Elasticsearch, Logstash, and Kibana (ELK)

Dwight Beaver dsbeaver@cert.org
Sean Hutchison shutchison@cert.org

January 2015

This material is based upon work funded and supported by Department of Homeland Security under Contract No. FA8721-05-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center sponsored by the United States Department of Defense.

References herein to any specific commercial product, process, or service by trade name, trade mark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by Carnegie Mellon University or its Software Engineering Institute.

NO WARRANTY. THIS CARNEGIE MELLON UNIVERSITY AND SOFTWARE ENGINEERING INSTITUTE MATERIAL IS FURNISHED ON AN “AS-IS” BASIS. CARNEGIE MELLON UNIVERSITY MAKES NO WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED, AS TO ANY MATTER INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR PURPOSE OR MERCHANTABILITY, EXCLUSIVITY, OR RESULTS OBTAINED FROM USE OF THE MATERIAL. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WARRANTY OF ANY KIND WITH RESPECT TO FREEDOM FROM PATENT, TRADEMARK, OR COPYRIGHT INFRINGEMENT.

DM-0001883

Who are we and what do we do?

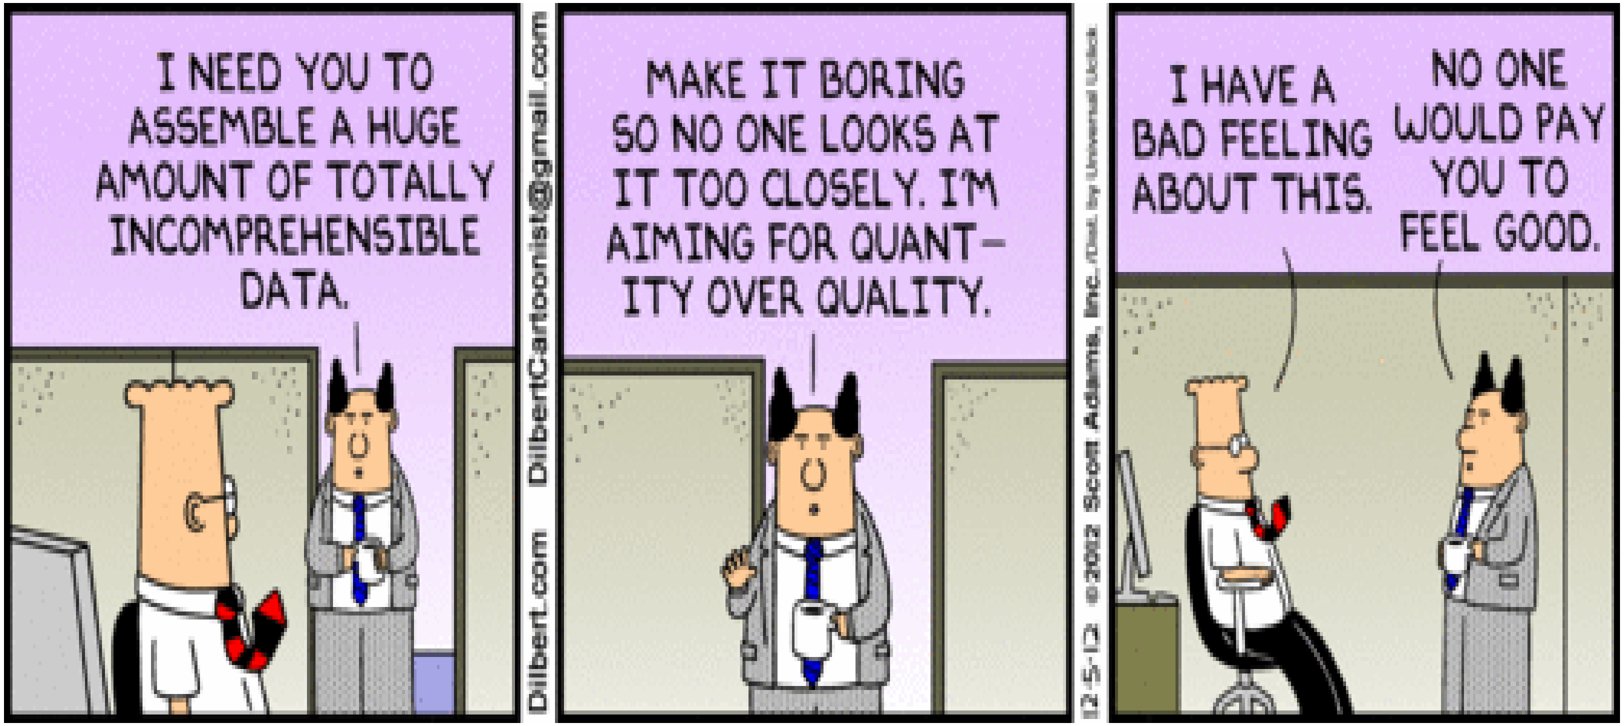
I NEED YOU TO ASSEMBLE A HUGE AMOUNT OF TOTALLY INCOMPREHENSIBLE DATA.
MAKE IT BORING SO NO ONE LOOKS AT IT TOO CLOSELY. I'M AIMING FOR QUANT-ITY OVER QUALITY.
I HAVE A BAD FEELING ABOUT THIS.
NO ONE WOULD PAY YOU TO FEEL GOOD.
Dilbert.com DilbertCartoonist@gmail.com
12-5-12 © 2002 Scott Adams, Inc./Dist. by Universal Uclick

What’s our problem?

- Small team
- Lots of users (all untrusted)
- Lots of systems
- Lots of logs
- Luckily, no “sensitive” information BUT

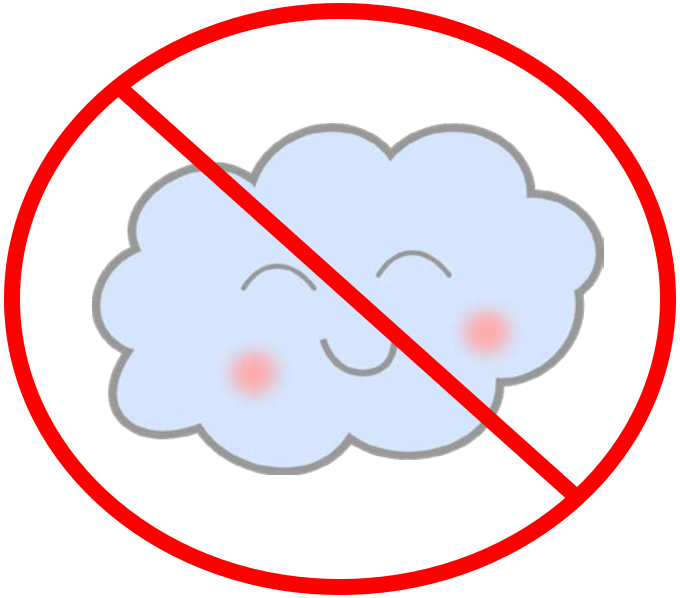

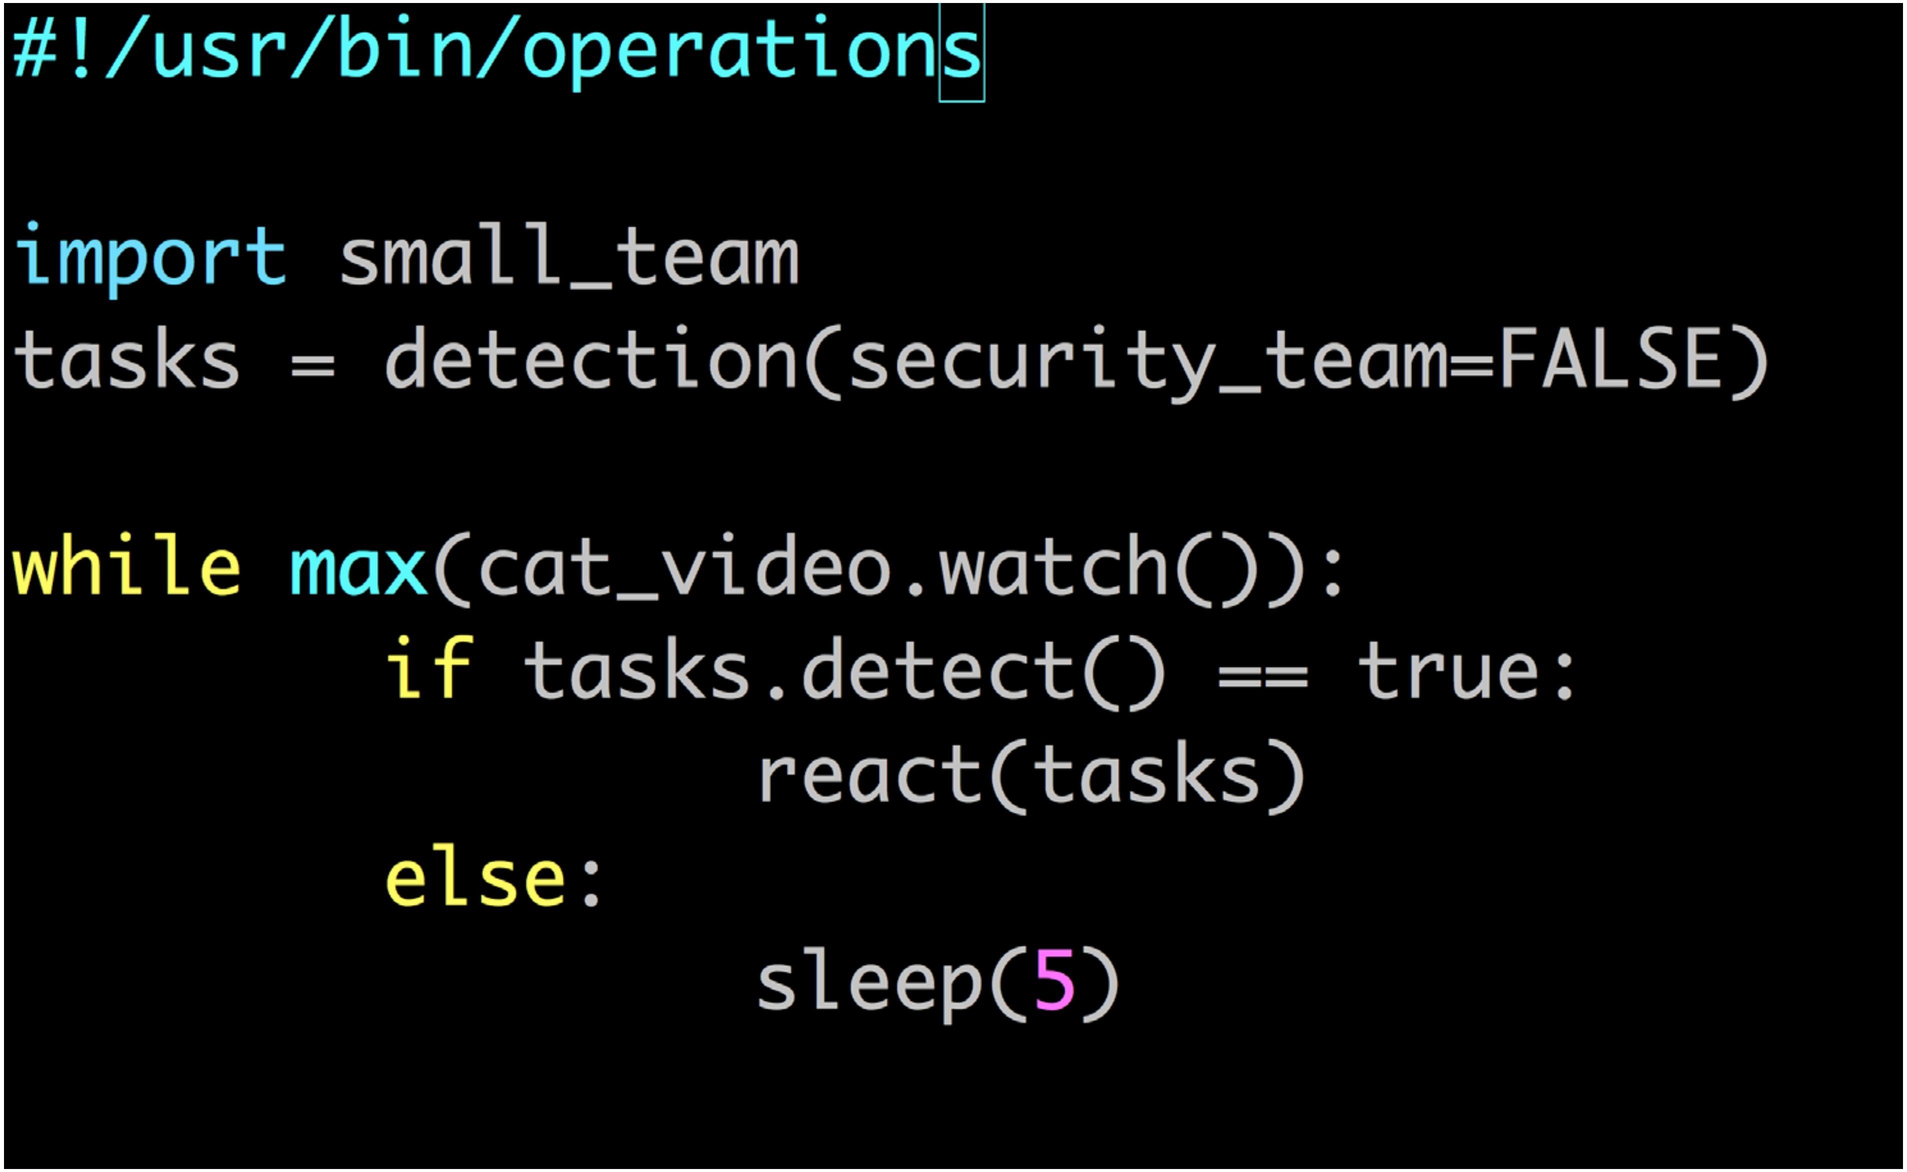

```
#!/usr/bin/operations

import small_team
tasks = detection(security_team=FALSE)

while max(cat_video.watch()):
        if tasks.detect() == true:
                react(tasks)
        else:
                sleep(5)
```

Why Elasticsearch

- Easy to deploy (minimum configuration)
- Scales vertically and horizontally
- Easy to use API
- Modules for most programming/scripting languages
- Actively developed with good online documentation
- It’s free

How Elasticsearch Works in 25 seconds

Shards

- Single instance of Lucene on a node
- Can be primary or replica

Index

- Mapping of shards to nodes
- Like a database within a relational database

Nodes

- Keeps a copy of the index
- Maintain primary and replica shards

Hardware and Infrastructure

- Blades
- Network attached storage – NFS
- Aggregate TAP, SPAN off switches (physical and virtual)
- Virtualization (VMware)
- Puppet

Nodes

8 x Nodes – virtualized

- 4x Cores
- 16 GB ram
- 500 GB data partition (NFS->NAS)

Deployed/Configured using Puppet modules.
https://forge.puppetlabs.com/

Software

(Data Collection)

(Queuing)

(Glue/Integration)

(Storage, index, search)

(Visualization)

Data Sources

- Windows Event Logs
- Syslog
- Bro (session data/dpi)
- SiLK (flow)
- SNMP
- PCAP (stored on disk, index information in ES)

Can I see a diagram with boxes and arrows?

