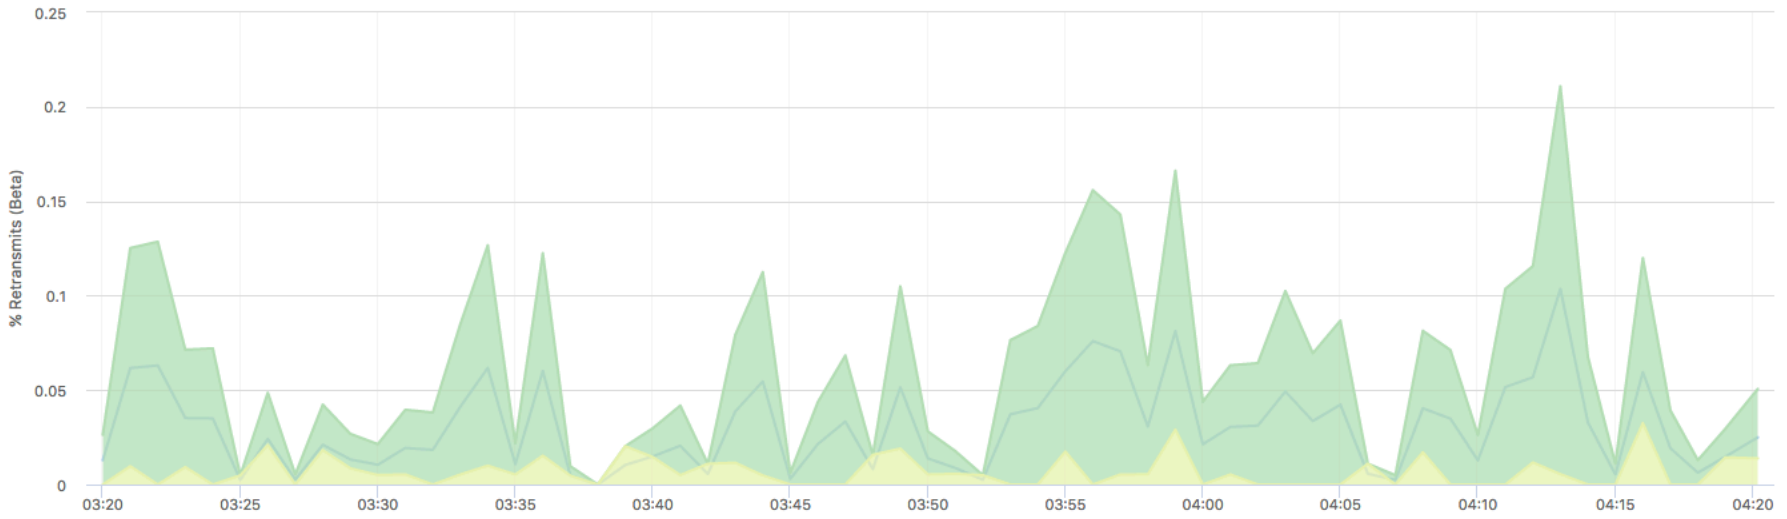

Left +Y Axis

	% Retransmits				Retransmits/s				Traffic			
key	Avg	p98th	Max	Last Datapoint	Avg/sec		p98th	Max/sec	Last Datapoint	Avg pkts/s	p98th pps	Avg mbps
Total	0.030	0.079	0.104	0.025	71.255		169.600	236.800	44.800	235,921	273,776	1,319.47
--- : --- (9277)	0.055	0.147	0.205	0.037	63.173	(88.7%)	163.200	230.400	32.000	114,877	133,549	670.74
--- : --- (0)	0.007	0.025	0.032	0.014	8.083	(11.3%)	28.800	32.000	12.800	121,045	140,589	648.73

kentik | Retransmits enhanced flow: rexmits / ASN

Total, Dest AS Number by Average % Retransmits (Beta)

