PNDA.io: when BGP meets Big-Data

THE LINUX FOUNDATION

26th April 2017

Let’s go back in time…

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?

SNAS Architecture

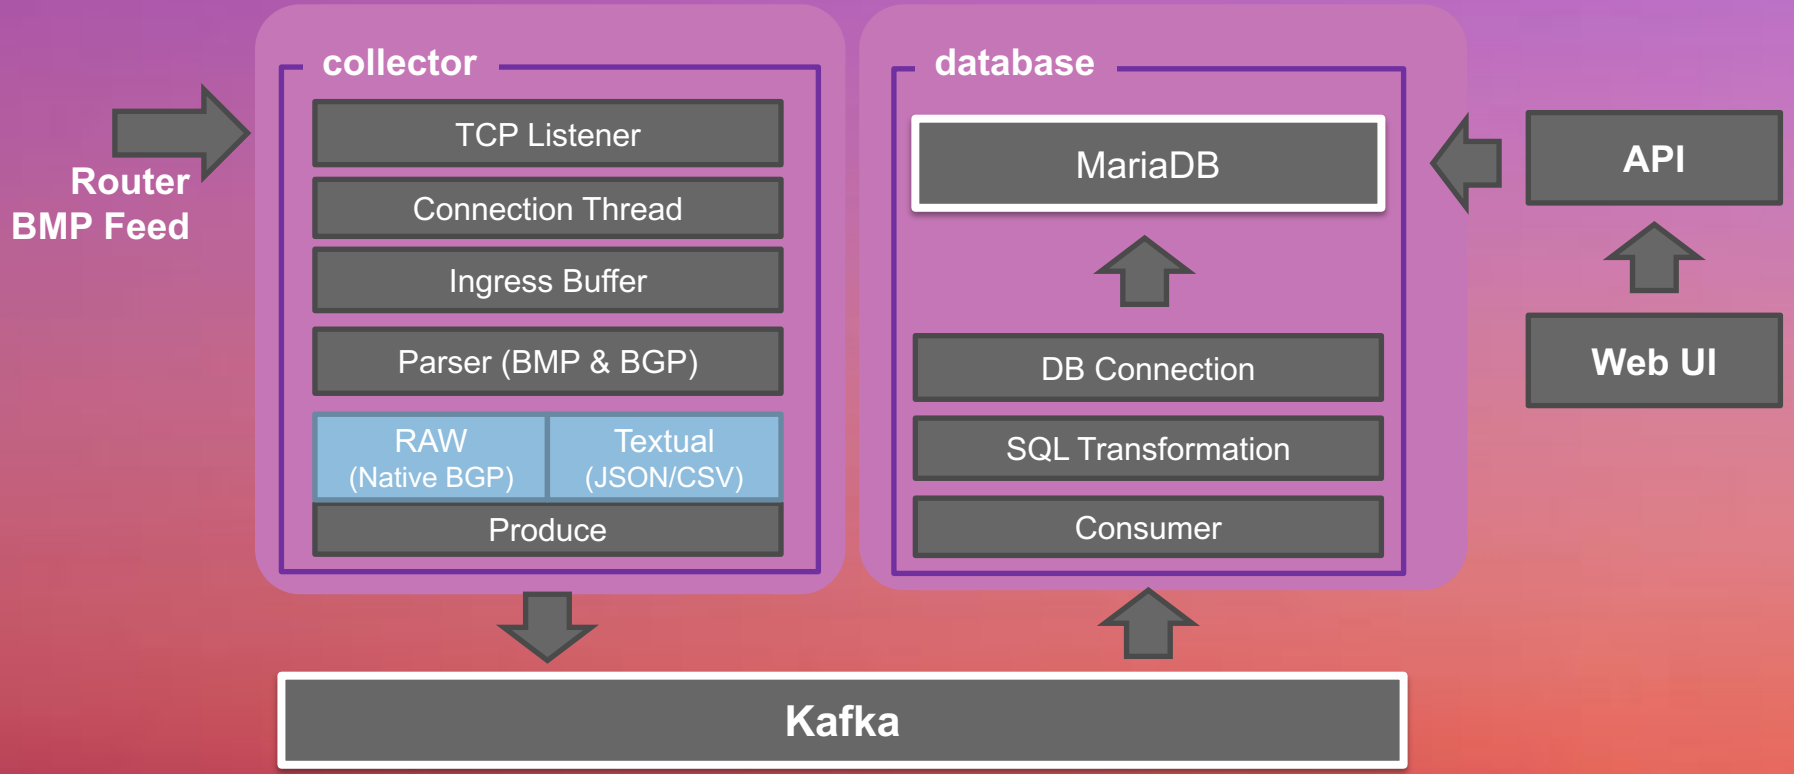

E2E architecture

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

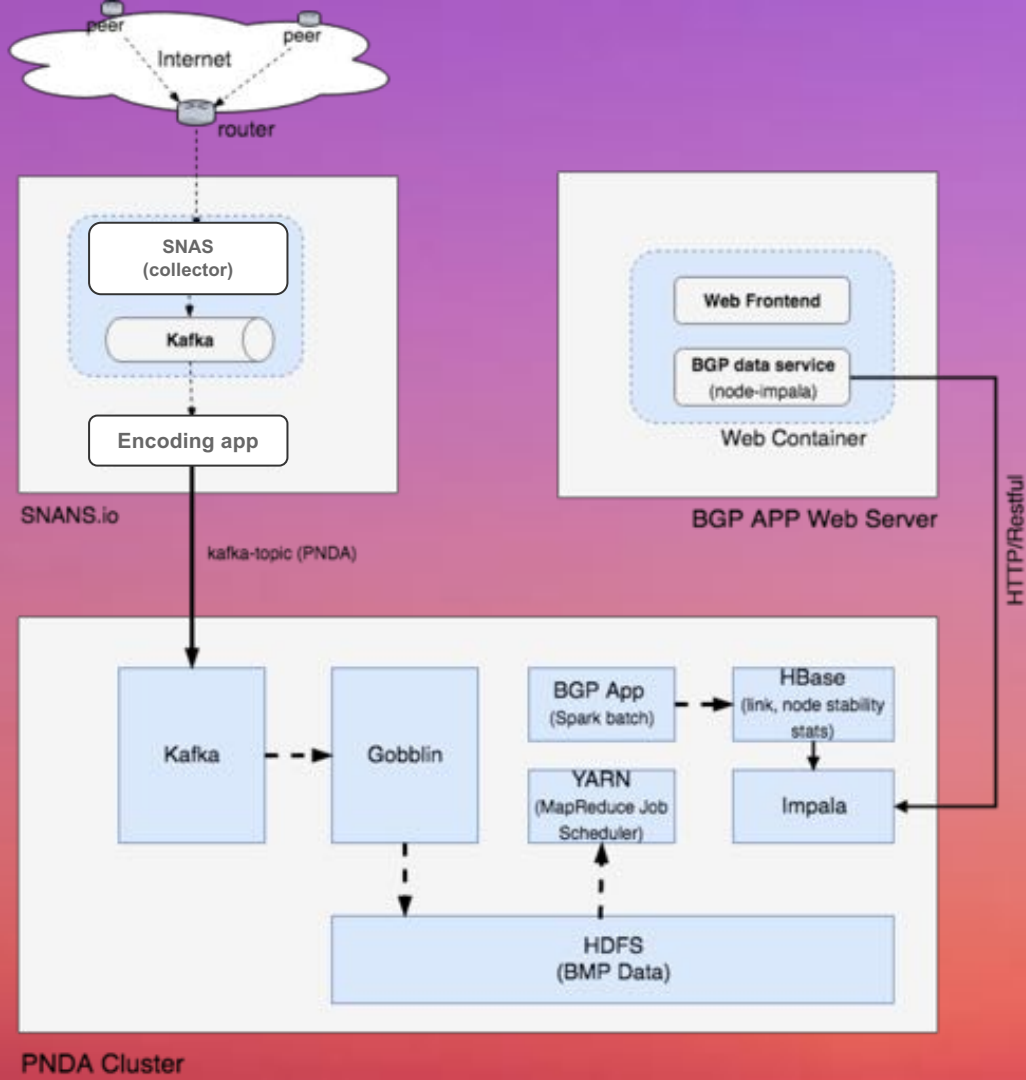

What does this give us?

