When BGP meets Big-Data

THE LINUX FOUNDATION

The Internet is very much ‘alive’

- Millions of BGP events occurring every day
  - 15 Routers Monitored
  - 410 active peers (both IPv4 and IPv6)
  - ~120,000,000 Prefixes Advertised
  - ~950,000 events per day from a single transit peer
  - ~202,000,000 changes per day
  - ~6,000,000,000 changes per month
- How do we extract ‘signal’ from ‘noise’?
- Can we apply techniques from other domains in this pursuit?

The Internet is very much ‘alive’

- If we know the questions we want to ask, how do we ask them?
- Enhance traditional dampening and suppression with analytics

Five Monitoring Points in BGP

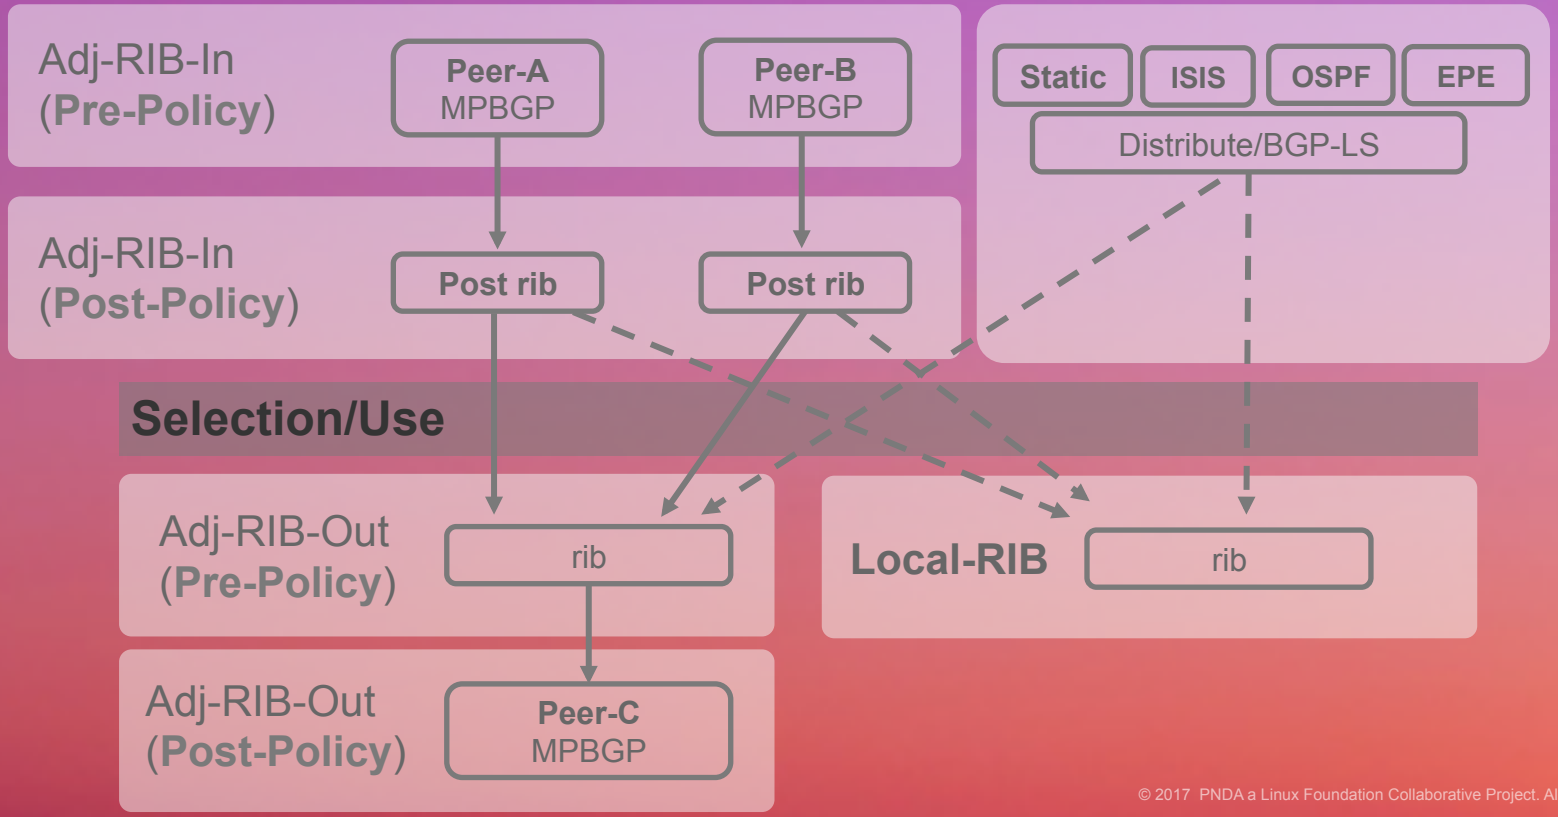

Five Monitoring Points in BGP

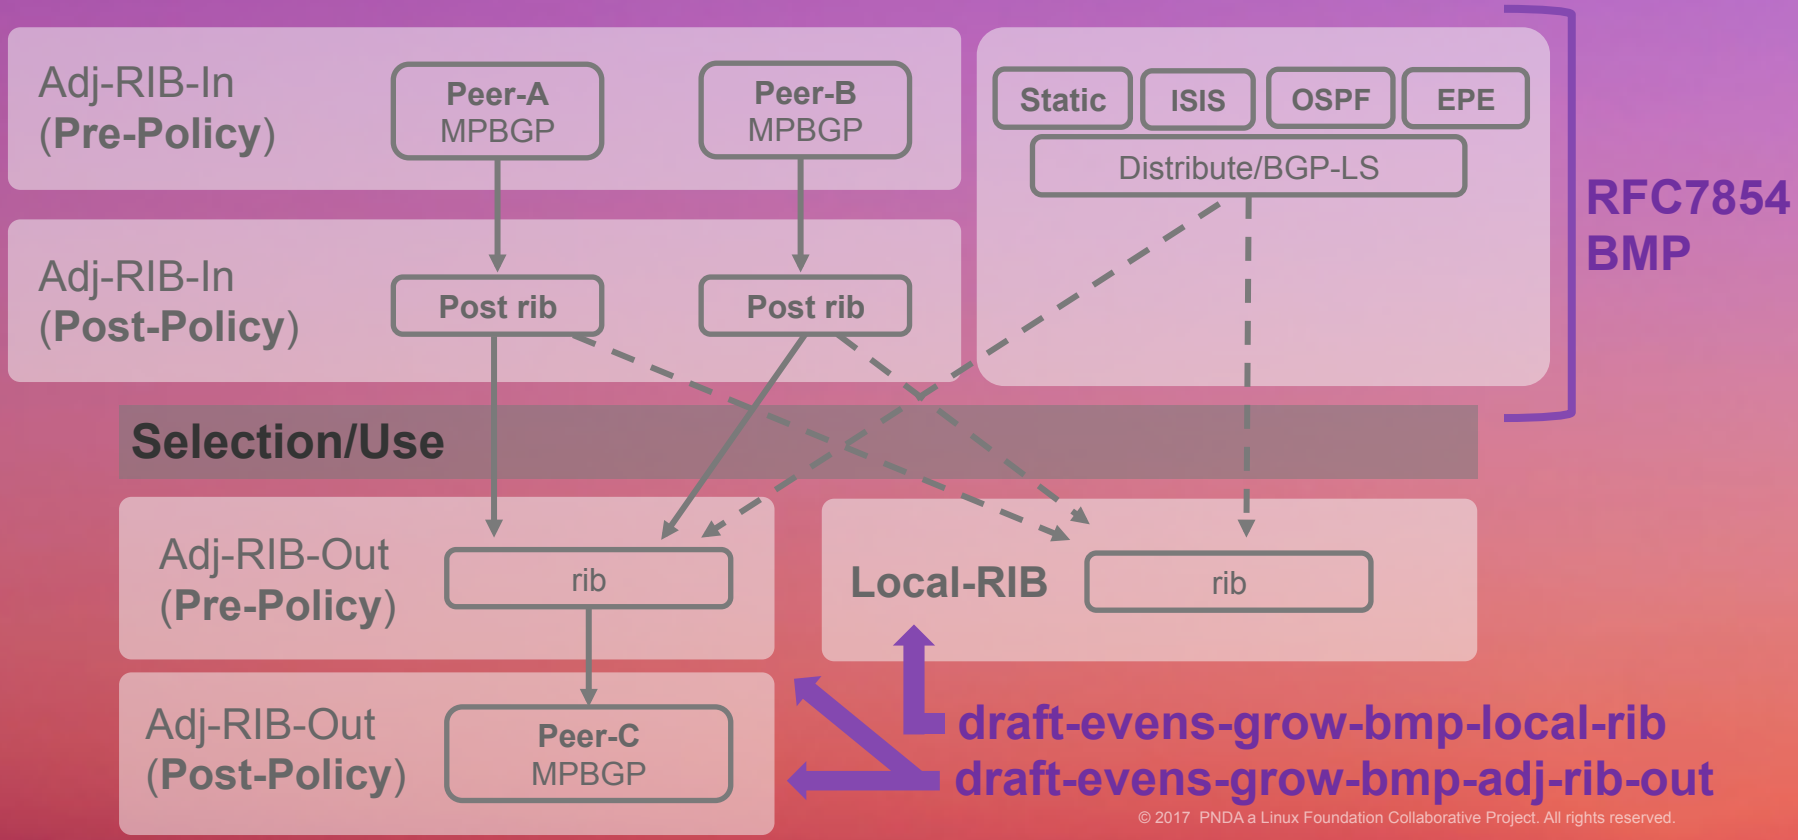

SNAS Architecture

