

Thales PunchPlatform



PunchPlatform team

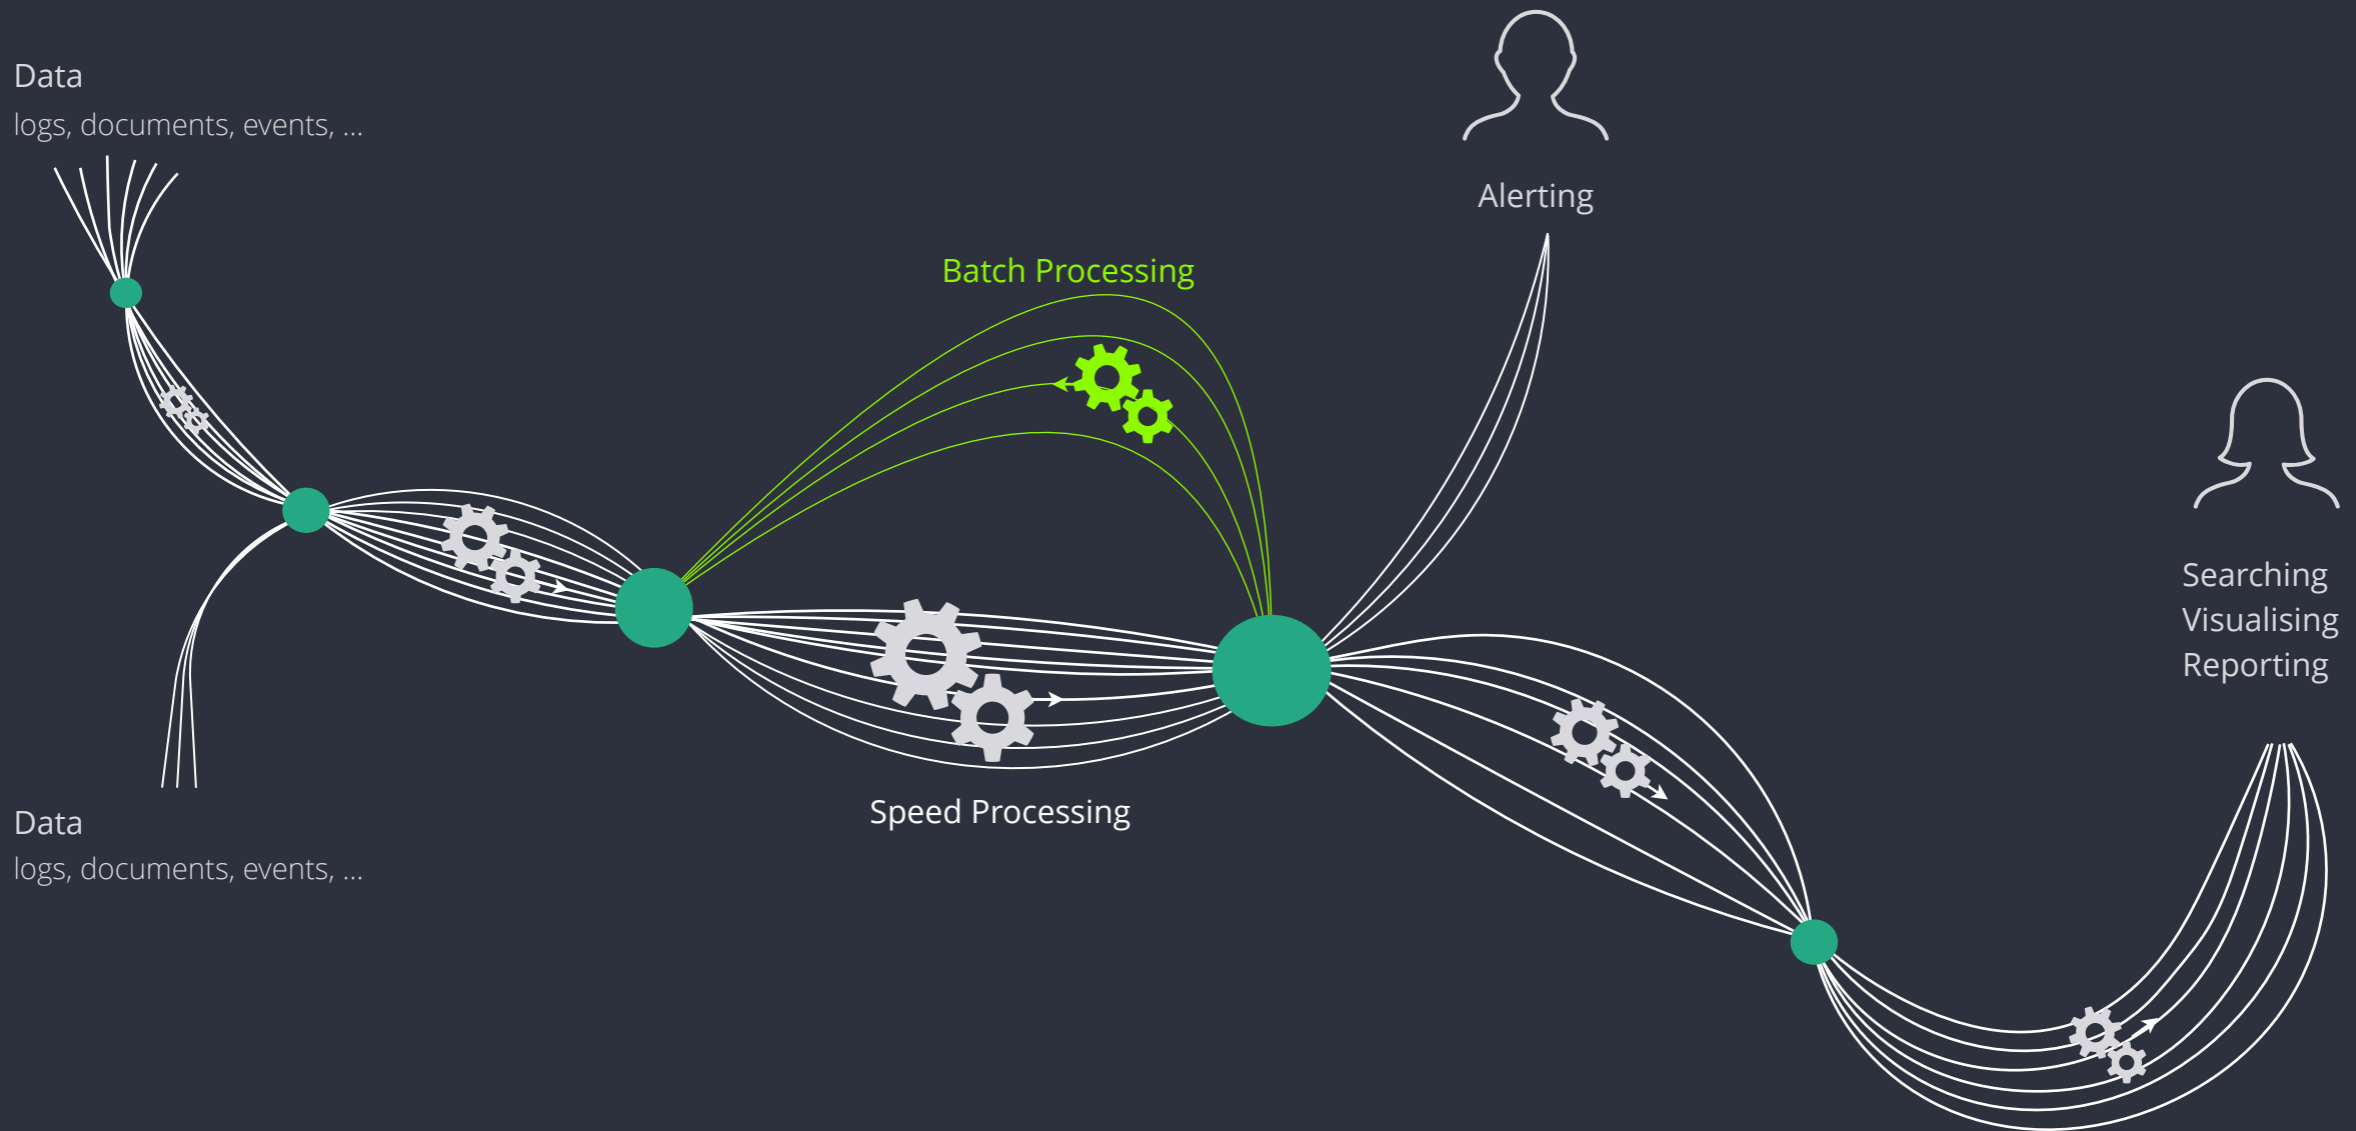
Agenda


- What It Does
- Building Blocks
- Deployment & Operations
- Typical Setups
- Customers and Use Cases
- RoadMap



What It Does

Compose Arbitrary Industrial Data Processing Channels





DataFlow

Collect and Transport Your Data
 Multi-nodes, -racks, -rooms, -sites
 Scalable, Resilient, Reliable



DataVisualization

Create your Dashboards to dig
 months/years of data
 Pick what suits best your needs:
 Kibana, Grafana, Zeppelin
 Data Extraction and Reporting
 Multi-Tenant, Secure

DataProcessing

Parse, Normalize, Enrich
 Store, Archive, Index
 Detect, Raise Alarms
 Search
 Reprocess, Replay




DataAnalytics

Plug in Arbitrary Processing
 Storm, Spark-Streaming, Flink
 Join the Thales Big Data and
 Analytics Community