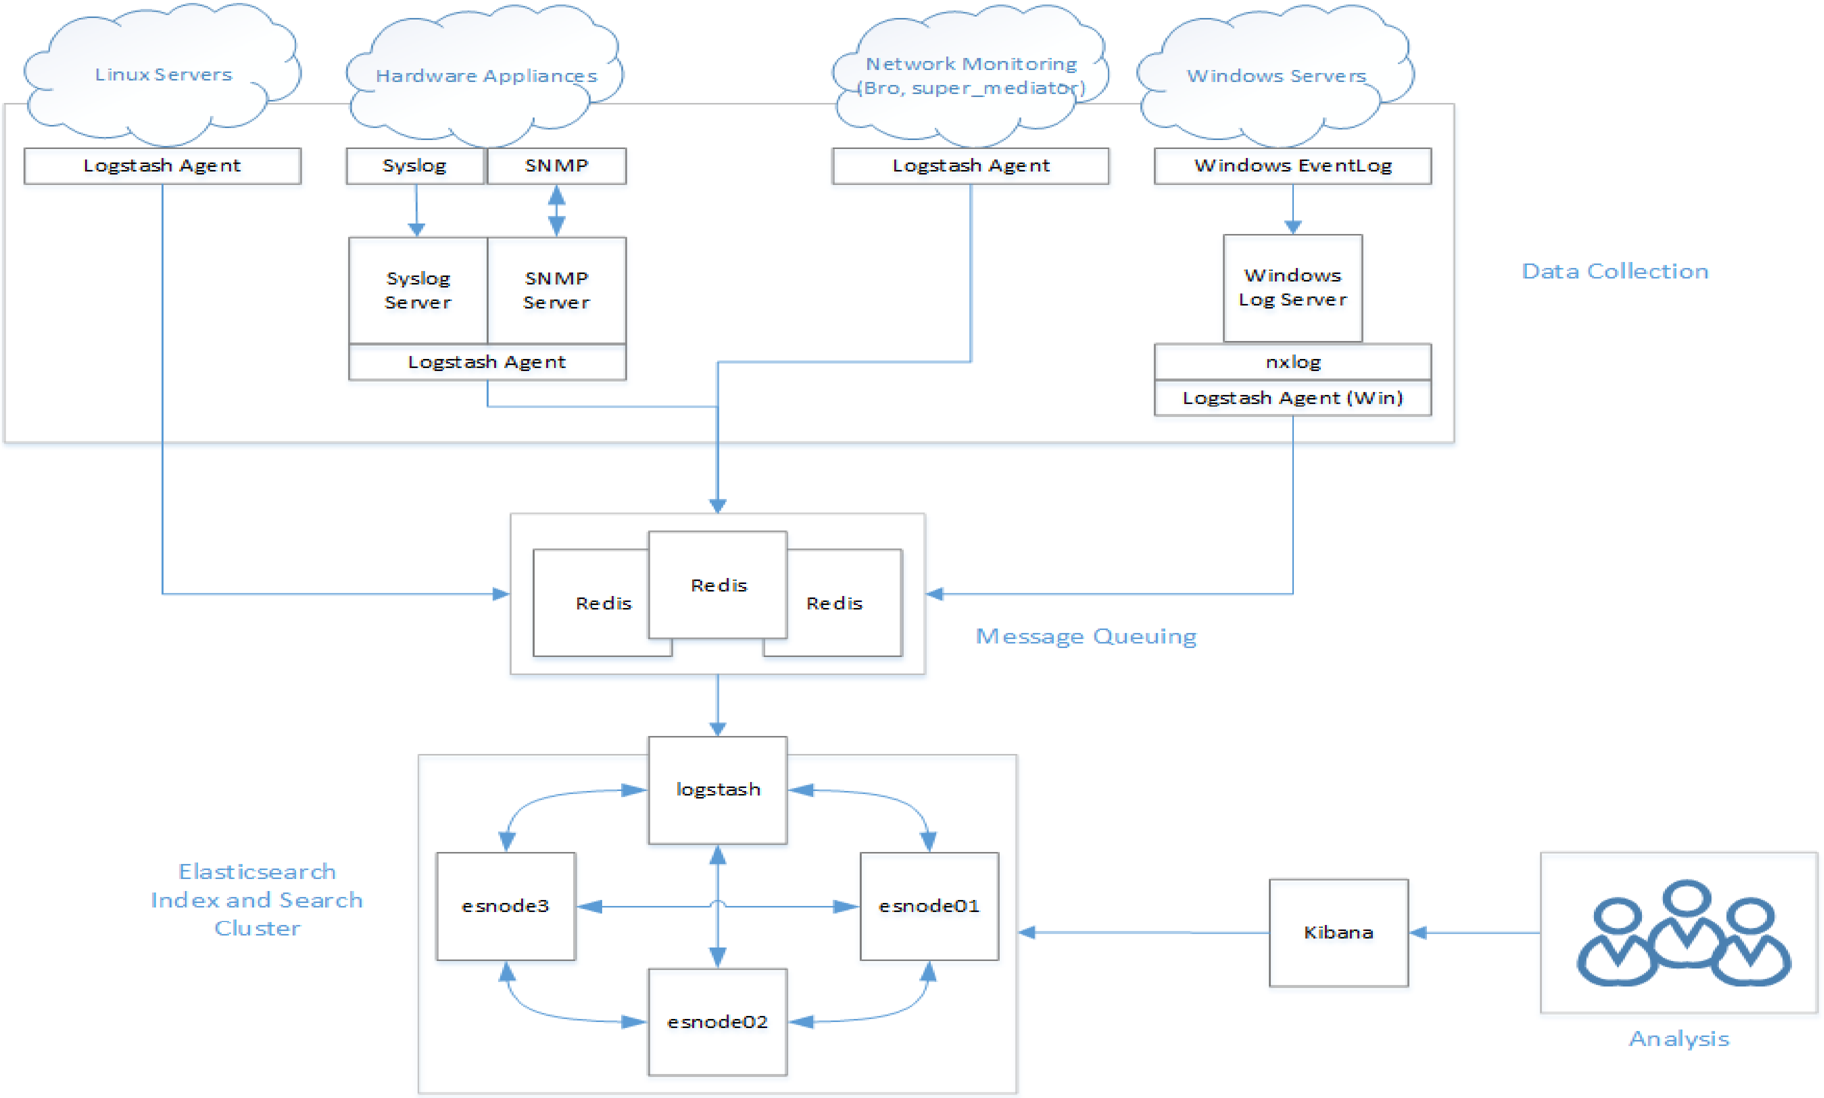

Things we can do

- Batch analysis (retrospective)
- Correlation between data sets
- Make pretty graphs for displaying on TVs – Kibana
- Alerting – Python/R

Where we want to do

Puppet / Applications containers (ie, Docker)

Our environment is defined in software.

Can we use this to automate auditing?

Batch/Retrospective Analysis

- Say we saw some interesting traffic coming from one of our servers – we want to know which processes were run around that time on that host…
- Set a simple filter in Kibana like…

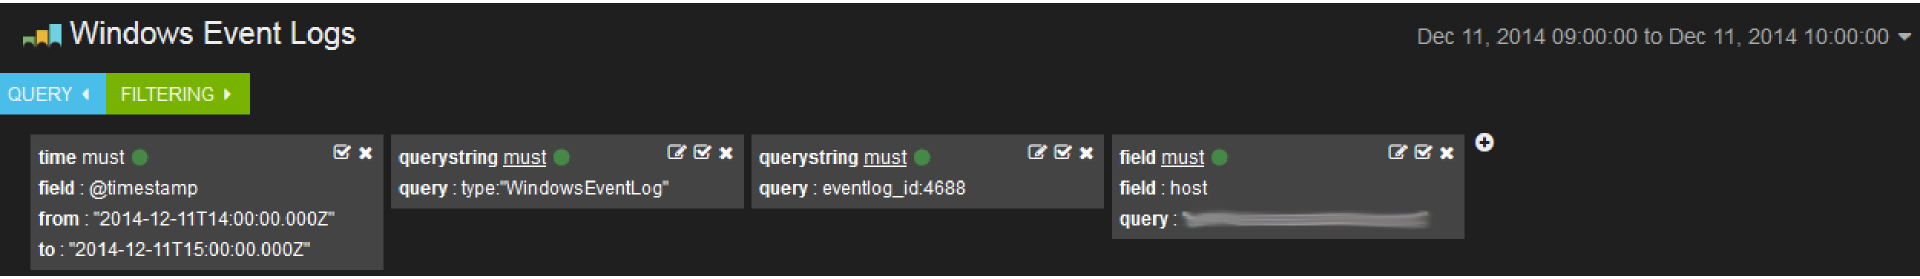

Windows Event Logs

Dec 11, 2014 09:00:00 to Dec 11, 2014 10:00:00

time must
field : @timestamp
from : "2014-12-11T14:00:00.000Z"
to : "2014-12-11T15:00:00.000Z"

querystring must
query : type:"WindowsEventLog"

querystring must
query : eventlog_id:4688

field must
field : host
query : [redacted]

