GridPP
UK Computing for Particle Physics

Security & Traceability

David Crooks

GridPP35, Liverpool, 11th September 2015

Context

- Security incidents are an operational reality
- When security incidents occur, need information about the chain of events:
  - *Who, What, When, Where*
- Well established procedures for current process, practised through SSC events
- Cloud and VM environment presents new challenges to gather and monitor the necessary information, using new tools and methods.

WLCG Cloud Traceability WG

- October 2014 call for volunteers to do practical work
- Face to face meetings at CERN February & June
- Many sites and all main LHC VOs represented, with good UK representation
- Areas of interest identified

Areas of interest

- Syslog
- Quarantine
- Externally observable behaviour
- Log management
- VO logging
- Policy

Syslog

- Logging from inside VMs
- Requires trust relationship between sites and VM providers (VOs)
- Evaluate merits of using job features vs. contextualisation to pass syslog configuration info

RAL, Andrew McNab

Quarantine

- Quarantining of VMs towards investigation of bad actors and forensics.

- Implementation on different cloud platforms

- Volume can be a problem with large image rate

RAL, Glasgow

Externally observable behaviour

- Hypervisor and Netflow
- Potentially strong uses in correlating events
- Might be new for many sites

Raul, David

Log management

- Increasingly high volumes of logs need work to manage appropriately

- Shared issue with “standard” operations

- Elasticsearch

RAL and others

VO logging

- All of this work will require closer links with the VO security teams

- Need to test the information logged by VOs

- SSC challenges

Sven Gabriel

Policy

- Policy evolution tracking work
- Follow up depending upon outcomes of other actions

Dave Kelsey

GridPP35, Liverpool, 11th September 2015

CERN SOC

- Liviu Vâlsan talked at HEPSYSMAN and elsewhere about the CERN SOC project

- Suite of tools to ingest, analyse and potentially act on a range of data sources - (*see http://lvalsan.web.cern.ch/lvalsan/presentations/workshops/UK_HEP_SYSMAN_2015/#/title*)

- Telemetry Capture Layer, Data Bus (Transport), Stream Processor, Long-Term Data Store, Real-Time Index and Search, and Visualisation