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

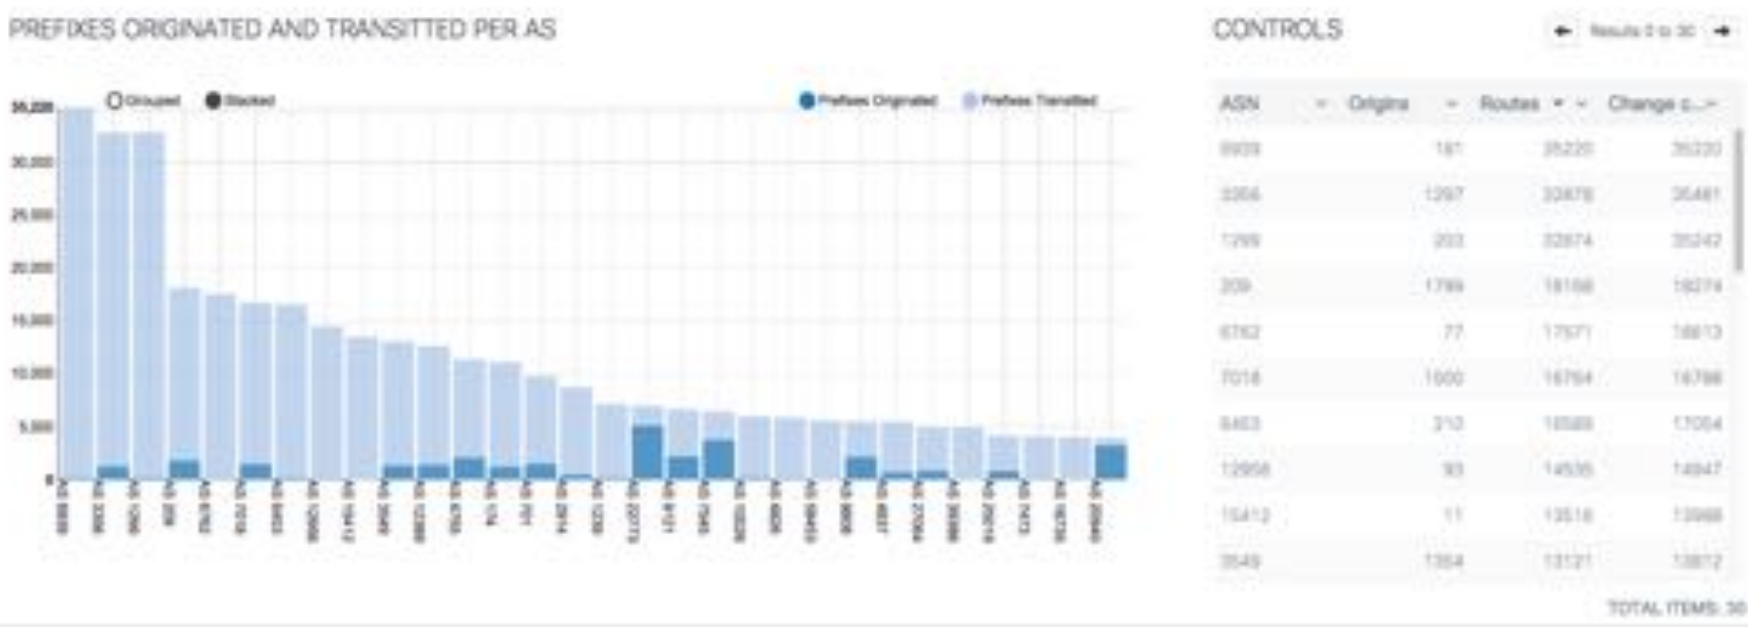

Path stability

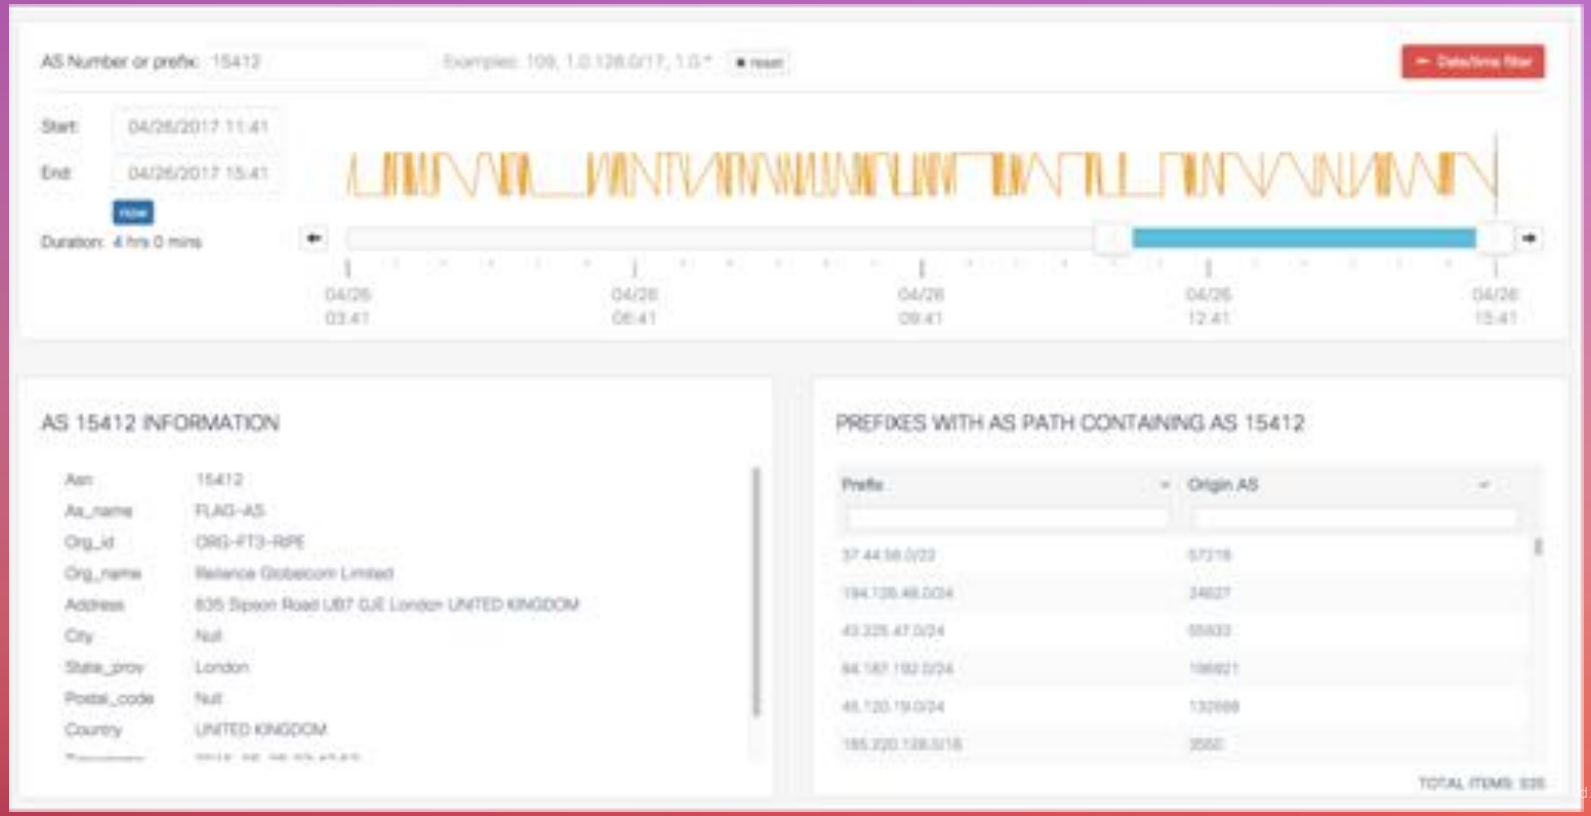

AS Connectivity - FLAG

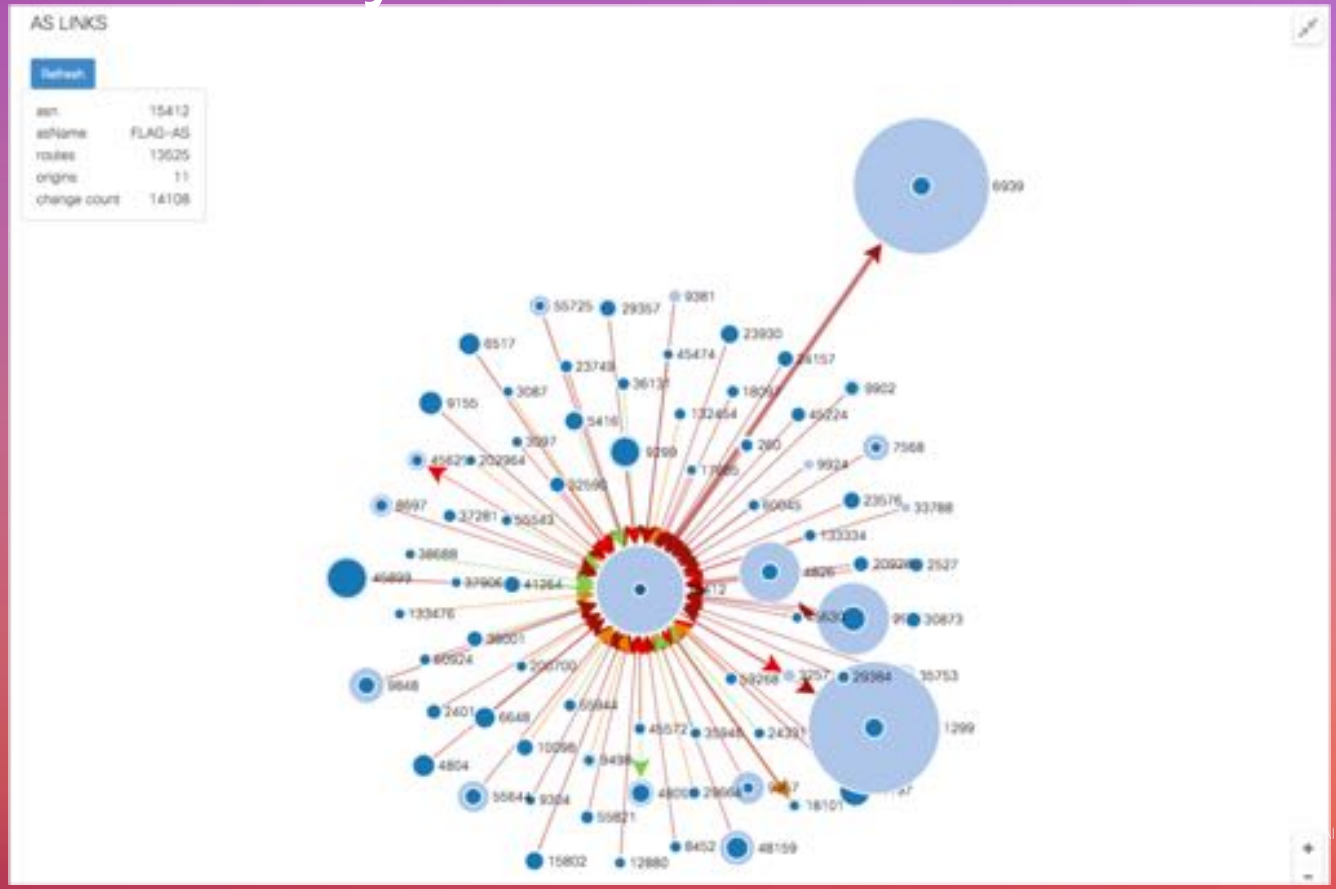