- Kibana queries ES and returns…

Batch/Retrospective Analysis

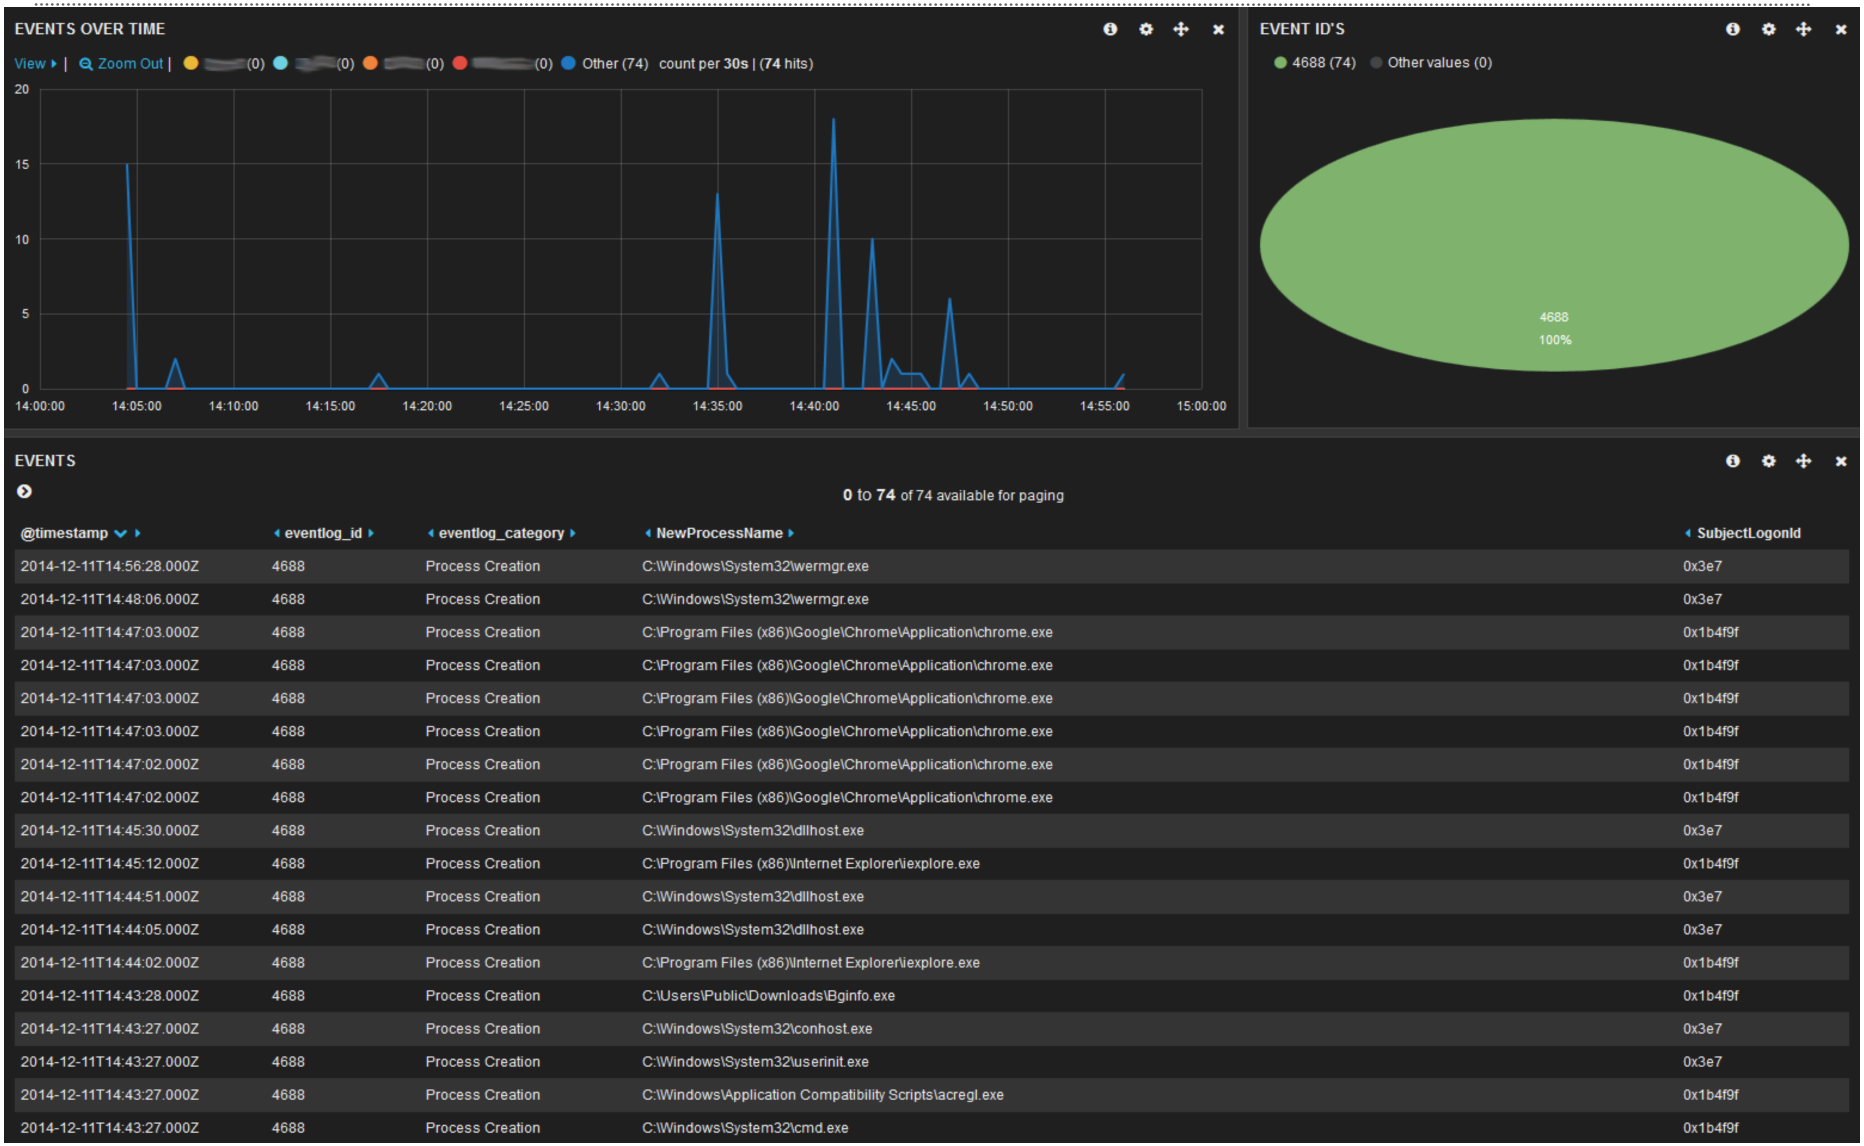

EVENTS OVER TIME
View ▸ | Zoom Out | ● [redacted] (0) ● [redacted] (0) ● [redacted] (0) ● [redacted] (0) ● Other (74) count per 30s | (74 hits)
20
15
10
5
0
14:00:00
14:05:00
14:10:00
14:15:00
14:20:00
14:25:00
14:30:00
14:35:00
14:40:00
14:45:00
14:50:00
14:55:00
15:00:00
EVENT ID'S
● 4688 (74) ● Other values (0)
4688
100%
EVENTS

0 to 74 of 74 available for paging

@timestamp	eventlog_id	eventlog_category	NewProcessName	SubjectLogonId
2014-12-11T14:56:28.000Z	4688	Process Creation	C:\Windows\System32\wermgr.exe	0x3e7
2014-12-11T14:48:06.000Z	4688	Process Creation	C:\Windows\System32\wermgr.exe	0x3e7
2014-12-11T14:47:03.000Z	4688	Process Creation	C:\Program Files (x86)\Google\Chrome\Application\chrome.exe	0x1b4f9f
2014-12-11T14:47:03.000Z	4688	Process Creation	C:\Program Files (x86)\Google\Chrome\Application\chrome.exe	0x1b4f9f
2014-12-11T14:47:03.000Z	4688	Process Creation	C:\Program Files (x86)\Google\Chrome\Application\chrome.exe	0x1b4f9f
2014-12-11T14:47:03.000Z	4688	Process Creation	C:\Program Files (x86)\Google\Chrome\Application\chrome.exe	0x1b4f9f
2014-12-11T14:47:02.000Z	4688	Process Creation	C:\Program Files (x86)\Google\Chrome\Application\chrome.exe	0x1b4f9f
2014-12-11T14:47:02.000Z	4688	Process Creation	C:\Program Files (x86)\Google\Chrome\Application\chrome.exe	0x1b4f9f
2014-12-11T14:45:30.000Z	4688	Process Creation	C:\Windows\System32\dllhost.exe	0x3e7
2014-12-11T14:45:12.000Z	4688	Process Creation	C:\Program Files (x86)\Internet Explorer\iexplore.exe	0x1b4f9f
2014-12-11T14:44:51.000Z	4688	Process Creation	C:\Windows\System32\dllhost.exe	0x3e7
2014-12-11T14:44:05.000Z	4688	Process Creation	C:\Windows\System32\dllhost.exe	0x3e7
2014-12-11T14:44:02.000Z	4688	Process Creation	C:\Program Files (x86)\Internet Explorer\iexplore.exe	0x1b4f9f
2014-12-11T14:43:28.000Z	4688	Process Creation	C:\Users\Public\Downloads\Bginfo.exe	0x1b4f9f
2014-12-11T14:43:27.000Z	4688	Process Creation	C:\Windows\System32\conhost.exe	0x3e7
2014-12-11T14:43:27.000Z	4688	Process Creation	C:\Windows\System32\userinit.exe	0x3e7
2014-12-11T14:43:27.000Z	4688	Process Creation	C:\Windows\Application Compatibility Scripts\acregl.exe	0x1b4f9f
2014-12-11T14:43:27.000Z	4688	Process Creation	C:\Windows\System32\cmd.exe	0x1b4f9f

Batch/Retrospective Analysis

- You can also use ES Python API to perform queries – http://elasticsearch-py.rtfd.org/
- Lots of query and filter options; JSON syntax; more flexibility and control
- Good for…
  - Running queries on-demand over any period of time
  - Checking on important events that are too cumbersome to alert on
  - Daily review of logs
  - Investigation

Batch/Retrospective Analysis

- Example query bodies

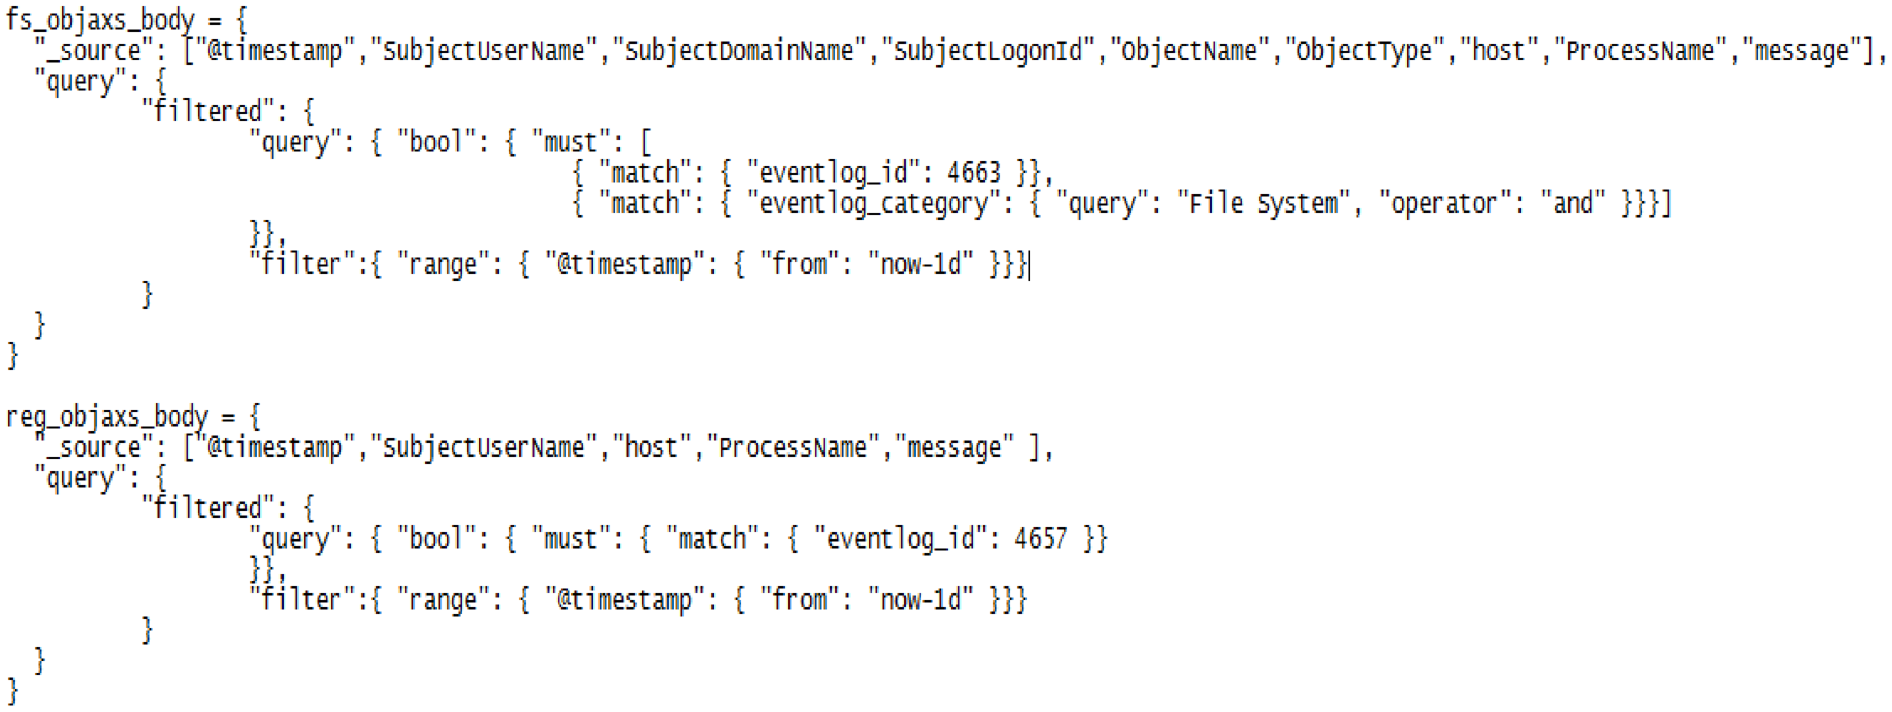

```
fs_objaxs_body = {
  "_source": ["@timestamp","SubjectUserName","SubjectDomainName","SubjectLogonId","ObjectName","ObjectType","host","ProcessName","message"],
  "query": {
          "filtered": {
                  "query": { "bool": { "must": [
                                          { "match": { "eventlog_id": 4663 }},
                                          { "match": { "eventlog_category": { "query": "File System", "operator": "and" }}}]
                  }},
                  "filter":{ "range": { "@timestamp": { "from": "now-1d" }}}
          }
  }
}

reg_objaxs_body = {
  "_source": ["@timestamp","SubjectUserName","host","ProcessName","message" ],
  "query": {
          "filtered": {
                  "query": { "bool": { "must": { "match": { "eventlog_id": 4657 }}
                  }},
                  "filter":{ "range": { "@timestamp": { "from": "now-1d" }}}
          }
  }
}
```

