Network Threat Detection at the University of Chicago

David Jordan
on behalf of the UChicago MWT2 Team

MANIAC LAB

10/26/2023

Who we are

- Midwest Tier 2: A tier 2 research computing site for the US ATLAS grid
  - Provide compute and storage resources for the CERN ATLAS project
  - We also host an analysis facility for interactive user resources
  - Offices in MCP and data center for MWT2 at 6045 Kenwood Pod-C
- Under the Worldwide LHC Computing Grid (WLCG) umbrella
  - All ATLAS computing grids and CERN projects (ATLAS, CMS, etc.)
- WLCG has their own Security Operations Center (SOC)
  - Focus on security at CERN, but also plans and models for WLCG at large
  - UChicago MWT2 team has primarily worked with

Reasoning

- Security becoming more important
  - University of Chicago has been clamping down on IT security and we can collaborate with them
- WLCG Security Operations Center at CERN wants to build a collaborative security group within the WLCG & research communities
- Work with the SOC on building a standardized model for new sites to join with as little effort as possible
- Share our build so others can use it as a reference

Project Goals

- Analyze **all** network traffic in & out of our ATLAS computing facility at Uchicago
- **Develop** specific, relevant **threat intelligence**
- Cooperate and **communicate with ITS Security** more easily about network attacks against our infrastructure
- **Share** any threat intelligence with WLCG SOC
- Use intelligence **other sites have gathered** to better protect ourselves
- **Coordinate on global attacks** within the research & education community
- Provide input for **traceability and response policies** within the research community

Equipment Setup at UChicago

- 2x bare metal machines for running Zeek w/ 2x 100Gb BlueField-2 NICs, 2x AMD EPYC 74F3 24-Core Processor, and 132GB RAM
- 1x 60/40 LC duplex fiber taps
- 2x Juniper QFX10002 in a virtual router each with 2x100Gb links to the WAN
  - In an active/backup configuration
- One machine per Juniper QFX10k switch in our racks in 6045 S Kenwood Pod-C. Each using half the 60/40 tap
- We purchased the nodes with enough CPU and memory to have multiple processes per interface capturing network traffic and writing output to logs
- Virtual machine running MISP (Malware Information Sharing Platform) on our VM stack