AS Connectivity – Deutsche Telekom

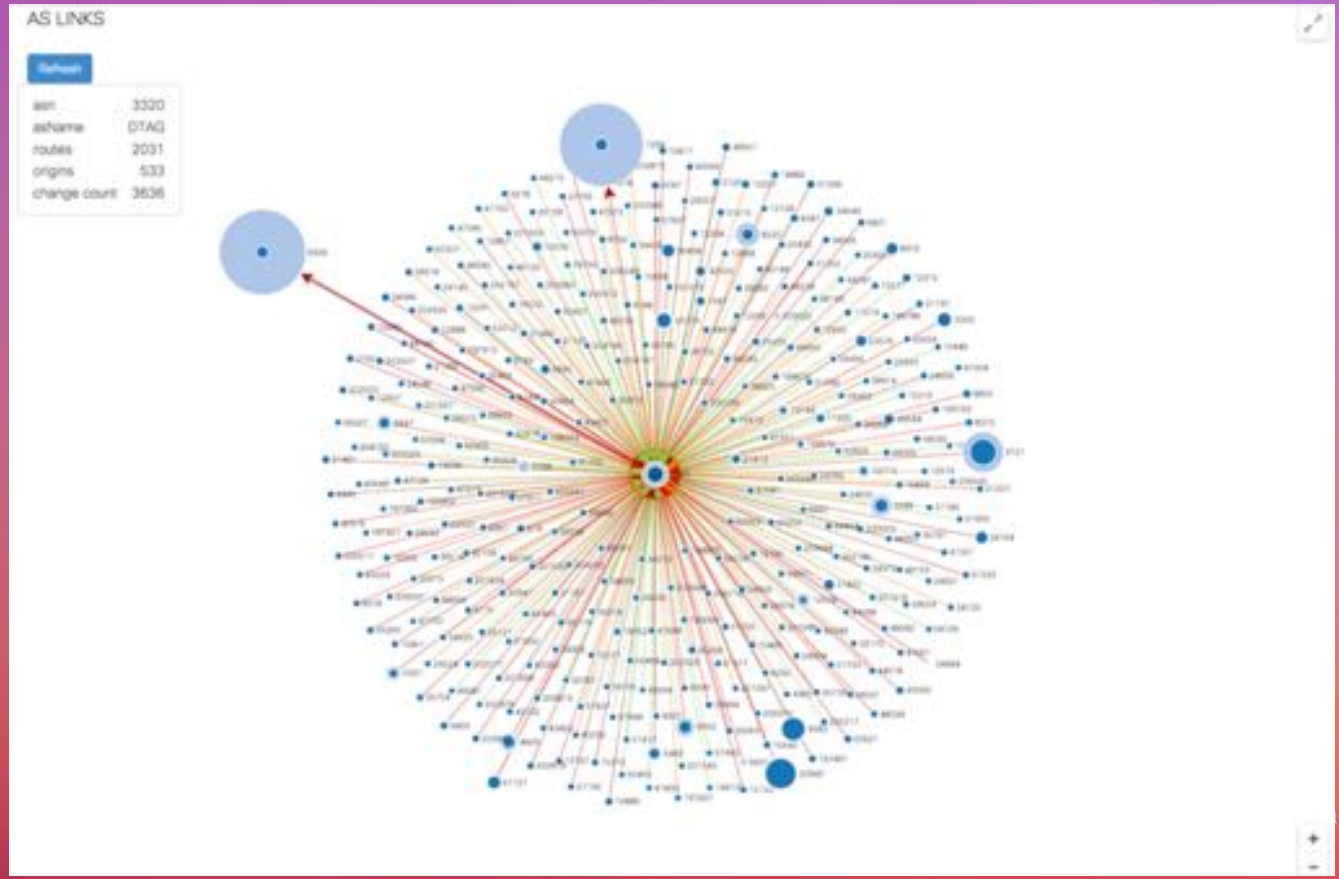

AS Path variance – 6939 to 8386

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

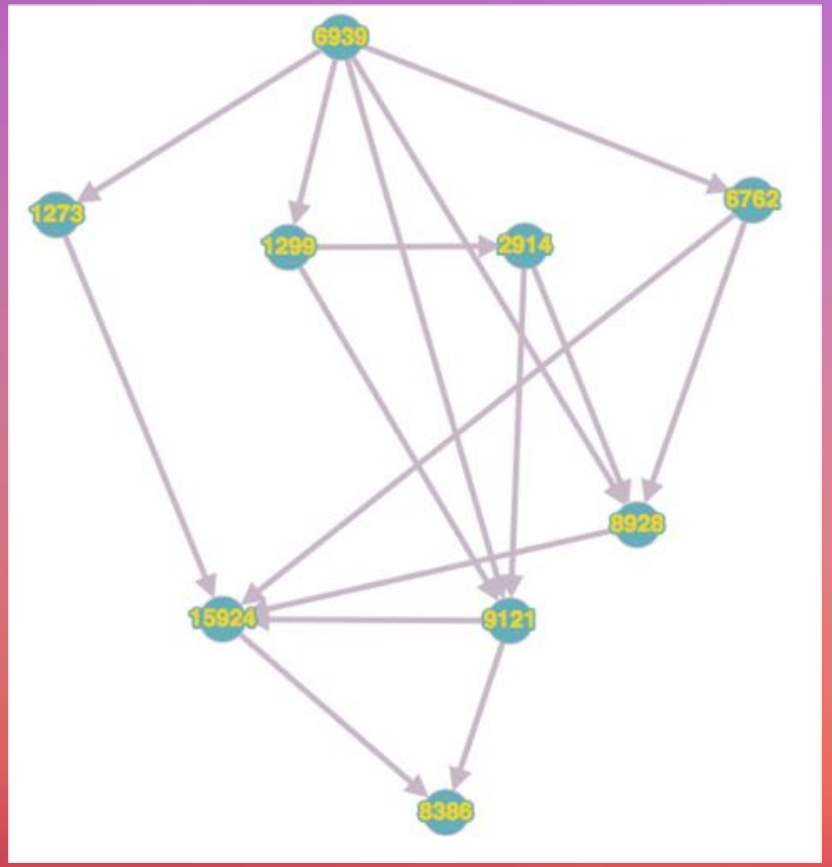

AS Path variance – 6939 to 8386