- And get…

Batch/Retrospective Analysis

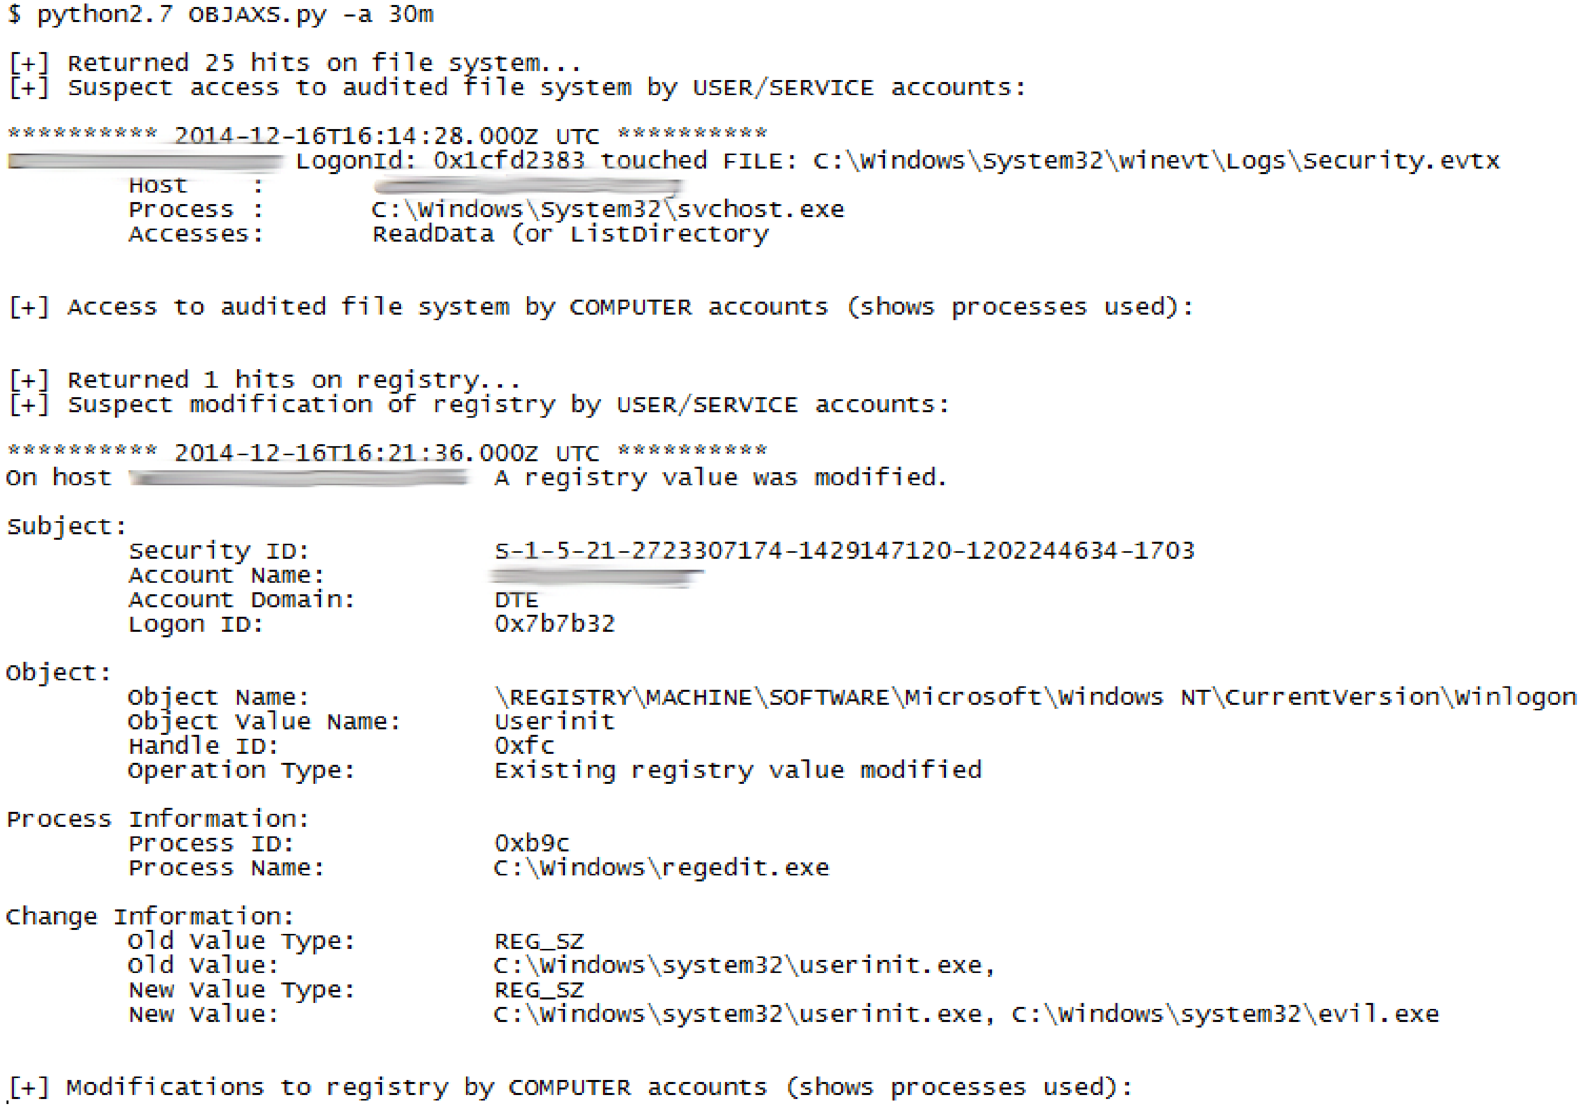

```
$ python2.7 OBJAXS.py -a 30m

[+] Returned 25 hits on file system...
[+] Suspect access to audited file system by USER/SERVICE accounts:

*********** 2014-12-16T16:14:28.000Z UTC ***********
[redacted] LogonId: 0x1cfd2383 touched FILE: C:\Windows\System32\winevt\Logs\Security.evtx
        Host    :       [redacted]
        Process :       C:\Windows\System32\svchost.exe
        Accesses:       ReadData (or ListDirectory


[+] Access to audited file system by COMPUTER accounts (shows processes used):


[+] Returned 1 hits on registry...
[+] Suspect modification of registry by USER/SERVICE accounts:

*********** 2014-12-16T16:21:36.000Z UTC ***********
On host [redacted] A registry value was modified.

Subject:
        Security ID:            S-1-5-21-2723307174-1429147120-1202244634-1703
        Account Name:           [redacted]
        Account Domain:         DTE
        Logon ID:               0x7b7b32

Object:
        Object Name:            \REGISTRY\MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon
        Object Value Name:      Userinit
        Handle ID:              0xfc
        Operation Type:         Existing registry value modified

Process Information:
        Process ID:             0xb9c
        Process Name:           C:\Windows\regedit.exe

Change Information:
        Old Value Type:         REG_SZ
        Old Value:              C:\Windows\system32\userinit.exe,
        New Value Type:         REG_SZ
        New Value:              C:\Windows\system32\userinit.exe, C:\Windows\system32\evil.exe


[+] Modifications to registry by COMPUTER accounts (shows processes used):
```

Correlation of Data Sets & Visualization

- Correlate events within and between data sets to gain context
- Visualizing data with Kibana facets…
  - Makes aspects of data more readily apparent
  - Aids perspective and understanding of data
  - Looks cool
- Typically…
  - Attach one or more Queries to individual facets
  - Drill down on specific data using Filters (whole page)
  - Plethora of info with just one or two filters

Correlation of Data Sets & Visualization

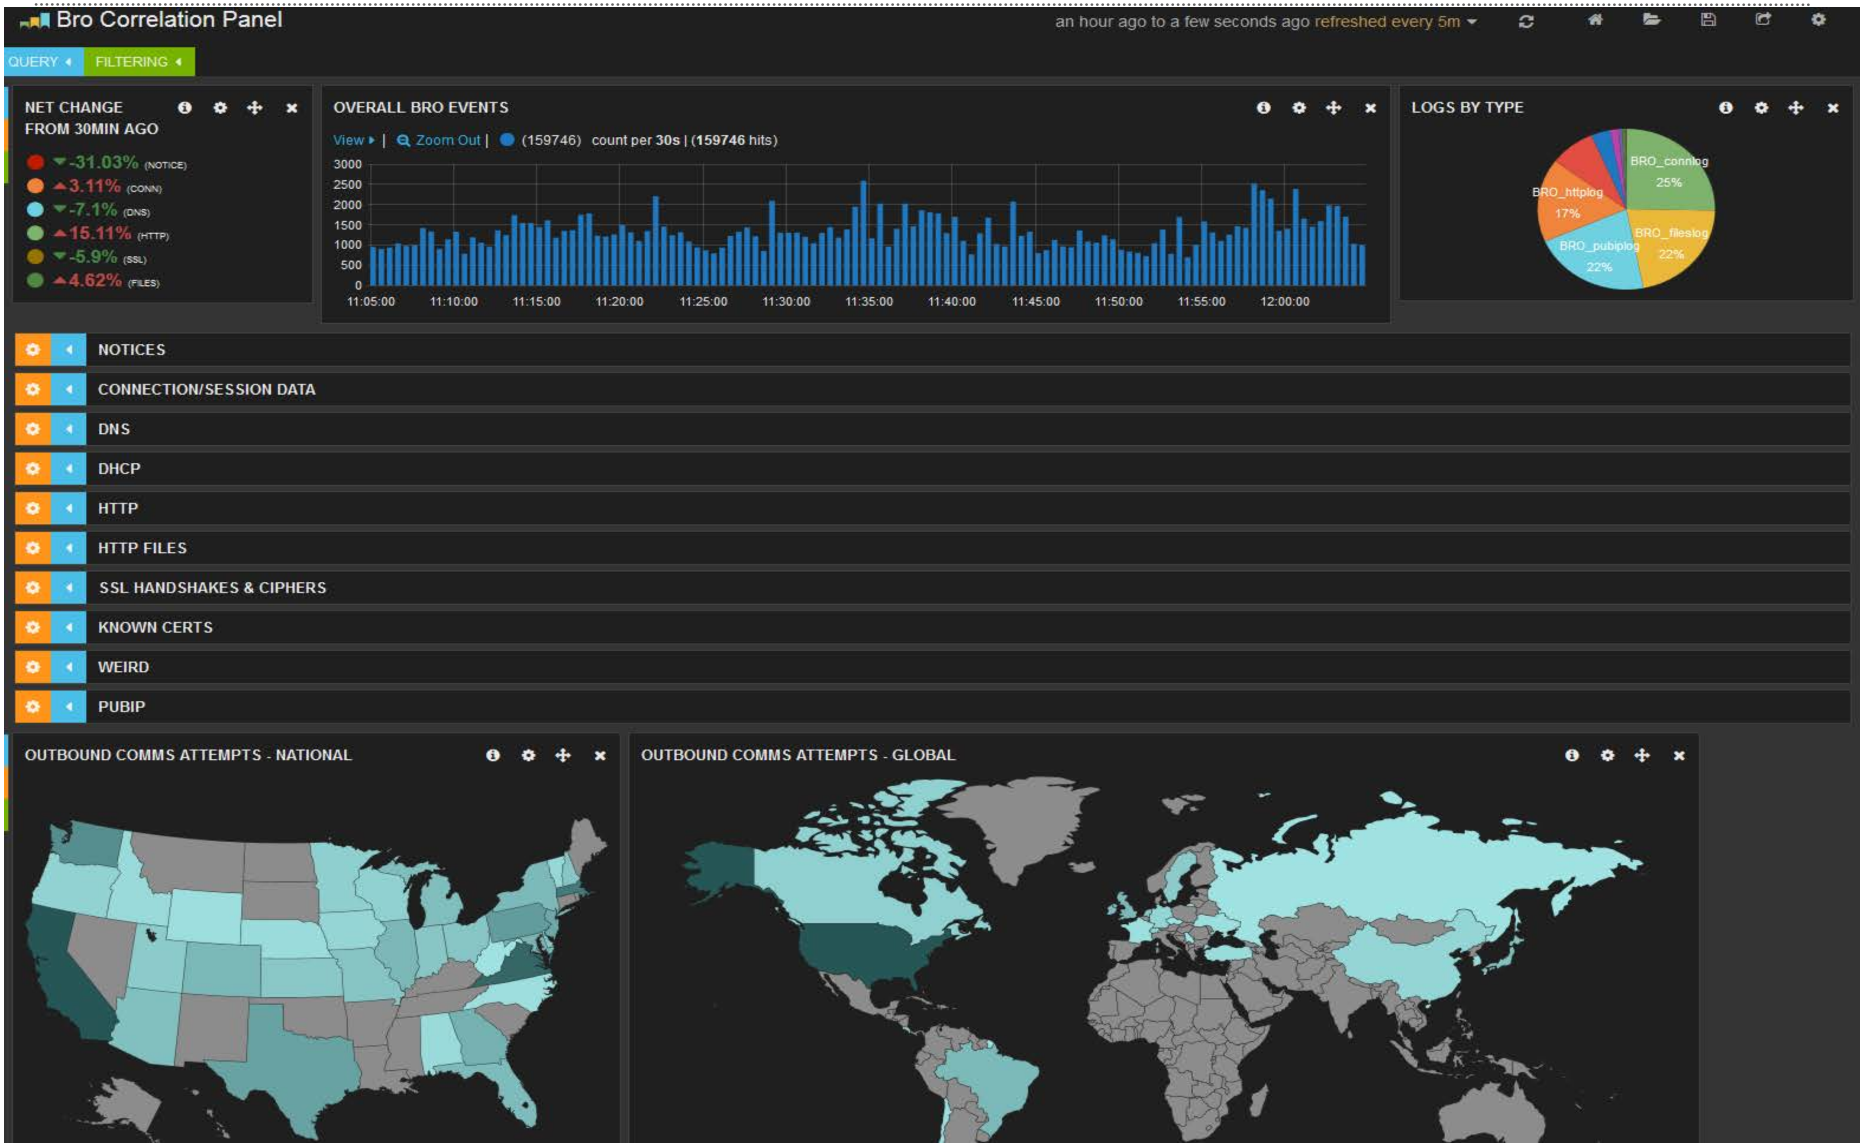

Bro Correlation Panel

an hour ago to a few seconds ago refreshed every 5m

QUERY

FILTERING

NET CHANGE FROM 30MIN AGO

- ● ▼-31.03% (NOTICE)
- ● ▲3.11% (CONN)
- ● ▼-7.1% (DNS)
- ● ▲15.11% (HTTP)
- ● ▼-5.9% (SSL)
- ● ▲4.62% (FILES)

OVERALL BRO EVENTS

View ▸ | Zoom Out | ● (159746) count per 30s | (159746 hits)

