

LEVERAGING ZEEK LOGS IN MOLOCH WITH FILEBEAT AND LOGSTASH



CISA
CYBER+INFRASTRUCTURE

What is Zeek?

- Extensible, open-source passive network analysis framework
- More than just an Intrusion Detection System
 - Live and post-capture traffic capture and inspection
 - Intrusion and anomaly detection
 - Flow tracking
 - Scripting framework
- See zeek.org



What is Zeek?

- Q: Doesn't Moloch already do a bunch of that stuff?
- A: Moloch is a **great** Full Packet Capture (FPC) system, but there are instances where FPC may not be possible or advisable
 - Storage or bandwidth limitations
 - Sensitive data (security/compliance limitations)

What is Zeek?

- Zeek has analyzers for some protocols Moloch doesn't (yet)
- Community for additional protocol parsers and plugins
- Write your own custom protocol parsers with BinPAC
 - High-level language for describing protocol parsers
 - Generates C++ code for Zeek





+



=

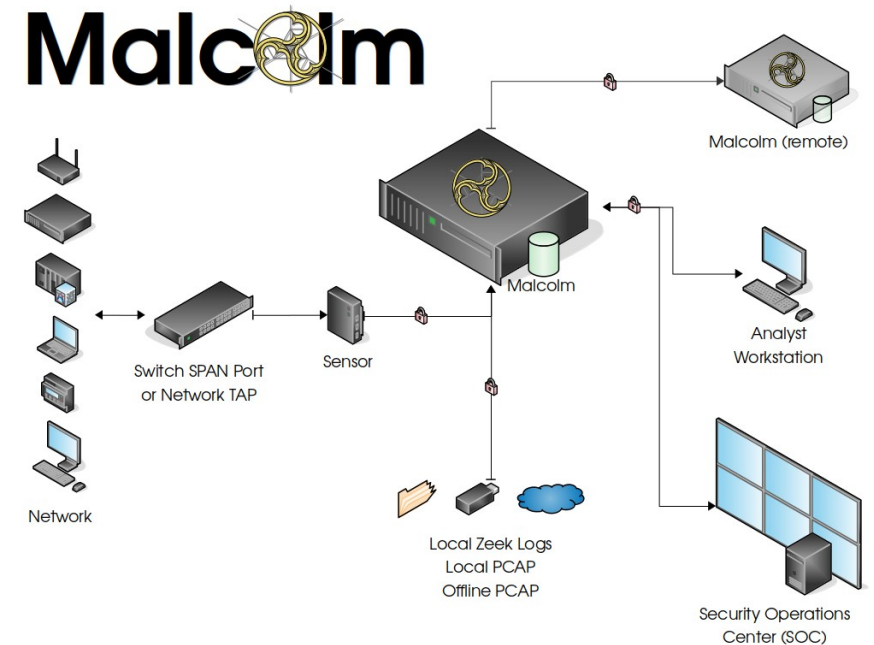


- Use `moloch-capture` for FPC and to provide integration with Moloch viewer's excellent payload analysis tools
- Use Zeek in instances where FPC is not an option, or to leverage parsers or plugins unavailable in Moloch
- Will it blend?
 - Ingest Zeek logs into Elasticsearch with Filebeat
 - Enrich the logs and map them to Moloch's schema with Logstash
 - Use WISE to define a data source to make browsing Zeek data in Moloch seamless
 - Use common fields to correlate Moloch sessions and Zeek logs during analysis



CISA
CYBER+INFRASTRUCTURE

- Docker-based framework of open-source network traffic analysis tools
 - Elastic (Filebeat, Logstash, Elasticsearch, Kibana, Curator), Moloch, Zeek, nginx, ElastAlert, ClamAV, ...
- Easy to deploy (`docker-compose`)
- Suitable for IT and OT networks
 - Upcoming development will emphasize creating new Zeek analyzers for industrial control systems (ICS) protocols
- See github.com/idaholab/Malcolm