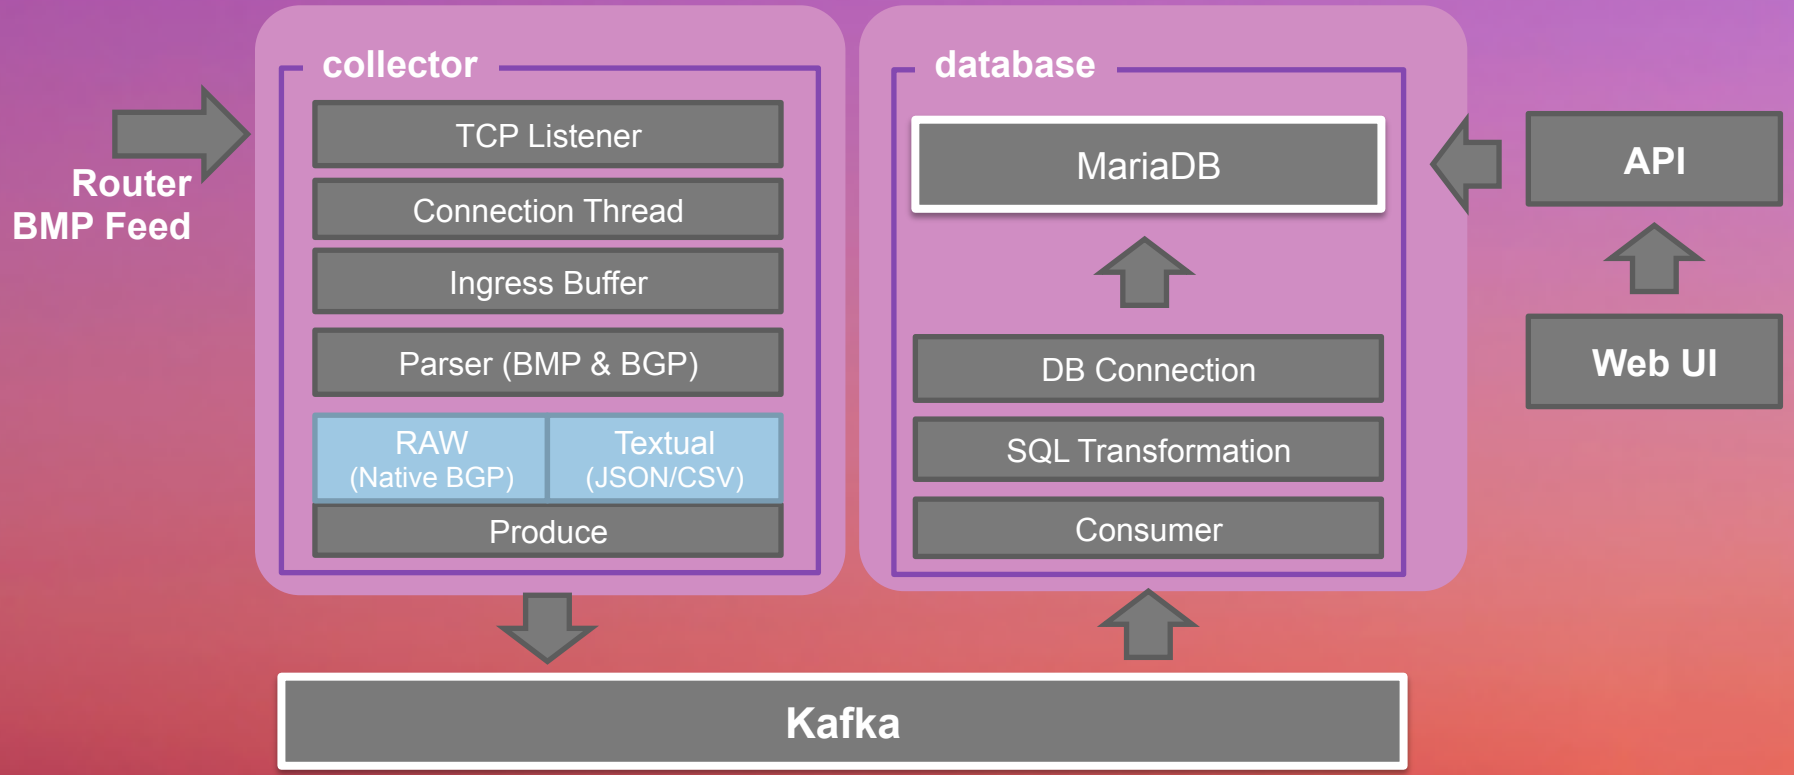

SNAS Architecture

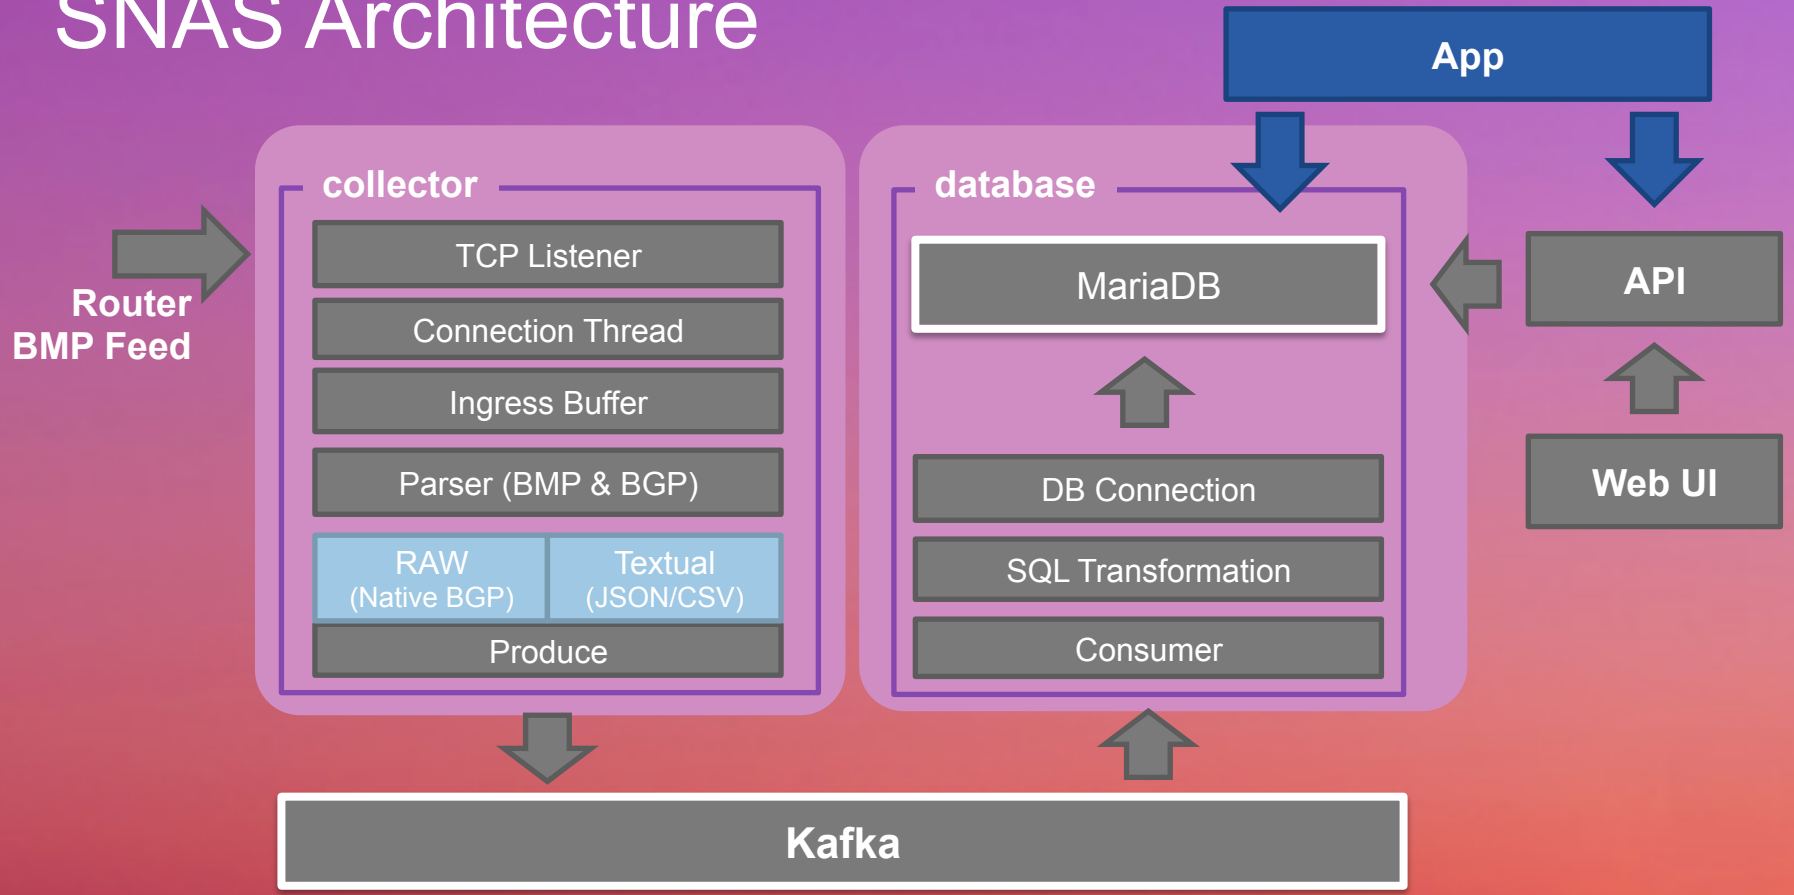

SNAS Architecture

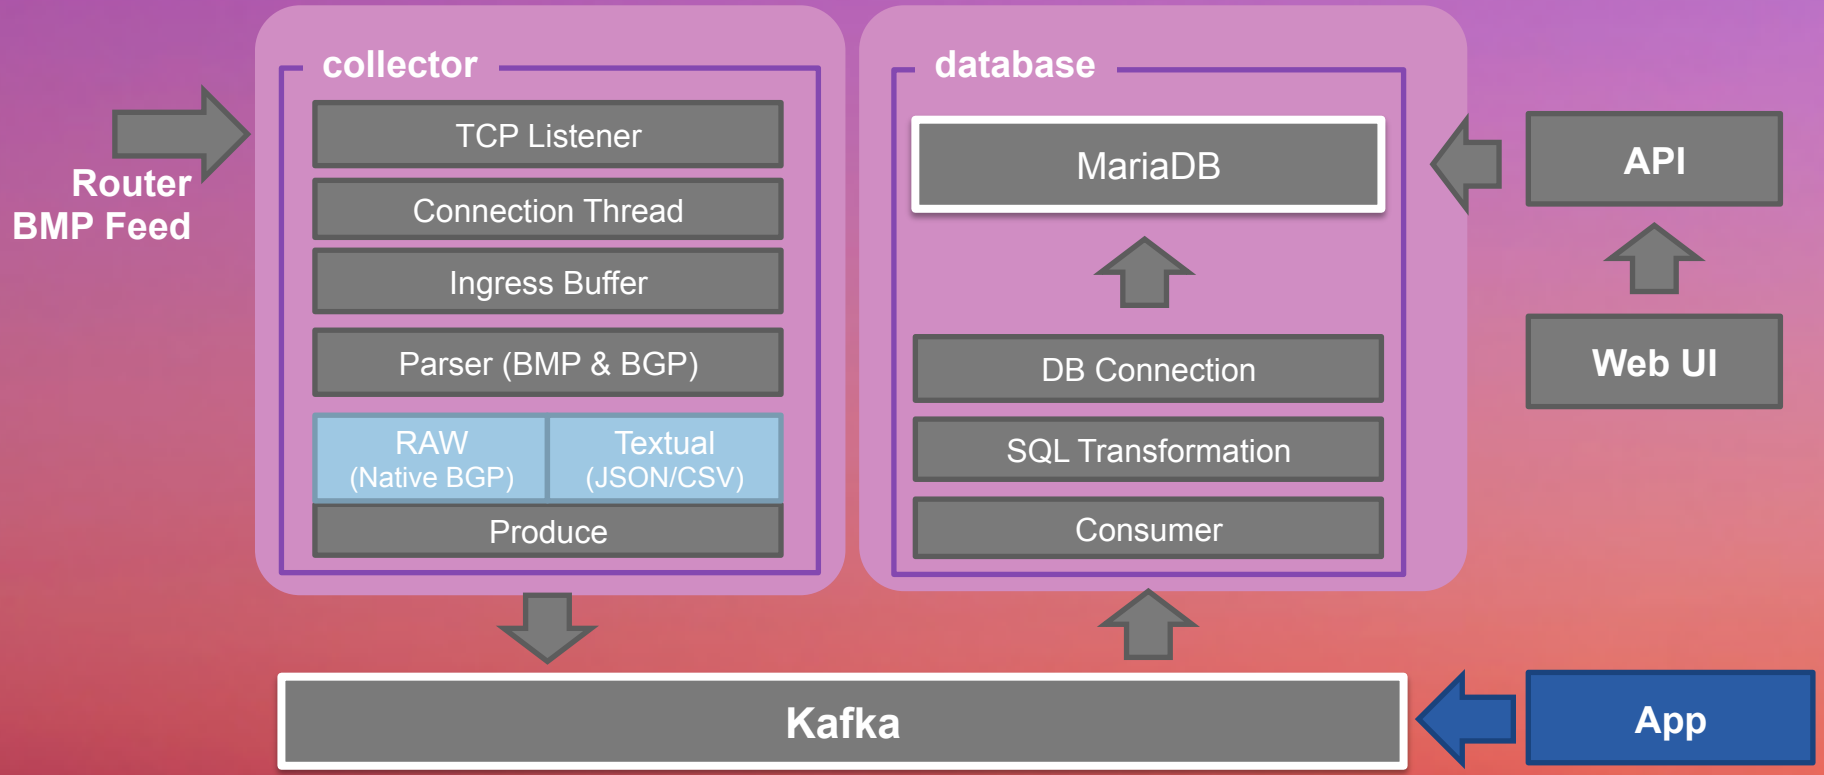

E2E architecture

- Encoding app required to perform ‘avro’ encoding of BMP data
- BGP App runs as Spark batch job, running periodically
- Can be converted to a Spark ‘streaming’ application for near-real-time processing

