PROCEEDINGS OF SCIENCE

Harnessing the Power of Threat Intelligence in Grids and Clouds: WLCG SOC Working Group

David Crooks*
University of Glasgow
E-mail: david.crooks@glasgow.ac.uk

Liviu Vâlsan
CERN
E-mail: liviu.valsan@cern.ch

Kashif Mohammad
University of Oxford
E-mail: kashif.mohammad@physics.ox.ac.uk

Mihai Carabas
Politehnica University of Bucharest
E-mail: mihai.carabas@cs.pub.ro

Shawn McKee
University of Michigan
E-mail: smckee@umich.edu

Jon Trinder
University of Glasgow
E-mail: Jon.Trinder@glasgow.ac.uk

On behalf of the WLCG SOC Working Group

PoS(ISGC 2018 & FCDD)012

The modern security landscape affecting grid and cloud sites is evolving to include possible threats from a range of avenues, including social engineering as well as more direct approaches. An effective strategy to defend against these risks must include cooperation between security teams in different contexts. It is essential that sites have the ability to share threat intelligence data with confidence, as well as being able to act on this data in a timely and effective manner.
As reported at ISGC 2017, the Worldwide LHC Computing Grid (WLCG) Security Operations Centres (SOC) Working Group (WG) [1] has been working with sites across the WLCG to develop a model for a Security Operations Centre reference design. This work includes not only the technical aspect of developing a security stack appropriate for sites of different sizes and topologies, but also the more social aspect of sharing data between groups of different kinds. In particular, since many Grid and Cloud sites operate as part of larger University or other Facility networks, collaboration between Grid and Campus / Facility security teams is an important aspect of maintaining overall security.
We discuss recent work on sharing threat intelligence, particularly involving the WLCG MISP [2] instance hosted at CERN. In addition, we examine strategies for the use of this intelligence, as well as considering recent progress in the deployment and integration of the Bro Intrusion Detection System (IDS) at contributing sites.
An important part of this work is a report on the first WLCG SOC WG workshop / hackathon which took place in December 2017. This workshop provided an opportunity to assist participating sites in the deployment of these security tools as well as giving attendees the opportunity to share experiences and consider site policies as a result. This workshop also played a substantial role in shaping the future goals of the working group, as well as shaping future workshops.

International Symposium on Grids and Clouds (ISGC) 2018 in conjunction with Frontiers in Computational Drug Discovery
16-23 March 2018
Academia Sinica, Taipei, Taiwan

*Speaker.

1. Introduction

The modern security landscape for grid and cloud sites continues to evolve to present threats from a range of sources. As new computing models using virtualised environments become more prevalent, the task of effectively monitoring these sites in a security context becomes both more important and more complex. The ability to aggregate important security data in a common framework is an important one, leading to the use of Security Operations Centres (SOCs). As reported on previously [1], the WLCG SOC Working Group has a mandate to investigate different models for SOCs and give advice to the WLCG community on recommended configurations. In this paper we report in particular on the outcomes of the first SOC WG workshop held at CERN [3] on December 11th - 12th 2017.

2. Context

The work of the SOC WG follows in many respects the work of the CERN Security Team in developing a large scale SOC able to process 5 TB of security logs / day. An operational diagram is shown below in figure 1.