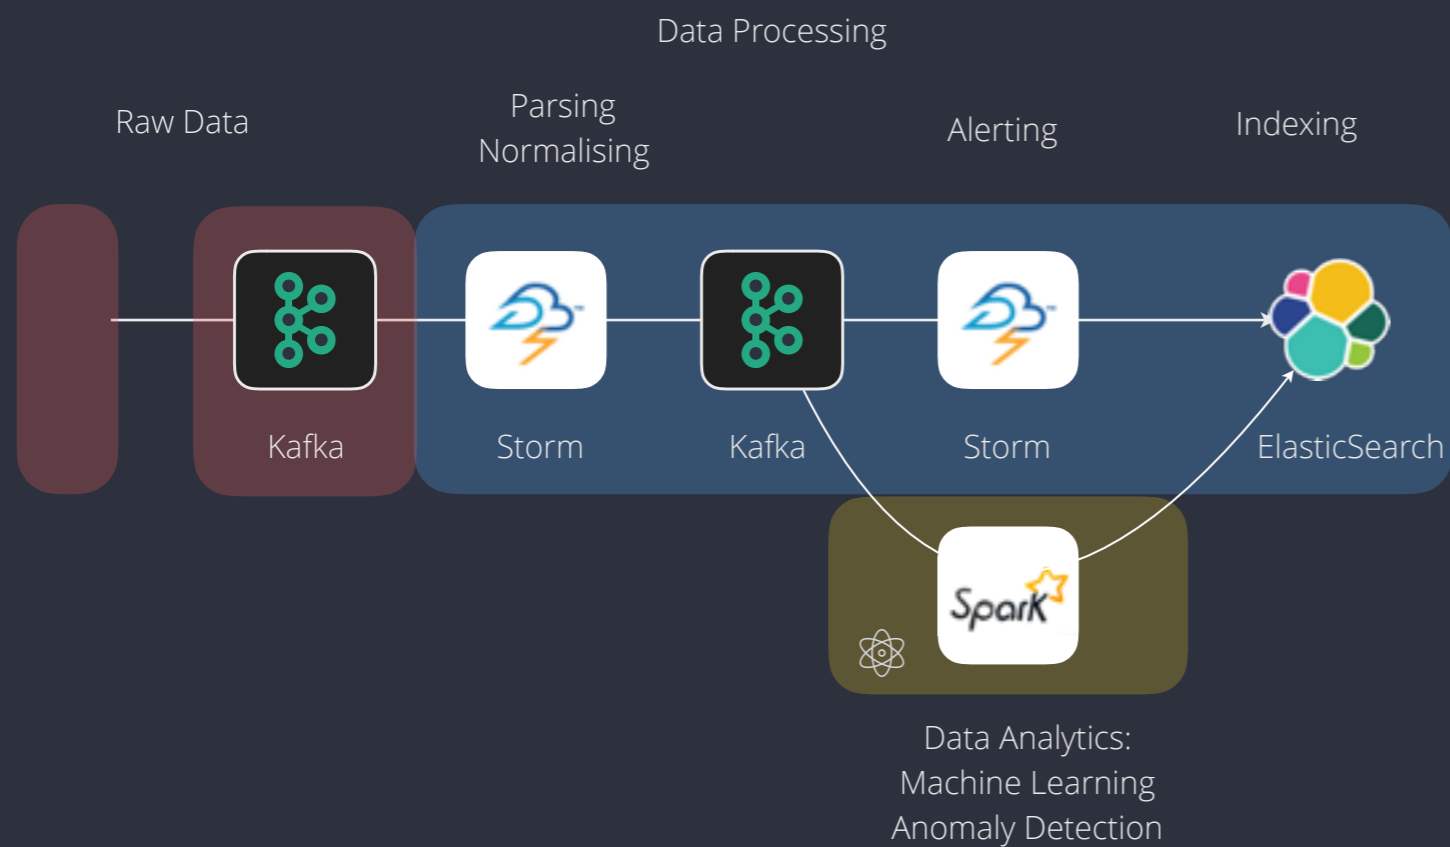
Understanding the (CyberSecurity) Data Pipeline






Log Management

Typical Setup




Searching
Visualising
Reporting
Alerting



Log Management + Analytics

Typical Setup



A vertical stack of four screenshots. The top one is a scatter plot with many colored dots. The second is a line graph with multiple blue lines showing trends over time. The third is a search interface with a 'Log Search' bar and various filters. The bottom one shows the logos for 'kibana' and 'Grafana'.