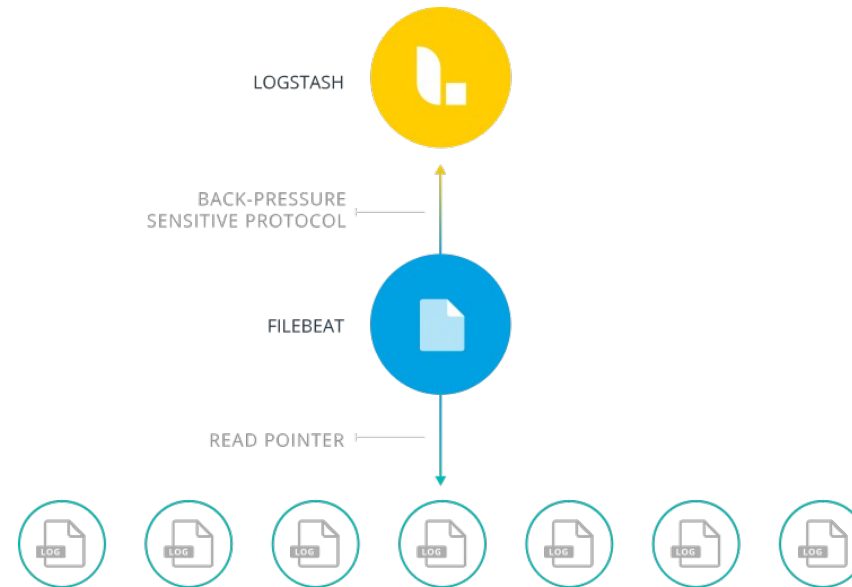


Zeek configuration

- Tweak Zeek's local site policy `local.bro`
- Enable fields not captured by default (VLANs, MACs)
- Make use of community plugins (CommunityID, JA3, HASSH)
- See [Malcolm/moloch/zeek/local.bro](https://github.com/Malcolm/moloch/zeek/local.bro)

Zeek log forwarding

- Filebeat is a lightweight log shipper
- Monitor Zeek log directories and forward to Logstash over SSL
- See [Malcolm/filebeat/filebeat.yml](#)



Logstash: parsing

- Default (tab-delimited) Zeek logs
 - **dissect** Logstash filter outperforms **grok** and **csv** filters
 - Watch out for log variations due to config or Zeek updates
 - See Malcolm/logstash/pipeline-main/11_zeek_logs.conf
- JSON Zeek logs
 - @load policy/tuning/json-logs.bro in **local.bro**
 - **json** Logstash filter
 - See Security Onion's approach
securityonion-elastic/configfiles/1100_preprocess_bro_conn.conf

```
conn.log
1 #separator \x09
2 #set_separator_
3 #empty_field_(empty)
4 #unset_field_
5 #path_conn
6 #open_2019-09-23-17-32-02
7 #fields ts uid id.orig_h id.orig_p id.resp_h id.resp_p proto service duration
8 #types time string addr port addr port enum string interval count count
9 1569259210.576120 Cy7x6f2BFqPtD7DAU 172.16.10.189 49174 52.37.243.173 443 tcp
10 1569259213.265169 CLui7t4SIRSxWg9cG4 172.16.10.189 59290 216.58.217.35 443 tcp
11 1569259224.265777 CF55Ht193an6e4a5Yl 172.16.10.189 54608 172.217.14.206 443 tc
12 1569259224.265846 CWE7ajMGftFPMqK9d 172.16.10.189 33240 172.217.3.206 443 tcp
13 1569259227.266467 C3XylD3Bvt2KzxHyJc 172.16.10.189 48014 172.217.14.234 443 tc
14 1569259227.266511 CAImKR2BBDcun6yiUb 172.16.10.189 48894 216.58.193.74 443 tcp
15 1569259741.335062 Cd0u38C07PxYj79Zj fe80::15:7bb1:c2aa:90d1 5353 ff02::fb 5353
16 1569259741.335062 CaixCa5NFrFc4jSve 172.16.10.136 5353 224.0.0.251 5353 udp dn
17 1569259741.354061 Cd6AvC3DS6dR37Y0Mj fe80::89c:48e7:846f:75a1 5353 ff02::fb
18 1569259741.354061 C81SFm4Di05VNYdVle 172.16.10.171 5353 224.0.0.251 5353 udp
19 1569259741.401418 CzBzG02AwJ67GZz2je 172.16.10.163 5353 224.0.0.251 5353 udp
20 1569259741.401418 ClEtU511VsamqmIJLg fe80::1c8b:7ac0:c497:8566 5353 ff02::fb
21 1569259238.400723 CW03441RmVRedZl0 172.16.10.189 5353 224.0.0.251 5353 udp dn
```