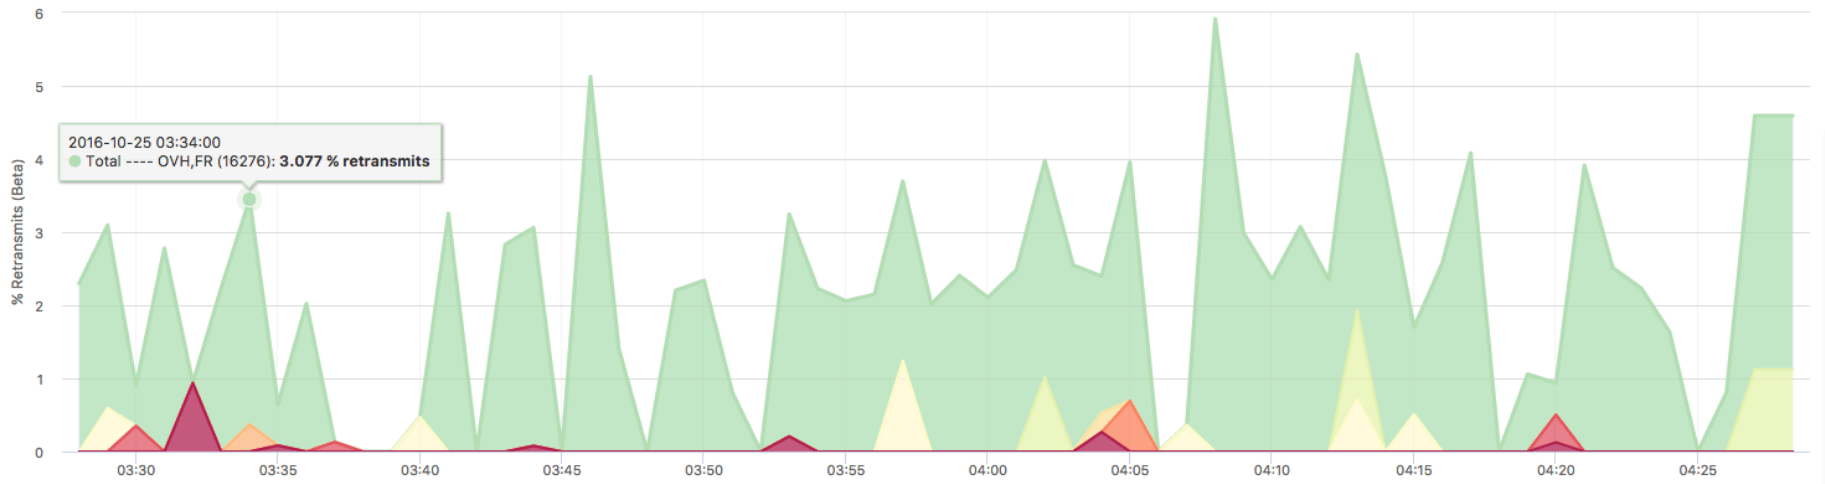

Left +Y Axis

	% Retransmits				Retransmits/s				Traffic		
name	Avg	p98th	Max	Last Datapoint	Avg/sec	p98th	Max/sec	Last Datapoint	Avg pkts/s	p98th pps	Avg mbps
Total ---- OVH,FR (16276)	2.397	4.605	5.917	3.468	55.582	150.400	217.600	76.800	2,319	6,886	18.51
Total ---- ATLANTIC-NET-1 - Atlantic.net, Inc.,US (6364)	0.150	1.124	1.220	1.124	0.431	6.400	6.400	6.400	289	1,350	1.34
Total ---- LATISYS-ASHBURN - Latisys-Ashburn, LLC,US (29944)	0.131	0.657	1.240	0.000	0.955	9.600	19.200	0.000	729	2,643	6.01

kentik | Retransmits enhanced flow: TCP latency / ASN

Top Dest AS Number by Average RTT/2 Client Latency (ms) (Beta)