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

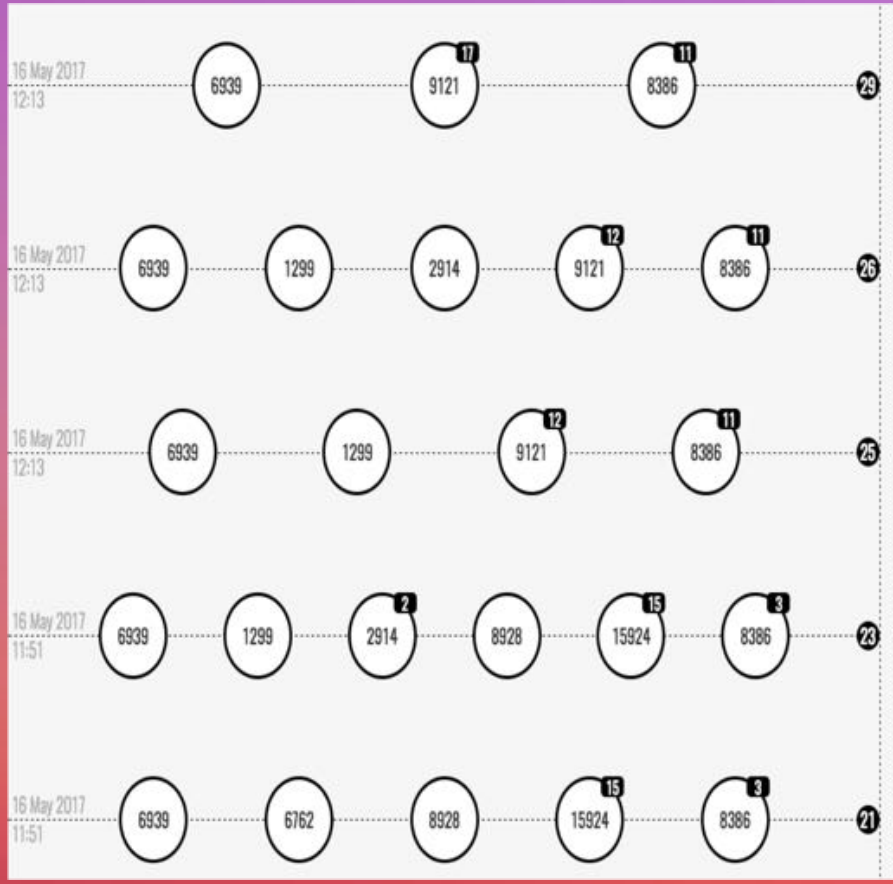

Security – Unallocated prefixes

Observed over a 12 hour period

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?