LOGS BY TYPE

NOTICES

CONNECTION/SESSION DATA

DNS

DHCP

HTTP

HTTP FILES

SSL HANDSHAKES & CIPHERS

KNOWN CERTS

WEIRD

PUBIP

OUTBOUND COMMS ATTEMPTS - NATIONAL

OUTBOUND COMMS ATTEMPTS - GLOBAL

Correlation of Data Sets & Visualization

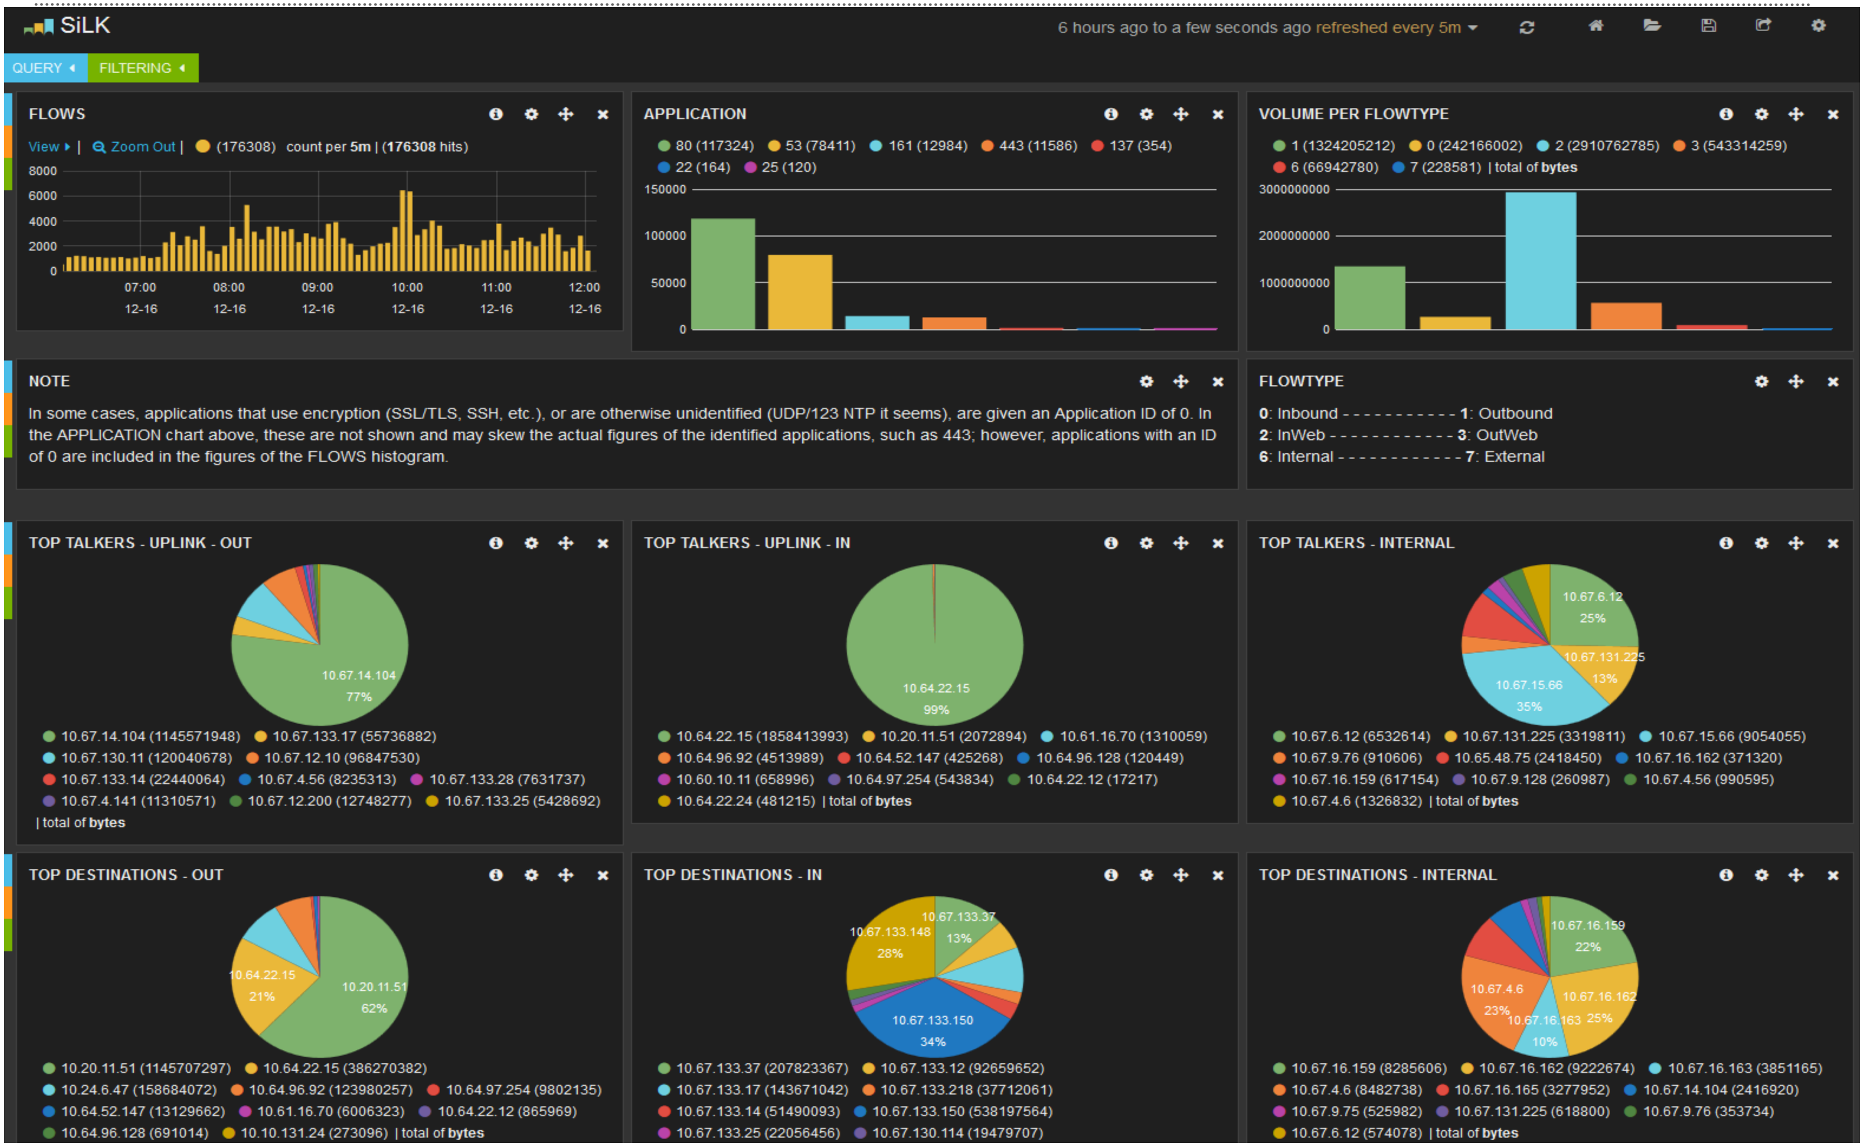

FLOWS

View | Zoom Out | ● (176308) count per **5m** | (**176308** hits)

APPLICATION

● 80 (117324) ● 53 (78411) ● 161 (12984) ● 443 (11586) ● 137 (354) ● 22 (164) ● 25 (120)

VOLUME PER FLOWTYPE

● 1 (1324205212) ● 0 (242166002) ● 2 (2910762785) ● 3 (543314259) ● 6 (66942780) ● 7 (228581) | total of **bytes**

NOTE

In some cases, applications that use encryption (SSL/TLS, SSH, etc.), or are otherwise unidentified (UDP/123 NTP it seems), are given an Application ID of 0. In the APPLICATION chart above, these are not shown and may skew the actual figures of the identified applications, such as 443; however, applications with an ID of 0 are included in the figures of the FLOWS histogram.

FLOWTYPE

0: Inbound - - - - - - - - - - - **1**: Outbound
2: InWeb - - - - - - - - - - - - **3**: OutWeb
6: Internal - - - - - - - - - - - - **7**: External

TOP TALKERS - UPLINK - OUT

● 10.67.14.104 (1145571948) ● 10.67.133.17 (55736882) ● 10.67.130.11 (120040678) ● 10.67.12.10 (96847530) ● 10.67.133.14 (22440064) ● 10.67.4.56 (8235313) ● 10.67.133.28 (7631737) ● 10.67.4.141 (11310571) ● 10.67.12.200 (12748277) ● 10.67.133.25 (5428692) | total of **bytes**

TOP TALKERS - UPLINK - IN

● 10.64.22.15 (1858413993) ● 10.20.11.51 (2072894) ● 10.61.16.70 (1310059) ● 10.64.96.92 (4513989) ● 10.64.52.147 (425268) ● 10.64.96.128 (120449) ● 10.60.10.11 (658996) ● 10.64.97.254 (543834) ● 10.64.22.12 (17217) ● 10.64.22.24 (481215) | total of **bytes**

TOP TALKERS - INTERNAL

● 10.67.6.12 (6532614) ● 10.67.131.225 (3319811) ● 10.67.15.66 (9054055) ● 10.67.9.76 (910606) ● 10.65.48.75 (2418450) ● 10.67.16.162 (371320) ● 10.67.16.159 (617154) ● 10.67.9.128 (260987) ● 10.67.4.56 (990595) ● 10.67.4.6 (1326832) | total of **bytes**

TOP DESTINATIONS - OUT

● 10.20.11.51 (1145707297) ● 10.64.22.15 (386270382) ● 10.24.6.47 (158684072) ● 10.64.96.92 (123980257) ● 10.64.97.254 (9802135) ● 10.64.52.147 (13129662) ● 10.61.16.70 (6006323) ● 10.64.22.12 (865969) ● 10.64.96.128 (691014) ● 10.10.131.24 (273096) | total of **bytes**

TOP DESTINATIONS - IN

● 10.67.133.37 (207823367) ● 10.67.133.12 (92659652) ● 10.67.133.17 (143671042) ● 10.67.133.218 (37712061) ● 10.67.133.14 (51490093) ● 10.67.133.150 (538197564) ● 10.67.133.25 (22056456) ● 10.67.130.114 (19479707)

TOP DESTINATIONS - INTERNAL

● 10.67.16.159 (8285606) ● 10.67.16.162 (9222674) ● 10.67.16.163 (3851165) ● 10.67.4.6 (8482738) ● 10.67.16.165 (3277952) ● 10.67.14.104 (2416920) ● 10.67.9.75 (525982) ● 10.67.131.225 (618800) ● 10.67.9.76 (353734) ● 10.67.6.12 (574078) | total of **bytes**

Alerting – Windows Event Monitoring

- Want to know about certain events as they occur
  - Administrator login
  - Local/SAM account login attempts
  - User account creation/re-enabling
  - Creation/Addition to Groups
  - Scheduled Task creation
  - Log cleared
- Uses ES Python API and CRON
  - Queries ES 15 times per hour
  - Every 4 minutes -> “from”: “now-4m”

Alerting – Windows Event Monitoring

- Example Alerts received via email notifications

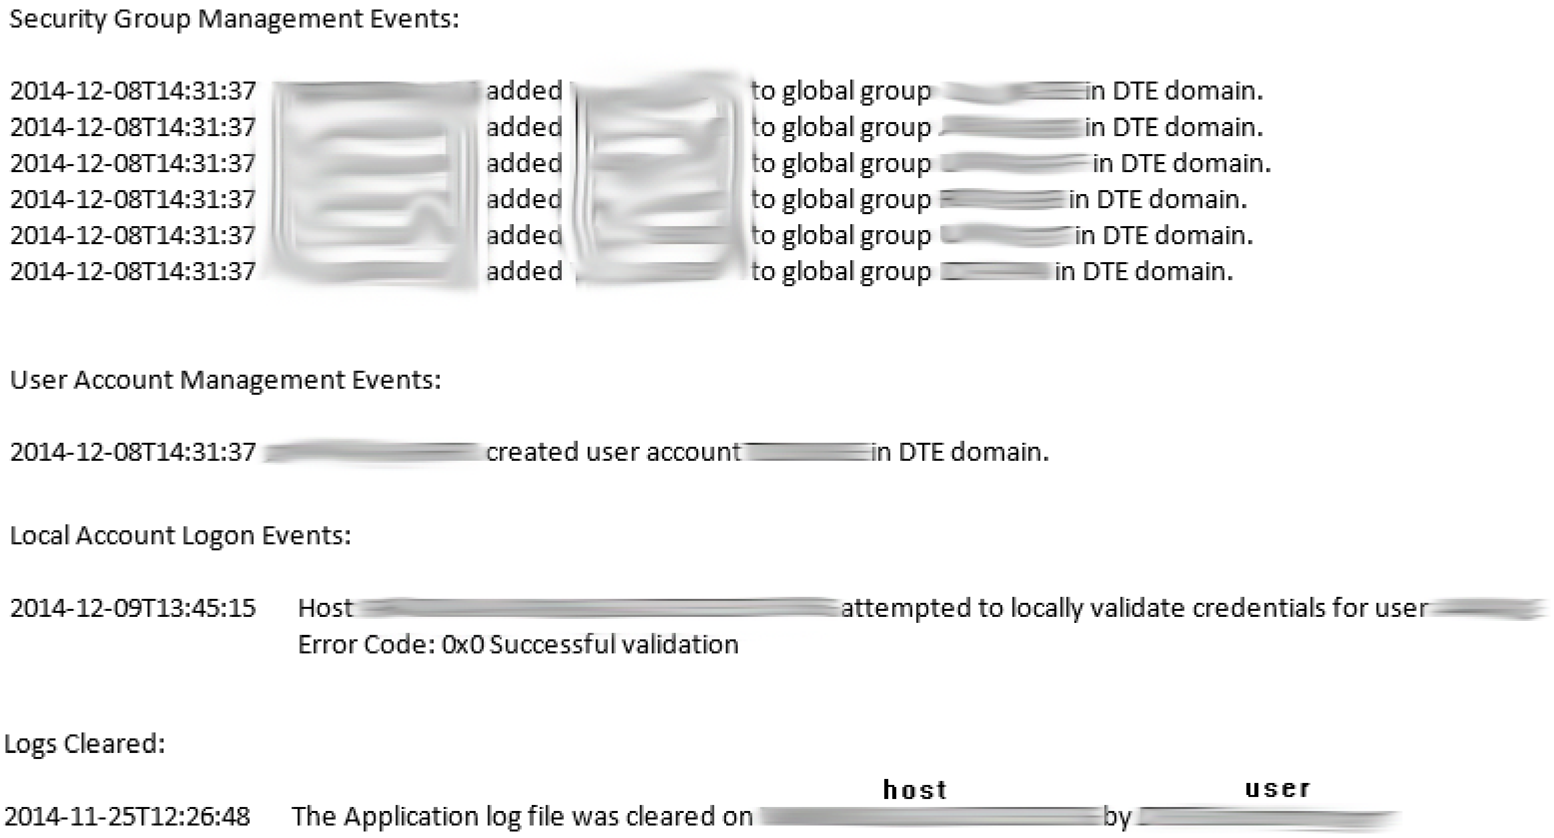

```
Security Group Management Events:

2014-12-08T14:31:37 [redacted] added [redacted] to global group [redacted] in DTE domain.
2014-12-08T14:31:37 [redacted] added [redacted] to global group [redacted] in DTE domain.
2014-12-08T14:31:37 [redacted] added [redacted] to global group [redacted] in DTE domain.
2014-12-08T14:31:37 [redacted] added [redacted] to global group [redacted] in DTE domain.
2014-12-08T14:31:37 [redacted] added [redacted] to global group [redacted] in DTE domain.
2014-12-08T14:31:37 [redacted] added [redacted] to global group [redacted] in DTE domain.


User Account Management Events:

2014-12-08T14:31:37 [redacted] created user account [redacted] in DTE domain.

Local Account Logon Events:

2014-12-09T13:45:15    Host [redacted] attempted to locally validate credentials for user [redacted]
                       Error Code: 0x0 Successful validation


Logs Cleared:
                                                                 host                  user
2014-11-25T12:26:48    The Application log file was cleared on [redacted] by [redacted]
```