Searching
Visualising
Reporting
Alerting

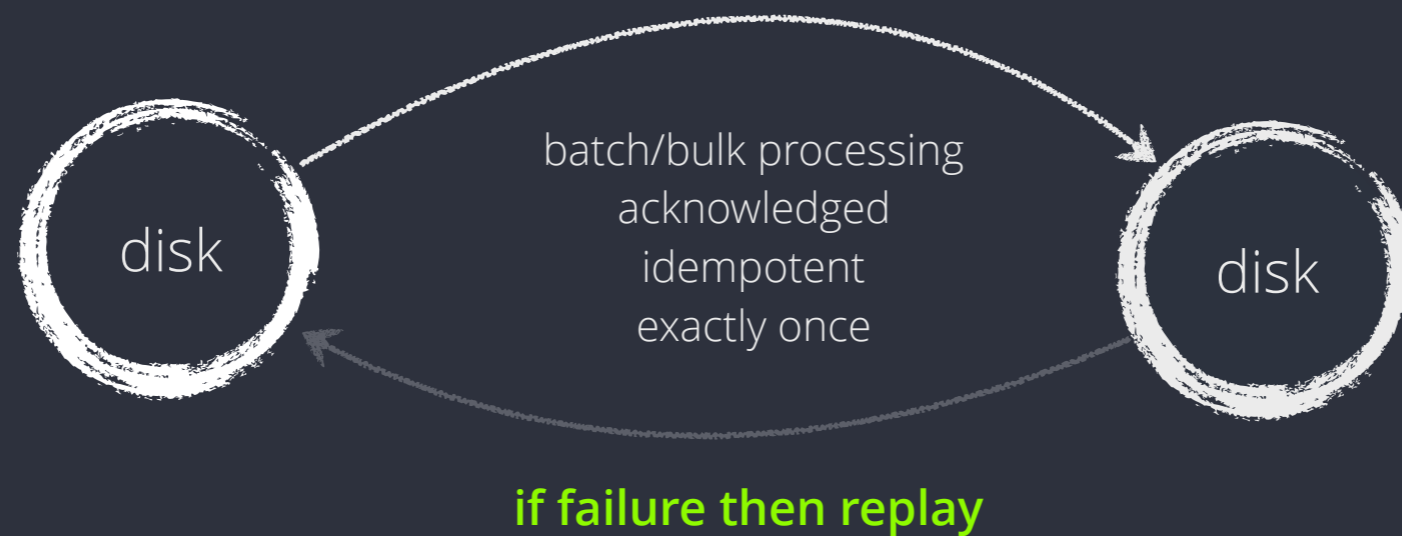




Stream or Batch Processing Architecture




In a nutshell : data processing is designed as follows



This requires : partition identifiers, timestamping, unique identifiers, batch identifiers, smart kafka offset handling, idempotent bulk file writing, on the fly efficient zero-copy compression, on the fly ciphering ... and of course real time supervision

... in a way manageable by the user. That is what the PunchPlatform provides.



 Deployment & Operations



PunchPlatform Benefits : Start from empty servers



What you need are plain linux servers, with local disks only.

Use dedicated hardware, VMs, Clouds. Whatever.

If you build your own infrastructure use the PunchPlatform infrastructure ansible libraries. Otherwise use Amazon, OpenStack, or any IaaS tool you have.