Looking for the needle using Jupyter Notebook

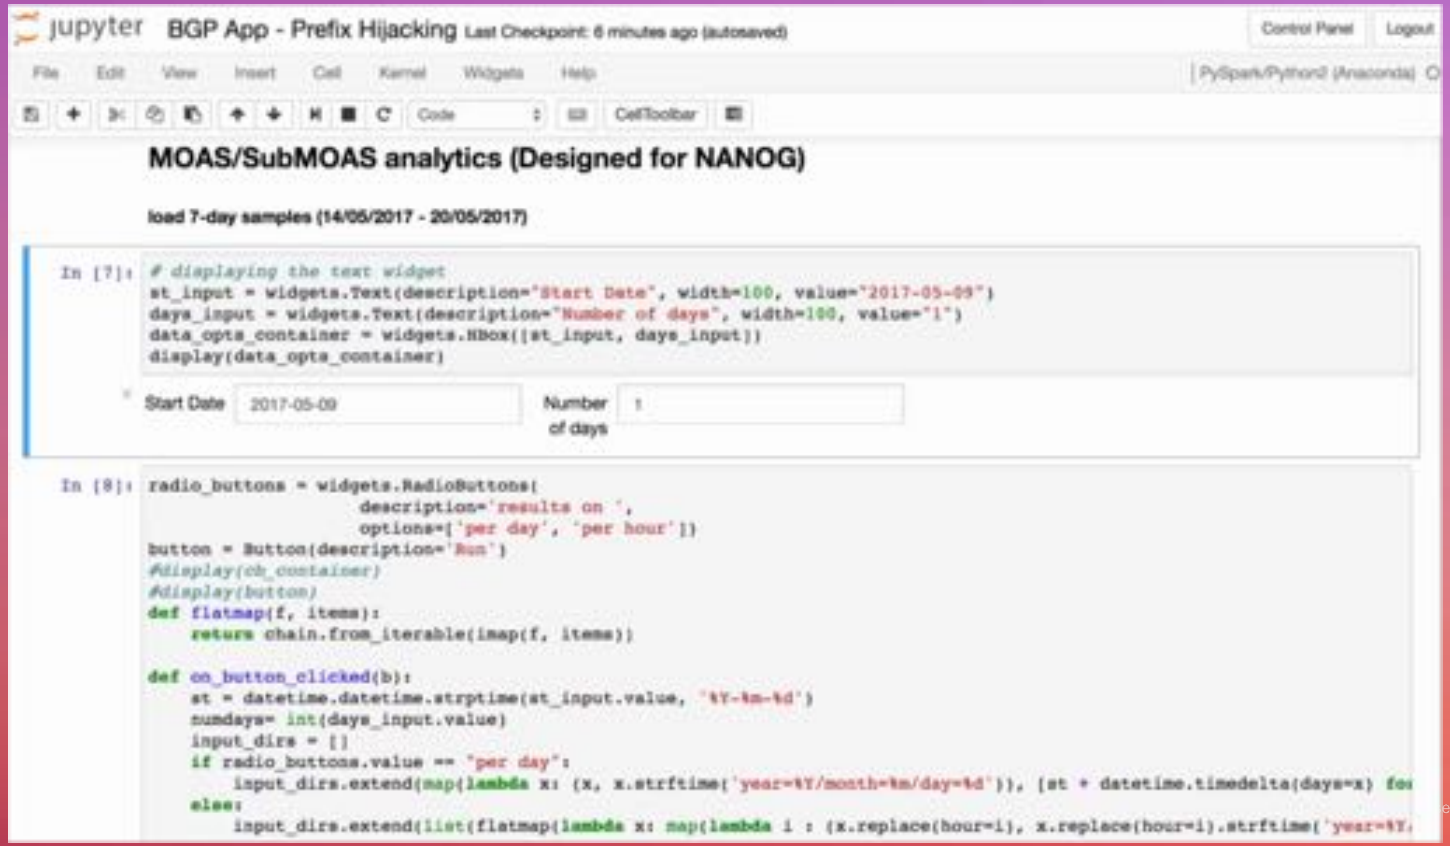

Multi-Origin AS prefixes ‘add’ detected – 24 hour period

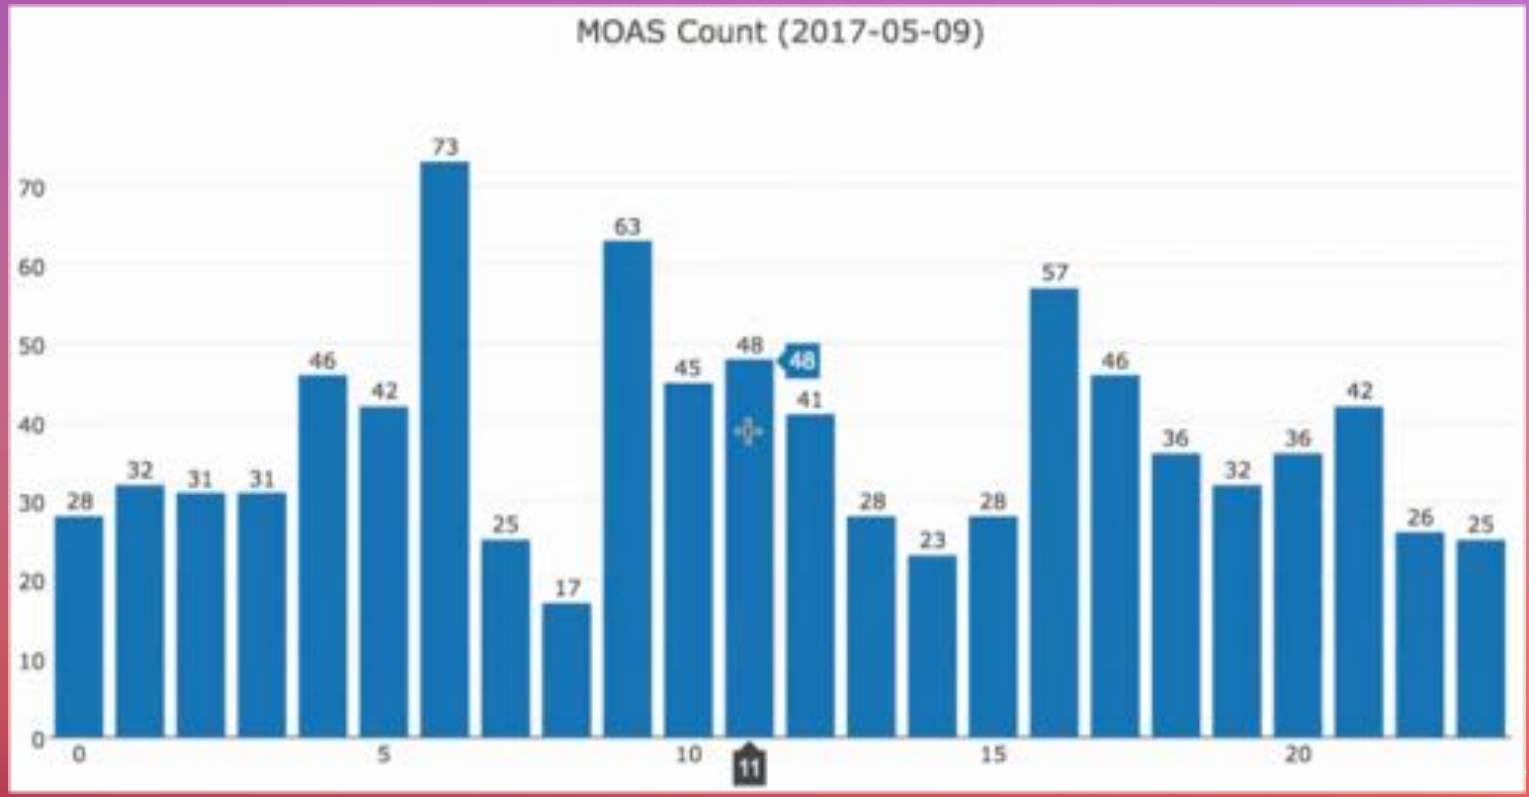