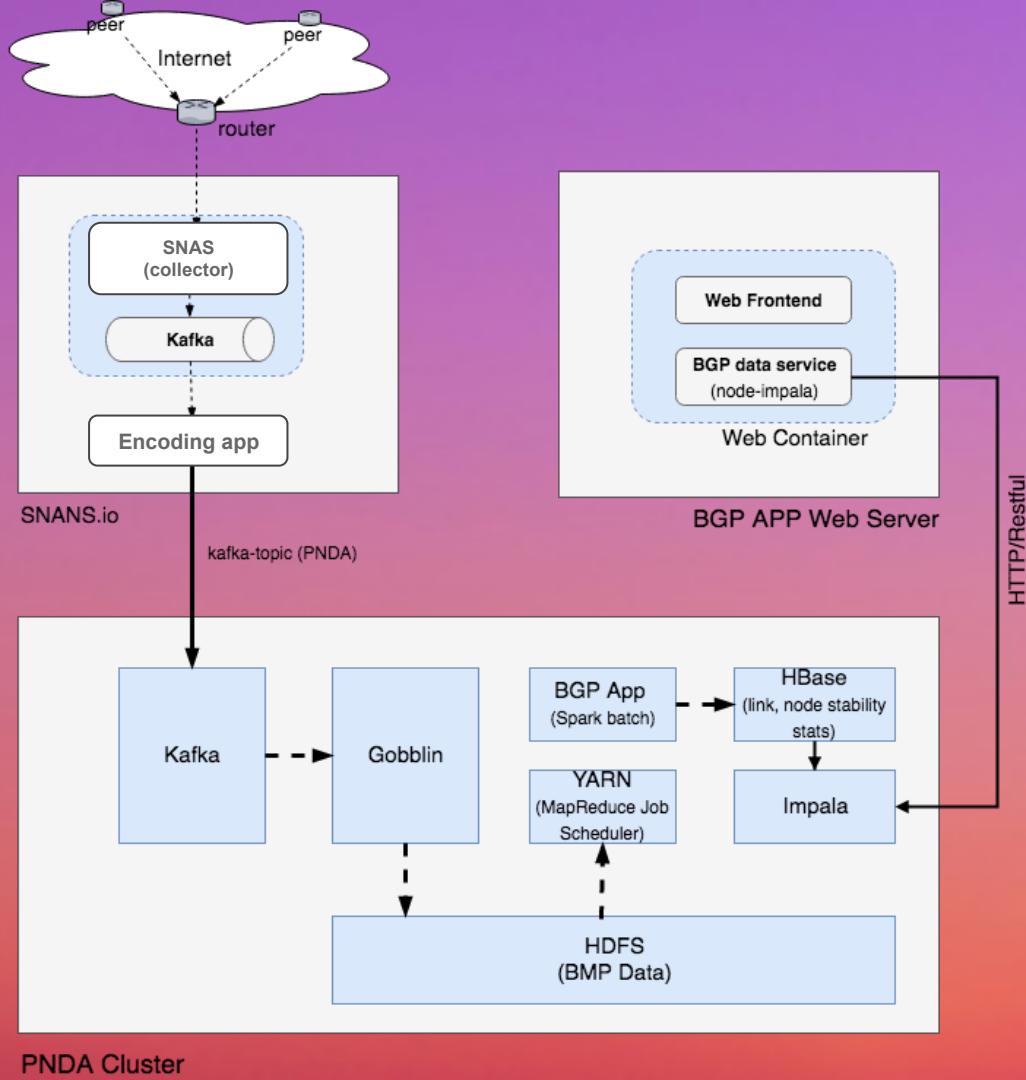

What does this give us?

SNAS.io gives us the ability to record the dynamics of the Internet

PNDA platform enables -

- ‘Raw’ event recording capability, with horizontal scaling (HDFS)
- Run analysis over very large data-sets with parallelism
- Ask questions of the aggregate data about the Internet
- Ask specific question
  - Per-prefix
  - Per-AS
  - Per AS-Path

Top-N analysis

PREFIXES ORIGINATED AND TRANSITTED PER AS

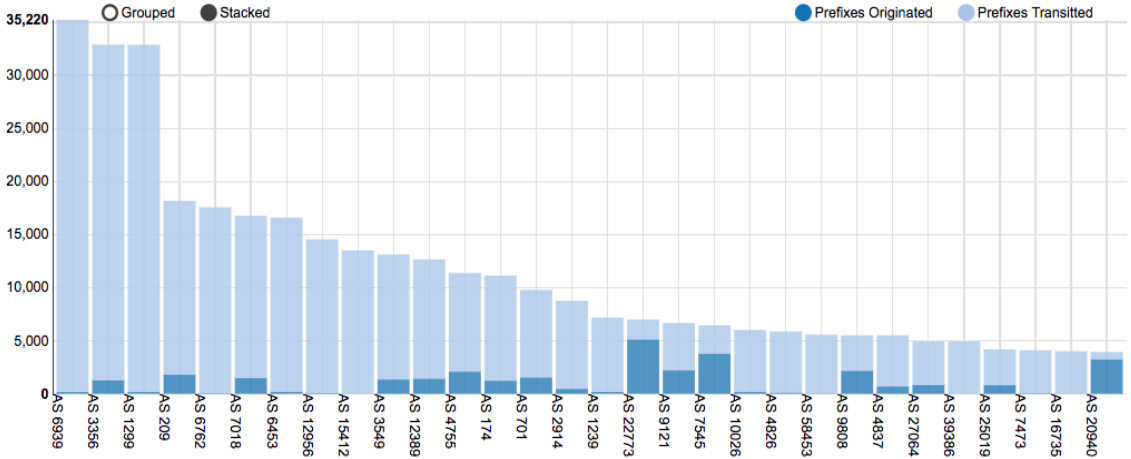

CONTROLS

Results 0 to 30

ASN	Origins	Routes	Change c...
6939	181	35220	35220
3356	1297	32878	35481
1299	203	32874	35242
209	1799	18168	18274
6762	77	17571	18613
7018	1500	16764	16788
6453	210	16589	17054
12956	93	14535	14947
15412	11	13516	13988
3549	1354	13121	13812