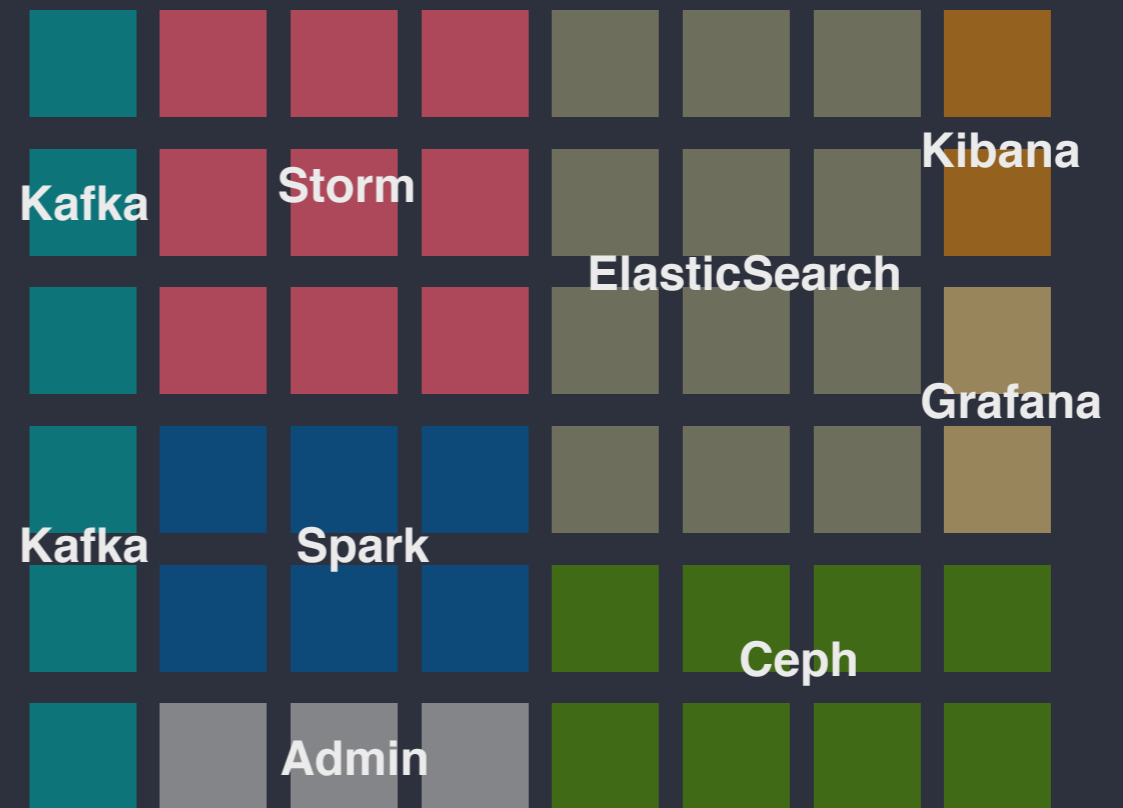




Deploy your services in minutes



Describe your setup in a configuration file. Use the PunchPlatform deployer to set it all.




Platform configuration file



Deploy your business logic in minutes



Describe your data channel in a configuration file.
Use the PunchPlatform channel command to set it all.
It can be a log parsing pipeline, a scada metric pipeline, whatever.



Channel configuration file




Resize Your Platform in minutes



Give more power to you platform, to scale up to you needs. Do that without service interruption.



Platform configuration file






Resize you Processing in seconds



Give more power to you data channel, to scale up to you needs. Do that without service interruption.