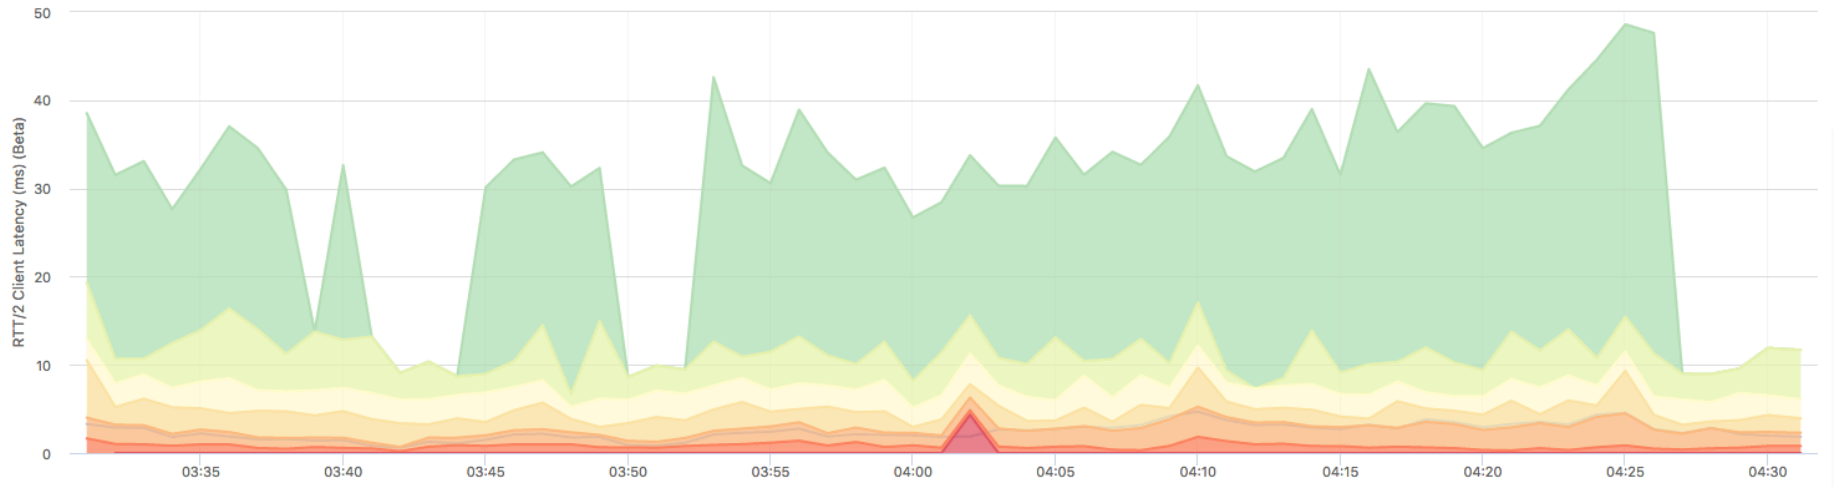

Left +Y Axis

key	Avg Latency (ms)	p98th Latency (ms)	Max Latency (ms)	Last Datapoint	p98th mbps	p98th pps
Total	2	4	5	2	1,494.62	262,166
AMAZON-02 - Amazon.com, Inc.,US (16509)	23	34	36	36	5.08	2,384
YAHOO-NE1 - Yahoo,US (36646)	4	7	9	6	20.79	4,173
YAHOO-3 - Yahoo!,US (26101)	3	4	4	2	28.63	6,502
RUBICONPROJECT - The Rubicon Project, Inc.,US (26667)	2	5	7	2	74.97	8,605

kentik | Building on top of the platform

You shouldn't have to stare at dashboards or watch logs to detect badness

Monitor top-x of any dimension combination (IP, ASN's, Geo, Interface)

Create baselines based on time of day

Be able to look at things beyond pps/bps such as retransmits, latency, logs

Detect shifts: did an ASN or IP on a particular interface suddenly move from top-x #200 to #2 and that is unusual for this time of day

This is available today (Open Source: Hadoop, Spark, Storm, Samza, Flink)

kentik | **Use case:** anomaly detection

Traffic from one ASN (network) unusually high. Operator notified at red line.

Unusual High BPS traffic from ASN per site (210)	AS_src: MICROSOFT-CORP-MSN-A... i_device_site_name: Chicago, IL	major Matches: 5	**757.22 Mbps**	488.28 Mbps CALCULA...
Unusual High BPS tra... CLEAR -> ALARM	AS_src: MICROSOFT-CORP-MSN-A... i_device_site_name: Ashburn, VA	major Matches: 17	**2,960.19 Mbps**	1,905.01 Mbps CALCULA...