TOTAL ITEMS: 30

Path stability

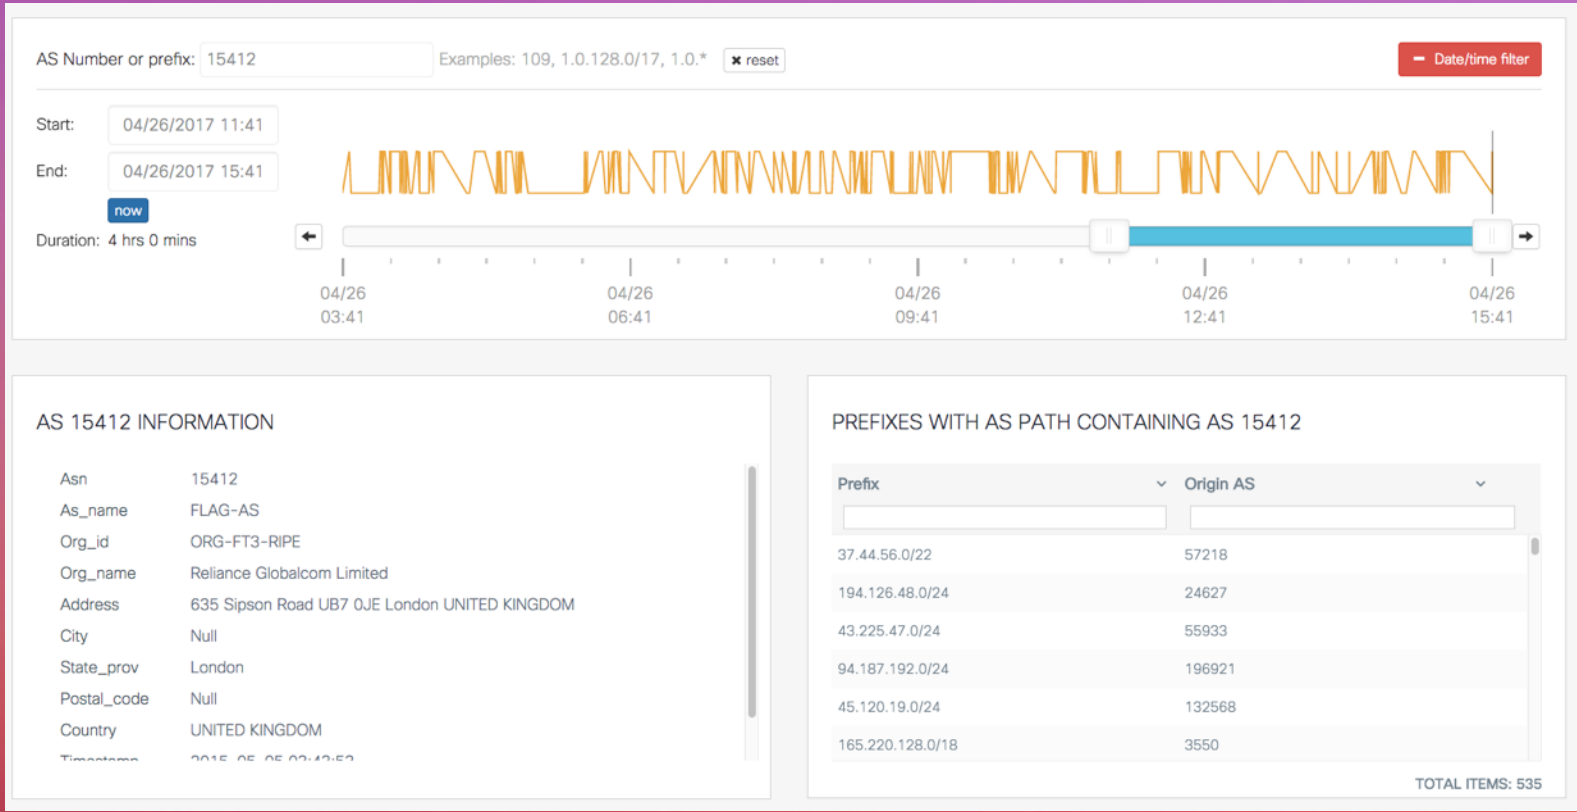

AS Connectivity - FLAG

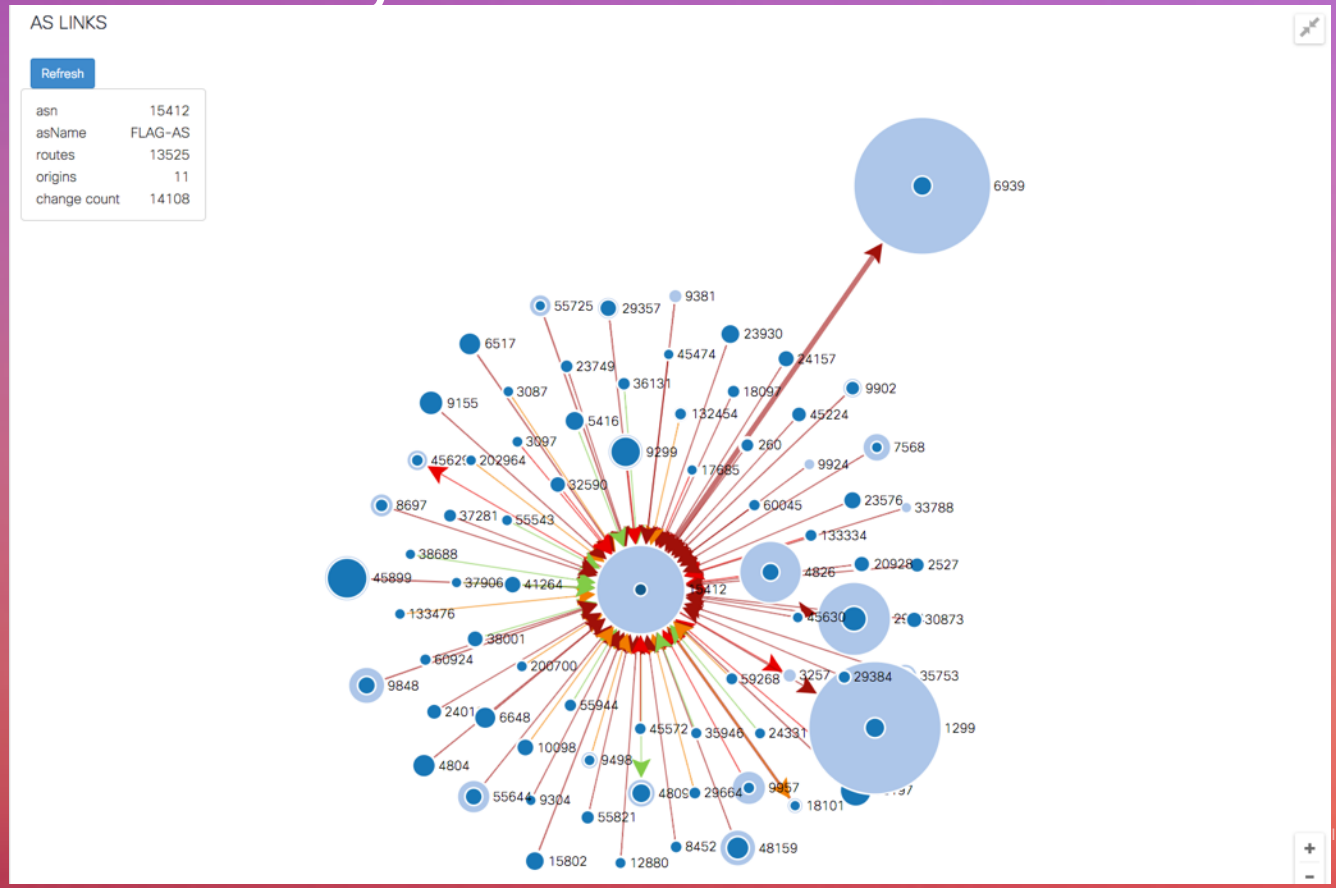

AS Connectivity – Deutsche Telekom

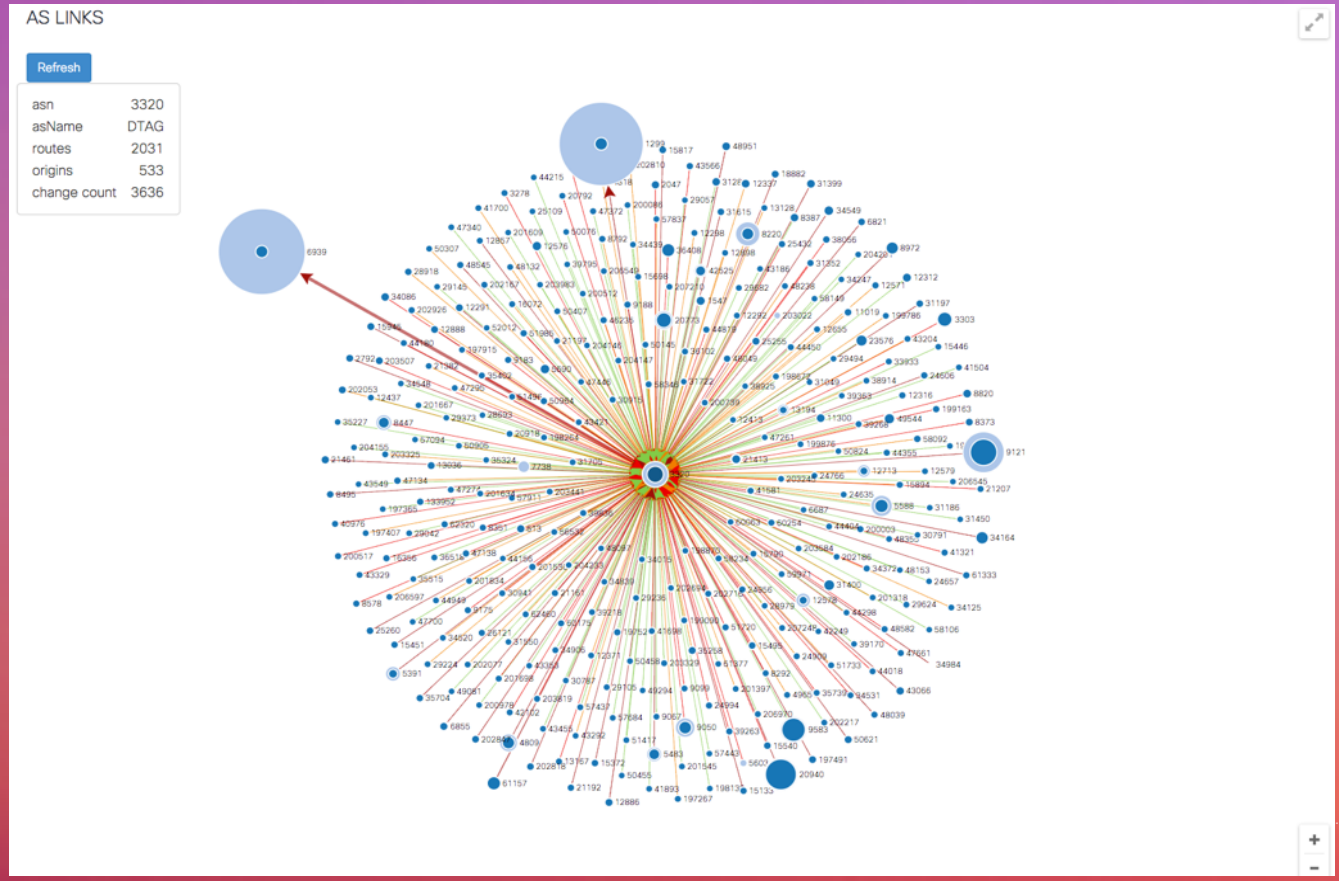

Prefix to Path history

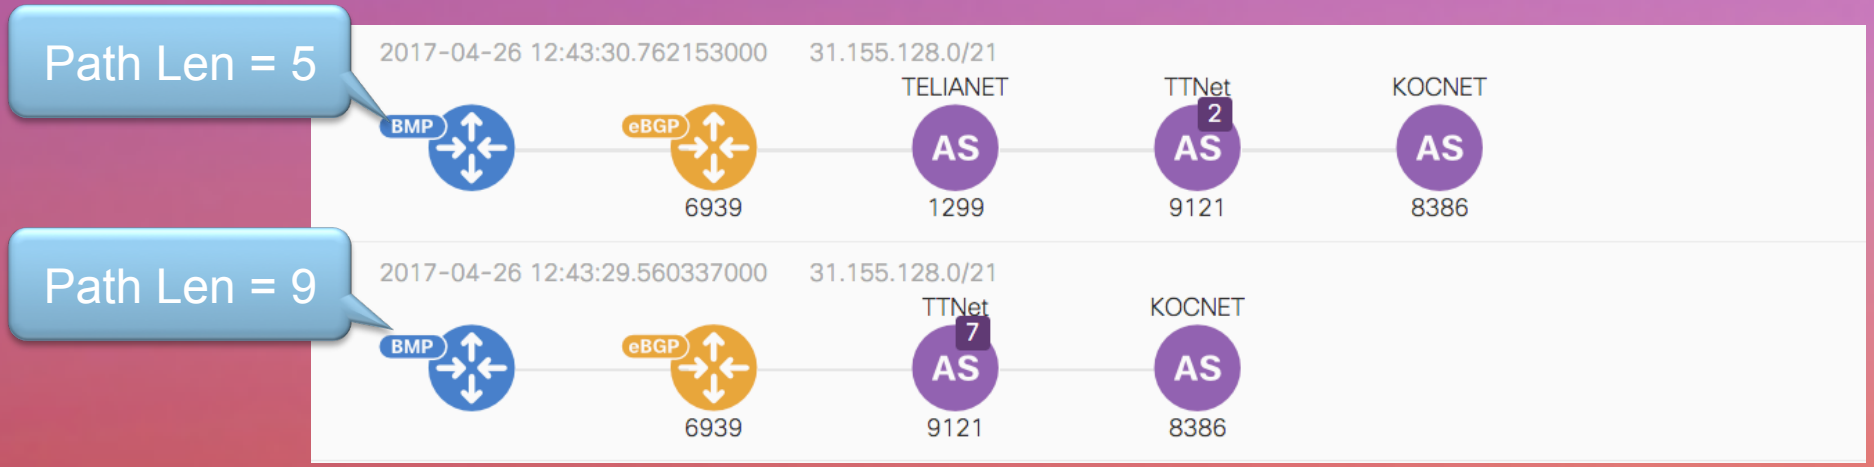

AS Path variance – 6939 to 8386

Shortest path − 3 hops
Longest path − 28 hops
Longest unique AS path − 5
Unique paths - 9
Largest prepend count − 17x
Prepend variation − [7-17]
Path with most updates − via AS1273

Data recorded in a 24hr period

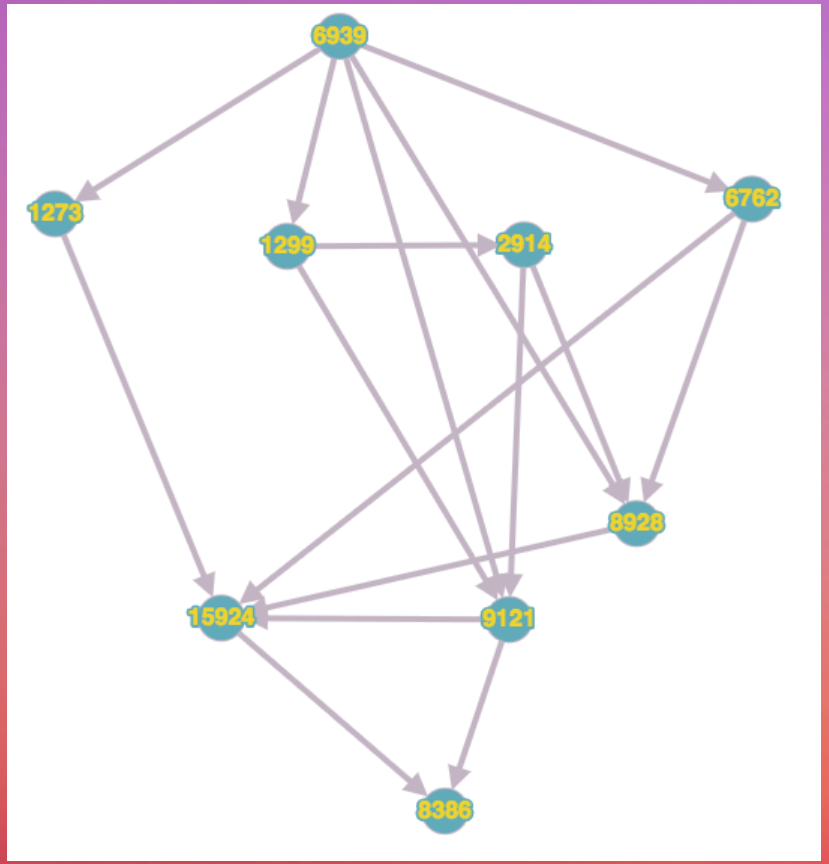

AS Path variance – 6939 to 8386

