# Cloud Traceability Working Group
# Face to Face Report

Ian Collier
ian.collier@stfc.ac.uk
GDB 10th June 2015

Agenda & notes: https://indico.cern.ch/event/396202/

# Areas of investigation from Feb F2F

- Externally observable behaviour - Hypervisor & netflow logging
- Logging from inside VMs (no update here)
- Better tools for managing large volumes of logs
- Deferred Deletion or Quarantining VMs
- VO Logging (to be driven by security service challenges)
- Policy evolution (to follow on from other work)

# Actions/Areas of investigation from Feb F2F

- Externally observable behaviour - Hypervisor & netflow logging
  - Raul Lopes & David Crooks
- Logging from inside VMs
  - RAL – connecting VMs to central loggers
  - Andrew MacNab – investigating methods for passing logger info – machine/job info vs cloudinit
- Better tools for managing large volumes of logs
  - RAL & others
- Deferred Deletion or Quarantining VMs
  - RAL investigating for OpenNebula & CEPH
- VO Logging
  - Service challenges – Sven Gabriel
- Policy evolution
  - Dave Kelsey – following other work.

# Externally Observable Behaviour – David Crooks

- Often neglected at our sites but esp in context of virtualisation only thing we can depend upon
- Presented survey of technologies (see slides for details)
- Network flow protocols – netflow & sflow
- Concerns about propriety tools, hardware dependencies (support in switches/routers, network capture cards etc.) and software tools
- Experience of several tools at sites surveyed:
  - Nfsen (Nikhef & Glasgow), ZNetS (IN2P3), Silk (JANET), CERN building Security Operation Centre along lines of OpenSOC – see later discussion
- Software tools for generating network flows
  - softflow (Glasgow), fprobe (CERN)
- Some discussion of aggregation/analytics tools picked up under management tools
- Some discussion of separating institutional vs grid data – not an issue everywhere
- Next steps
  - More detailed evaluation and practical investigation

# Log management tools – Ian Collier

- David Crooks already noted significant effort investigating ELK (elasticsearch, logstash, kibana) stack – particular reference to UK but widely seen elsewhere too
- RAL have ELK infrastructure (on back of castor work), migration to new hardware longer than planned. In place now. Next step to send logs from cloud into that.
  - RAL have done low level work to tune ELK to better scale for their purposes, this will be ready to share later this year
- Staff from RAL visited IBM Research recently
  - As a side point saw their cloud security analytics tools – use some machine learning techniques. Looking to see if we can work together. We can offer very diverse environments.

# Security Service Challenges – Ian Collier

- At Feb F2F agreed the way to test gaps in VO logging for traceability purposes should be by service challenges
- Sven Gabriel (EGI CSIRT) is organising SSCs for EGI Federated Cloud
  - Also a commitment to apply those to WLCG cloud instances
- Ian noted there may be some specific issues with such challenges for sites using VAC

# Deferred Deletion/Quarantine – Ian Collier

- Want images from VM instances to persist for traceability/forensics purposes after VM is detroyed
- At RAL have so far implemented this using HTCondor
  - Easy managed by HTCondor but only applies to Condor managed VMs
  - Working on building into storage service – tricky to get hooks to work cleanly with OpenNebula – will involve developers if necessary
- Openstack
  - David Crooks has it on his plan
  - Tim Bell noted OpenStack has the functions to defer deletion but causes resource problems and full chain to do it is not there.
  - 300-400 images/hr on Cern production cloud so problems with volume.
  - Upstream would probably be happy to talk about this.