15 Days Graph

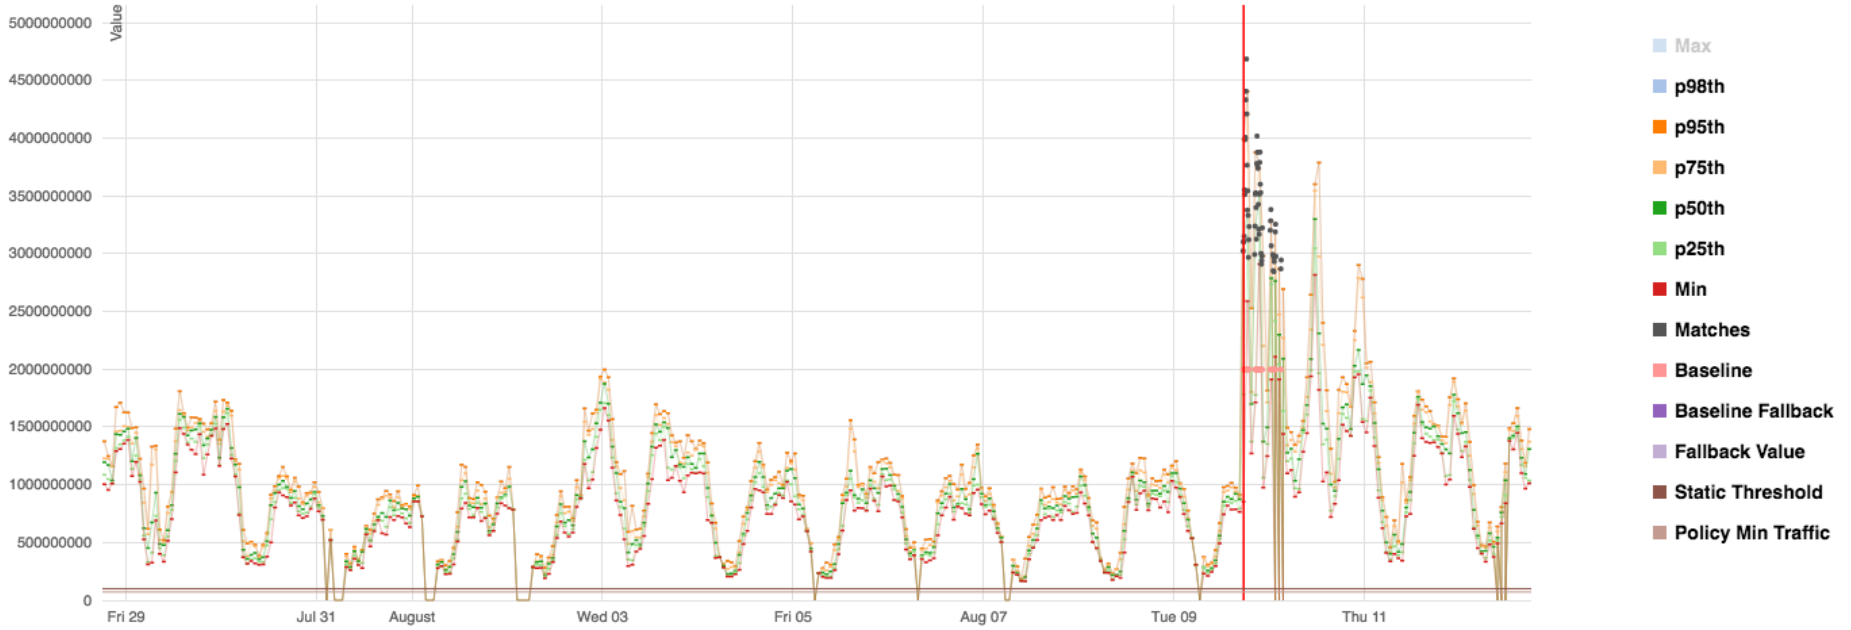

kentik | **Use case:** traffic anomaly detection & annotation

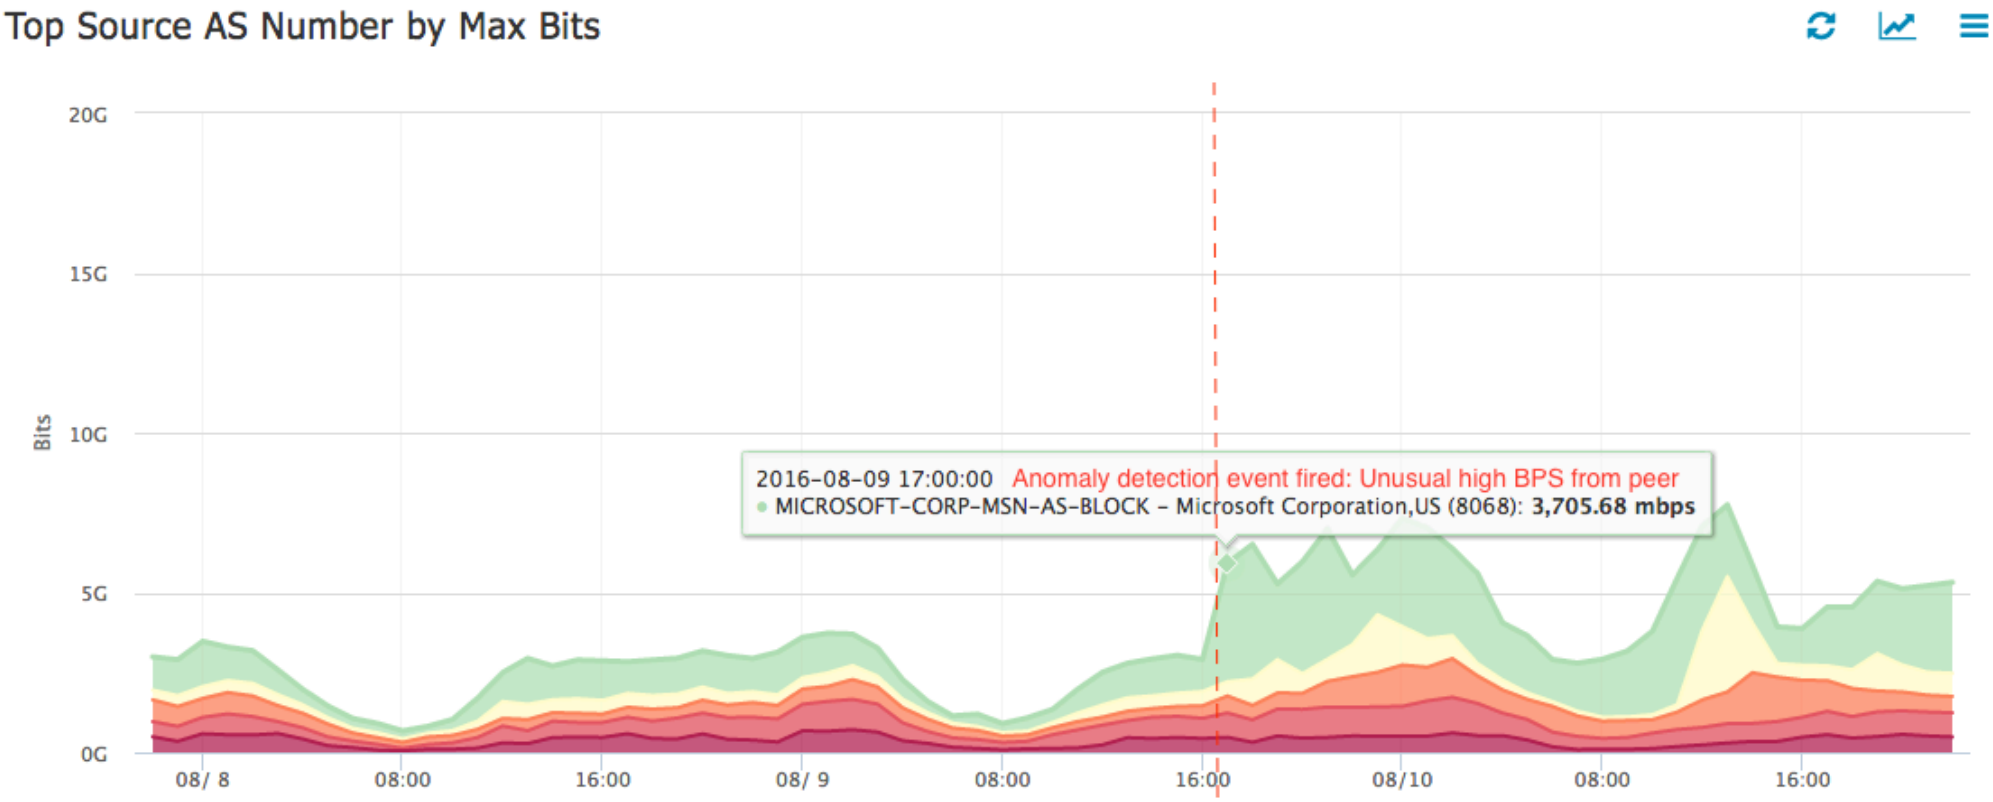

kentik | **Use case:** traffic annotated w/ multiple events

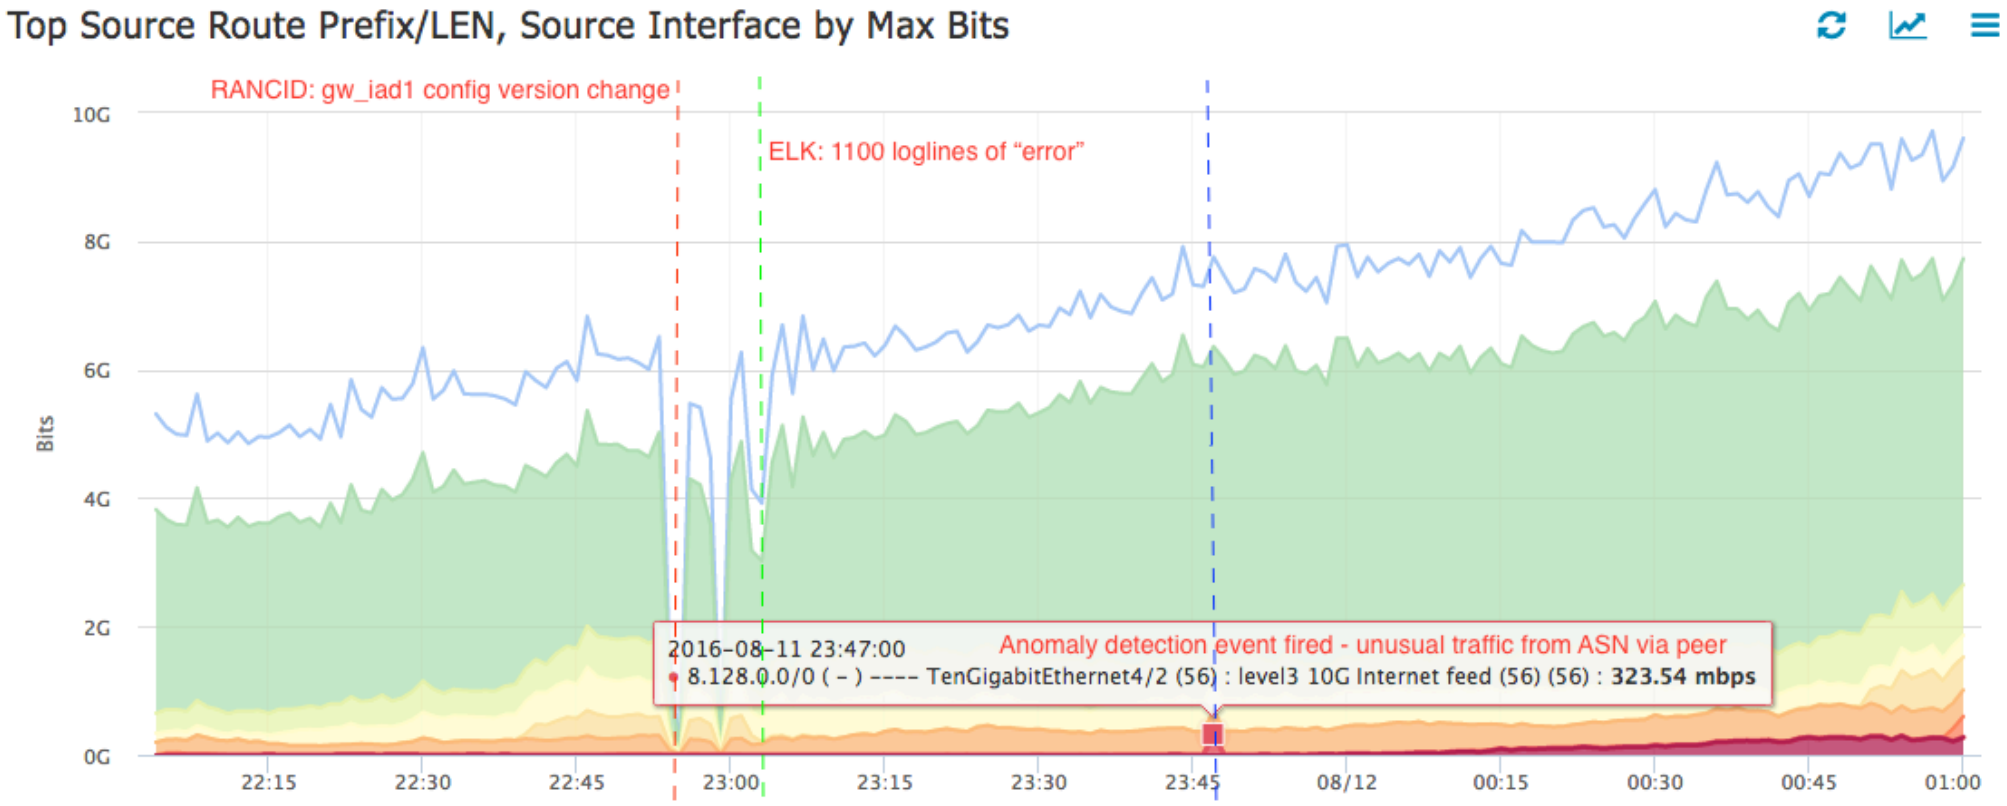

kentik | Anomaly detection: DDoS detection & characteristics

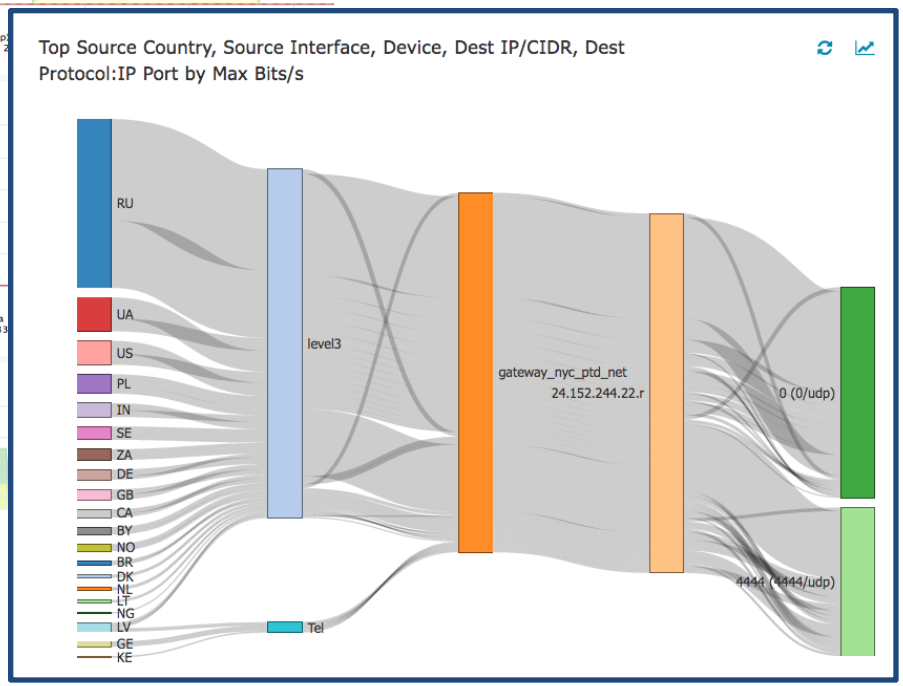

