Network flows

David Crooks
for WLCG Cloud Traceability TF

Introduction

- *Netflow* developed by CISCO to monitor network traffic

- *sflow* originally developed by InMon, used by several vendors

- Allow you to trace source and destination IPs, ports, packets and bytes transferred for network traffic

Benefits for traceability

- “External” source of data for work done in VMs
- Potential strong uses in correlating events

Concerns

- Vendor specific?
- Hardware sources?
- Software sources?

Tools

- nfsen (http://nfsen.sourceforge.net)
  - Nikhef (Dennis van Dok), Glasgow PPE (Andrew Pickford)
  - Frontend for nfdump (http://nfdump.sourceforge.net)
- ZNetS (http://www.znets.net)
  - IN2P3 (Jerome Bernier)
  - free for public institutions in France
- SOC stacks
  - CERN (Liviu Valsan), OpenSOC (http://opensoc.github.io)
- Silk (http://tools.netsa.cert.org/silk/index.html)
  - JANET (John Green)

Software flow generators

- softflow (https://code.google.com/p/softflowd/)
    - Glasgow PPE
- fprobe (http://fprobe.sourceforge.net)

nfsen

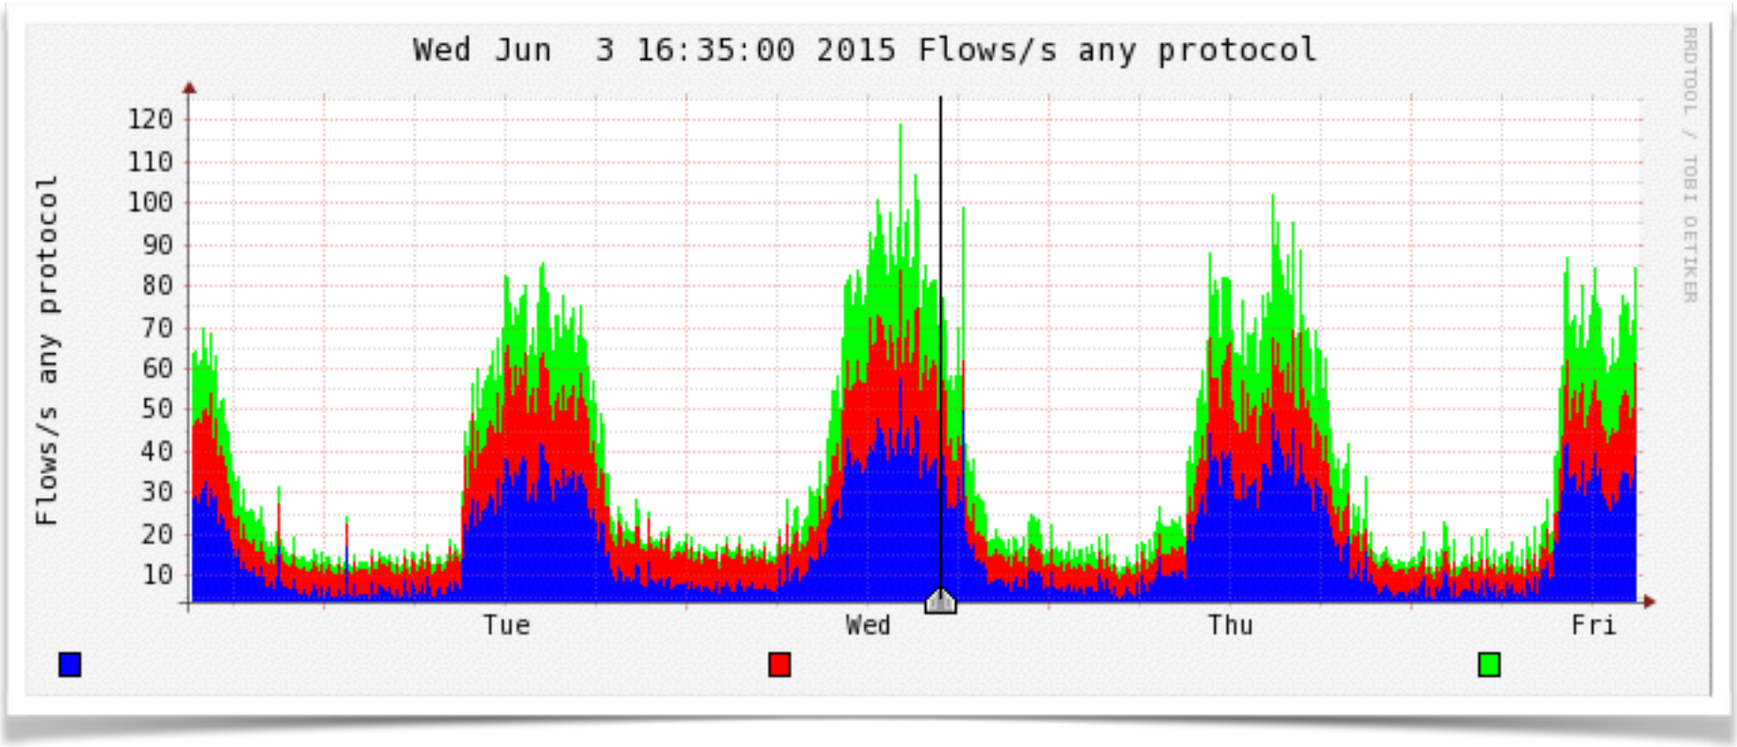

Netflow Processing

Source: **Filter:** **Options:**

All Sources **and** <none>

List Flows Stat TopN

Top: 10

Stat: Any IP Address **order by** flows

Limit: Packets > 0 -

Output: / IPv6 long

Clear Form process

Glasgow PPE

OpenSOC

- Open sourced by CISCO last year
- “an extensible and scalable advanced security analytics tool”
    - Apache Flume 1.4.0 +
    - Apache Kafka 0.8.1+
    - Apache Storm 0.9 +
    - Apache Hadoop 2.x (any distribution)
    - Apache Hive 12 + (13 recommended)
    - Apache Hbase 0.94+
    - Elastic Search 1.1 +
    - MySQL 5.6+