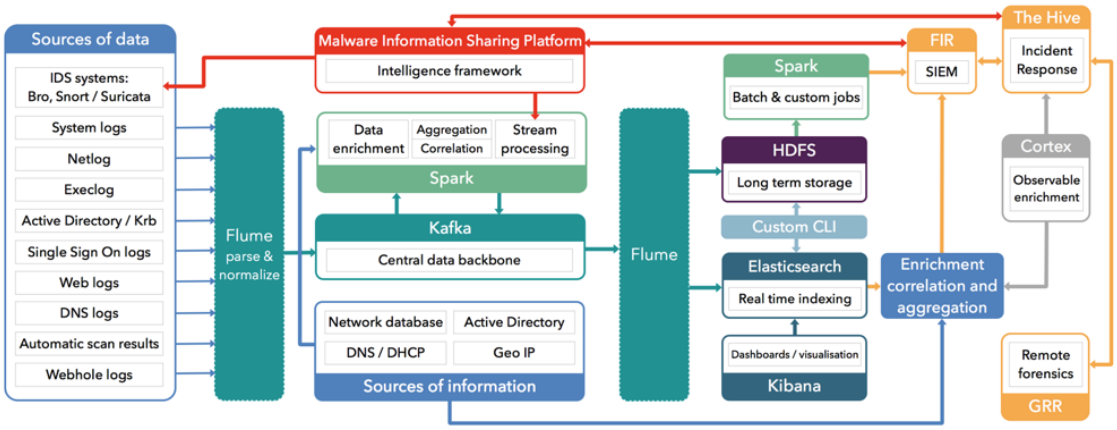

Figure 1: Operational diagram of the CERN SOC

More details on this work can be found in *Building a large scale Intrusion Detection System using Big Data technologies*, also in this issue [4]. In the SOC WG in general and in this paper in specific, we focus on the Intrusion Detection System (IDS) source of data and the Malware Information Sharing Platform (MISP) intelligence framework.

3. Areas of work

As the working group has progressed, two main areas of work have emerged. These are the technology stack necessary to perform the function of a SOC, and the cultural change necessary to allow the best use of threat intelligence.

3.1 Technology stack

To help with defining the technology stack for this work, we can ask two questions:

- What is happening in my cluster?
- What events are taking place that I need to know about?

To address these, we use the IDS Bro [5] to monitor what is happening in a cluster, and the threat intelligence platform MISP to determine what events are taking place. These form the foundation for our stack, which will be added to in due course.

3.2 Social & cultural aspects

A key element identified as part of this work is the need for collaboration between different security teams, for example those belonging to grid sites and campuses. Although traditionally these teams have often worked separately, the nature of the types of security threats seen (including a preponderance of phishing attacks, for example) means that different security teams need to collaborate in order to harness the available expertise, knowledge and resources available.

One good example is that of CERN, who played a founding role in this working group and where there is one single Computer Security Team handling the operational security aspects across the entire organisation, including:

- Office computing (i.e. campus).
- Data centre computing (i.e. grid).
- SCADA and control systems (i.e. CERN's accelerator complex).

This is in contrast with the way Computer Security activities are being organised at many WLCG sites, having dedicated teams for campus security and grid security.

4. December 2017 workshop

The workshop took place at CERN on the 11th and 12th of December 2017. The structure was one afternoon of introduction (including a demonstration of the CERN SOC) along with a full day of hands-on technical work, with half a day for MISP and half a day for Bro. The workshop had 27 registered attendees, of which 17 were in person. 19 institutes were represented, from 8 countries. Following the work described in the previous paper [1], focus was given to the following broad areas.

- Deploying MISP (section 5.2)
- Syncing MISP to WLCG instance hosted in CERN (section 5.2)
- Discussing the formation of trust groups (Section 5.3)
- Deploying Bro (section 6.1)

5. Threat Intelligence & MISP

5.1 CERN & WLCG instance

At the core of the CERN SOC lies threat intelligence data, structured information on various ongoing and past security events. This includes Indicators of Compromise (IoC), e.g. malicious IP addresses, domains or file hashes (signatures) of various malware samples. This information is constructed from the results of the investigations of security incidents discovered at CERN, but also received from partner organisations. Through participation in vetted trust groups the CERN Computer Security Team is automatically exchanging threat intelligence information with peer organisations, spanning different fields (academia, industry, etc). The data exchanged covers Indicators of Compromise as well as Tactics, Techniques and Procedures used by the various threat actors or groups of threat actors.

As of March 2018 the CERN main MISP instance contains more than 600 000 Indicators of Compromise. All the threat intelligence information that CERN is allowed to share with third parties is being made available to participating sites in the WLCG SOC WG via the central WLCG MISP instance. The primary access method is eduGAIN [6], of which the SIRTFI [7] certificated subset is considered eligible. As of March 2018 that instance contains more than 310 000 Indicators of Compromise.

5.2 MISP Deployment

All WLCG sites present at the workshop were able to deploy MISP, following the documentation provided by CERN and working with their respective provisioning and configuration management systems. CERN Puppet modules were used as a reference which could be used either as part of a Puppet master / agent configuration or standalone. Most newly deployed MISP instances have been configured to sync with the central WLCG MISP instance.