Setup Diagram

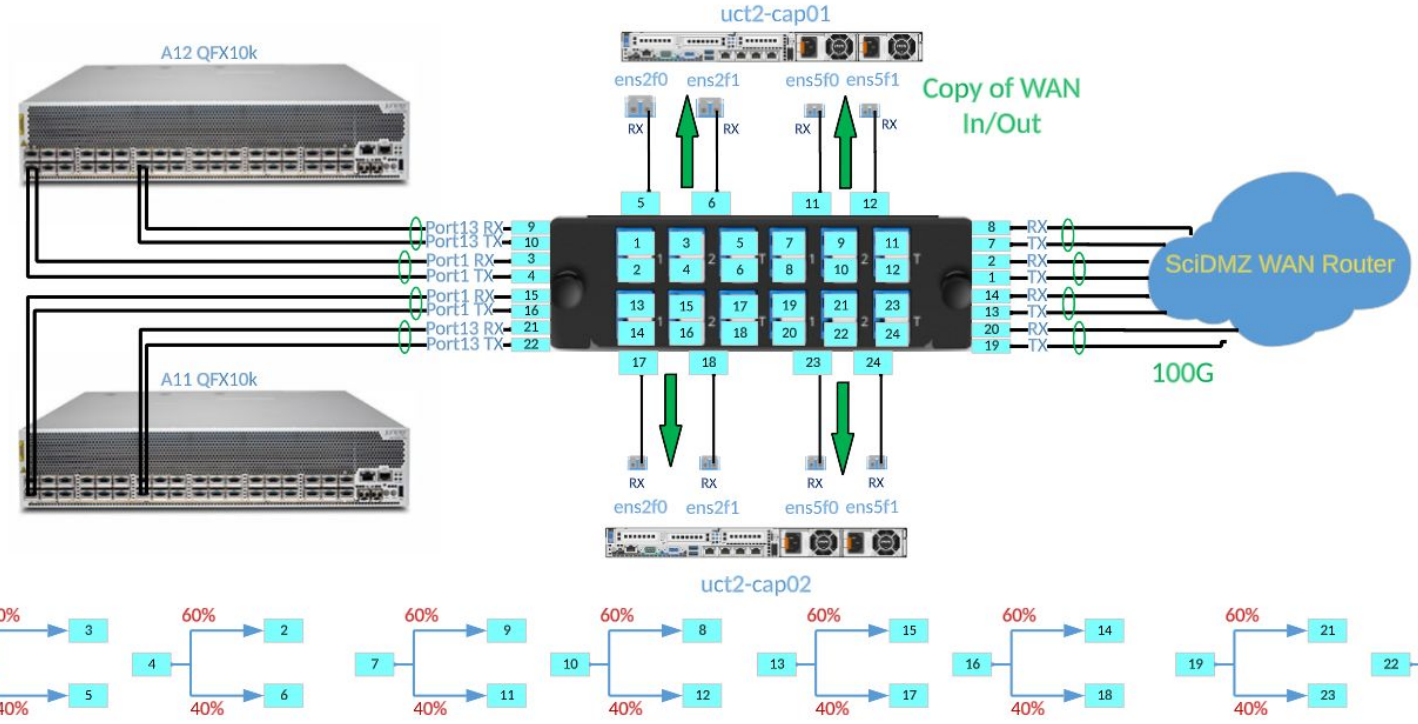

Setup Pictures

Original 70/30 splitters

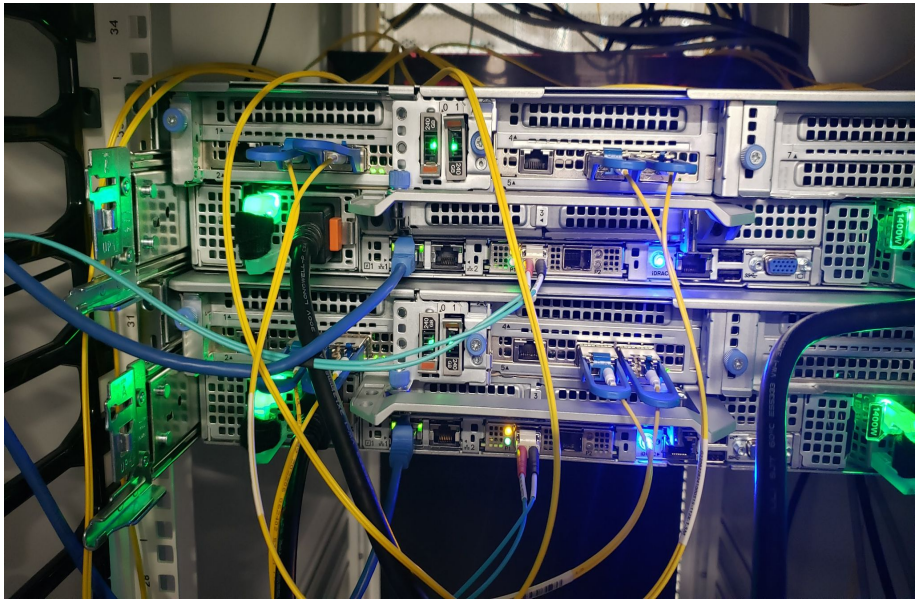

Network Traffic Analysis w/ Zeek

- Open Source monitoring tool used to capture and parse network activity.
- **Passive capture** and will only read packets routed to it
- Performs monitoring and logging at the network and application layers
  - Monitors network traffic and can display performance data
- Customizable and scalable, so adjustments can be made to fit different needs
- Interprets network traffic and generates log files based on that traffic
- Can be configured to alert and notify when it spots potential compromise
- Chosen as the standard for the WLCG SOC network security model

Zeek Cont.

- Zeek cannot read the payload of the packet, only the headers (ips, port, etc.)
- Zeek saves packet information in various logs, dividing up with a best guess as to what type of traffic the packet was (ssh, dns, etc.)
- Zeek can create a cluster of multiple nodes
  - The manager process runs on the head node where zeek is configured
  - Manager node must have root access to other machines in the cluster
- Zeek has built-in scripts
  - An intelligence framework that can be fed information to generate hits
  - Email notifications and alerts based on intel hits
  - Loaded on startup from local.zeek

Configuring Zeek

- Three main config files
  - node.cfg
    - Describes processes and the nodes they run on
  - networks.cfg
    - List of networks considered yours
  - zeekctl.cfg
    - Zeek directory locations, logs settings, etc.
- local.zeek
  - Where scripts are loaded

```
[uct2-cap01-logger]
type=logger
host=uct2-cap01.mwt2.org

[uct2-cap02-logger]
type=logger
host=uct2-cap02.mwt2.org

[uct2-cap01-manager]
type=manager
host=uct2-cap01.mwt2.org

[uct2-cap01-proxy]
type=proxy
host=uct2-cap01.mwt2.org

[uct2-cap01-ens2f0]
type=worker
host=uct2-cap01.mwt2.org
interface=ens2f0
lb_method=af_packet
af_packet_fanout_id=1
lb_procs=12
pin_cpus=0,1,2,3,4,5,6,7,8,9,10,11

[uct2-cap01-ens2f1]
type=worker
host=uct2-cap01.mwt2.org
interface=ens2f1
lb_method=af_packet
af_packet_fanout_id=2
lb_procs=12
pin_cpus=12,13,14,15,16,17,18,19,20,21,22,23
```

```
[root@uct2-cap01 ~]# cat /opt/zeek/etc/networks.cfg
# List of local networks in CIDR notation, optionally followed by a descriptive
# tag. Private address space defined by Zeek's Site::private_address_space set
# (see scripts/base/utils/site.zeek) is automatically considered local. You can
# disable this auto-inclusion by setting zeekctl's PrivateAddressSpaceIsLocal
# option to 0.

# UC Subnets
192.168.0.0/16          Internal Network
192.170.240.0/23        Pod-C SciDMZ
192.170.209.0/25        Pod-C Enterprise
192.170.231.128/25      Hinds SciDMZ
128.135.20.128/25       Hinds Enterprise
192.170.236.0/24        River SciDMZ
192.170.231.0/26        Hinds SciDMZ

2605:9a00:10:200a::/64  Pod-C v6 SciDMZ
2605:9a00:10:200b::/64  Hinds v6 SciDMZ
2605:9a00:10:200c::/64  River v6 SciDMZ

# IU Subnets
149.165.224.0/23        IU MWT2 Network
2001:18e8:c02:5::/64    IU v6 MWT2 Network

# UIUC Subnets
72.36.96.0/24           UIUC Network
2620:0:e01:4800::/64    UIUC v6 Network
[root@uct2-cap01 ~]#
```

```
[root@uct2-cap01 ~]# grep = /opt/zeek/etc/zeekctl.cfg
MailTo = djordan66@uchicago.edu
MailConnectionSummary = 1
MinDiskSpace = 5
MailHostUpDown = 1
LogRotationInterval = 86400
LogExpireInterval = 30day
StatsLogEnable = 1
StatsLogExpireInterval = 0
StatusCmdShowAll = 0
CrashExpireInterval = 0
SitePolicyScripts = local.zeek
LogDir = /opt/zeek/logs
SpoolDir = /opt/zeek/spool
BrokerDBDir = /opt/zeek/spool/brokerstore
CfgDir = /opt/zeek/etc
[root@uct2-cap01 ~]#
```