UDP_BADPORTS **CLEAR -> ALARM**	IP_dst: **70.44.101.7**	major Matches: 9	**2,445.91 Mbps** 8,370 unique_src_ip	381.47 Mbps DEFAULT...
UDP_HIGHBPS **CLEAR -> ALARM**	IP_dst: **70.44.101.7**	major2 Matches: 9	**3,914.95 Mbps** 10,340 unique_src_ip	95.37 Mbps DEFAULT...

kentik | Building on top of the platform

Once you have a platform, what's next?

✓Augmented flow (retransmits, latency, URL, DNS)

✓Anomaly detection

❑ **Multi-hop exit determination**

Challenging to map traffic from ingest to exit point, multi-hop

❑ **BGP-path congestion detection**

Detect individual congested paths within a circuit that isn't congested

kentik | Summary

Networks can produce large amounts of data that will make your life easier

Big Data platforms are able to consume this data

Specific tools for Network Operators are beginning to appear (free & paid)

- Paid tools are more specific to network use (UI, easy setup, etc)
- Free tools have the "power" but require cobbling together pieces
- Much work to be done re fusing data such as logs, changes, alerts, DNS

SaaS providers will provide community views and enable data-sharing

QUESTIONS ?

Dan Ellis
dan@kentik.com

CREDITS ntop PMACCT kibana elasticsearch qb Grafana Project Herbert ARBOR NETWORKS Dropbox kentik