CERN SOC

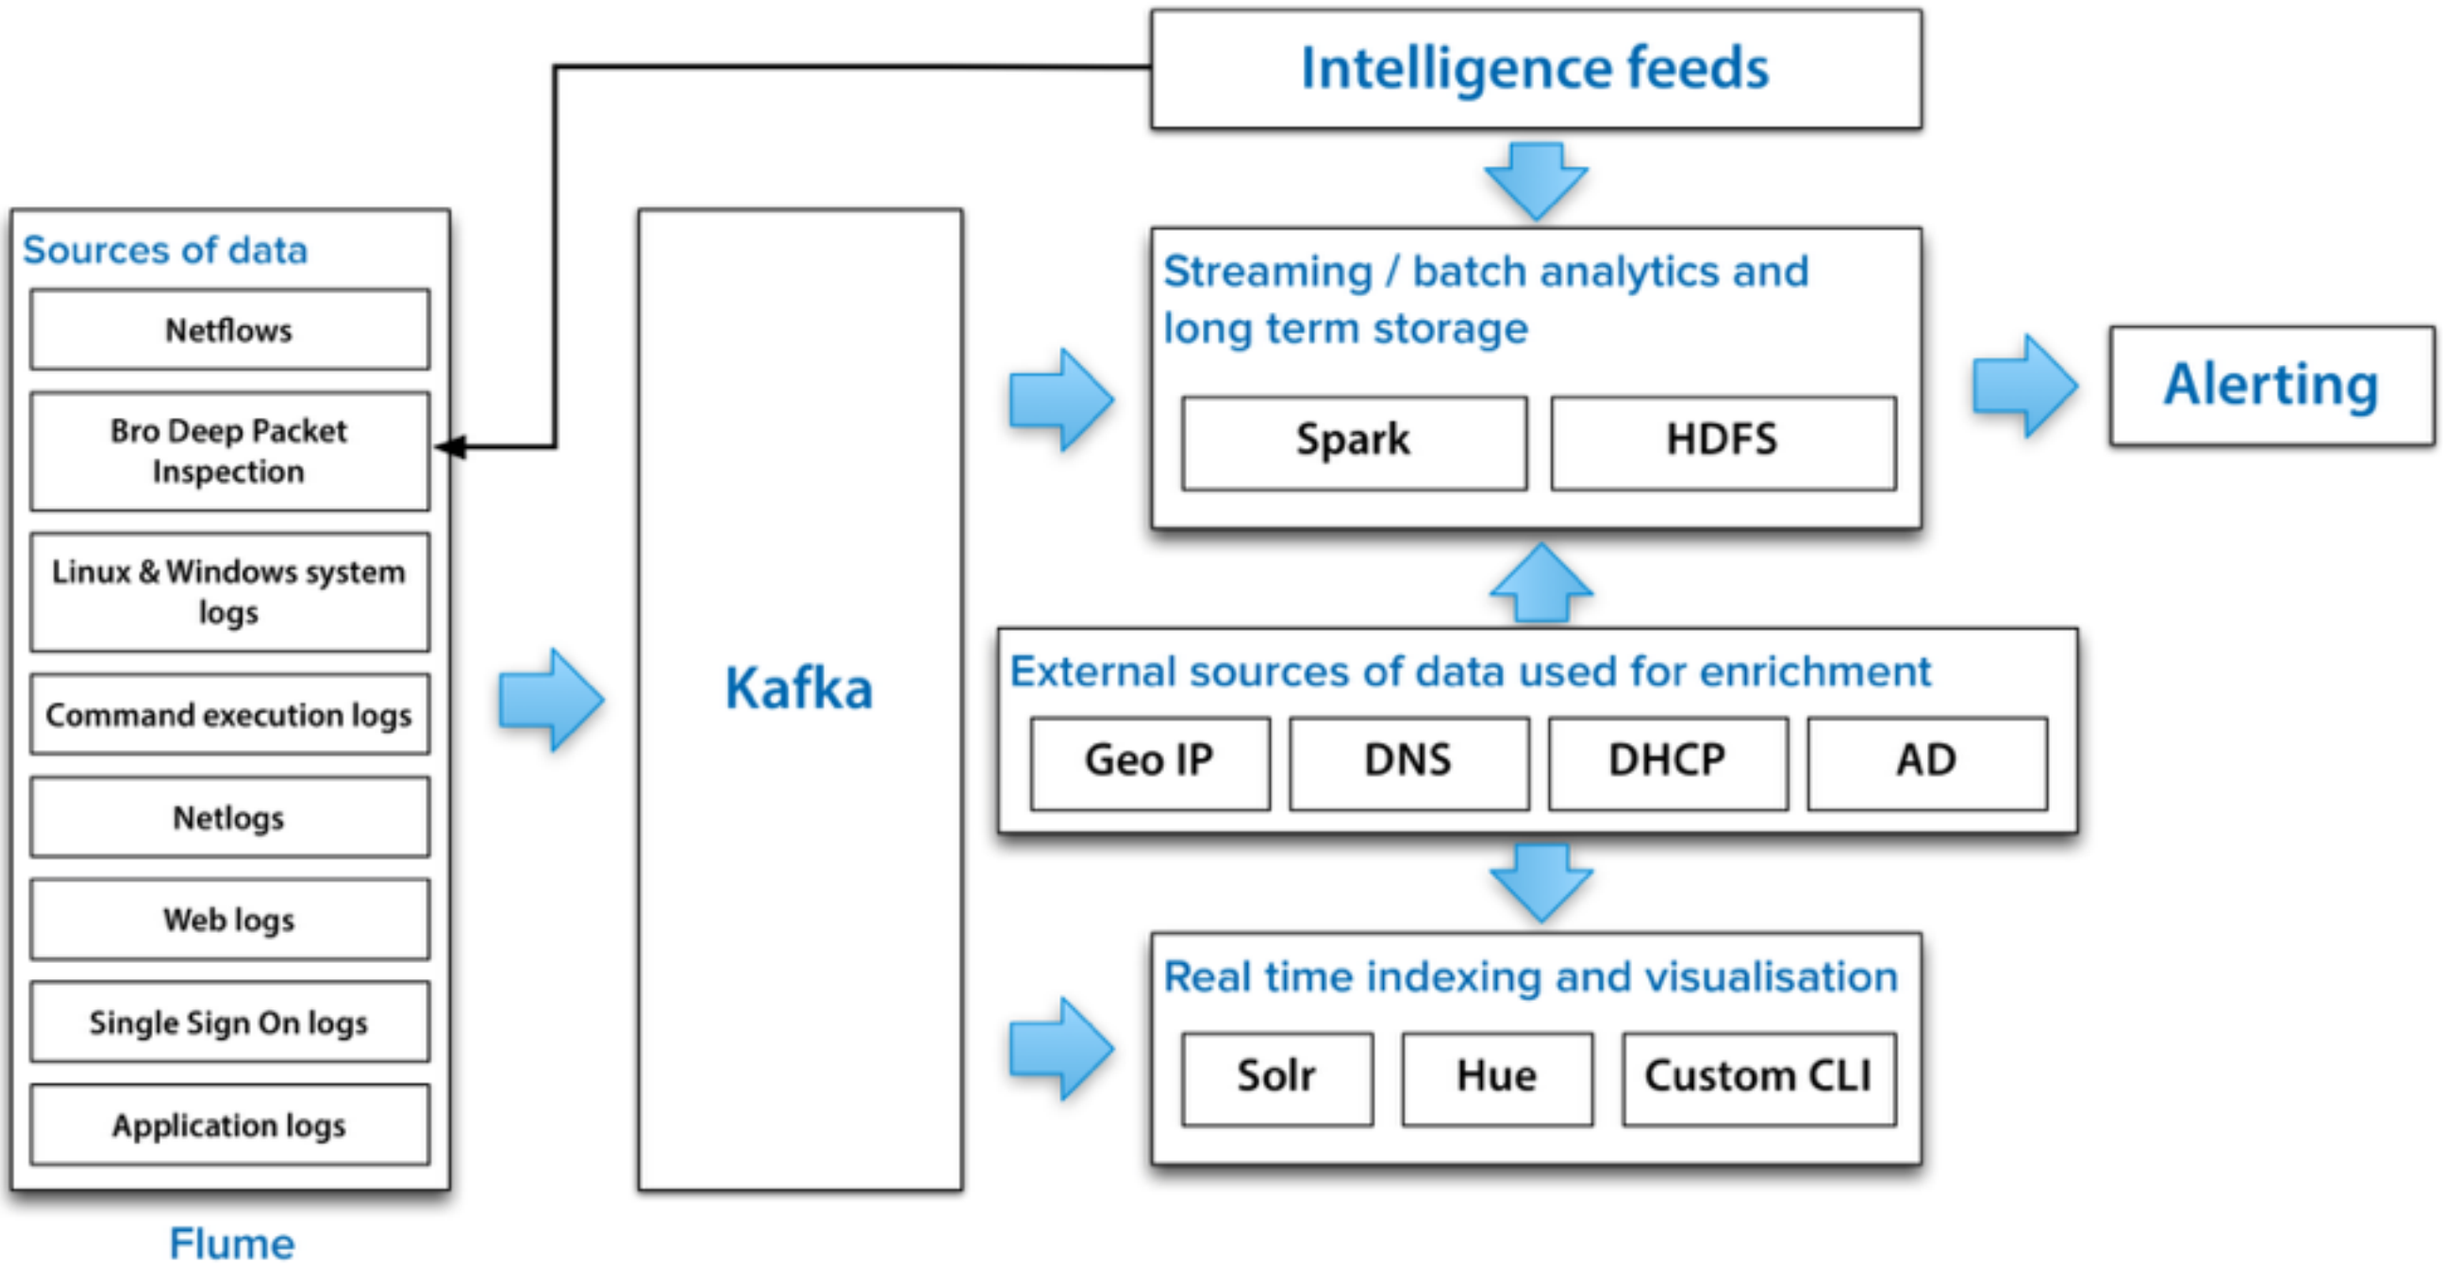
Intelligence feeds
Sources of data
Netflows
Bro Deep Packet Inspection
Linux & Windows system logs
Command execution logs
Netlogs
Web logs
Single Sign On logs
Application logs
Flume
Kafka
Streaming / batch analytics and long term storage
Spark
HDFS
Alerting
External sources of data used for enrichment
Geo IP
DNS
DHCP
AD
Real time indexing and visualisation
Solr
Hue
Custom CLI