Zeek Logs

```
[root@uct2-cap01 ~]# ll /opt/zeek/logs/current/
total 1811369
-rw-r--r-- 1 root zeek 123677682 Sep 25 13:57 analyzer.log
-rw-r--r-- 1 root zeek     46582 Sep 25 13:39 broker.log
-rw-r--r-- 1 root zeek     19917 Sep 25 13:57 capture_loss.log
-rw-r--r-- 1 root zeek     46065 Sep 25 12:27 cluster.log
-rw-r--r-- 1 root zeek 574739737 Sep 25 13:57 conn.log
-rw-r--r-- 1 root zeek     24621 Sep 25 13:57 dhcp.log
-rw-r--r-- 1 root zeek 237455773 Sep 25 13:57 dns.log
-rw-r--r-- 1 root zeek       721 Sep 25 13:57 dpd.log
-rw-r--r-- 1 root zeek   2794346 Sep 25 13:57 files.log
-rw-r--r-- 1 root zeek 191868310 Sep 25 13:57 http.log
-rw-r--r-- 1 root zeek       816 Sep 25 13:24 kerberos.log
-rw-r--r-- 1 root zeek       211 Sep 25 12:26 known_hosts.log
-rw-r--r-- 1 root zeek    106587 Sep 25 13:57 known_services.log
-rw-r--r-- 1 root zeek     33203 Sep 25 12:26 loaded_scripts.log
-rw-r--r-- 1 root zeek     86146 Sep 25 13:57 notice.log
-rw-r--r-- 1 root zeek   2157665 Sep 25 13:57 ntp.log
-rw-r--r-- 1 root zeek      6766 Sep 25 13:46 ocsp.log
-rw-r--r-- 1 root zeek       209 Sep 25 12:26 packet_filter.log
-rw-r--r-- 1 root zeek      1331 Sep 25 13:57 radius.log
-rw-r--r-- 1 root zeek    150123 Sep 25 13:55 sip.log
-rw-r--r-- 1 root zeek      2194 Sep 25 13:01 smtp.log
-rw-r--r-- 1 root zeek    711324 Sep 25 13:57 snmp.log
-rw-r--r-- 1 root zeek    163620 Sep 25 13:56 software.log
-rw-r--r-- 1 root zeek  21106288 Sep 25 13:57 ssh.log
-rw-r--r-- 1 root zeek  25781633 Sep 25 13:57 ssl.log
-rw-r--r-- 1 root zeek    139940 Sep 25 13:56 stats.log
-rw-r--r-- 1 root zeek         0 Sep 25 12:26 stderr.log
-rw-r--r-- 1 root zeek       204 Sep 25 12:26 stdout.log
-rw-r--r-- 1 root zeek 336462574 Sep 25 13:57 syslog.log
-rw-r--r-- 1 root zeek  70099795 Sep 25 13:57 telemetry.log
-rw-r--r-- 1 root zeek     31069 Sep 25 13:56 tunnel.log
-rw-r--r-- 1 root zeek 266134673 Sep 25 13:57 weird.log
-rw-r--r-- 1 root zeek     53330 Sep 25 13:47 x509.log
```