Alerting – Windows Event Monitoring

- Example Alerts received via email notifications

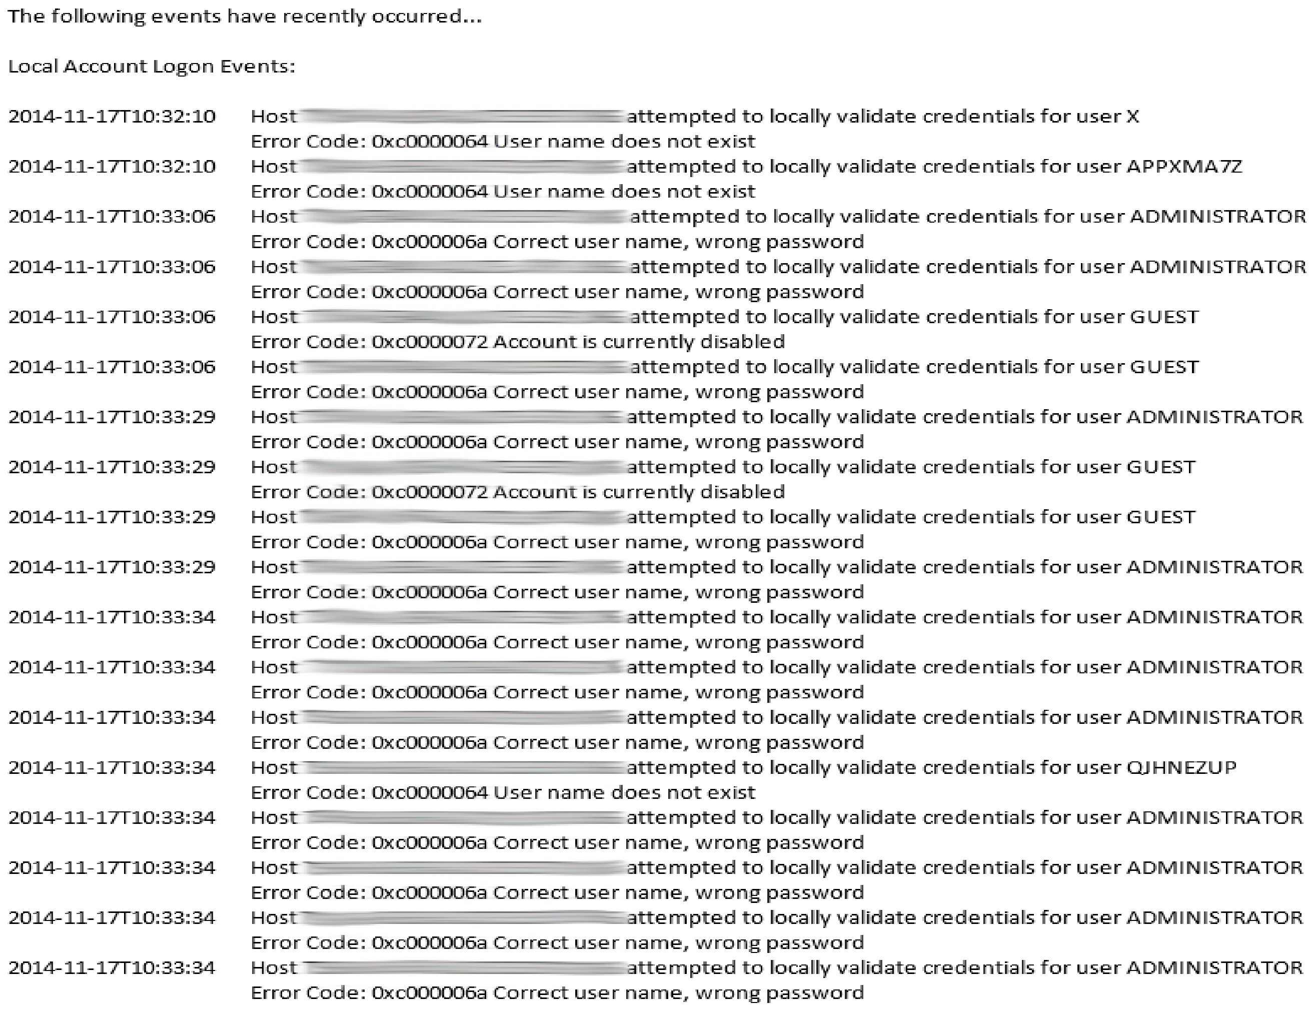

The following events have recently occurred...

Local Account Logon Events:

```
2014-11-17T10:32:10    Host [redacted] attempted to locally validate credentials for user X
                       Error Code: 0xc0000064 User name does not exist
2014-11-17T10:32:10    Host [redacted] attempted to locally validate credentials for user APPXMA7Z
                       Error Code: 0xc0000064 User name does not exist
2014-11-17T10:33:06    Host [redacted] attempted to locally validate credentials for user ADMINISTRATOR
                       Error Code: 0xc000006a Correct user name, wrong password
2014-11-17T10:33:06    Host [redacted] attempted to locally validate credentials for user ADMINISTRATOR
                       Error Code: 0xc000006a Correct user name, wrong password
2014-11-17T10:33:06    Host [redacted] attempted to locally validate credentials for user GUEST
                       Error Code: 0xc0000072 Account is currently disabled
2014-11-17T10:33:06    Host [redacted] attempted to locally validate credentials for user GUEST
                       Error Code: 0xc000006a Correct user name, wrong password
2014-11-17T10:33:29    Host [redacted] attempted to locally validate credentials for user ADMINISTRATOR
                       Error Code: 0xc000006a Correct user name, wrong password
2014-11-17T10:33:29    Host [redacted] attempted to locally validate credentials for user GUEST
                       Error Code: 0xc0000072 Account is currently disabled
2014-11-17T10:33:29    Host [redacted] attempted to locally validate credentials for user GUEST
                       Error Code: 0xc000006a Correct user name, wrong password
2014-11-17T10:33:29    Host [redacted] attempted to locally validate credentials for user ADMINISTRATOR
                       Error Code: 0xc000006a Correct user name, wrong password
2014-11-17T10:33:34    Host [redacted] attempted to locally validate credentials for user ADMINISTRATOR
                       Error Code: 0xc000006a Correct user name, wrong password
2014-11-17T10:33:34    Host [redacted] attempted to locally validate credentials for user ADMINISTRATOR
                       Error Code: 0xc000006a Correct user name, wrong password
2014-11-17T10:33:34    Host [redacted] attempted to locally validate credentials for user ADMINISTRATOR
                       Error Code: 0xc000006a Correct user name, wrong password
2014-11-17T10:33:34    Host [redacted] attempted to locally validate credentials for user QJHNEZUP
                       Error Code: 0xc0000064 User name does not exist
2014-11-17T10:33:34    Host [redacted] attempted to locally validate credentials for user ADMINISTRATOR
                       Error Code: 0xc000006a Correct user name, wrong password
2014-11-17T10:33:34    Host [redacted] attempted to locally validate credentials for user ADMINISTRATOR
                       Error Code: 0xc000006a Correct user name, wrong password
2014-11-17T10:33:34    Host [redacted] attempted to locally validate credentials for user ADMINISTRATOR
                       Error Code: 0xc000006a Correct user name, wrong password
2014-11-17T10:33:34    Host [redacted] attempted to locally validate credentials for user ADMINISTRATOR
                       Error Code: 0xc000006a Correct user name, wrong password
```

Alerting – Windows Event Monitoring

- Example Alerts received via email notifications

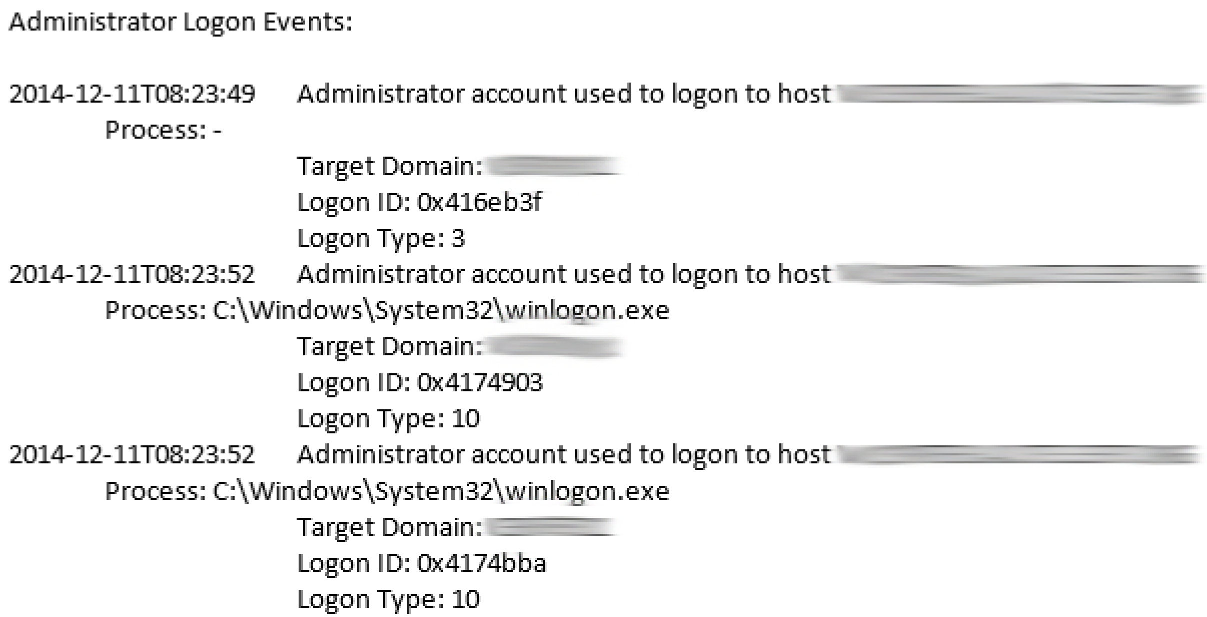

```
Administrator Logon Events:

2014-12-11T08:23:49     Administrator account used to logon to host [redacted]
          Process: -
                        Target Domain: [redacted]
                        Logon ID: 0x416eb3f
                        Logon Type: 3
2014-12-11T08:23:52     Administrator account used to logon to host [redacted]
          Process: C:\Windows\System32\winlogon.exe
                        Target Domain: [redacted]
                        Logon ID: 0x4174903
                        Logon Type: 10
2014-12-11T08:23:52     Administrator account used to logon to host [redacted]
          Process: C:\Windows\System32\winlogon.exe
                        Target Domain: [redacted]
                        Logon ID: 0x4174bba
                        Logon Type: 10
```

Alerting – Windows Event Monitoring

- Example Alerts received via email notifications

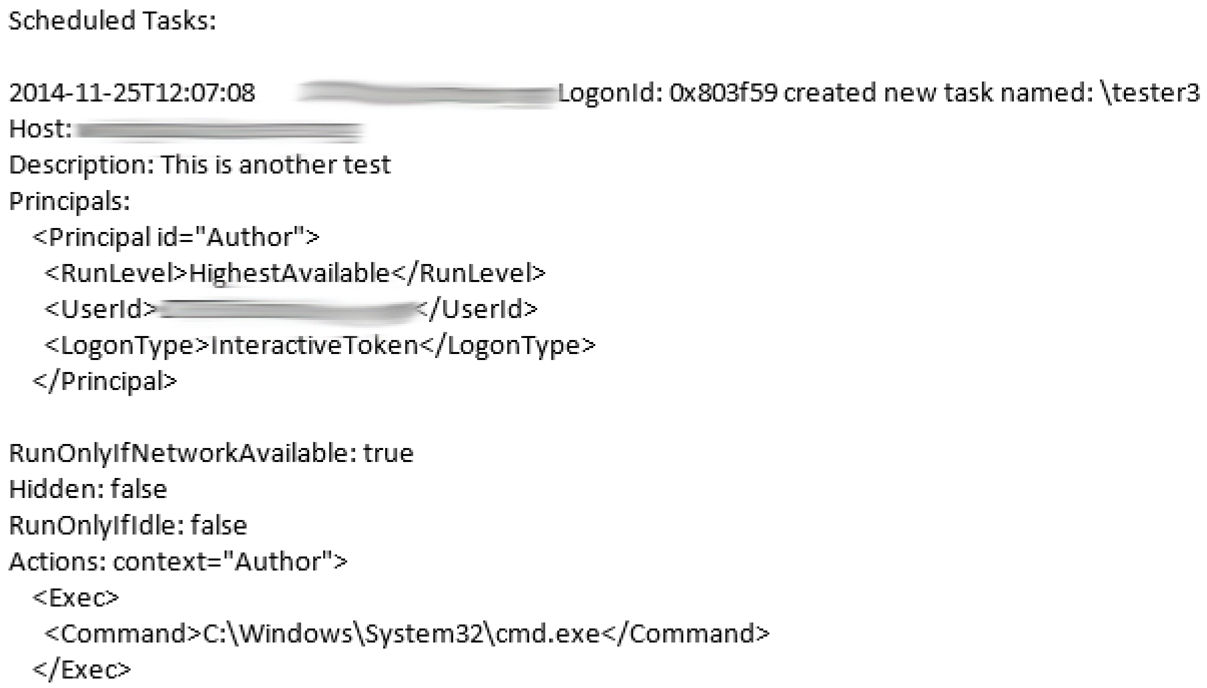

```
Scheduled Tasks:

2014-11-25T12:07:08     [redacted] LogonId: 0x803f59 created new task named: \tester3
Host: [redacted]
Description: This is another test
Principals:
  <Principal id="Author">
   <RunLevel>HighestAvailable</RunLevel>
   <UserId>[redacted]</UserId>
   <LogonType>InteractiveToken</LogonType>
  </Principal>

RunOnlyIfNetworkAvailable: true
Hidden: false
RunOnlyIfIdle: false
Actions: context="Author">
  <Exec>
   <Command>C:\Windows\System32\cmd.exe</Command>
  </Exec>
```