Logstash: field mapping

- Get Moloch field definitions

```
$ curl -XGET "http://localhost:9200/fields/_search?pretty&size=1000"
```

```
...
```

```
{  
  "_index" : "fields_v3",  
  "_type" : "field",  
  "_id" : "mac.src",  
  "_score" : 1.0,  
  "_source" : {  
    "friendlyName" : "Src MAC",  
    "group" : "general",  
    "help" : "Source ethernet mac addresses set for session",  
    "dbField2" : "srcMac",  
    "type" : "lotermfield",  
    "transform" : "dash2Colon"  
  }  
},
```

```
...
```



Logstash: field mapping

- Where possible, map Zeek fields to Moloch fields

```
...
mutate { add_field => { "[srcIp]"           => "%{[zeek][orig_h]}" } }
mutate { add_field => { "[srcPort]"        => "%{[zeek][orig_p]}" } }
mutate { add_field => { "[dstIp]"          => "%{[zeek][resp_h]}" } }
mutate { add_field => { "[dstPort]"        => "%{[zeek][resp_p]}" } }
mutate { add_field => { "[communityId]"    => "%{[zeek][community_id]}" } }
mutate { merge =>      { "[srcMac]"         => "[zeek][orig_l2_addr]" } }
...
```

- Map `zeek.uid` to Moloch's `rootId` to tie Zeek sessions together across logs
- See `Malcolm/logstash/pipeline-main/11_zeek_logs.conf`