dns.log

```
#separator \x09
#set_separator  ,
#empty_field    (empty)
#unset_field    -
#path   dns
#open   2023-09-25-12-26-32
#fields ts      uid     id.orig_h       id.orig_p       id.resp_h       id.resp_p       proto   trans_id        rtt     query   qclass  qclass_name     qtyp
e       qtype_name      rcode   rcode_name      AA      TC      RD      RA      Z       answers TTLs    rejected
#types  time    string  addr    port    addr    port    enum    count   interval        string  count   string  count   string  bool    bool
        bool    bool    count   vector[string]  vector[interval]        bool    -       -       -       -       -       -       -       -
1695662792.928466       CwuxufaLB5i2Xhf2U       192.170.241.218 47554   128.135.247.50  53      udp     8404    0.000581        0.9.d.f.4.b.e.f.f.f.f.9.6.3.
2.a.a.0.0.2.0.1.0.0.0.0.a.9.5.0.6.2.ip6.arpa    1       C_INTERNET      12      PTR     0       NOERROR F       F       T       T       0       uct2-s28.mwt
2.org   216.000000      F
1695662792.691921       CPa4w4Ip7iLSwcC5b       192.170.241.148 43682   128.135.247.50  53      udp     24079   -       -       -       -       -       -
        3       NXDOMAIN        F       F       F       F       0       -       -       F
1695662792.727944       CKwF2w1uc0feCox0lb       192.170.241.102 45505   128.135.247.50  53      udp     24236   -       v4cvmfs.mwt2.org        -       -
        -       -       0       NOERROR F       F       F       T       0       192.170.241.192 103.000000      F
1695662792.776805       CdRjv326rDsXfTdy31      192.170.240.203 38792   128.135.247.50  53      udp     36542   -       v4cvmfs.mwt2.org        1       C_IN
TERNET  28      AAAA    -       -       F       F       T       F       0       -       -       F
1695662792.757535       C1GcML2IsnYY4fjy1i      192.170.241.170 32845   128.135.247.50  53      udp     1225    -       voatlasrucio-server-prod.cern.ch
        1       C_INTERNET      28      AAAA    -       -       F       F       T       F       0       -       -       F
1695662792.757807       C3xlhUR0xYHNNg7R2       192.170.241.170 40641   128.135.247.50  53      udp     12863   -       230.206.153.98.in-addr.arpa     1
        C_INTERNET      12      PTR     -       -       F       F       T       F       0       -       -       F
1695662792.757520       Cft6c23bDWLnHUwb53      192.170.241.170 32845   128.135.247.50  53      udp     1225    -       voatlasrucio-server-prod.cern.ch
        1       C_INTERNET      28      AAAA    -       -       F       F       T       F       0       -       -       F
1695662792.794395       CbCcgs2r5eyOir7BJ2      192.170.241.154 35640   128.135.247.50  53      udp     17635   -       v4a.mwt2.org    -       -       -
        -       0       NOERROR F       F       F       T       0       2605:9a00:10:200a:e63d:1aff:fed9:a850   117.000000      F
1695662792.818410       CzDjuM22M0zbgjcM6d      192.170.241.49  56513   128.135.247.50  53      udp     39938   -       pandaserver.cern.ch     1       C_IN
TERNET  1       A       -       -       F       F       T       F       0       -       -       F
1695662792.816937       Ch36w42Z5kkFnIb5db      192.170.241.184 56675   128.135.247.50  53      udp     57516   -       134.128.35.189.in-addr.arpa     1
        C_INTERNET      12      PTR     -       -       F       F       T       F       0       -       -       F
1695662792.846178       C3pJzR1djv6WiXsFva      192.170.241.154 38013   128.135.247.50  53      udp     53443   -       v4a.mwt2.org    1       C_INTERNET
        1       A       -       -       F       F       T       F       0       -       -       F
1695662792.846213       C8G9wC389pJskU43Wh      192.170.241.154 38013   128.135.247.50  53      udp     11473   -       v4a.mwt2.org    1       C_INTERNET
        28      AAAA    -       -       F       F       T       F       0       -       -       F
1695662792.854651       CoAHjJ1woRaGestti5      192.170.240.75  38342   128.135.247.50  53      udp     14349   -       -       -       -       -       -
        3       NXDOMAIN        F       F       F       F       0       -       -       F
1695662792.855203       CBqy7L1l2Oj6IBob8h      192.170.240.75  34262   128.135.247.50  53      udp     25209   -       21.99.253.201.telecom.com.ar.iu.edu
        1       C_INTERNET      1       A       -       -       F       F       T       F       0       -       -       F
/opt/zeek/logs/current/dns.log
```

Zeek actively adds to the logs as capture data comes in

Using Zeek Intel

```
[root@uct2-cap01 ~]# ll /opt/zeek/intel/
total 11765
-rw-r--r-- 1 root root      90 Sep 25 14:00 certhash.txt
-rw-r--r-- 1 root root    7251 Sep 25 14:00 domain.txt
-rw-r--r-- 1 root root      90 Sep 25 14:00 email.txt
-rw-r--r-- 1 root root 7512330 Sep 25 14:00 filehash.txt
-rw-r--r-- 1 root root    3589 Sep 25 14:00 filename.txt
-rw-r--r-- 1 root root   17948 Sep 25 14:00 ip.txt
-rw-r--r-- 1 root root      90 Sep 25 14:00 software.txt
-rw-r--r-- 1 root zeek     207 Sep 25 08:56 test2.txt
-rw-r--r-- 1 root root     215 Sep 25 11:15 test.txt
-rw-r--r-- 1 root root 4302417 Sep 25 14:00 url.txt
[root@uct2-cap01 ~]#
```

local.zeek

```
redef Intel::read_files += {
# MISP feeds
        "/opt/zeek/intel/domain.txt",
        "/opt/zeek/intel/email.txt",
        "/opt/zeek/intel/filehash.txt",
        "/opt/zeek/intel/filename.txt",
        "/opt/zeek/intel/ip.txt",
        "/opt/zeek/intel/url.txt",
        "/opt/zeek/intel/test.txt",
};
```

```
[root@uct2-cap01 ~]# cat /opt/zeek/intel/test.txt
#fields indicator       indicator_type  meta.source     meta.desc       meta.url        meta.do_notice  meta.if_in
67.176.240.185  Intel::ADDR     Manual  Test intel      test-url.mwt2.org       T       -
192.170.231.213 Intel::ADDR     Manual  Test intel      test-url.mwt2.org       T       -
```