Sub-prefix injections over a 24 hour period

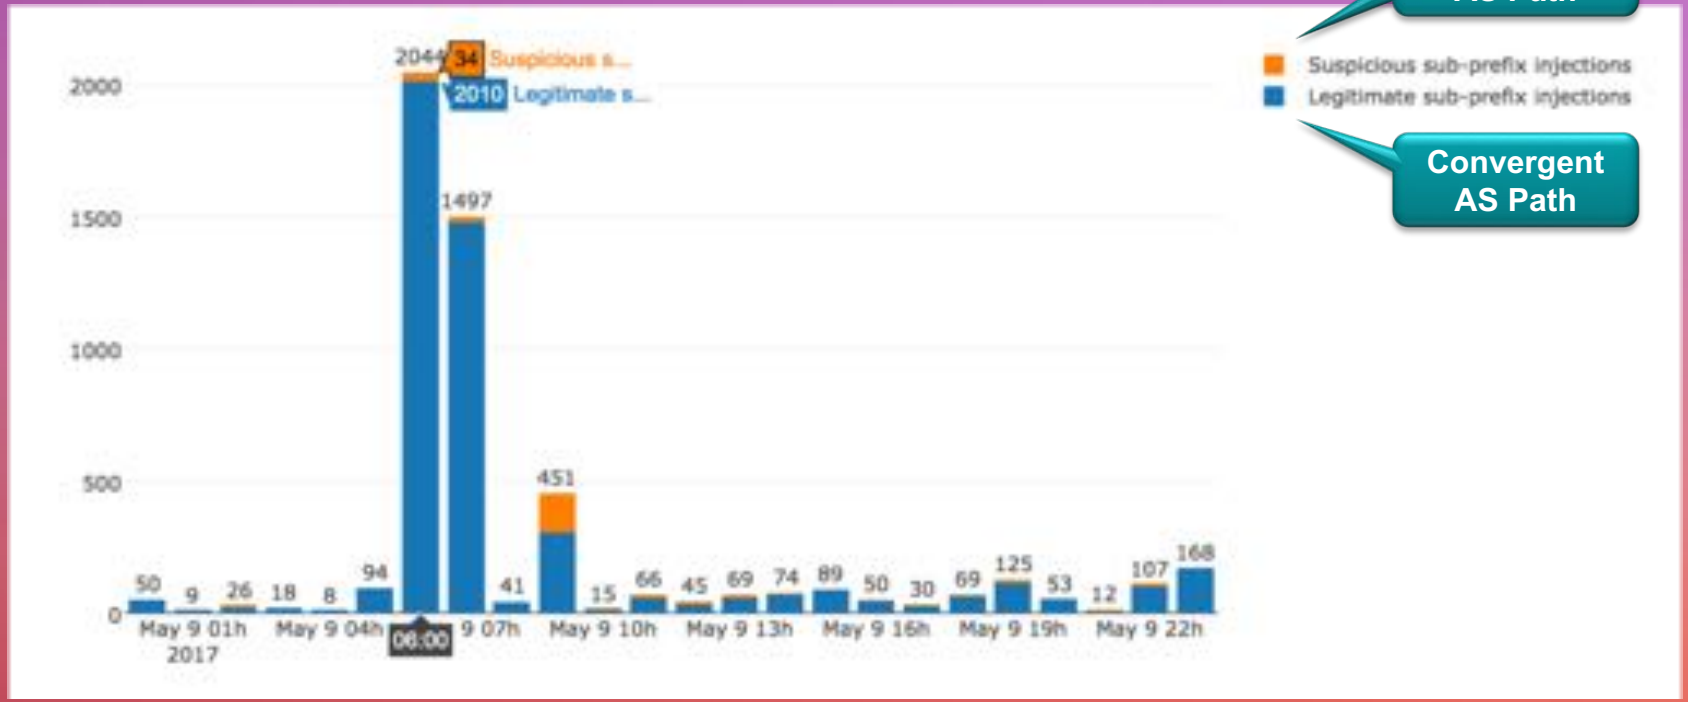

Sub-prefix injection victims – 24 hour period

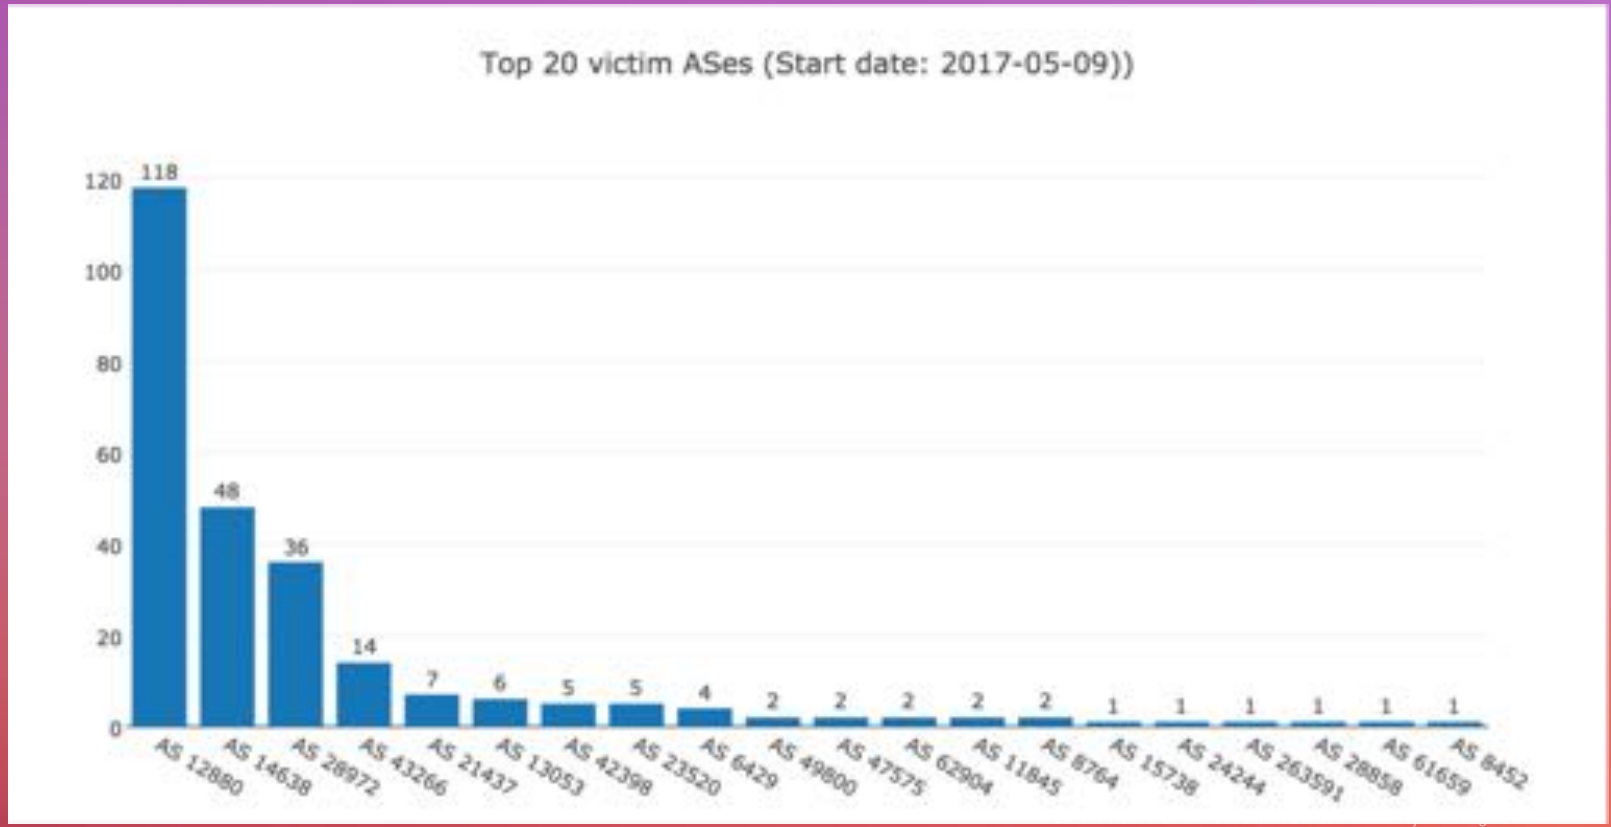

Sub-prefix injection attackers – 24 hour period