Shortest path − 4 hops
Longest path − 29 hops
Longest unique AS path − 6
Unique paths - 9
Largest prepend count − 17
Prepend variation − [7-17]
Path with most updates − via AS1273

Data recorded in a 24hr period

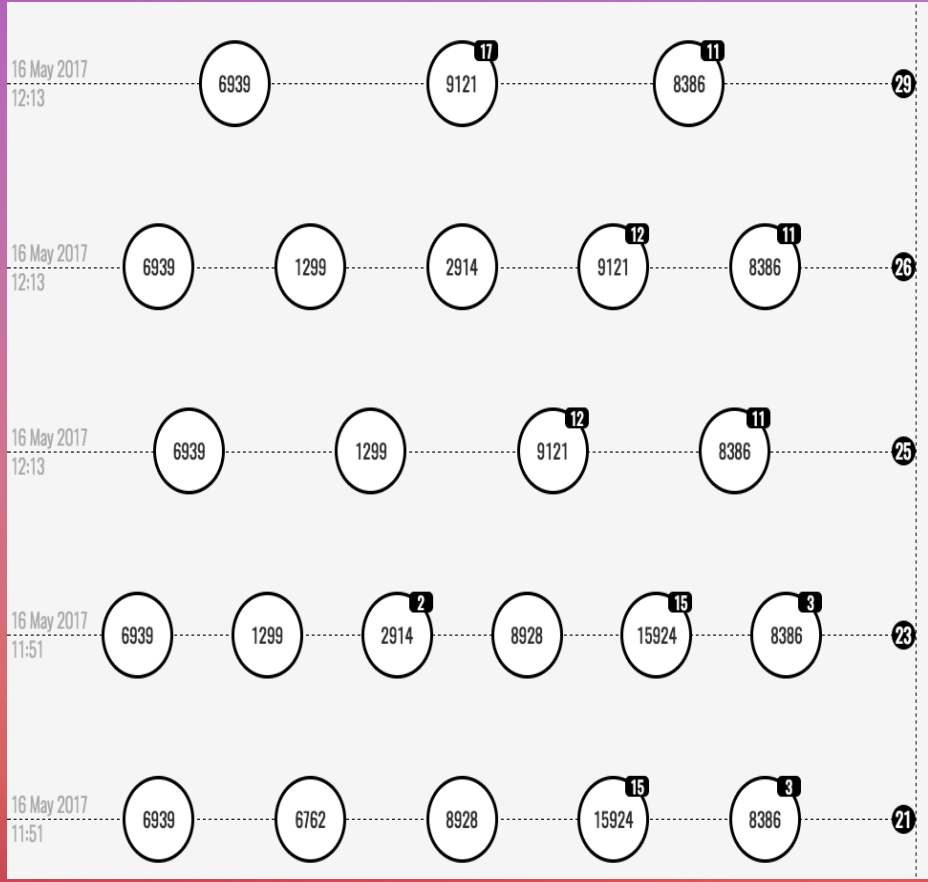

Security – Short prefix / long prefix detection

Martian anomalies

390 -

Prefix length anomalies

2 -

PREFIX LENGTH ANOMALIES

Download data: JSON

Prefix	Origin AS	Peer AS	AS Path	Advertising Router	Type	Timestamp
0.0.0.0/0	6939	11017	11017 6939	192.133.197.1	IPv4	2017-04-26 13:38:10
216.66.32.160/28	6939	11017	11017 6939	192.133.197.1	IPv4	2017-04-26 13:38:10