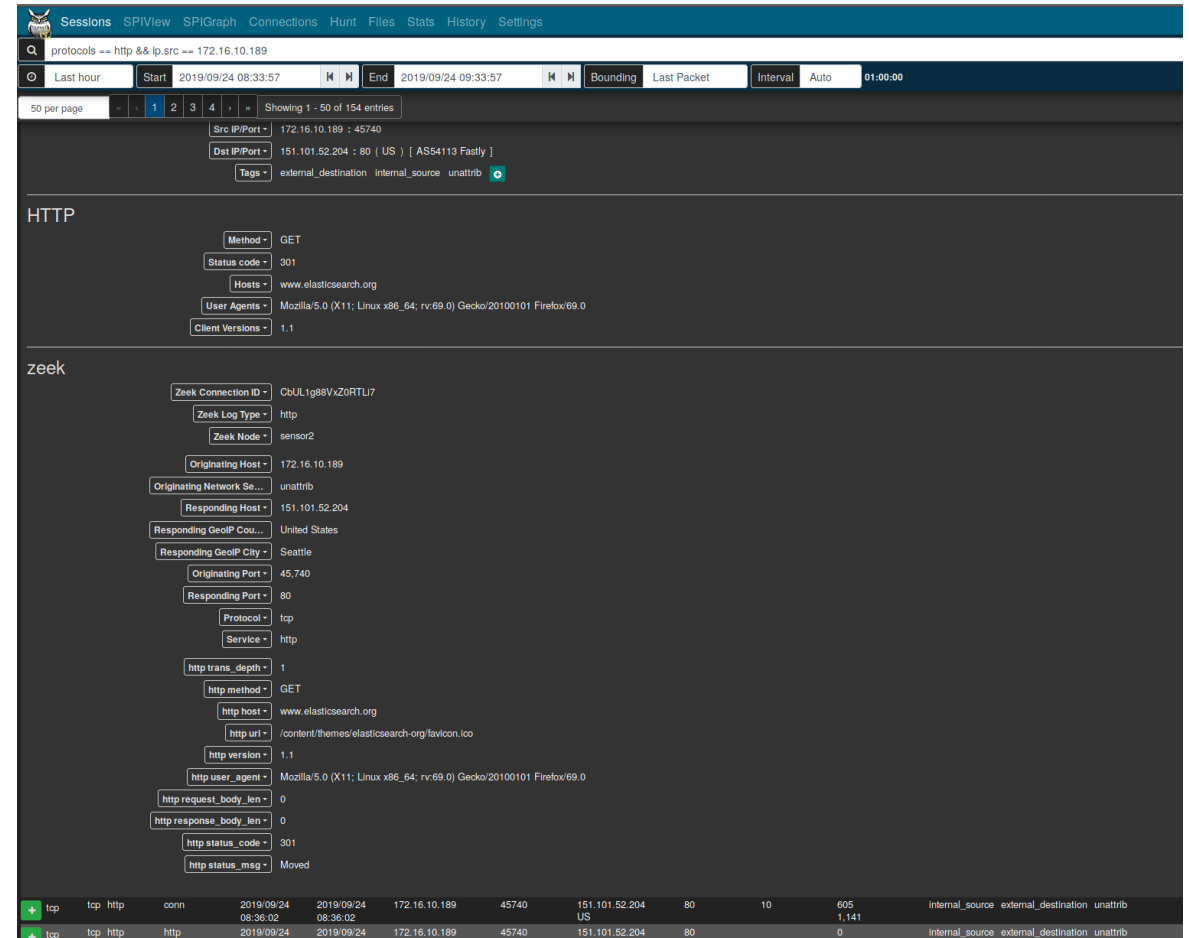
Logstash: enrichment

- Logstash filters to enrich your data (similar to WISE)
 - `geoip` for GeoIP and ASN lookups
 - `ieee_oui` for MAC address OUI lookups
 - `translate` for custom mappings
 - `add_tag` for custom tagging
 - `ruby` to roll your own



WISE: Make Zeek logs “native” in Moloch

- Create a Zeek WISE data source definition for Moloch
 - Defines templates for Zeek fields not mapped to Moloch
 - Defines fields' GUI elements for Moloch viewer
- See [Malcolm/moloch/wise/source.zeeeklogs.js](#) and [wise.ini](#)