5.3 Trust groups

A key area of the deployment of MISP is that of establishing trust groups. MISP represents the technical means of forming a trust relationship between different groups - between the WLCG central instance and a grid site, for example, or between different institutes or infrastructures.

Following the workshop, two examples of a discussion around the establishment of institutional trust group arose.

IN2P3

A distributed laboratory in France, the intention was to investigate how MISP could be incorporated into sharing threat intelligence between the component sites.

Glasgow

Additionally, the College of Science and Engineering at the University of Glasgow has been discussing how to share security information between member Schools, which could involve the use of MISP.

This collaboration is between 5 schools (Engineering, Computing Science, Chemistry, Psychology, Physics and Astronomy). Each has its own unique perspective on how security issues are handled. The participants have widely varying levels of knowledge and experience in operational security, threat intelligence sharing and importantly the language and terminology. It is believed that collaborating on a project around MISP (and other tools such as OpenVAS [8]) will provide incentive for frequent discussions around security and that this will provide opportunities to collaborate, share knowledge, learn from each other and gain a shared understanding of the terminology - this is an embodiment of a *community of practice* [9].

Beyond the support teams many computer users neglect security considerations until there is a problem locally or a high profile case reported in the media. A further aspiration is that by making it widely known to users that vulnerability scans and threat intelligence is being shared, users will be encouraged to inform their local support of suspicious activity, phishing attempts and increase communication. What has been a pleasant surprise is that pen testing scans inside schools have been welcomed (so far) by users when any shortcomings of their systems have been identified.

UK

In the UK, the GridPP [10] project (which provides the UK component of the WLCG computing resources) is discussing how a central MISP instance with additional site instances could work, potentially using a hierarchical deployment model. It is anticipated that contributing sites may use a mix of some with local instances and some using remote (API) access to the WLCG MISP instance.

6. Network Monitoring: Bro

Several of the WLCG sites present at the workshop deployed Bro, at least to the extent of having Bro workers running and logs being generated. More detailed site summaries are given in the next section.

6.1 Site configurations & Deployment

6.1.1 Oxford

Oxford is a medium sized WLCG Tier 2 site with 3 000 logical CPU cores and 1 PB of storage. The site is running a set of services typical for a WLCG site.

All of the service nodes are virtual machines (VMs) running on oVirt [11] virtualization hosts. Worker nodes and storage pool nodes are physical machines, with the bulk of the network traffic going to these machines. But the service nodes have more interesting traffic as they are exposed to the internet and running many web services so we decided, in the first instance, to mirror only the oVirt host port which gives insight into all service VMs. A simple schematic diagram of our setup is shown in figure 2.

All the traffic which passes through oVirt is mirrored to the Bro server. The logs produced by Bro are sent to the Elasticsearch, Logstash, Kibana (ELK) stack [12] through the use of FileBeat [13]. The data collected from Bro is further enriched by GeoIP tagging during Logstash parsing. Currently we are running using a single Elasticsearch node and it seems capable to process the

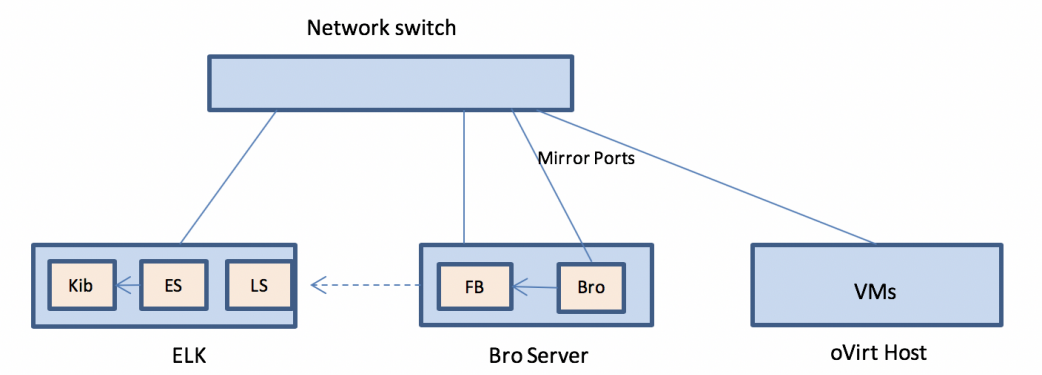

Figure 2: Oxford network monitoring setup

data generated by around ten VMs. Once the data is in Elasticsearch, we can use Kibana to query various combinations and look for interesting patterns.