Default Route and Long prefix injection detected

Security – Unallocated prefixes

Download data: JSON ground truth

Prefix	Origin AS	Peer AS	AS Path	Advertising Route..	Type	Timestamp	Last Seen	Still ..	Category
202.181.6.0/24	134943	11017	11017 6939 9498 134943	192.133.197.1	IPv4	2017-04-26 13:38:10	2017-04-26 11:16:35	true	unallocated
202.181.6.0/24	134943	6939	6939 3491 9498 134943	192.133.197.1	IPv4	2017-04-26 13:38:10	2017-04-26 11:16:35	true	unallocated
116.199.203.0/24	38521	6939	6939 3491 58552 38521	192.133.197.1	IPv4	2017-04-26 13:38:09	2017-04-25 16:14:40	true	unallocated
202.181.6.0/24	134943	6939	6939 1299 5511 9498 134943	192.133.197.1	IPv4	2017-04-26 13:38:09	2017-04-26 11:16:35	true	unallocated
103.207.91.0/24	63969	11017	11017 6939 3491 9498 58715 63969	192.133.197.1	IPv4	2017-04-26 13:38:09	2017-04-25 16:14:39	true	unallocated
103.247.31.0/24	132122	6939	6939 1299 5511 9498 9730 132122	192.133.197.1	IPv4	2017-04-26 13:38:09	2017-04-26 05:14:54	true	unallocated
103.243.8.0/22	133676	6939	6939 5511 9498 133676	192.133.197.1	IPv4	2017-04-26 13:38:09	2017-04-25 16:14:40	true	unallocated

TOTAL ITEMS: 957

Observed over a 12 hour period

Security – Prefix drill-down

Download data: JSON ground truth

Prefix	Origin AS	Peer AS	AS Path	Advertising Route..	Type	Timestamp	Last Seen	Still ..	Category
103.212.									
103.212.178.0/24	56124	6939	6939 15412 18101 55410 56124	192.133.197.1	IPv4	2017-04-26 13:37:47	2017-04-25 16:13:41	true	unallocated
103.212.178.0/24	56124	6939	6939 1273 55410 56124	192.133.197.1	IPv4	2017-04-26 13:37:46	2017-04-25 16:13:41	true	unallocated
103.212.178.0/24	56124	6939	6939 3356 55410 55410 56124	192.133.197.1	IPv4	2017-04-26 13:37:46	2017-04-25 16:13:41	true	unallocated
103.212.178.0/24	56124	6939	6939 3356 55410 55410 56124	192.133.197.1	IPv4	2017-04-26 13:37:46	2017-04-25 16:13:41	true	unallocated
103.212.178.0/24	56124	6939	6939 1299 2914 15412 18101 55410 56124	192.133.197.1	IPv4	2017-04-26 13:37:46	2017-04-25 16:13:41	true	unallocated
103.212.178.0/24	56124	6939	6939 3209 55410 55410 55410 56124	192.133.197.1	IPv4	2017-04-26 13:37:40	2017-04-25 16:13:41	true	unallocated

AS PATHS

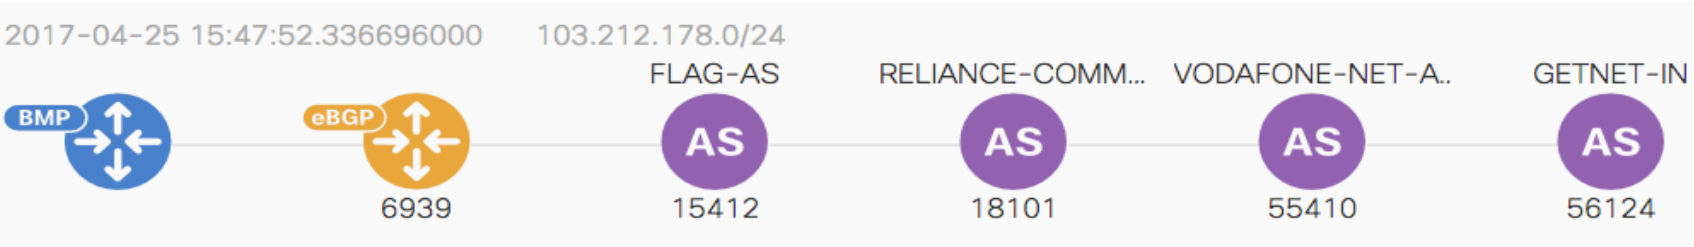

Security – drill-down

Download data: JSON ground truth

Prefix	Origin AS	Peer AS	AS Path	Advertising Route..	Type	Timestamp	Last Seen	Still ..	Category
	58934								
191.37.252.0/24	58934	11017	11017 6939 12389 48066 58271 58934	192.133.197.1	IPv4	2017-04-26 13:37:50	2017-04-25 16:13:49	true	unallocated
138.59.180.0/23	58934	6939	6939 1273 12389 48066 58271 58934	192.133.197.1	IPv4	2017-04-26 13:37:50	2017-04-25 16:13:48	true	unallocated
200.3.10.0/23	58934	11017	11017 6939 12389 48066 58271 58934	192.133.197.1	IPv4	2017-04-26 13:37:50	2017-04-25 16:13:48	true	unallocated
177.154.93.0/24	58934	6939	6939 12389 48066 58271 58934	192.133.197.1	IPv4	2017-04-26 13:37:49	2017-04-25 16:13:50	true	unallocated
200.0.202.0/23	58934	6939	6939 12389 48066 58271 58934	192.133.197.1	IPv4	2017-04-26 13:37:49	2017-04-25 16:13:48	true	unallocated
177.73.253.0/24	58934	6939	6939 12389 48066 58271 58934	192.133.197.1	IPv4	2017-04-26 13:37:49	2017-04-25 16:13:48	true	unallocated
177.154.93.0/24	58934	6939	6939 12389 48066 58271 58934	192.133.197.1	IPv4	2017-04-26 13:37:49	2017-04-25 16:13:50	true	unallocated

AS PATHS

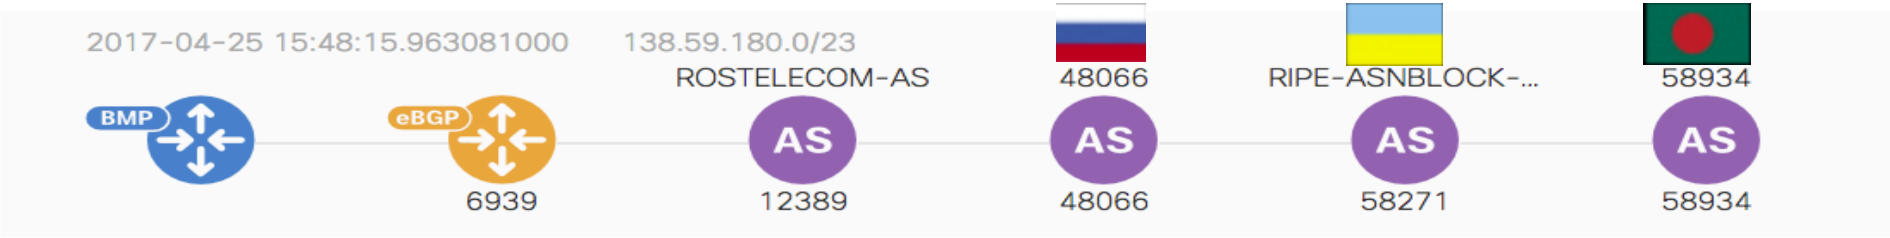

More specific prefix detection

- AS 12345 originates 100.100.0.0/18
- Hijacker originates 100.100.63.0/24
- Basically a needle in a large haystack, does anyone notice?
- What does RPKI show?
- Do the origin ASNs match?
- Does the less specific share the same transit set or similar as_paths?
- Does RIR have the same organization name or contacts for both origins?
- Anything out of the norm for the new originating ASN?

Potential

What can we do with large-scale collection of historical event information?

- Event impact analysis –
  - Stability
  - Security
  - Misconfiguration
  - Forensics
- Application of ML/DL to data-set
- Pattern-detection and network ‘weather forecasting’

PNDA.io – the platform

What is PNDA?

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

Linux Foundation Collaborative Project based on the Apache ecosystem

Where is PNDA today?