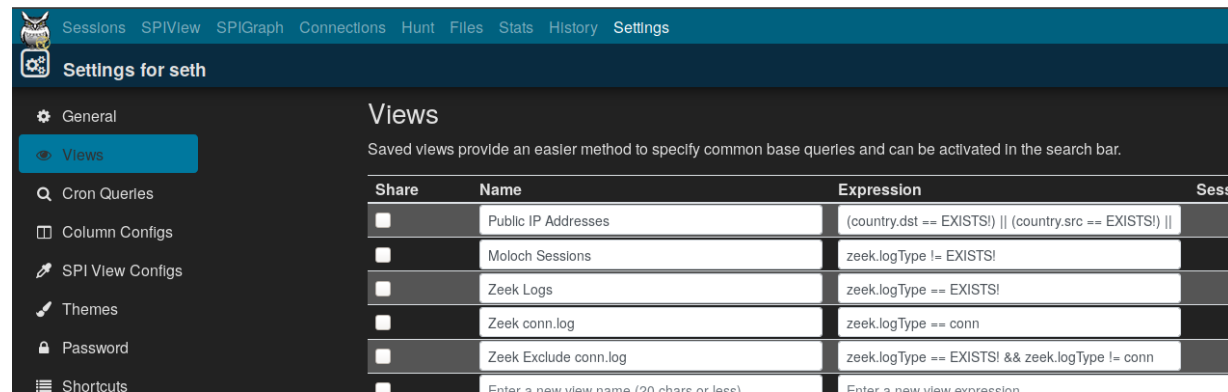


The screenshot displays the Moloch viewer interface. At the top, there are navigation tabs: Sessions, SPIView, SPIGraph, Connections, Hunt, Files, Stats, History, and Settings. A search bar contains the query `protocols == http && ip.src == 172.16.10.189`. Below the search bar, there are controls for time range (Start: 2019/09/24 08:33:57, End: 2019/09/24 09:33:57), bounding (Last Packet), interval (Auto), and a refresh button. The main content area shows a Zeek log entry for an HTTP GET request. The entry is expanded to show details for both the HTTP and Zeek protocols. The HTTP section shows Method: GET, Status code: 301, Hosts: www.elasticsearch.org, User Agents: Mozilla/5.0 (X11; Linux x86_64; rv:69.0) Gecko/20100101 Firefox/69.0, and Client Versions: 1.1. The Zeek section shows Zeek Connection ID: CbUL1g88VvZ0RTU7, Zeek Log Type: http, Zeek Node: sensor2, and various host and port information. At the bottom, a table lists the log entries with columns for protocol, type, connection, timestamps, IP addresses, ports, and other fields.

Protocol	Type	Conn	Start	End	Src IP	Src Port	Dst IP	Dst Port	Len	Offset	Internal Source	External Destination	Unattrib	
tcp	tcp	http	conn	2019/09/24 08:36:02	2019/09/24 08:36:02	172.16.10.189	45740	151.101.52.204	80	10	605	internal_source	external_destination	unattrib
tcp	tcp	http	http	2019/09/24 08:36:02	2019/09/24 08:36:02	172.16.10.189	45740	151.101.52.204	80	0	1,141	internal_source	external_destination	unattrib

Putting it all together

- Zeek logs and Moloch sessions coexist in Moloch viewer
 - (Though not exactly “apples to apples”)
- Views can help differentiate when needed