6.1.2 University Politehnica of Bucharest

All production services at University Politehnica of Bucharest are virtualised, including network routers and firewalls. There is a virtual environment in each data center in University Politehnica of Bucharest comprised of:

- Nodes running hypervisors (Microsoft Hyper-V).
- Shared Storage Area Network (SAN) connected via FibreChannel or iSCSI.
- High Availability management component (Microsoft Failover Cluster Manager).

In the case where MISP and Bro have been deployed:

- One Lenovo compute chassis with 14 compute nodes.
- One SAN from EMC2 connected via iSCSI 10 Gbps to the compute nodes.

Microsoft Windows 2012 R2 Datacenter edition is running on each node, using Microsoft Failover Cluster in order to ensure high availability. The above setup is used for running different services, including MISP and Bro.

In order to increase throughput in virtualised environments, there are two approaches:

- Using the synthetic Microsoft network drivers and virtual queues (VMQ) [14].
- Using Single Root I/O Virtualisation (SR-IOV) for exposing hardware NIC queues to a VM.

The latter is preferred for performance reasons as it does not load the host, but is poses problems for live migrations of VMs.

In a virtual environment traffic mirroring can be performed in two ways:

- Mirror traffic inside the virtual switch [15].

- Dedicate a special secondary port for traffic mirroring on the host and enable SR-IOV on that port. This basically means that the mirroring will be done from the hardware switch.

These are the approaches used at University Politehnica of Bucharest, with Microsoft Hyper-V. The same can be achieved with VMWare or other virtualisation software stacks.

6.2 University of Michigan

The ATLAS Great Lakes Tier-2 (AGLT2) center at the University of Michigan and Michigan State University hosts over 7 PB of disk storage and more than 10,000 job slots for ATLAS and Open Science Grid (OSG) jobs. Almost all AGLT2 services run on VMware v5.5 virtual machines. The University of Michigan site is connected to the wide-area network (WAN) with two 40Gbit/s links in a load-balanced configuration. While there are many security best practices implemented in the service and systems deployment for AGLT2, the ability to monitor traffic coming into the site was previously missing.

The challenge for AGLT2 was how to best split out the network traffic for analysis by Bro. It turned out that the AGLT2 router, a Juniper EX9208, was unable to be configured to use SPAN or port-mirroring to get all the traffic. Because of this, options were investigated to insert passive optical splitters, many of which were costly (> 4000 USD). Fortunately a much more cost effective option from Fibrestore [16] was found as shown in the following parts list:

- Optical splitter module (handles 4 fiber-pairs) with customisable splitting ratio $200.
- Shelf to host up to 4 optical splitter modules $107.
- Various single mode patch cables (ST-LC, LC-LC) from 1m - 3m lengths ($2-$4 each)
- LR4 40G Optics (need two to monitor incoming WAN traffic) $300 each.
- Dual-port 40G NIC (had existing Mellanox ConnectX-3 card).
- Dedicated Bro analysis host (re-purposed existing Dell R630 node).

Total cost to monitor the incoming site traffic, including shipping and handling, was less then 1100 USD. Note that each module can have its splitting tuned from the default of 50% /50% to either 60% /40% or 70% /30% (the numbers represent the percentage of the light sent on the primary path versus the split-out path). For AGLT2 we selected a 60% / 40% split.

Shown in Figure 3 is the deployed optical splitter configuration. Because of the redundancy in the 2x40G LACP WAN connection, we were able to insert the optical splitter without a site downtime. Unfortunately the first attempt failed because the vendor failed to provide an accurate map of the fibers in the splitter module. Through subsequent testing we verified that the vendor incorrectly mapped the fibers, swapping ports 1 and 3 and ports 12 and 14, which prevented the use of duplex cables. The work-around was to purchase simplex patch cables to map the connections as required.

We measured the incoming light on the Juniper EX9208 both before and after introducing the optical tap. Table 1 summarizes the changes.