Using Zeek Intel Cont.

```
[root@uct2-cap01 ~]# rm -f /root/test-zeek-logs/*
[root@uct2-cap01 ~]# ll test-zeek-logs/
total 0
[root@uct2-cap01 ~]# zeek -r /root/dump.trace Log::default_logdir=/root/test-zeek-logs local
[root@uct2-cap01 ~]# ll test-zeek-logs/
total 236
-rw-r--r-- 1 root root   1061 Sep 25 15:32 analyzer.log
-rw-r--r-- 1 root root    321 Sep 25 15:32 capture_loss.log
-rw-r--r-- 1 root root 130765 Sep 25 15:32 conn.log
-rw-r--r-- 1 root root    617 Sep 25 15:32 dhcp.log
-rw-r--r-- 1 root root    429 Sep 25 15:32 dpd.log
-rw-r--r-- 1 root root    839 Sep 25 15:32 intel.log
-rw-r--r-- 1 root root  32371 Sep 25 15:32 loaded_scripts.log
-rw-r--r-- 1 root root    899 Sep 25 15:32 notice.log
-rw-r--r-- 1 root root    278 Sep 25 15:32 packet_filter.log
-rw-r--r-- 1 root root   5080 Sep 25 15:32 ssh.log
-rw-r--r-- 1 root root    840 Sep 25 15:32 stats.log
-rw-r--r-- 1 root root  12273 Sep 25 15:32 syslog.log
-rw-r--r-- 1 root root  22908 Sep 25 15:32 telemetry.log
```

Zeek Intel Hits

```
[root@uct2-cap01 ~]# cat /root/test-zeek-logs/intel.log
#separator \x09
#set_separator  ,
#empty_field    (empty)
#unset_field    -
#path   intel
#open   2023-09-25-15-32-32
#fields ts      uid     id.orig_h       id.orig_p       id.resp_h       id.resp_p       seen.indicator  seen.indicator_type     seen.w
here      seen.node  matched  sources fuid    file_mime_type  file_desc
#types  time    string  addr    port    addr    port    string  enum    enum    string  set[enum]       set[string]     string  string
  string
1695652764.576800       Cr8m04kmfcPNEfuZ        192.170.231.213 60224   192.170.240.8   22      192.170.231.213 Intel::ADDR     Conn::
IN_ORIG   zeek    Intel::ADDR Manual  -       -       -
1695652782.301889       CGFleT1wWN1u6fscF9      192.170.231.213 60226   192.170.240.8   22      192.170.231.213 Intel::ADDR     Conn::
IN_ORIG   zeek    Intel::ADDR Manual  -       -       -
1695652783.772364       CAESIn1m9ciZtKKbpd      192.170.231.213 60228   192.170.240.8   22      192.170.231.213 Intel::ADDR     Conn::
IN_ORIG   zeek    Intel::ADDR Manual  -       -       -
#close  2023-09-25-15-32-33
```

Hits on 192.170.231.213 based on the test.txt file fed into Zeek with Intel::read_files

Zeek Notify

```
[root@uct2-cap01 ~]# cat /root/test-zeek-logs/notice.log
#separator \x09
#set_separator  ,
#empty_field    (empty)
#unset_field    -
#path   notice
#open   2023-09-25-15-32-32
#fields ts      uid     id.orig_h       id.orig_p       id.resp_h       id.resp_p       fuid    file_mime_type  file_desc       proton
ote     msg     sub     src     dst     p       n       peer_descr      actions email_dest      suppress_for    remote_location.countr
y_code  remote_location.region  remote_location.city    remote_location.latitude        remote_location.longitude
#types  time    string  addr    port    addr    port    string  string  string  enum    enum    string  string  addr    addr    port c
ount    string  set[enum]       set[string]     interval        string  string  string  double  double
1695652764.576800       Cr8m04kmfcPNEfuZ        192.170.231.213 60224   192.170.240.8   22      -       -       -       tcp     Intel:
:Notice Intel hit on 192.170.231.213 at Conn::IN_ORIG   Indicator = 192.170.231.213     192.170.231.213 192.170.240.8   22      -    N
otice::ACTION_LOG,Notice::ACTION_EMAIL  djordan66@uchicago.edu,(empty)  43200.000000    -       -       -       -       -
#close  2023-09-25-15-32-33
```

Only one notification per indicator hit per 12hrs to limit spam (configurable)

Our Zeek

- Two bare metal nodes with AlmaLinux 9
  - uct2-cap01.mwt2.org & uct2-cap02.mwt2.org
- Configured as a Zeek cluster with cap01 as the manager
- We are monitoring only our own Ingress and Egress traffic
- BlueField-2 NICs only used for capturing data
  - A separate 10Gb link is used for management
  - Configuration management via Puppet
- Zeek Release 6.0.0
- 12 processes per interface capturing data (48 per node)
  - Each capture process has own dedicated core

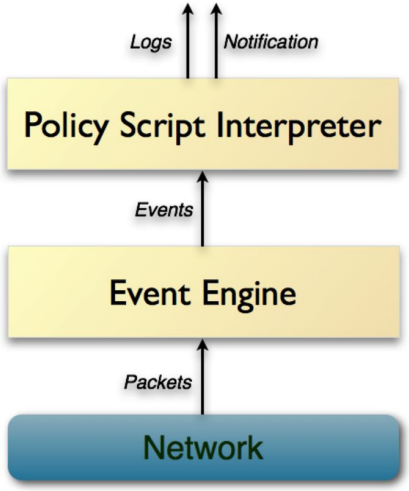

MISP (Malware Information Sharing Platform)

- An open source threat sharing software with contributors including the Belgian Ministry of Defense, NATO Cyber Security Centre, and Computer Incident Response Center Luxembourg
- There is a web portal portion to the MISP software that is the main way to interact with the software
- MISP does not interact with the network in any way. It is a list of threats and bad actors. Zeek pulls this intel
- We can define the scope of information shared to send more or less information about individual threats, compromises, etc. using TLP (Traffic Light Protocol) definitions