opensoc.github.io

OpenSOC

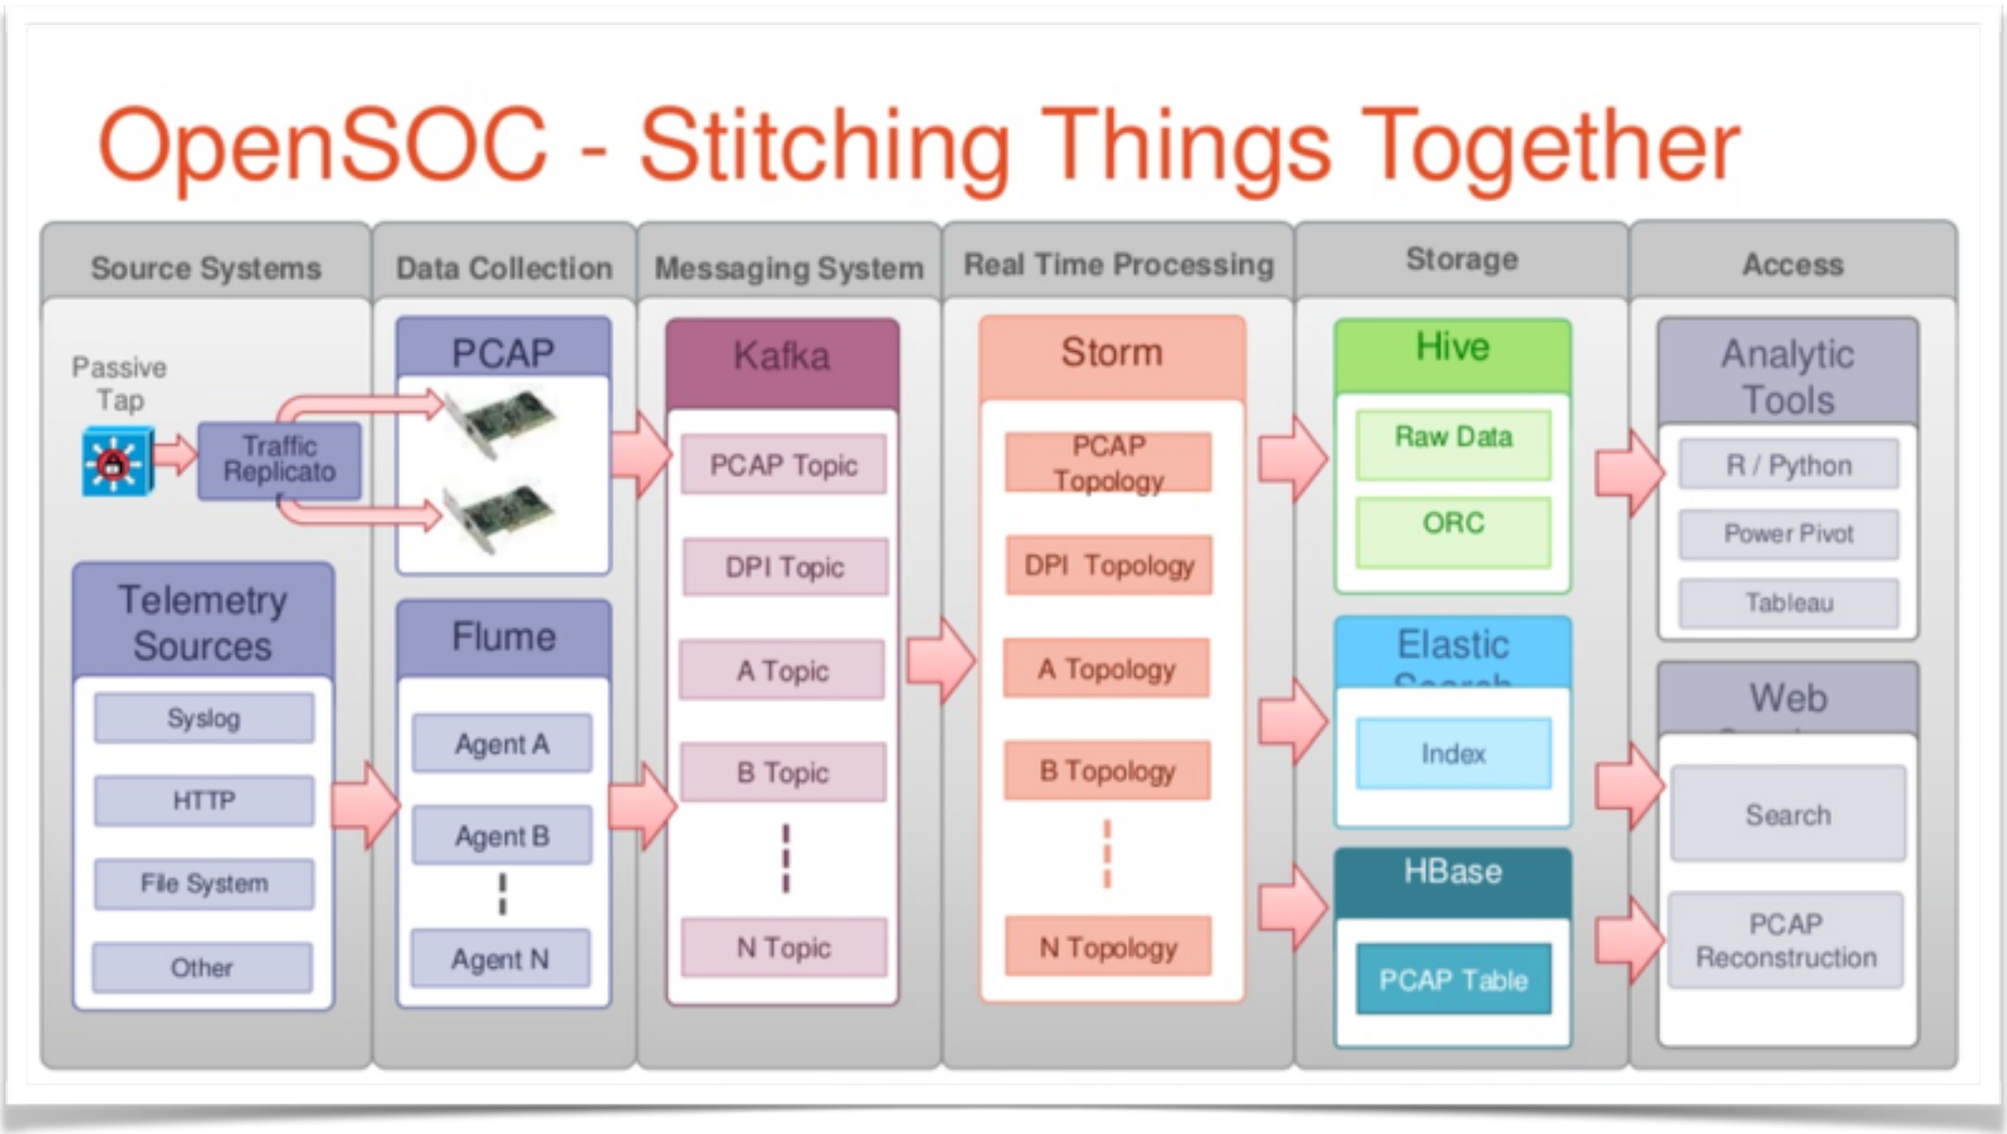

http://www.slideshare.net/JamesSirota/cisco-opensoc

Data rates

- Based on CERN SOC work (with thanks to Liviu; see later talk)
    - 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

(UK) side note

- Analytics
- Recent HEPSYSMAN meeting at RAL
- Considerable interest in, for example, ELK stack analytics

OSSEC & Analytics

- Glasgow has been been looking at OSSEC & ELK
    - Host IDS
- Complementary to this work; informs process

Next steps

- Testing
  - OpenSOC
  - nfsen, Silk
  - Software and Hardware sources
  - Common solutions
  - Data rates

Glasgow Cloud

- Dedicated cloud setup now in place after some delays
  - Controller, dedicated 12 TB storage, ~128 compute cores
- Software stack installation underway
  - OpenStack 2015.1.0 (Kilo)
- In this context will pursue work on syslog/netflow/containers