Data rates

- Based on CERN SOC work
  - Netflow: ~ 25 GB/day
  - System logs, execution logs, netlogs, application logs: ~ 450 GB/day
  - Logging of HTTP connections: ~ 5 GB / day
  - KDC logs, SSO logs: ~ 2 GB / day

Feeds

- One area of cooperation noted as potentially useful by CERN Security is in intelligence feeds; they intend to subscribe to these and could share with us (except commercial sources)

OpenSOC

- Open sourced by CISCO last year
- “an extensible and scalable advanced security analytics tool”
- Similar to CERN approach - differences are partly down to CERN specific concerns
- Romain Wartel noted that commercial providers converging on OpenSOC + BRO IDS
    - OpenSOC is still immature, considerable work to implement but long term for many sites only sensible option

opensoc.github.io

OpenSOC

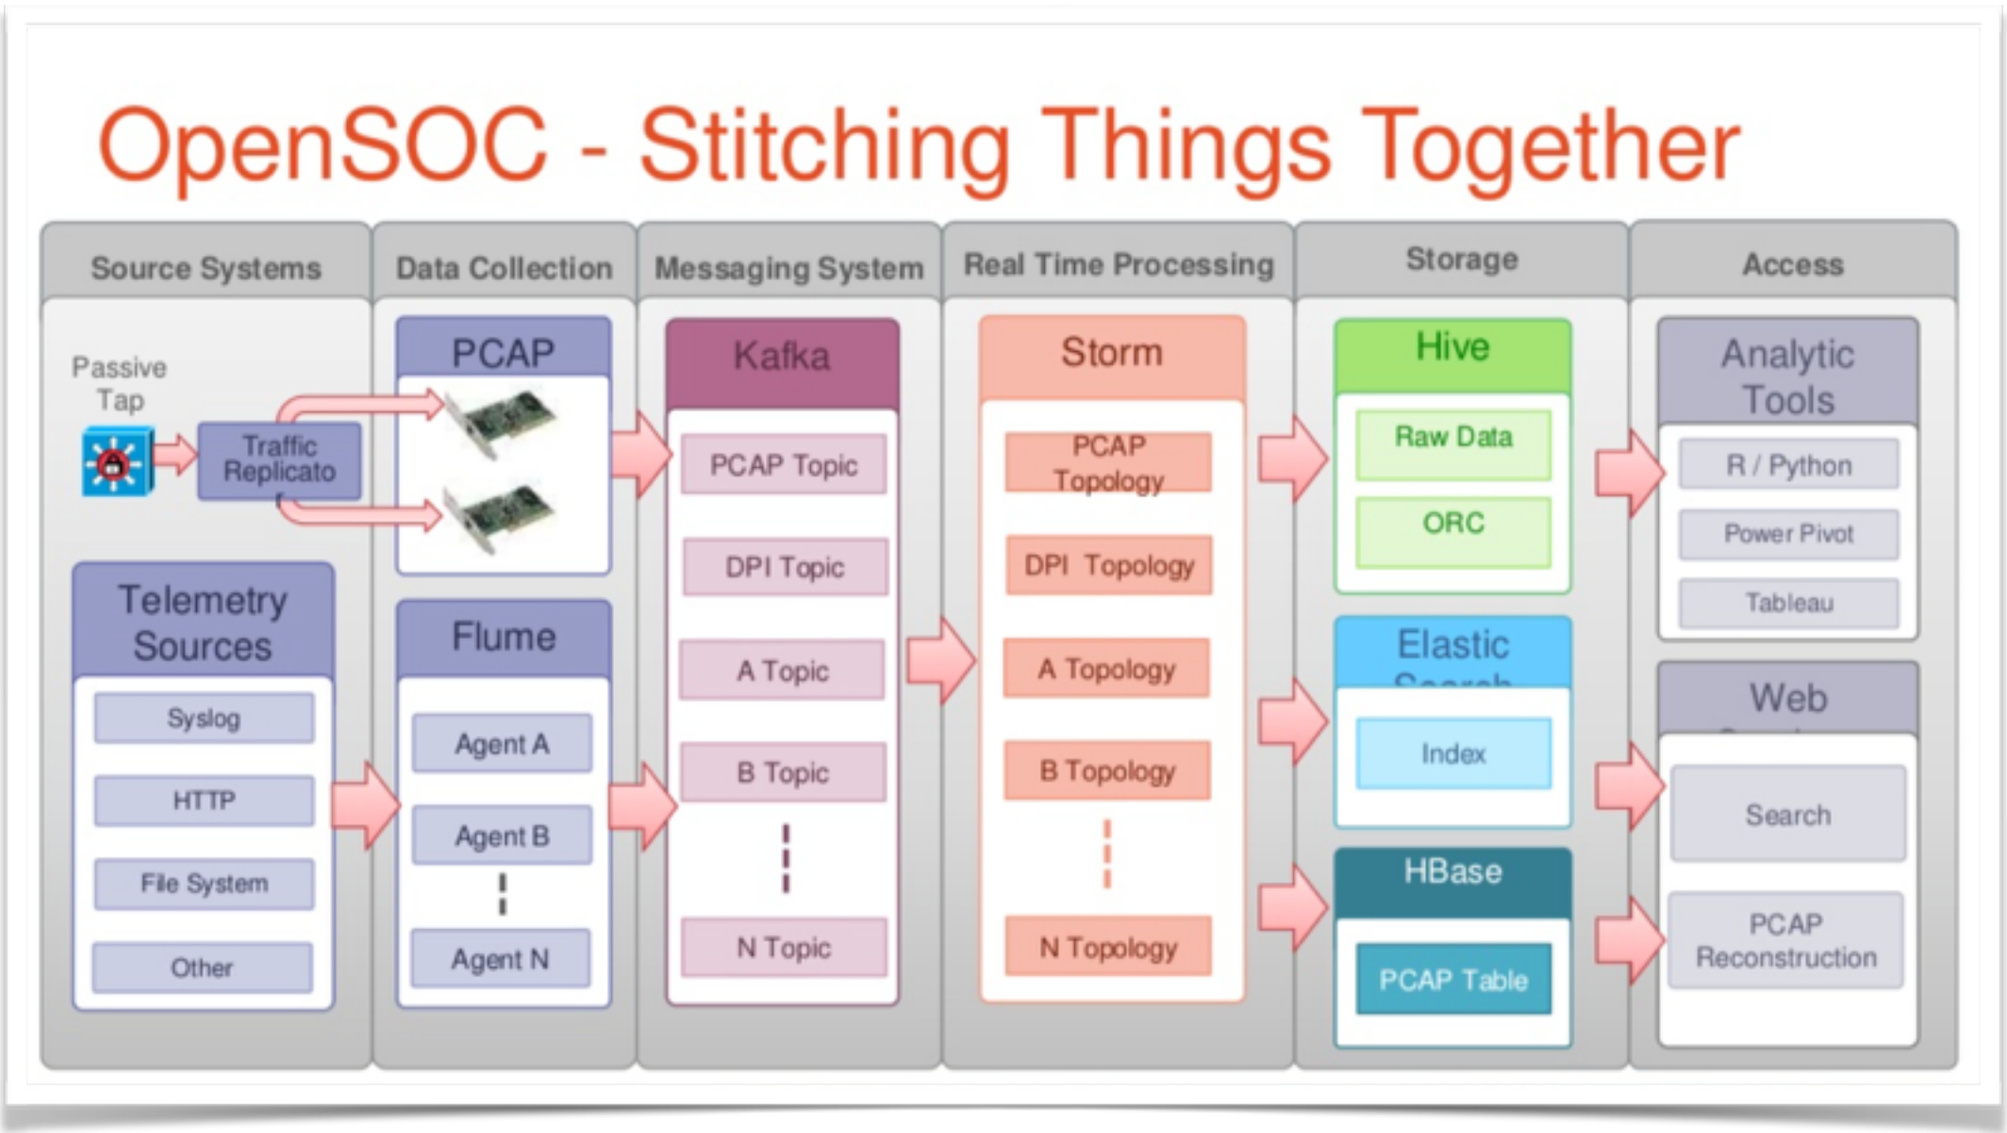
OpenSOC - Stitching Things Together
Source Systems
Data Collection
Messaging System
Real Time Processing
Storage
Access
Passive Tap
Traffic Replicato
PCAP
Kafka
Storm
Hive
Analytic Tools
PCAP Topic
PCAP Topology
Raw Data
R / Python
DPI Topic
DPI Topology
ORC
Power Pivot
Tableau
Telemetry Sources
Flume
A Topic
A Topology
Elastic
Web
Syslog
Agent A
B Topic
B Topology
Index
HTTP
Agent B
Search
File System
HBase
N Topic
N Topology
PCAP Reconstruction
Other
Agent N
PCAP Table

OpenSOC

- https://github.com/OpenSOC

- Includes a vagrant config for a test setup - currently evaluating this (somewhat complete stack though not entirely)

- In active development, but potentially not packaged for general installation at this time

- Participation welcome!

Site concerns

- Tools & strategies
- Scale
- Engagement with Central IT teams and policies

Summary

- Now is the time to have this discussion as cloud technologies are being more widely used
- Suite of tools needed to cover new environment
- Engagement with VOs critical to success
- More participation in the working group welcome - intending to have meeting before end of year