Our MISP Instance

- Runs on a VM on our OpenStack cluster
  - uct2-misp.mwt2.org
- The MWT2 instance connects with CERN's to pull threat intelligence they have from various organizations
  - Must have access to the other instance and use an API key
- We have accounts for only our team on our instance
  - MFA required
- Currently do not publicly publish any events or threats from the MWT2 instance
- Installed and running via docker containers
  - https://github.com/JiscCTI/misp-docker

MISP Web Portal

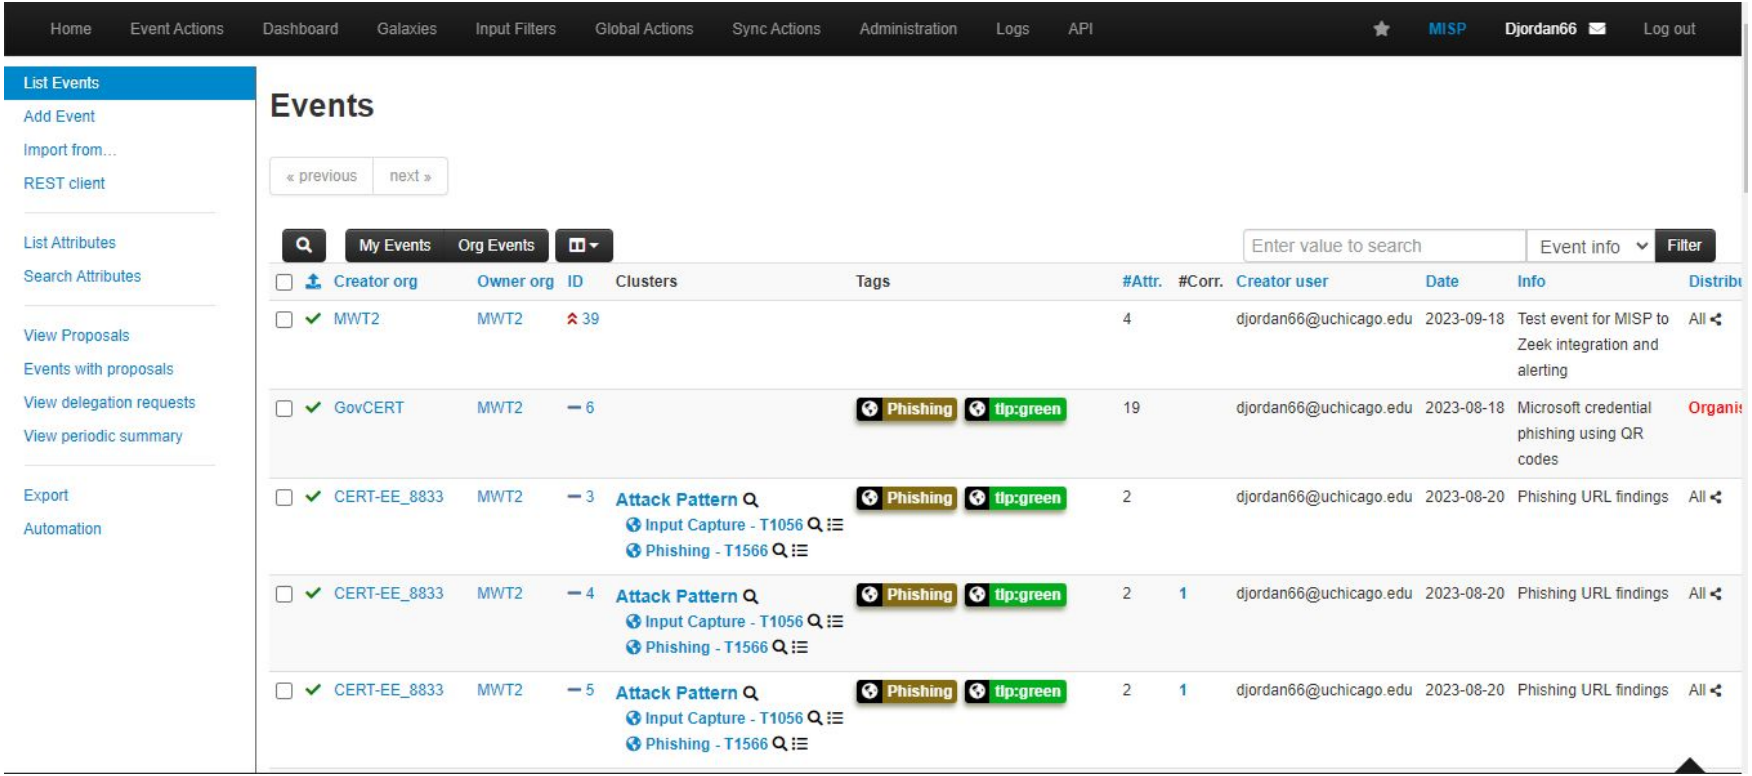

Elasticsearch (for Zeek)

- Distributed RESTful search and analytics engine
- Robust clusterization and easily scalable
  - Supported to run on the cloud (elastic cloud) or dedicated hardware
- Can enrich the data with processing and analytic tools (e.g. GeoIP)
- Has built in Zeek integration
  - Need to install an “elastic-agent” service on the host with the logs
- Extensive documentation