Alerting – Irregular Login Activity

- Want to keep an eye on privileged account use
- Want to know…
  - When users login to hosts they never or rarely ever login to
  - When users login from atypical source IPs
  - When user logins violate certain thresholds based on previous behavior
- Uses ES Python API, CRON, R, and sqlite3 DB
  - Delivers daily login stats
  - Updates weekly and expires old weeks
  - Checks against DB with 4 weeks of aggregated data

Alerting – Irregular Login Activity

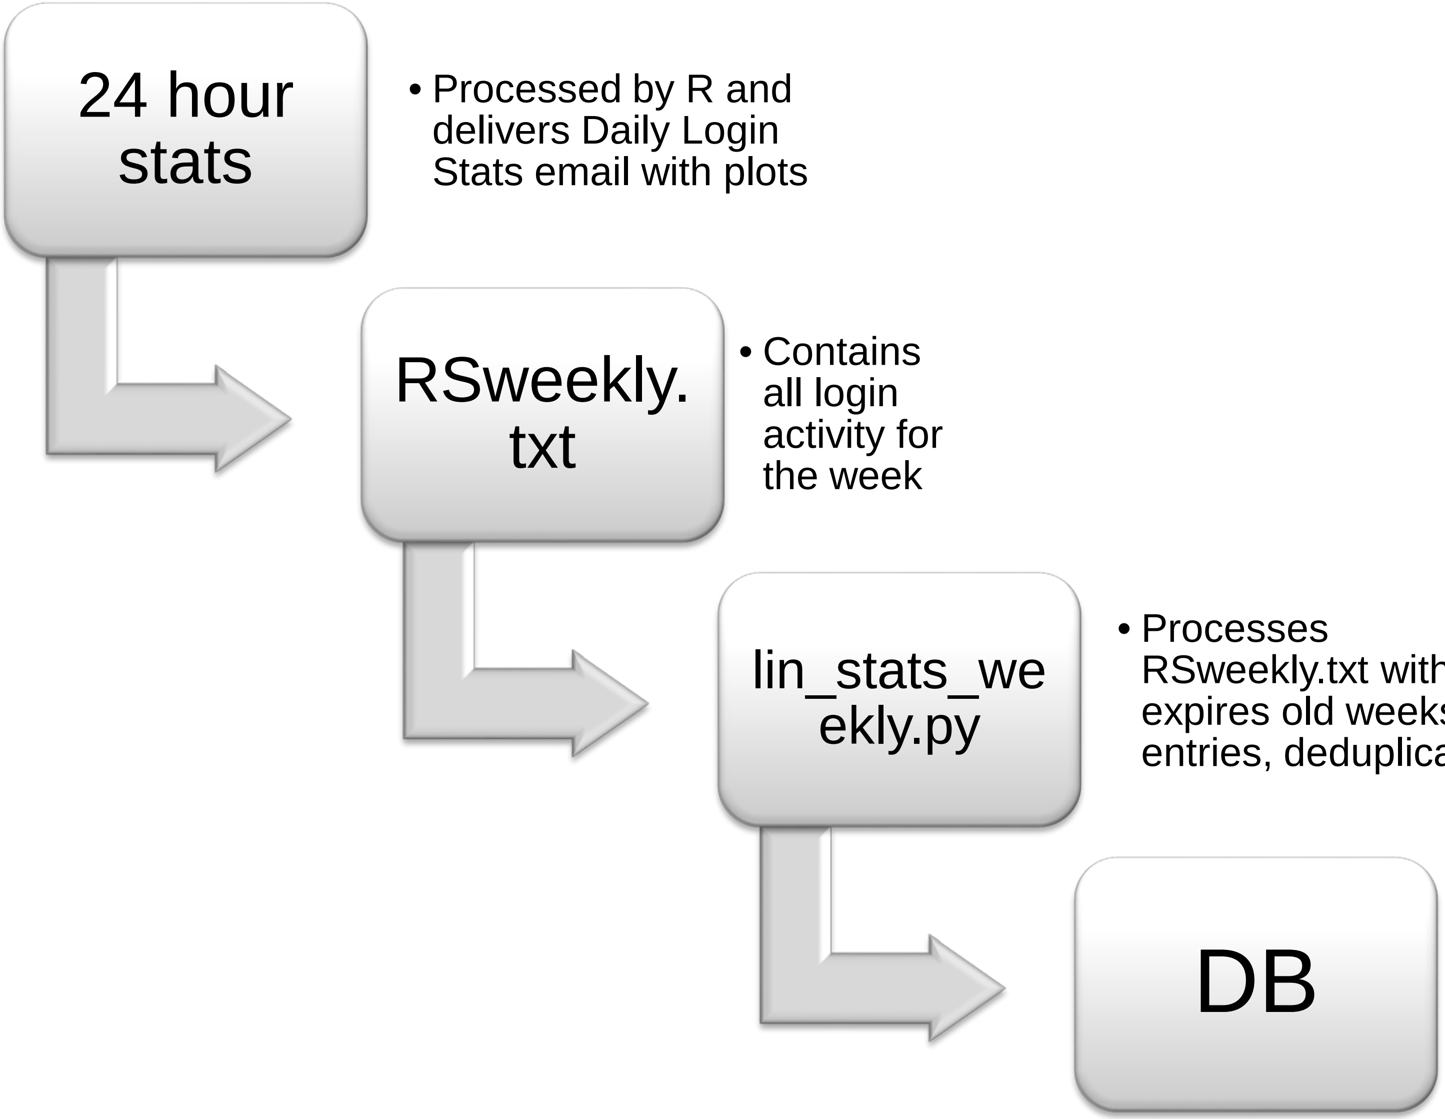

24 hour stats

- Processed by R and delivers Daily Login Stats email with plots

RSweekly.txt

- Contains all login activity for the week

lin_stats_weekly.py

- Processes RSweekly.txt with R, expires old weeks / entries, deduplicate

DB

Alerting – Irregular Login Activity

- Example Daily Login Stats with plots
- Email Message

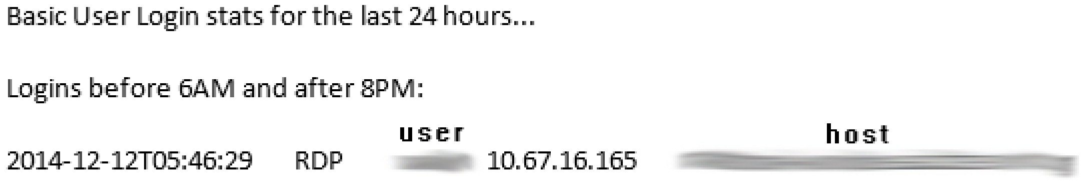

Basic User Login stats for the last 24 hours...

Logins before 6AM and after 8PM:

```
                                 user                                    host
2014-12-12T05:46:29     RDP     [redacted]  10.67.16.165     [redacted]
```