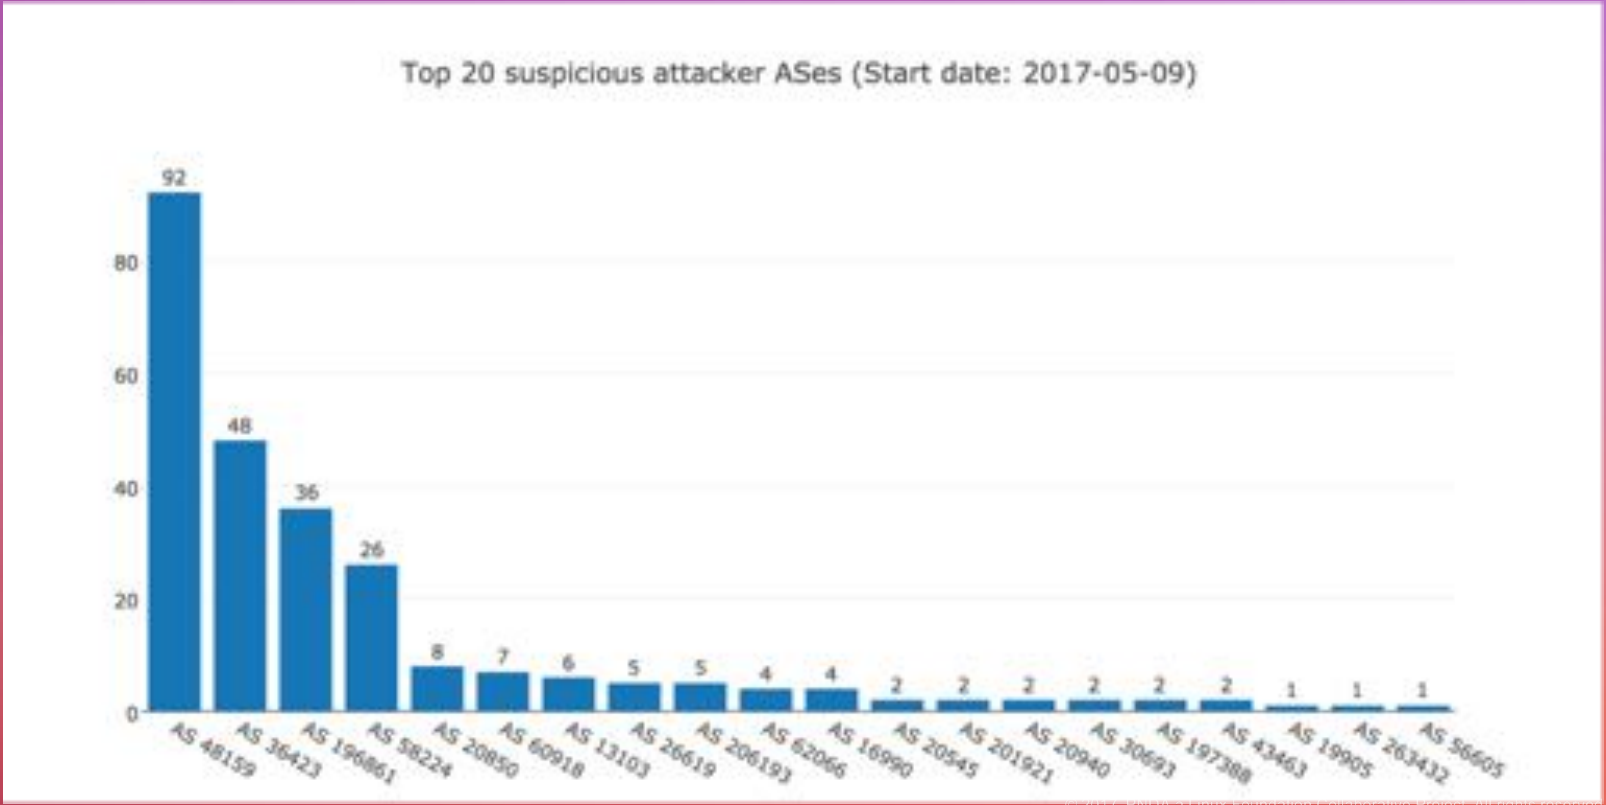

Move to real-time analytics

PNDA.io – the platform

THE LINUX FOUNDATION

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?