Our Elasticsearch (regarding Zeek)

- 29x nodes in the cluster
  - 22 data nodes for storage
  - 6 head nodes
  - 1 ingress and kibana host
- Elastic-agent installed on uct2-misp and uct2-cap01 to send MISP and Zeek data respectively to elasticsearch through the ingress host
- Processes the Zeek connection logs to break down % of connection types, most popular hosts, etc.
- Uses GeoIP to map the connections that come through our network on the dashboard

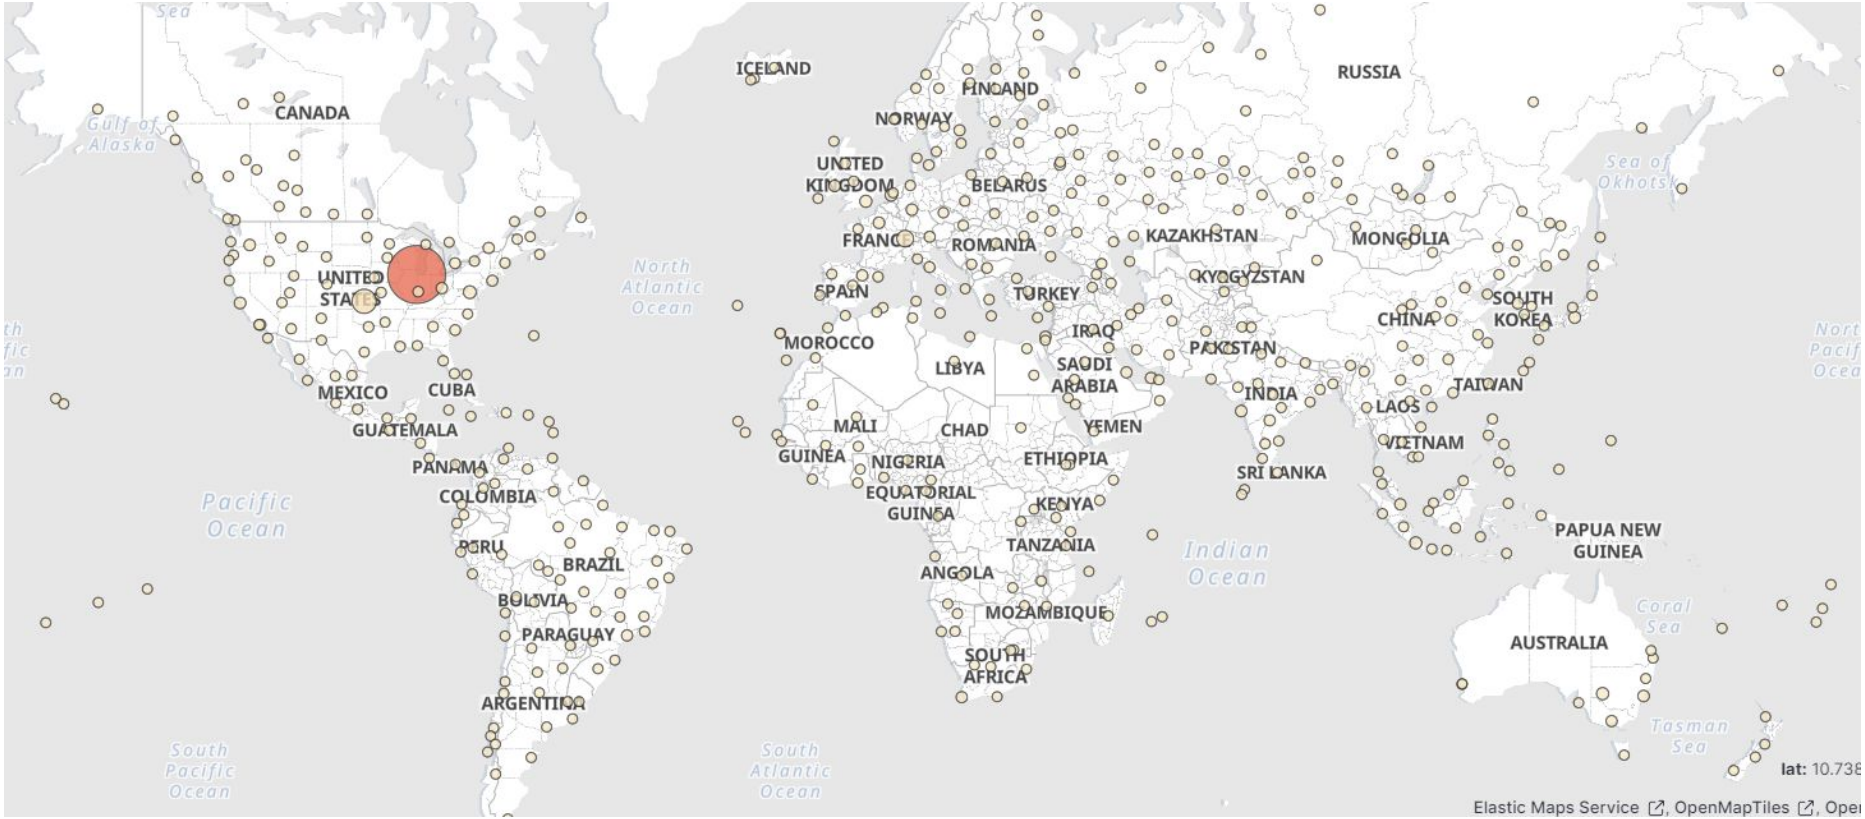

All Connection Destinations outside MWT2 networks 10/4/23–10/11/23

XrootD (port 1094) and WebDav Traffic (port 8443) 10/4/23–10/11/23

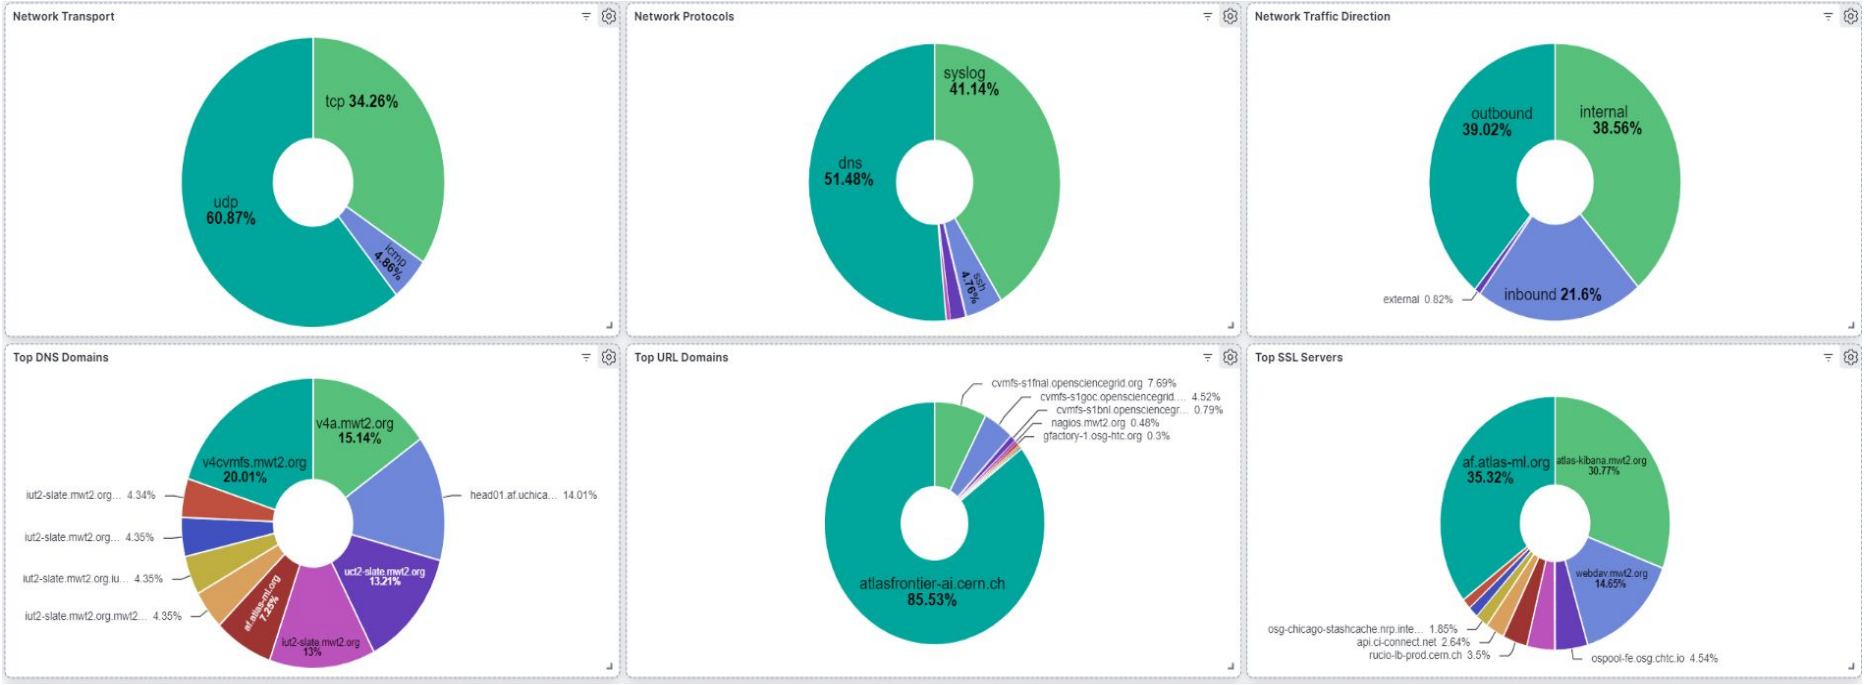

Graphs below the maps (all traffic)

Other Integrations

- Kafka
  - Data processing of Zeek logs and intel
- pDNSSOC
  - Collect DNS logs to send alerts
- We are not currently investigating either of these

Summary and Next Steps

- WLCG SOC wants to create a standardized model for a federated security framework within the WLCG community
- Zeek and MISP (Malware Information Sharing Platform) are installed and configured at UChicago
- Zeek can alert based on intelligence hits
- MISP pulls data from CERN's central instance
  - Zeek then grabs this data from the MWT2 MISP instance
- Elasticsearch copies Zeek logs and processes data
  - Continue work on elasticsearch dashboard and data enrichment

References

WLCG-SEC-OPS-COMBINED-GDB-DEC-2021 – WLCG presentation from 2021

Zeek Github

Zeek site

- What is Zeek? Docs

MISP

pDNSSOC

Questions?

Email: djordan66@uchicago.edu