Channel configuration file






Real Example Setups



Connecting Your Sites

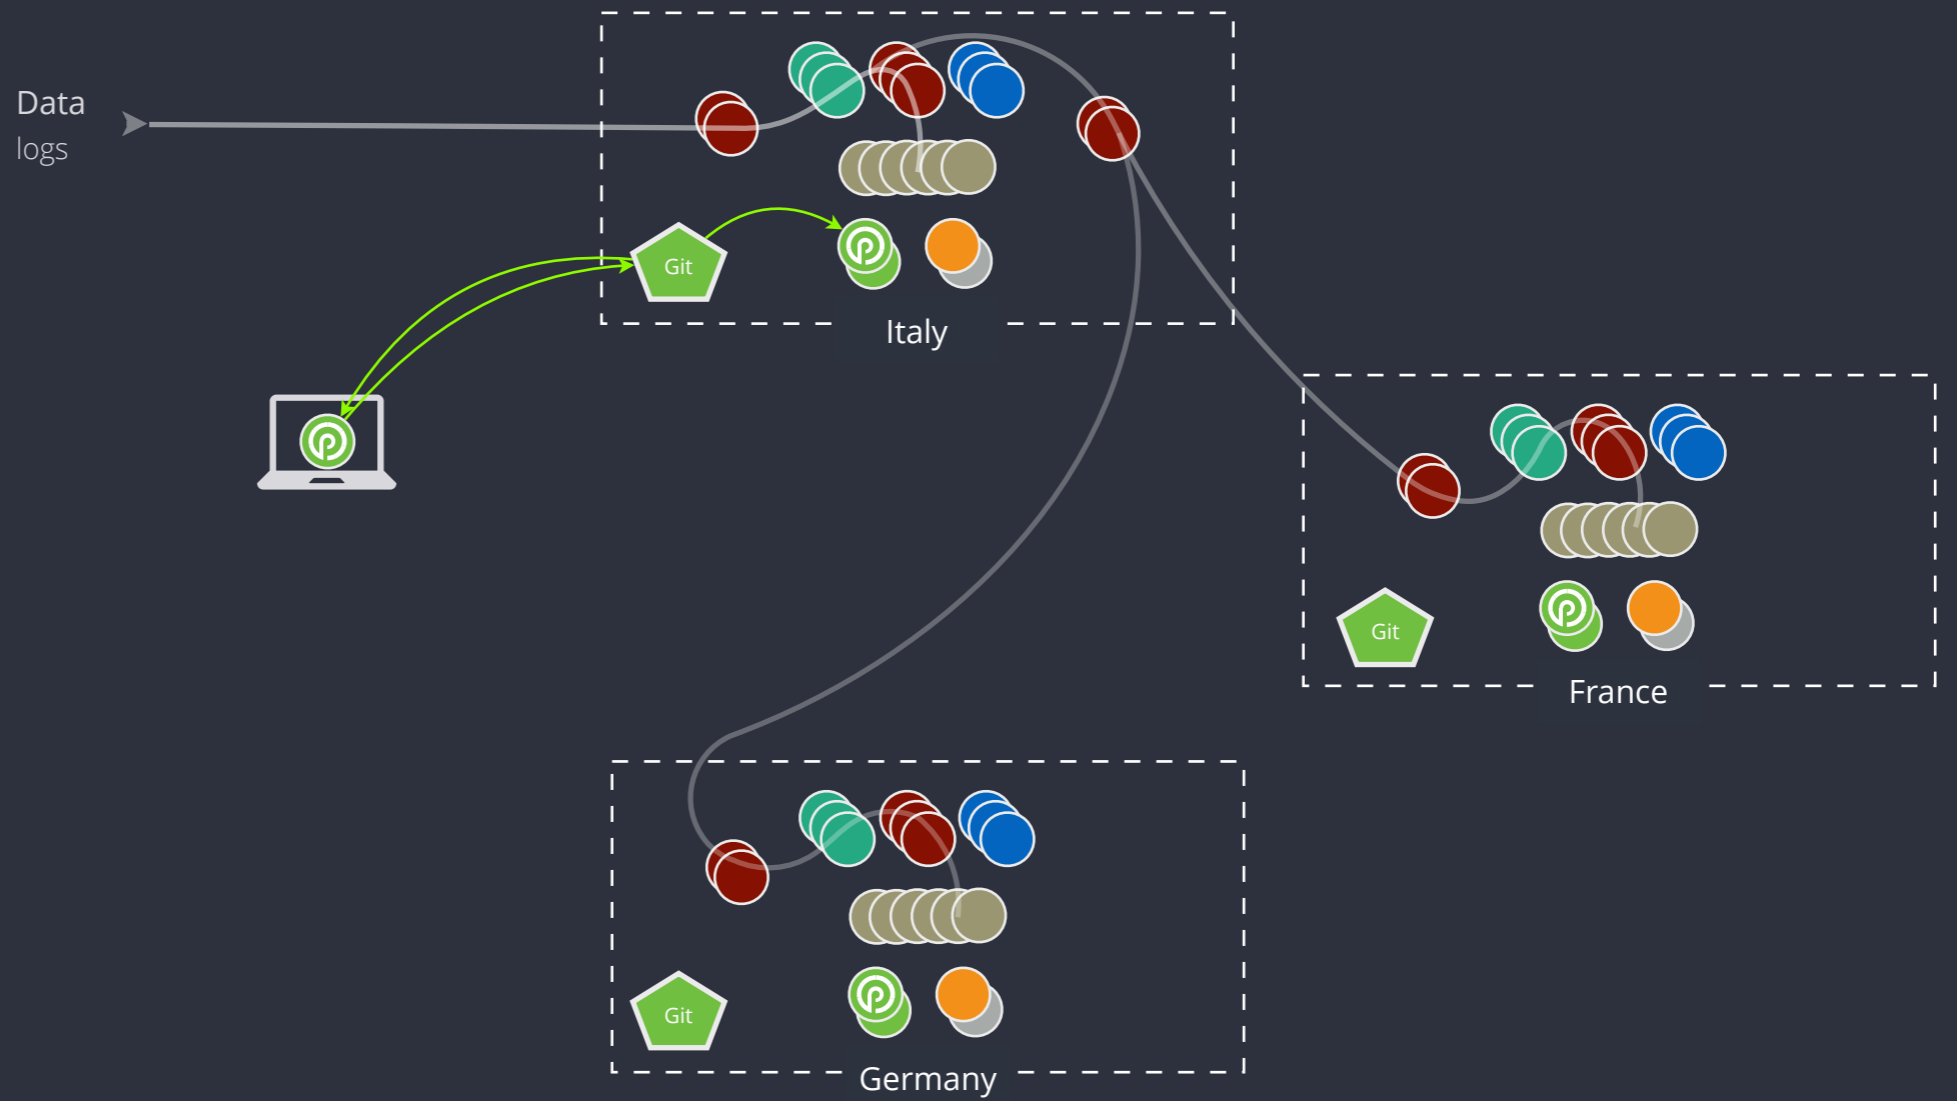
CyberSecurity Platforms : Holland, France, Honk Kong





Multi Sites Replication

CyberSecurity & Supervision Platform : Toulouse, Thales Avionics




dns/ldap/etc storm zookeeper kafka elasticsearch ceph



Small Scale Deployment

Transportation Monitoring Systems : Toronto



dns/ldap/etc storm zookeeper kafka elasticsearch



Customers and Use Cases



Data Agnostic :

- . Logs, XML, json, text
- . Snmp traps, Netflow metrics, time series

Open Platform : powered by

- . ElasticSearch
- . Cassandra
- . OpenTsdB
- . CEPH
- . Storm
- . Spark
- . Kibana/Grafana/Zeppelin

Wide Range of Business Cases :

- . CyberSecurity
- . Industry 4.0 : manufacturing, transportation, energy
- . Monitoring : IT infrastructures





Summary & Roadmap



PunchPlatform Stack



Parsers and Channels
Simple, straight, industrial

Stream Processing
Takes care of simple but key stream processing

Deployer/Updater
Install, Update. Fully documented.