Alerting – Irregular Login Activity

- Example Daily Login Stats with plots

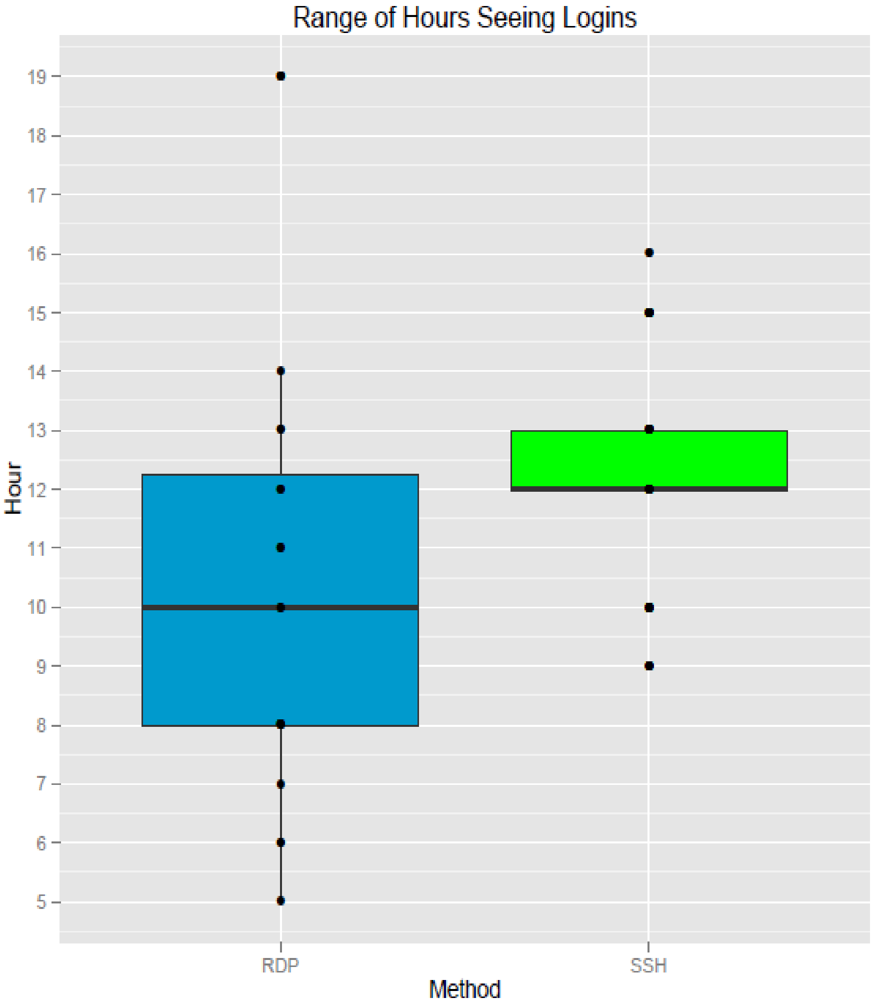

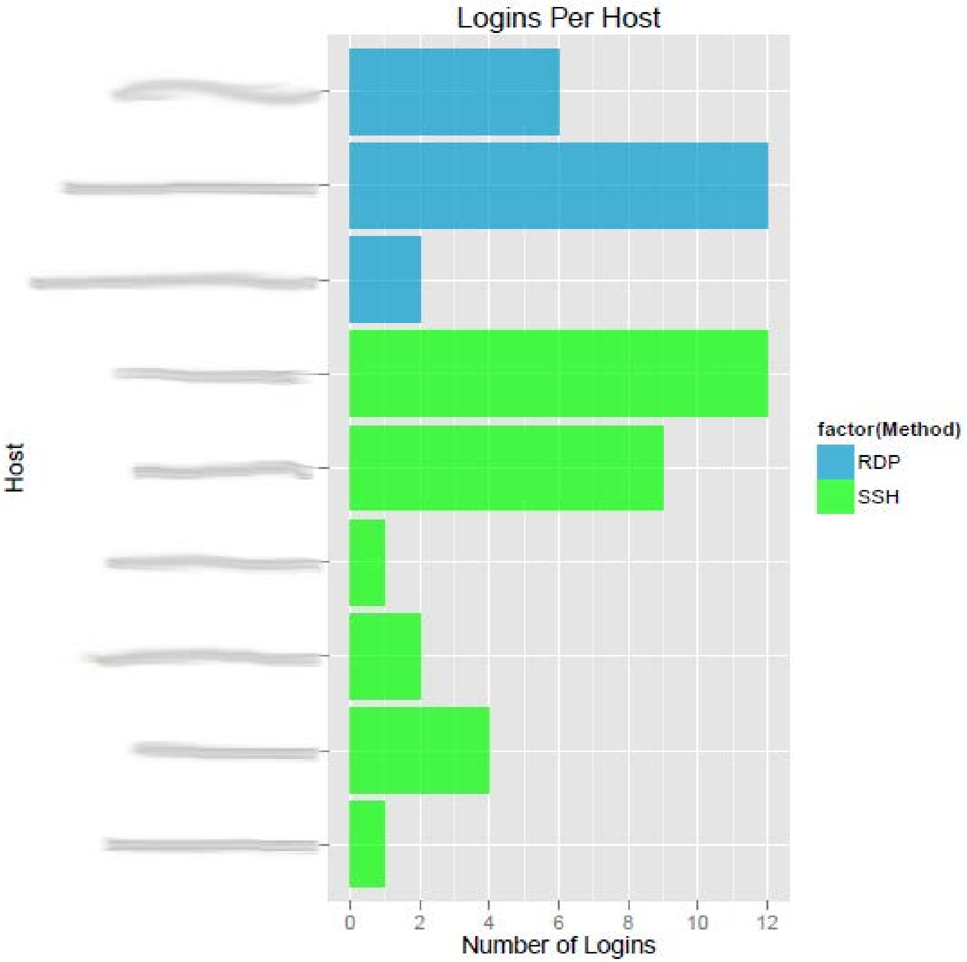

Alerting – Irregular Login Activity

- Example Daily Login Stats with plots

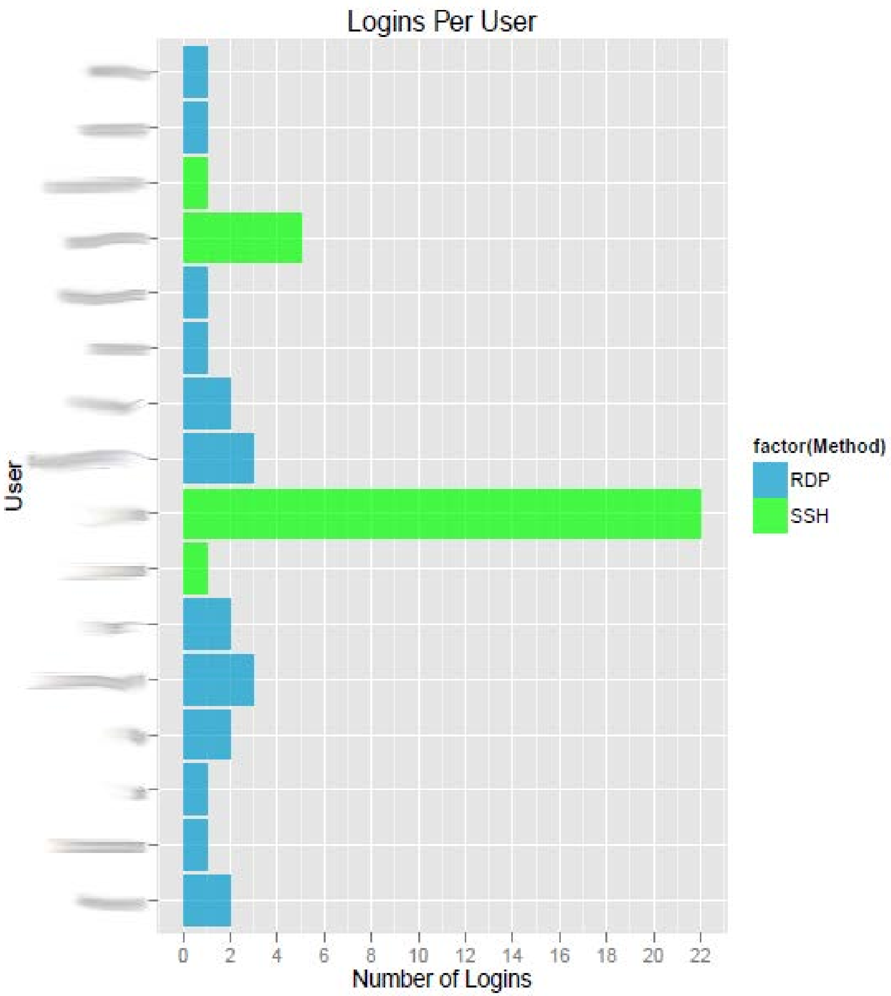

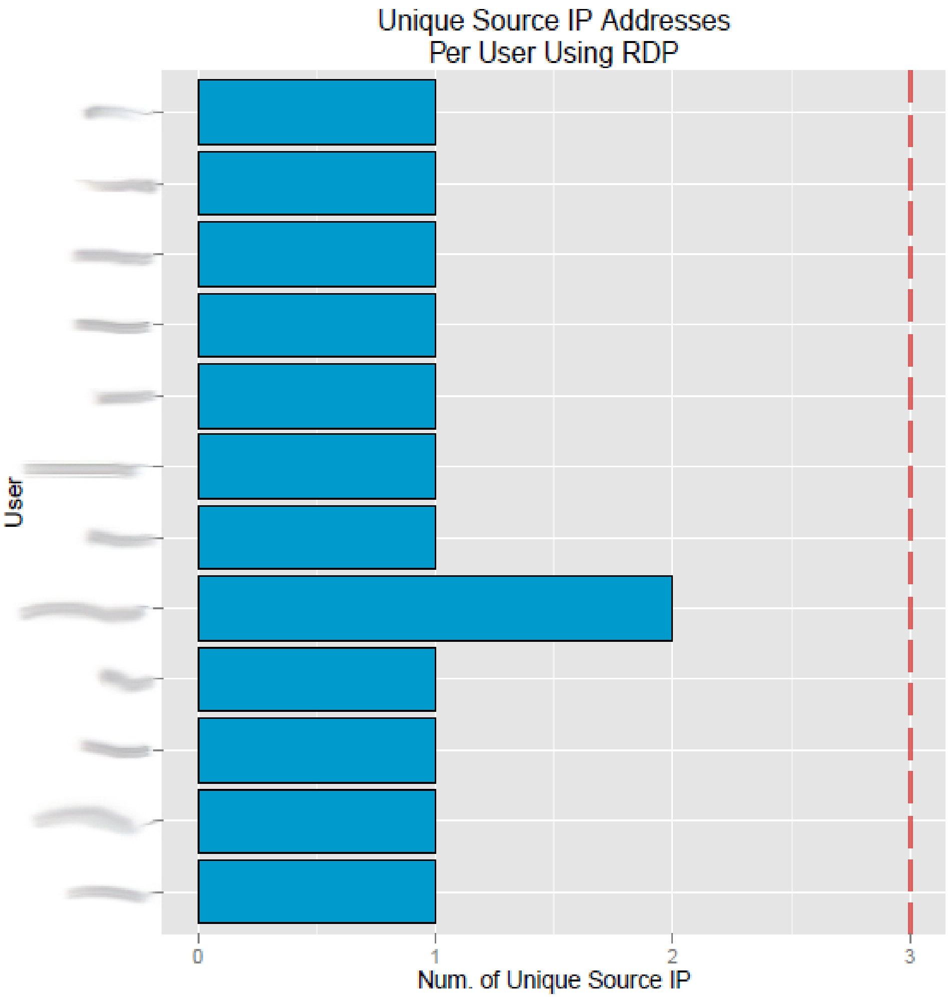

Alerting – Irregular Login Activity

- Example Daily Login Stats with plots

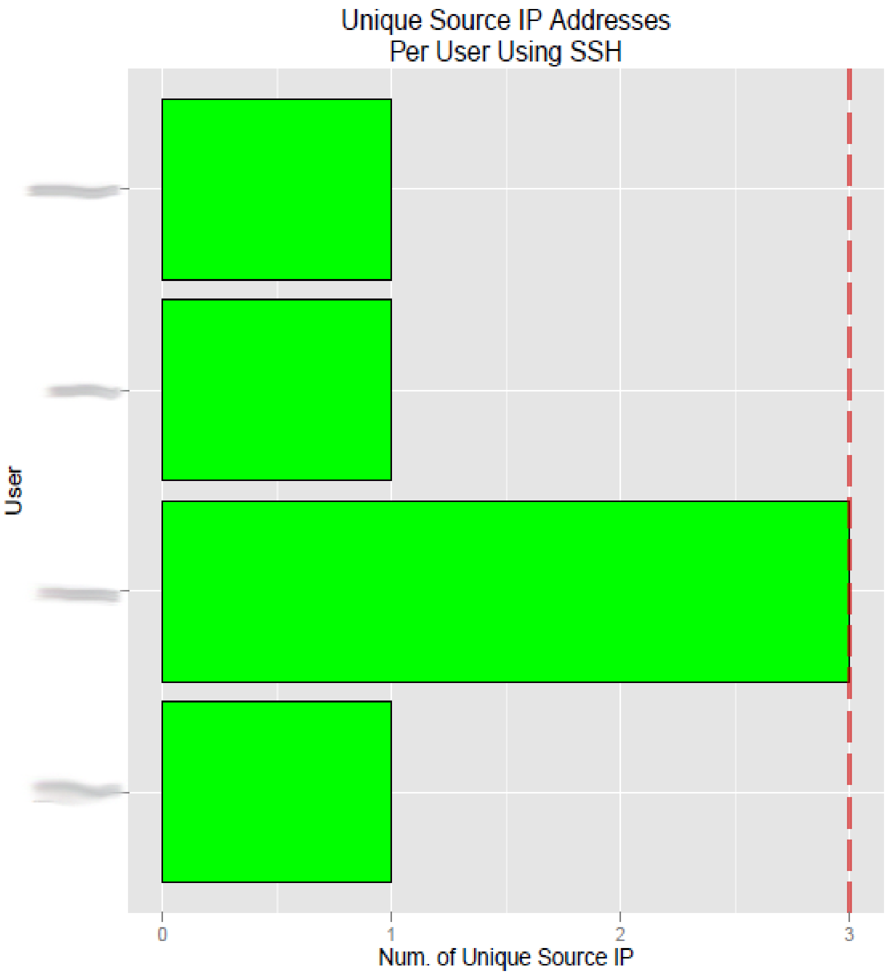

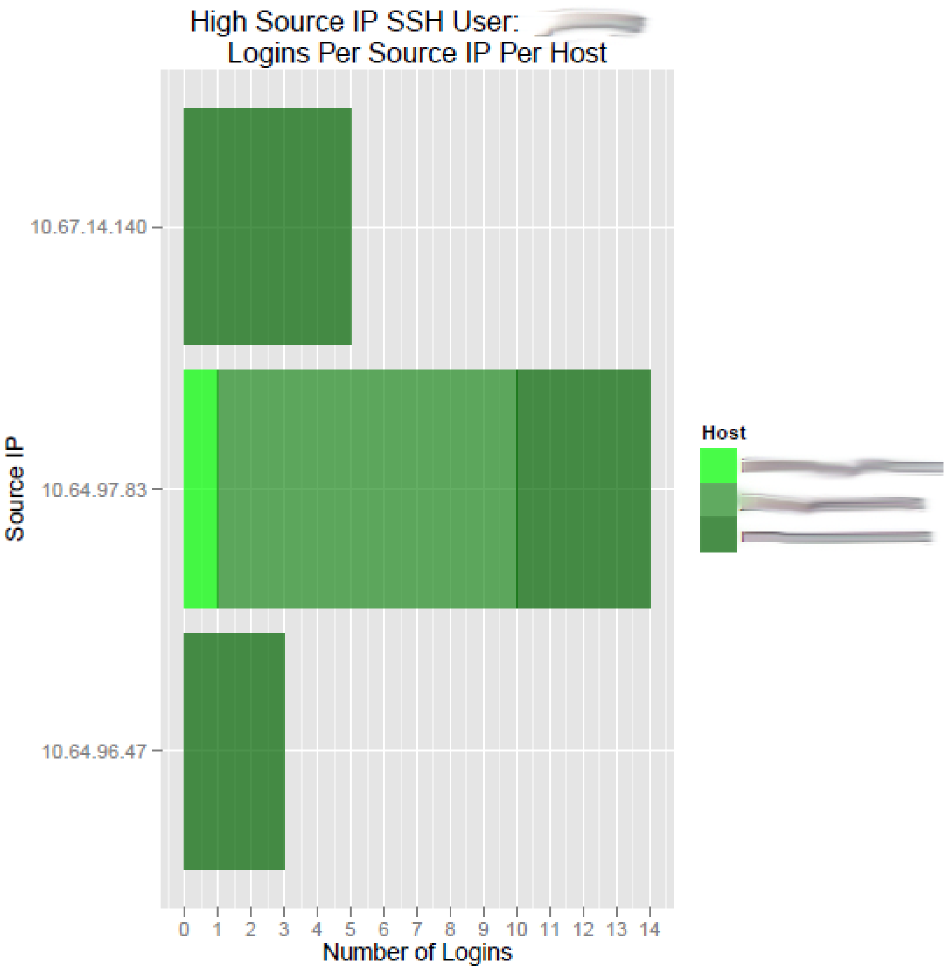

Alerting – Irregular Login Activity

- Example Alerts received via email notification

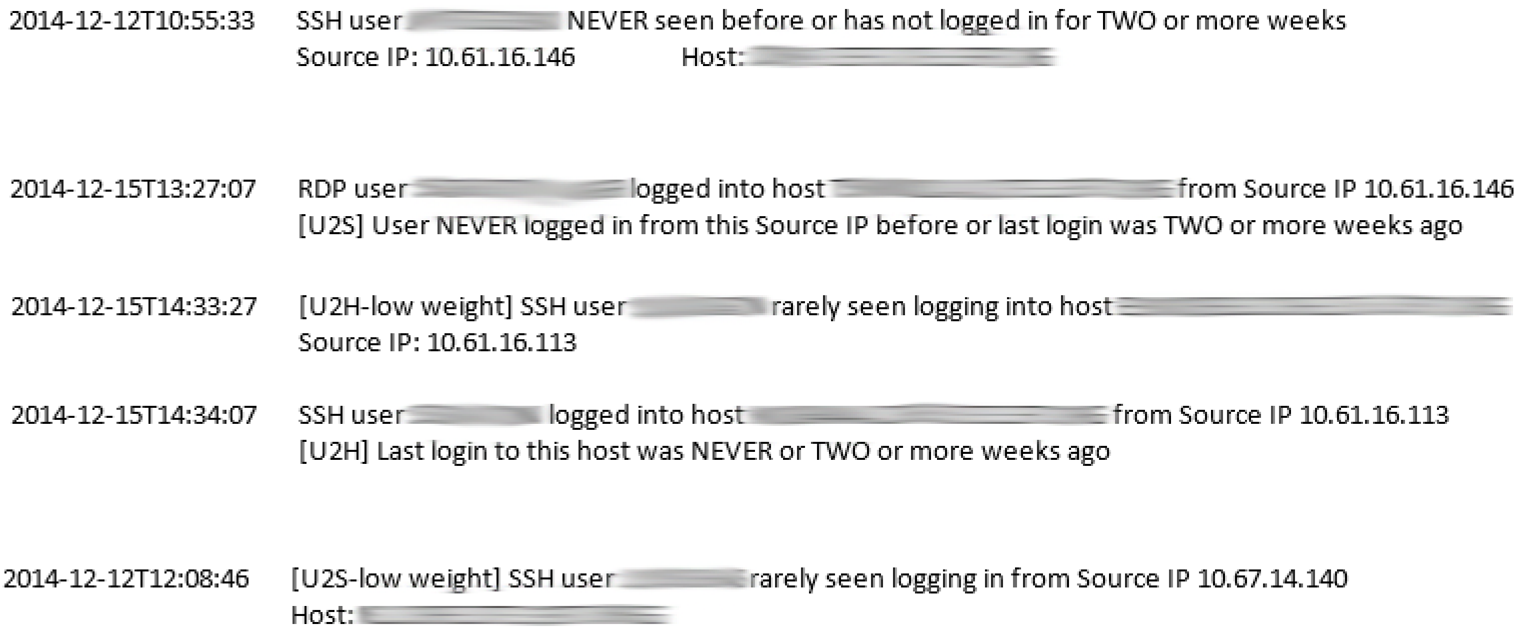

```
2014-12-12T10:55:33     SSH user [redacted] NEVER seen before or has not logged in for TWO or more weeks
                        Source IP: 10.61.16.146          Host: [redacted]


2014-12-15T13:27:07     RDP user [redacted] logged into host [redacted] from Source IP 10.61.16.146
                        [U2S] User NEVER logged in from this Source IP before or last login was TWO or more weeks ago

2014-12-15T14:33:27     [U2H-low weight] SSH user [redacted] rarely seen logging into host [redacted]
                        Source IP: 10.61.16.113

2014-12-15T14:34:07     SSH user [redacted] logged into host [redacted] from Source IP 10.61.16.113
                        [U2H] Last login to this host was NEVER or TWO or more weeks ago


2014-12-12T12:08:46     [U2S-low weight] SSH user [redacted] rarely seen logging in from Source IP 10.67.14.140
                        Host: [redacted]
```

CERT

Questions and Discussion