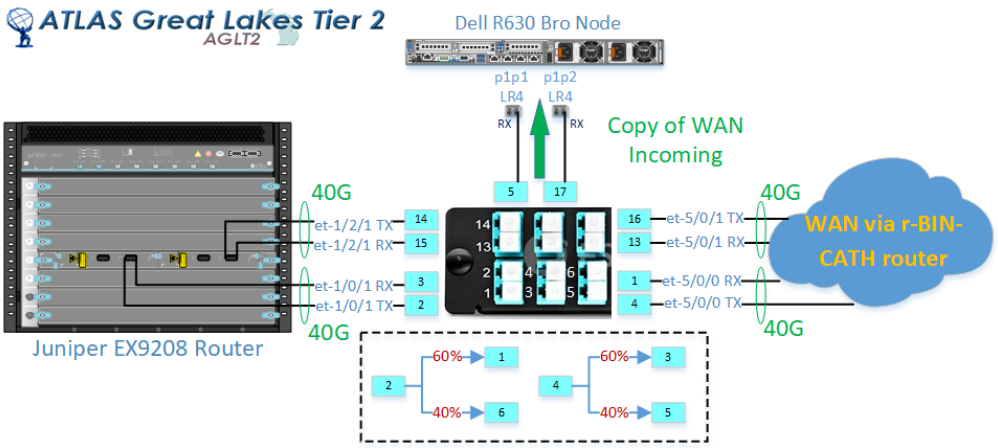

Figure 3: AGLT2 optical tap configuration splitting the incoming WAN 80 Gbps into a Bro monitoring node. Note input TX fibers are split 60/40 as shown in the two example maps in the bottom center.

Interface	Lane	Pre-tap Received	Post-tap Received
et-1/0/1	0	0.510 mW / -2.92 dBm	0.447 mW / -3.50 dBm
	1	0.502 mW / -2.99 dBm	0.426 mW / -3.71 dBm
	2	0.498 mW / -3.03 dBm	0.429 mW / -3.68 dBm
	3	0.439 mW / -3.58 dBm	0.360 mW / -4.44 dBm
et-1/2/1	0	0.445 mW / -3.52 dBm	0.291 mW / -5.36 dBm
	1	0.388 mW / -4.11 dBm	0.251 mW / -6.00 dBm
	2	0.319 mW / -4.96 dBm	0.205 mW / -6.88 dBm
	3	0.354 mW / -4.51 dBm	0.222 mW / -6.54 dBm

Table 1: Received light levels before and after optical tap insertion.

Once the incoming light was being split and sent to the two LR4 optics on the Bro system, the last step was to properly configure the network interfaces that are receiving the incoming network traffic. The system, a Dell R630 with 256 GB of RAM, dual E5-2680 v3 processors (12C/24T each) and a Mellanox ConnectX-3 dual 40G port NIC, is running CentOS 7.4 and has Bro installed following the WLCG SOC WG documentation. The Mellanox NIC was installed in slot 1 and the Linux network devices for the two 40G ports show up as p1p1 and p1p2. The Mellanox drivers and firmware was previously updated to the most current.

For CentOS 7, we need to create /etc/sysconfig/network-scripts/ifcfg-p1p1 and similarly for p1p2. Here is the configuration we used (p1p2 is identical except for the DEVICE and HWADDR lines):

```
DEVICE=p1p1
ONBOOT=yes
HWADDR=xx:xx:xx:xx:xx:xx
# Increase MTU for Bro captures
MTU=9100
BOOTPROTO=static
PROMISC=yes
```

```
USERCTL=no
# Persist ethtool settings for Bro use
# Increase receive ring to max,
# turn off pause frames; remove offloading
ETHTOOL_OPTS="-G \${DEVICE} rx 8192; -A \${DEVICE} rx off tx off;
-K \${DEVICE}  rx off tx off sg off tso off gso off lro off rxhash
off ntuple off txvlan off rxvlan off"
```

The ETHTOOL_OPTS defines the ethtool settings we want to persist. Specifically we are maximizing the receive buffer, turning off PAUSE frames and all the network offloading the NIC provides since we want to capture the original packets on the wire. The network interfaces can be brought up with 'ifup p1p1; ifup p1p2'.