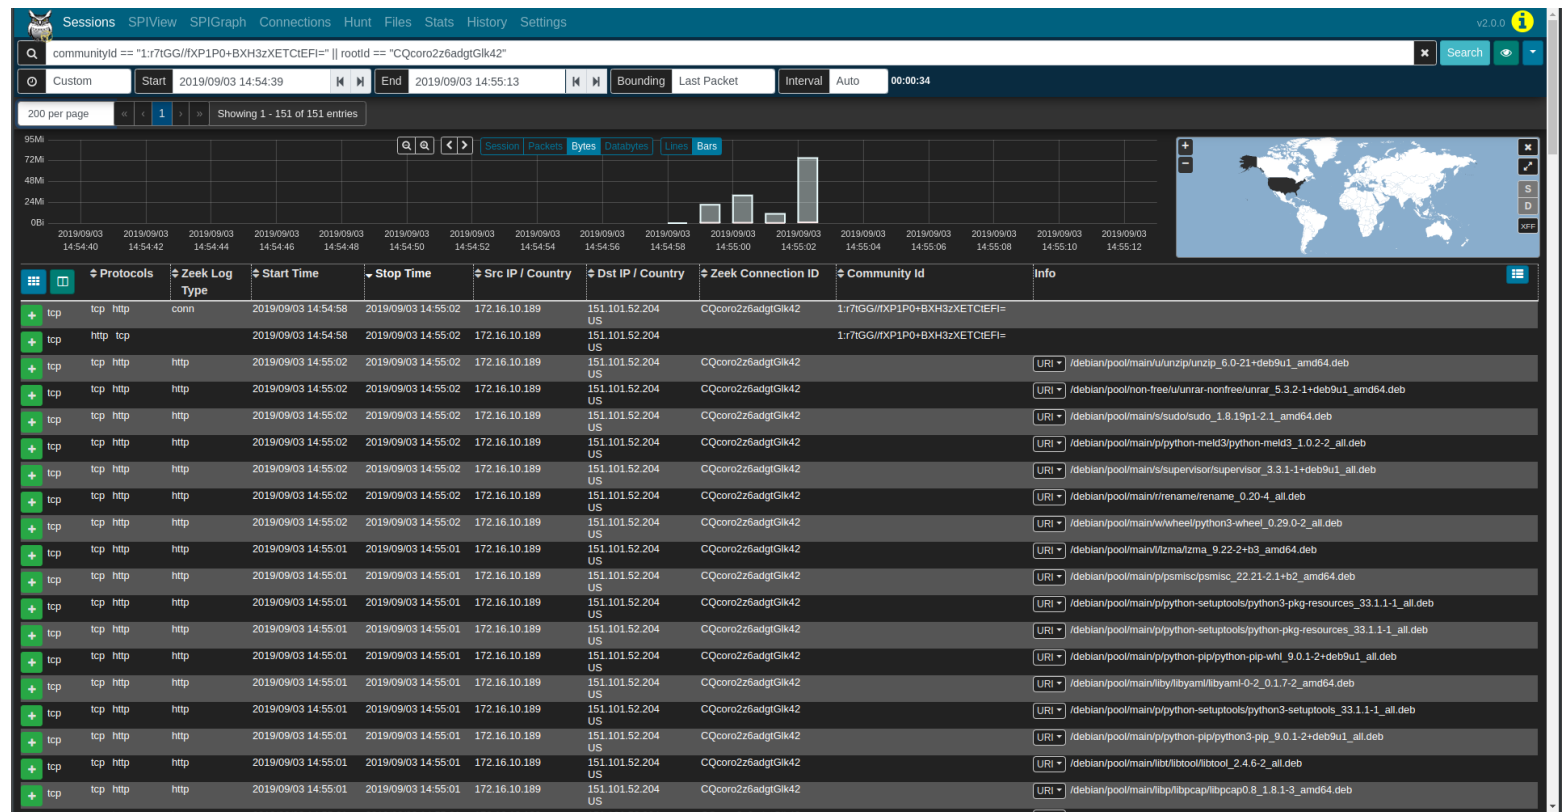


The screenshot shows the 'Settings for seth' page in the Moloch viewer, specifically the 'Views' section. The page has a dark theme and a sidebar on the left with navigation options: General, Views (selected), Cron Queries, Column Configs, SPI View Configs, Themes, Password, and Shortcuts. The main content area is titled 'Views' and includes a description: 'Saved views provide an easier method to specify common base queries and can be activated in the search bar.' Below this is a table with columns for 'Share', 'Name', 'Expression', and 'Sess'. The table lists several saved views with their corresponding expressions.

Share	Name	Expression	Sess
<input type="checkbox"/>	Public IP Addresses	(country.dst == EXISTS!) (country.src == EXISTS!)	
<input type="checkbox"/>	Moloch Sessions	zeek.logType != EXISTS!	
<input type="checkbox"/>	Zeek Logs	zeek.logType == EXISTS!	
<input type="checkbox"/>	Zeek conn.log	zeek.logType == conn	
<input type="checkbox"/>	Zeek Exclude conn.log	zeek.logType == EXISTS! && zeek.logType != conn	
<input type="checkbox"/>	Enter a new view name (20 chars or less)	Enter a new view expression	

Putting it all together

- Use common fields to correlate Moloch sessions and Zeek logs
 - `zeek.uid` (if mapped to Moloch's `rootId`) ties Zeek sessions together across logs
 - Community ID flow hashing (supported by both Moloch and Zeek) links Moloch sessions and the corresponding Zeek logs
 - Use both to see the Moloch sessions and Zeek logs for a particular network connection





CISA
CYBER+INFRASTRUCTURE

For more information:
github.com/idaholab/Malcolm

Questions?
Email: seth.grover@inl.gov



CISA
CYBER+INFRASTRUCTURE