- 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 Transportation use-cases

PNDA

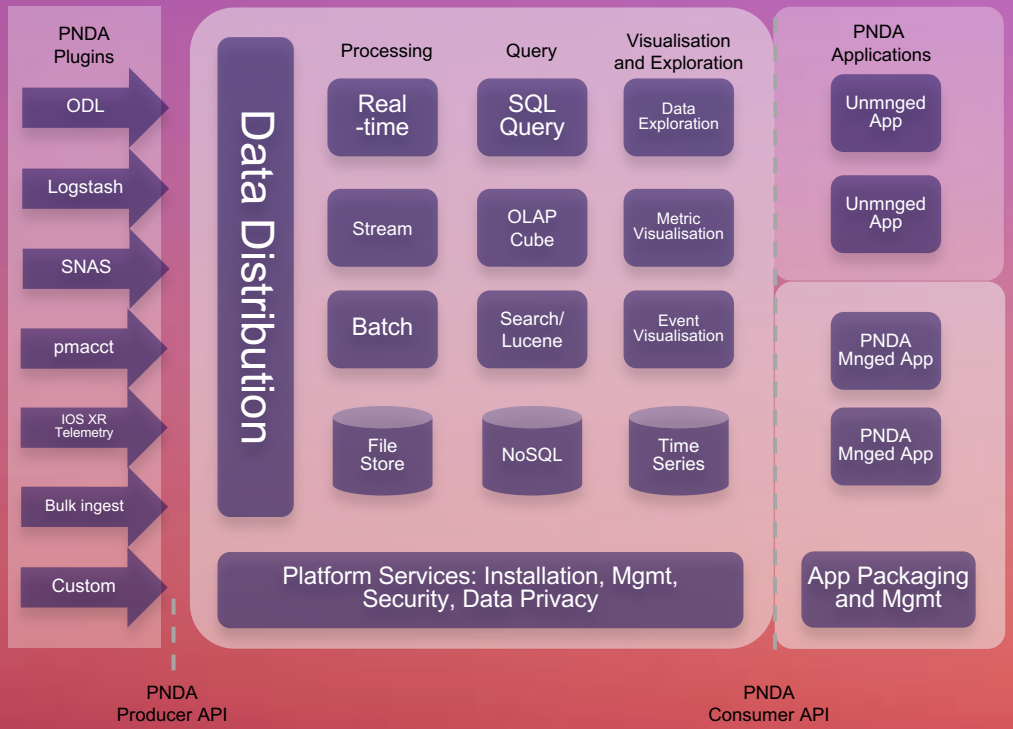

- Horizontally scalable platform for analytics and data processing applications
- Support for near-real-time stream processing and in-depth batch analysis on massive datasets
- Decouples data collection and aggregation from data analysis
- Consuming applications can be either platform apps developed 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

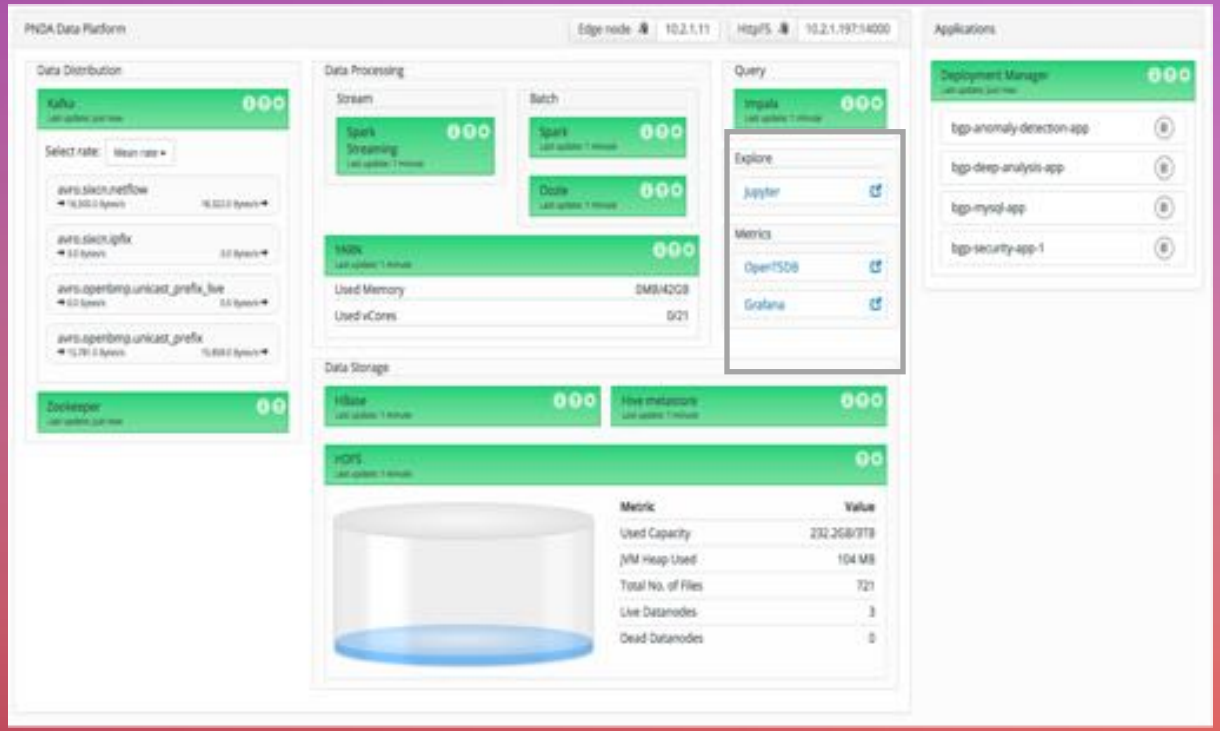

- Simple, scalable open data platform
- Provides a common set of services for developing analytics applications
- Accelerates the process of developing big data analytics applications whilst significantly reducing the TCO
- PNDA provides a platform for convergence of network data analytics

Red-PNDA

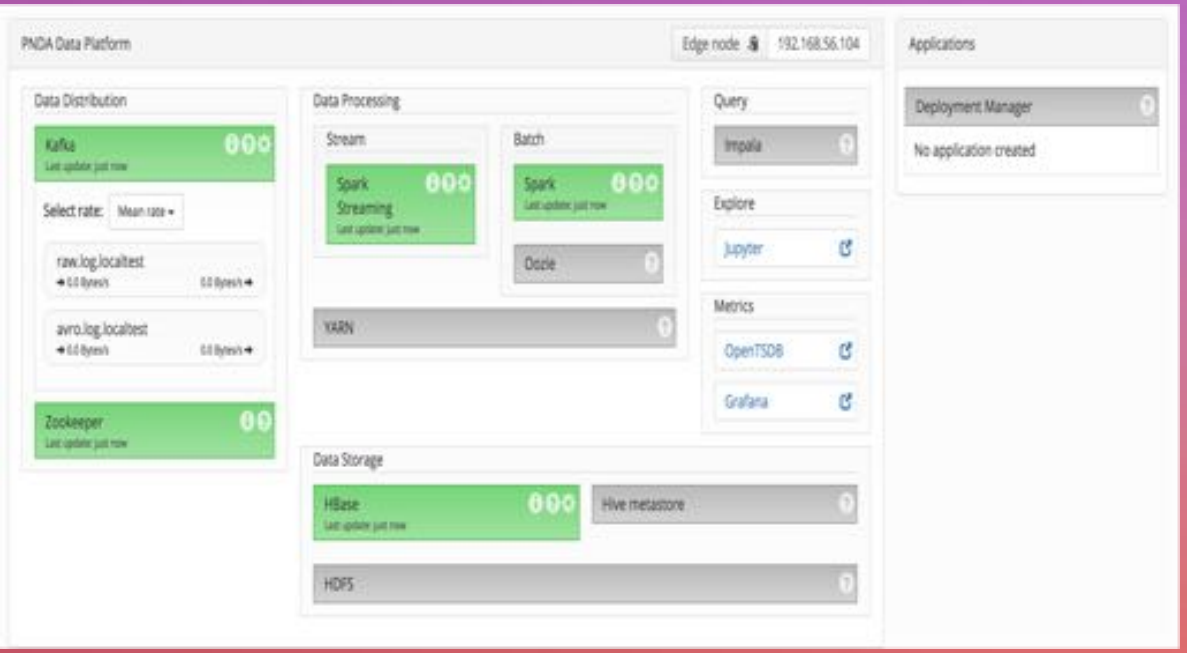

- A reduced set of components providing a PNDA-like environment for education and basic prototyping
- Miniature PNDA – fits your laptop
  - Lightweight simplified “Big Data” platform

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’

Where can I learn more?

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

THE LINUX FOUNDATION