The last thing we need to do is configure Bro to use these interfaces. Our Bro configuration files are in /opt/bro/etc. The node.cfg file is the one which needs to be configured to define the capture interfaces. We are using pf_ring and about 2/3 of our processors to capture and analyze traffic. Below is the relevant configuration snippet from node.cfg:

```
...
[bro.aglt2.org-p1p1]
type=worker
host=bro.aglt2.org
interface=p1p1
lb_method=pf_ring
lb_procs=8
pin_cpus=0,2,4,6,8,10,12,14

[bro.aglt2.org-p1p2]
type=worker
host=bro.aglt2.org
interface=p1p2
lb_method=pf_ring
lb_procs=8
pin_cpus=16,18,20,22,24,26,28,30
...
```

We should note that all the configuration details above are the result of our initial testing and will likely evolve as we gain experience with the system.

7. MISP and Bro integration

A script provided by CERN to generate Bro import data from MISP IOCs has been used. All workshop participants have tested exporting IoC data from MISP to Bro's intelligence framework format using the MISP API.

8. Next steps

8.1 Steps following the December 2017 workshop

A few next steps have been identified following the December 2017 workshop. These include the continued deployment of Bro at different sites, and also the tuning of these deployments. We start by

identifying a small set of traffic to monitor to begin with in order to tune the configuration, before adding additional traffic as the deployment is validated. This was established as a recommended practice during the workshop.

Secondly, an important step is to import IoCs from MISP into Bro, thus completing the integration of these components.

8.2 Next workshop

The second workshop took place during 27 - 29th of June 2018 at CERN and contained an expanded programme, including a number of topics beyond an introductory session. These included Elasticsearch and associated technologies, data access and visualisation. Advanced aggregation techniques, correlation and enrichment of generated alerts were also discussed, with presentations of different approaches in use at CERN. Four different WLCG sites presented four different network topologies, with an emphasis on different possible network tap points and strategies for traffic mirroring. Lastly, various Bro optimisations were investigated. Further discussion of this workshop will form a future publication.

9. Conclusion

The workshop marked a good level of progress for the working group, expanding the participants and leading to a useful improvement in the documentation available. Notable results included an increase in the installed instances of MISP. This clearly demonstrates that the basic deployment of MISP is straightforward, which leaves the way open for further developments in enabling sites to sync with the WLCG instance, as well as enabling work towards forming suitable trust frameworks for sharing threat intelligence both within the WLCG community and beyond. The deployment of Bro was also demonstrated, although this is more complex than MISP.

Further work remains, including a more complete set of guidelines on the type and rates of network traffic required. The addition of further SOC components, including visualisation and data enrichment, is another important area, which will be a subject of a future workshop, in addition to a deeper consideration of different network topologies, possible network tap points and strategies for traffic mirroring.

References

[1] D. Crooks and L. Vâlsan (2017) WLCG Security Operations Centres Working Group. International Symposium on Grids & Clouds 2017 (ISGC 2017), BHSS, Academia Sinica, Taipei, 5-10 March 2017

[2] Malware Information Sharing Platform, www.misp-project.org

[3] https://home.cern

[4] L. Vâlsan, et al. (2018) Building a large scale Intrusion Detection System using Big Data technologies. International Symposium on Grids & Clouds 2018 (ISGC 2018), BHSS, Academia Sinica, Taipei, 16-23 March 2018. *To Be Confirmed*

[5] http://bro.org

[6] https://www.geant.org/Services/Trust_identity_and_security/eduGAIN

[7] H. Short and R. Wartel (2016) International Symposium on Grids and Clouds (ISGC) 2016 (ISGC 2016), Academia Sinica, Taipei, 13-18 March 2016

[8] http://www.openvas.org

[9] Wenger, E. "Communities of practice: Learning, meaning, and identity", 1998, Cambridge Univ Press, Cambridge

[10] www.gridpp.ac.uk

[11] https://ovirt.org/

[12] https://www.elastic.co

[13] https://www.elastic.co/products/beats/filebeat

[14] https://charbelnemnom.com/2015/05/how-to-enable-configure-vmqdvmq-on-windows-server-2012-r2-with-below-ten-gig-network-adapters-hyperv-vmq-vrss/

[15] https://blogs.technet.microsoft.com/networking/2015/10/16/setting-up-port-mirroring-to-capture-mirrored-traffic-on-a-hyper-v-virtual-machine/

[16] https://fs.com