- Linux Foundation project
- Selected by MEF for Analytics function within Lifecycle Service Orchestration framework
- In service trials with two Service Providers
- One platform supporting a range of use-cases including
  - Network security – Apache Spot
  - 6CN
  - Virtualization infrastructure monitoring and analysis
  - Smart Cities
  - Smart Transportion use-cases

PNDA

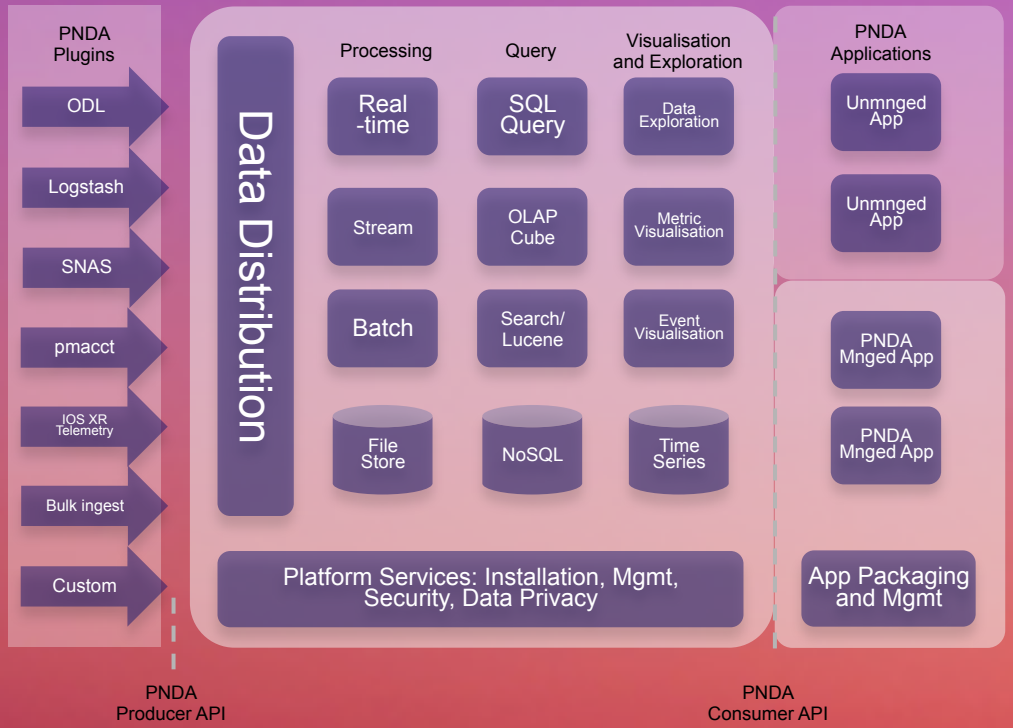

- Horizontally scalable platform for analytics and data processing applications
- Support for near-real-time stream processing and in-depth batch analysis on massive datasets
- Decouples data collection and aggregation from data analysis
- Consuming applications can be either platform apps developed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

PNDA

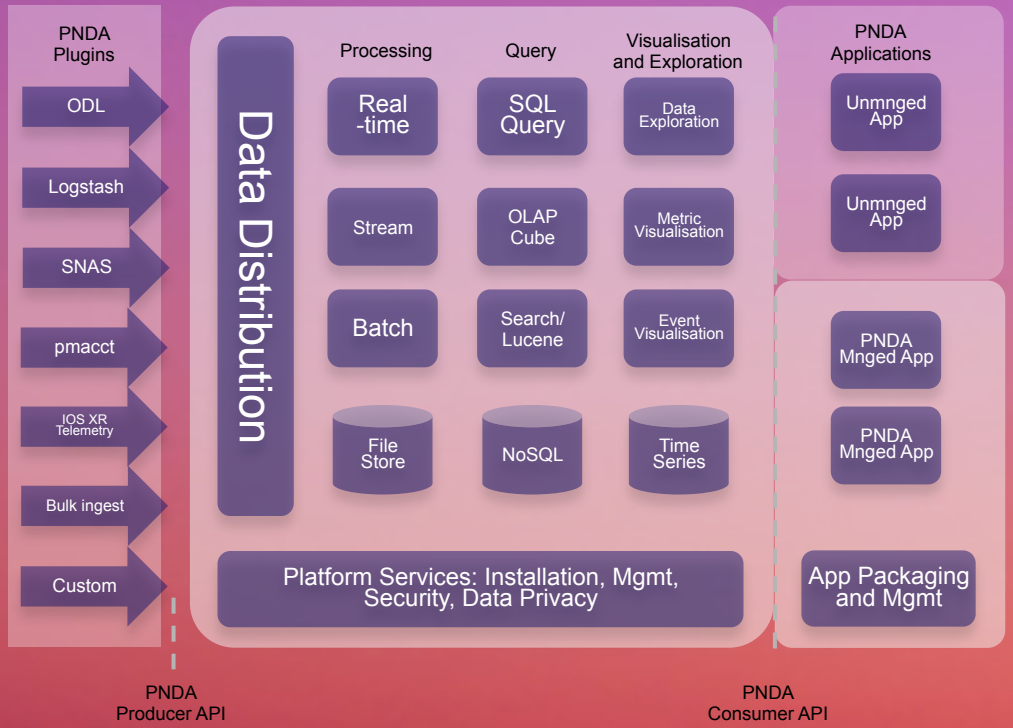

- Simple, scalable open data platform
- Provides a common set of services for developing analytics applications
- Accelerates the process of developing big data analytics applications whilst significantly reducing the TCO
- PNDA provides a platform for convergence of network data analytics

Why PNDA?

Innovation in the big data space is extremely rapid, but combining multiple technologies into an end-to-end solution can be extremely complex and time-consuming

PNDA removes this complexity and allows you to focus on developing the analytics applications, not on developing the pipeline – significantly reducing the effort required and time-to-value

PNDA Software Components

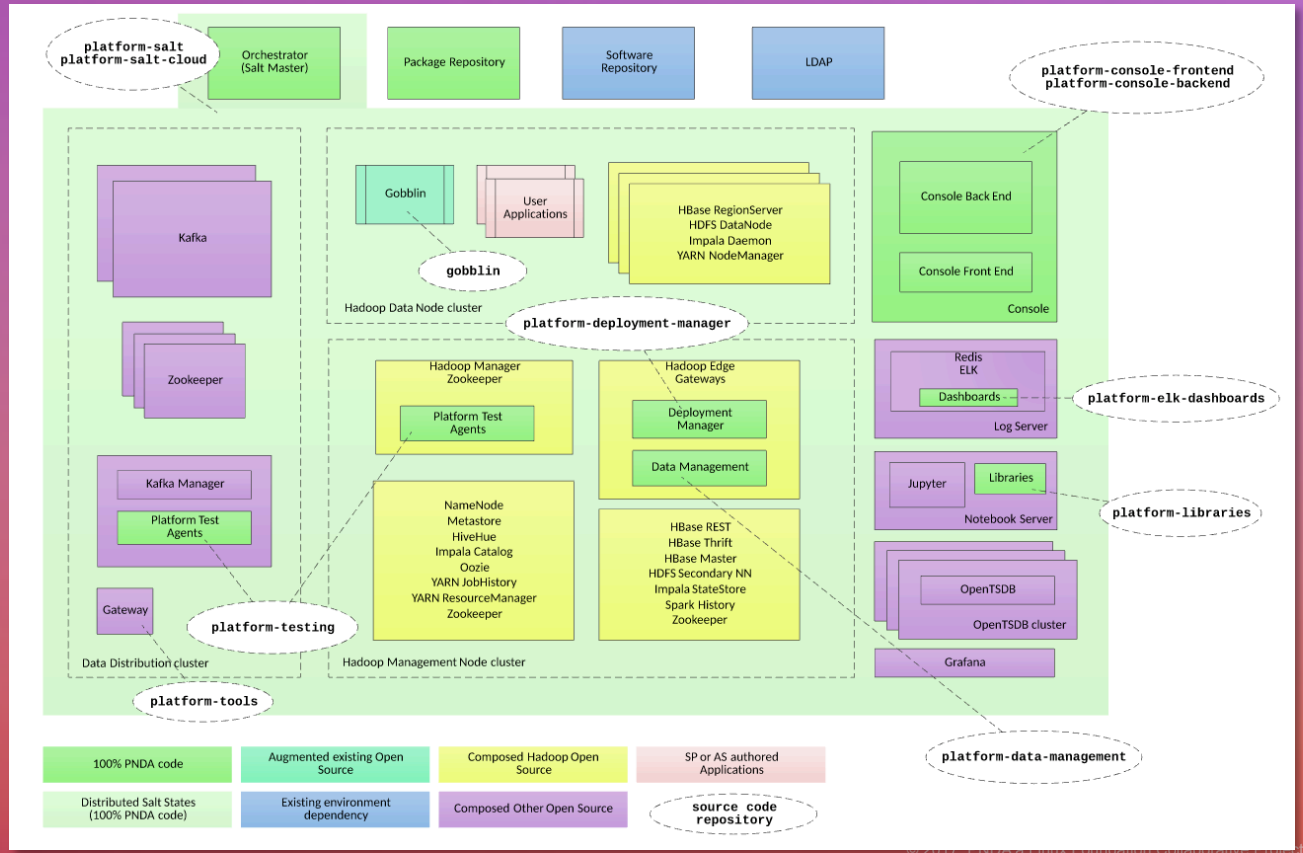

Where can I learn more?

- www.pnda.io
- https://github.com/pndaproject
- www.snas.io

THE LINUX FOUNDATION