# Security Operations Centres – Liviu Valsan

- CERN developing Security Operations Centre – much in common with Cisco's open sourced OpenSOC
  - See slides for details
- Aggregates range of logs, network flows, other intelligence sources in to ELK like framework
- Romain Wartell 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
- Opportunity to collaborate on sharing (some) intelligence sources

# OpenSOC - Stitching Things Together

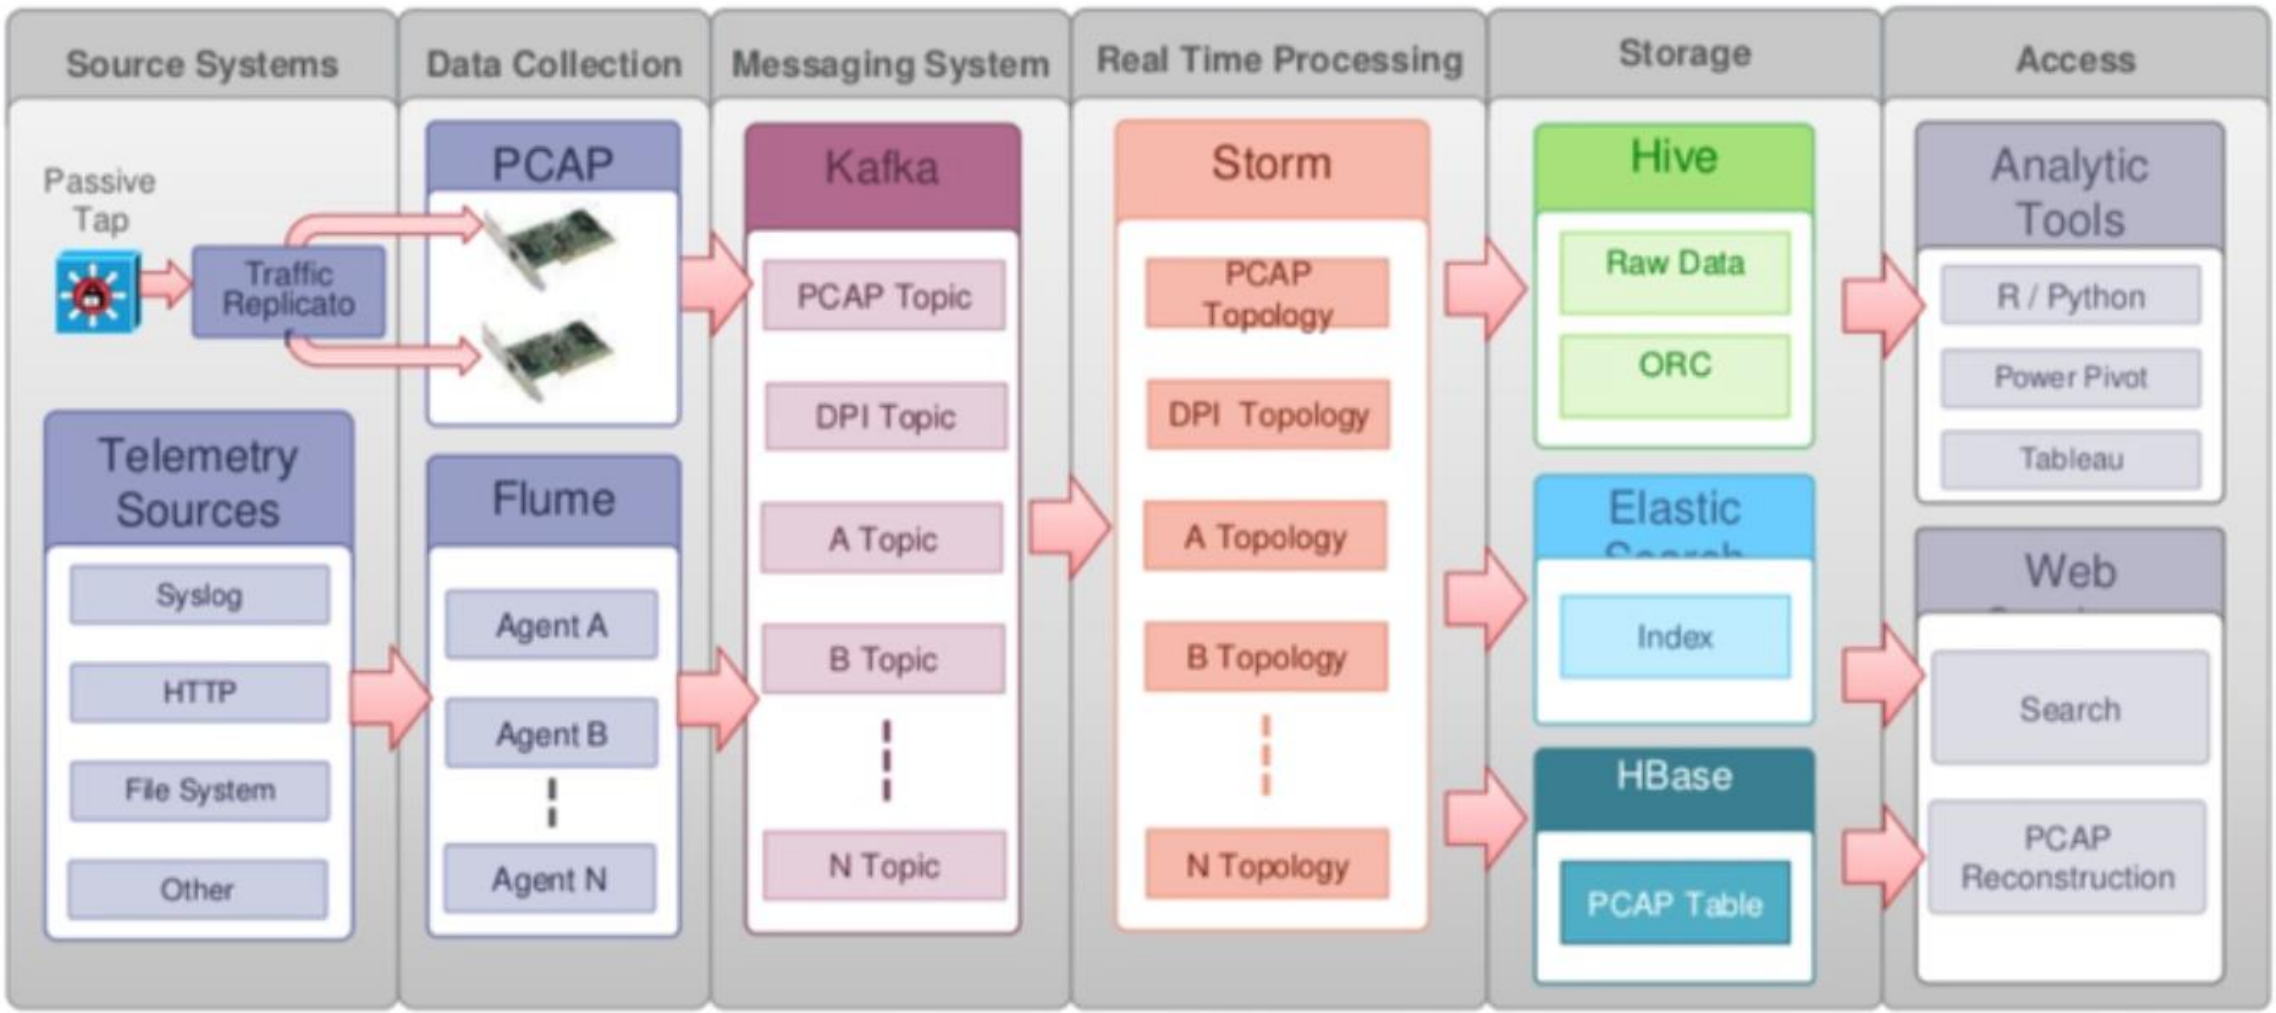


http://www.slideshare.net/JamesSirota/cisco-opensoc

# Summary

- Active work since February Face to Face
- Variable progress
  - Most in network flow/log management/SOC area
- Contact Ian Collier is you are interested in participating