Elastic Stack
all of it

Monitoring
end to end

Machine Learning
By configuration or by coding

Archiving
Long term storage. Secured



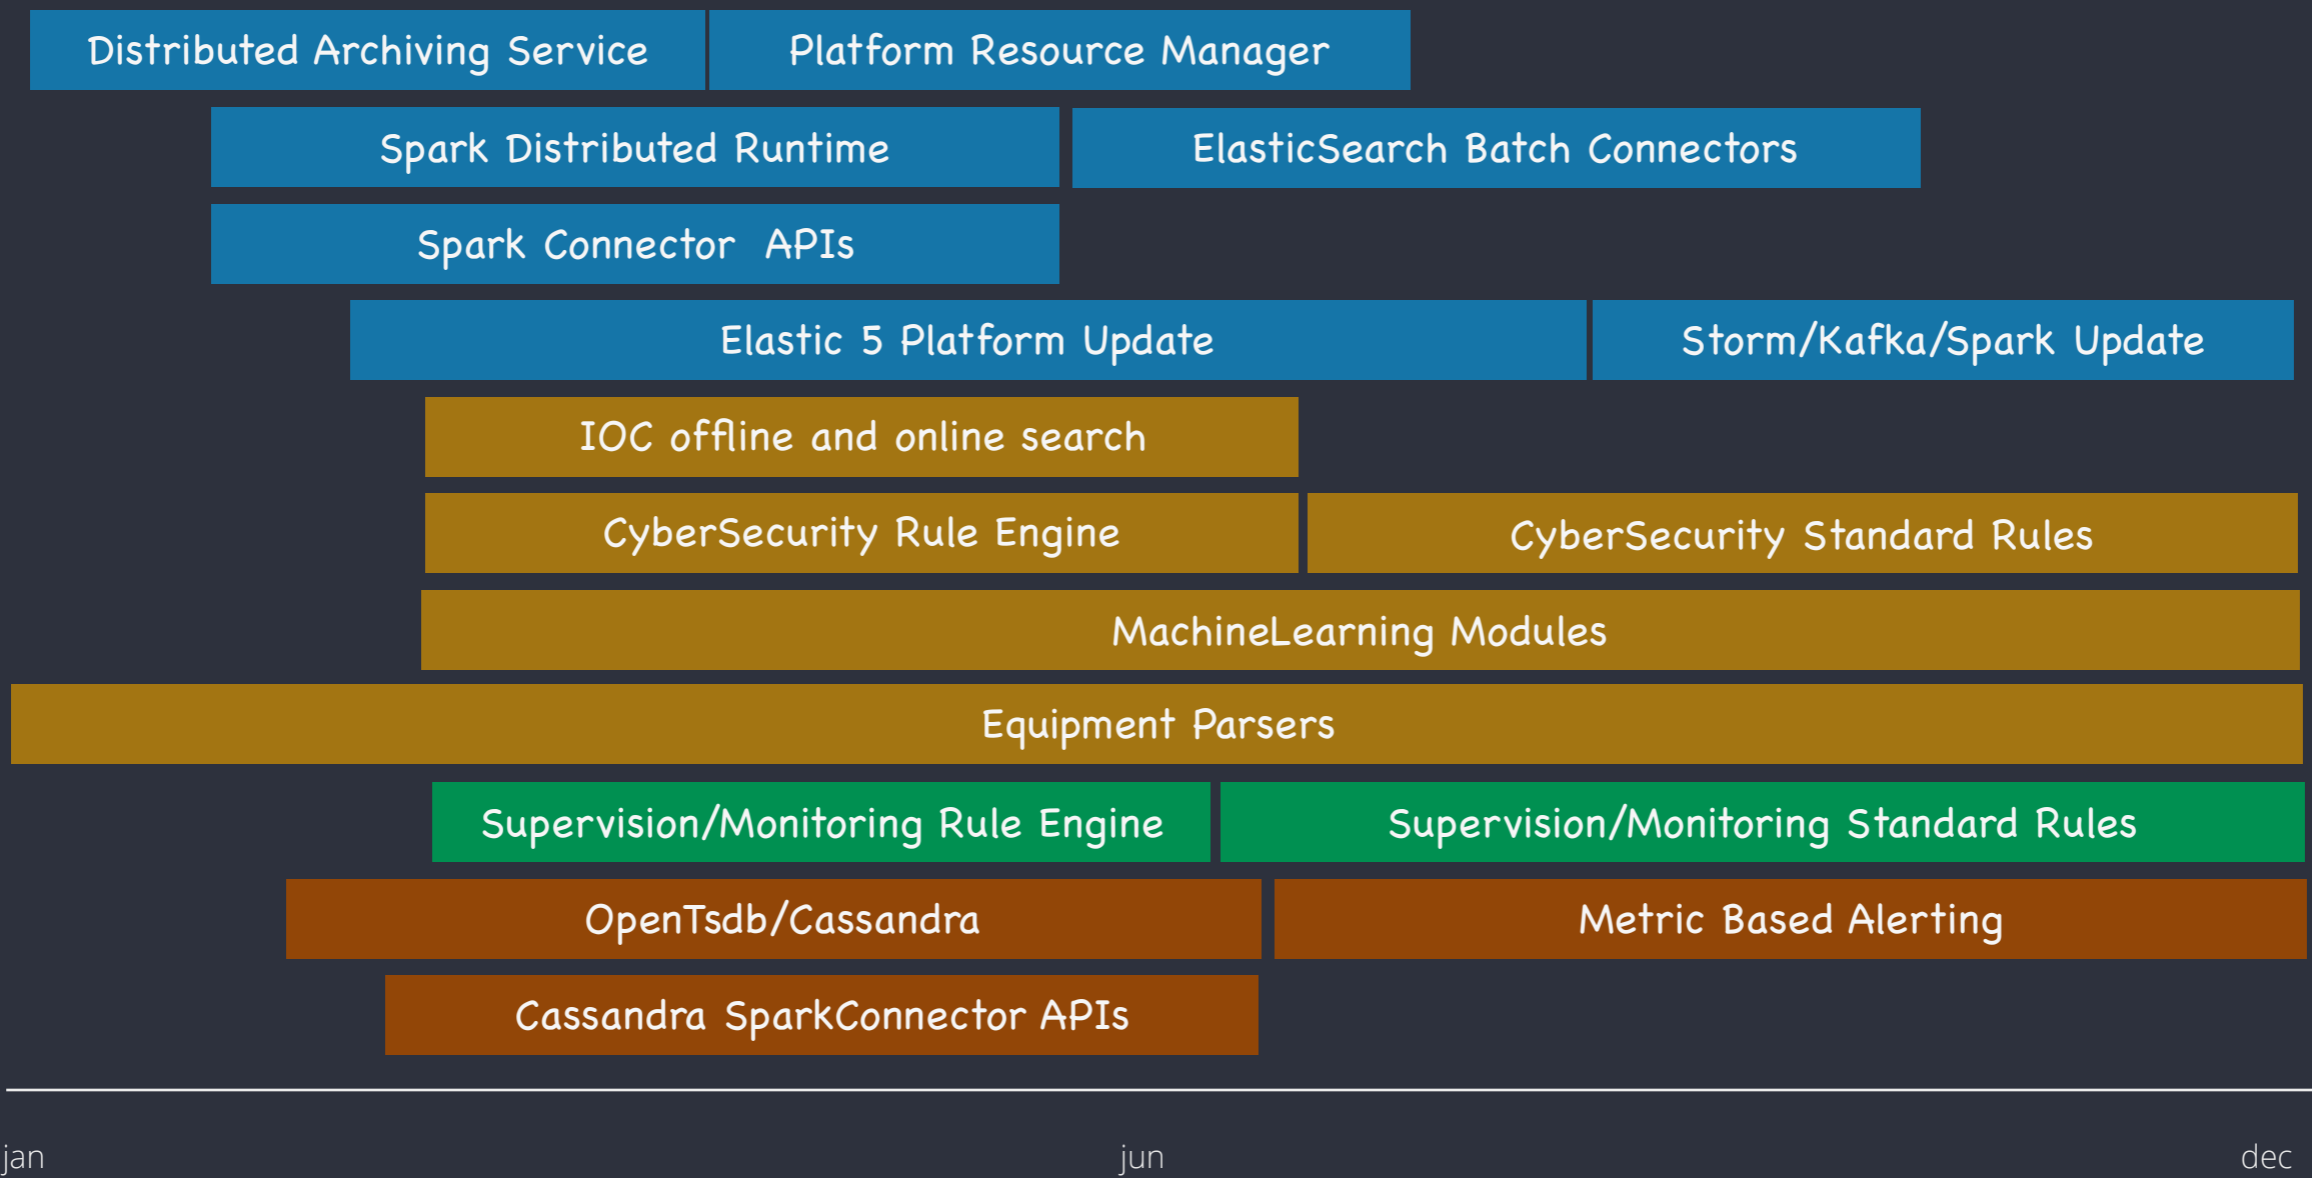
2017 RoadMap

Platform

CyberSecurity

Supervision

Industry





2018 RoadMap

Platform

CyberSecurity



Spark Distributed Runtime

Elastic 6 (7 8 ..) Platform Update

IOC/Replay/Extraction Kibana Plugins

Equipment Parsers

MachineLearning

jan

jun

dec



Thanks !