TREND MICRO
Securing Your Journey to the Cloud

Securing Hadoop with OSSEC

Vic Hargrave | vichargrave@gmail.com | @vichargrave

$ whoami

- Community Manager for the OSSEC Project
- Software Architect for Trend Micro Data Analytics Group
- Blogger for Trend Micro Security Intelligence and Simply Security
- Twitter: @vichargrave
- LinkedIn: www.linkedin.com/in/vichargrave
- Email: ossec@vichargrave.com

Outline

- Hadoop Security
- OSSEC in a Nutshell
- OSSEC for Hadoop
- Security Event Analysis
- Summing Up

Hadoop Security

Hadoop Security

Kerberos
User Authorization

HDFS encryption

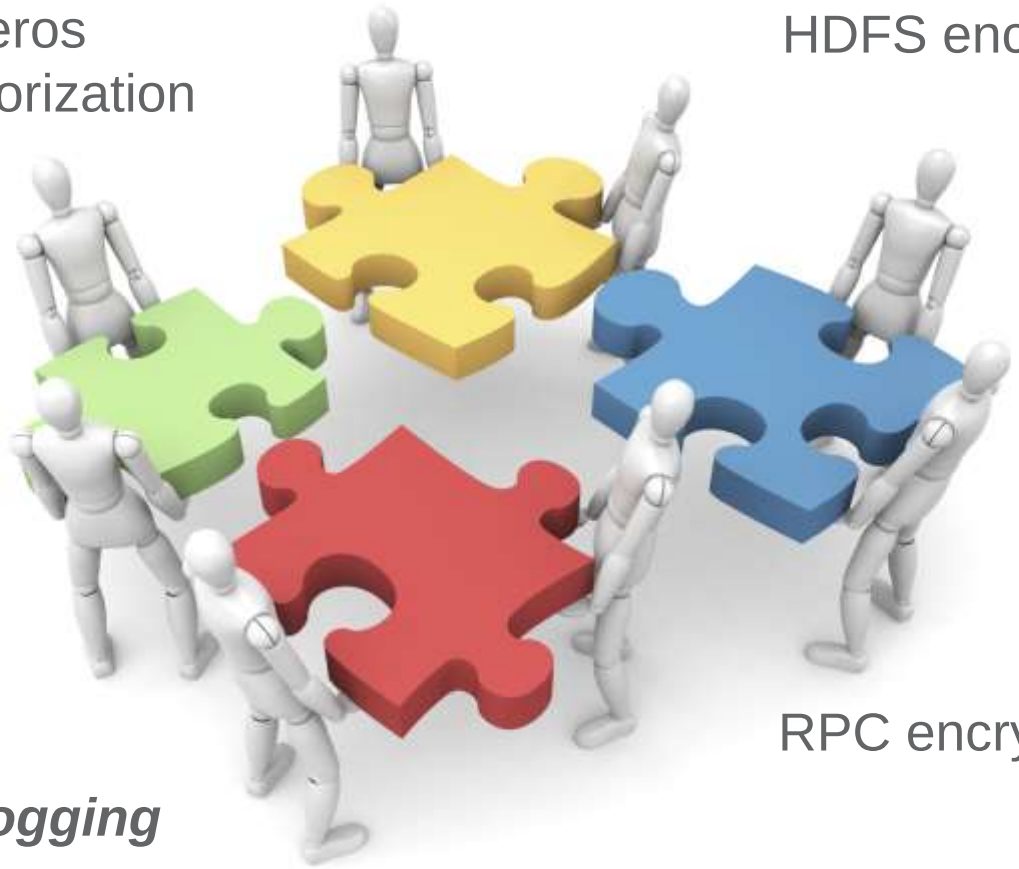

RPC encryption

Extensive Logging

Missing Pieces

Firewall protection

Network intrusion detection

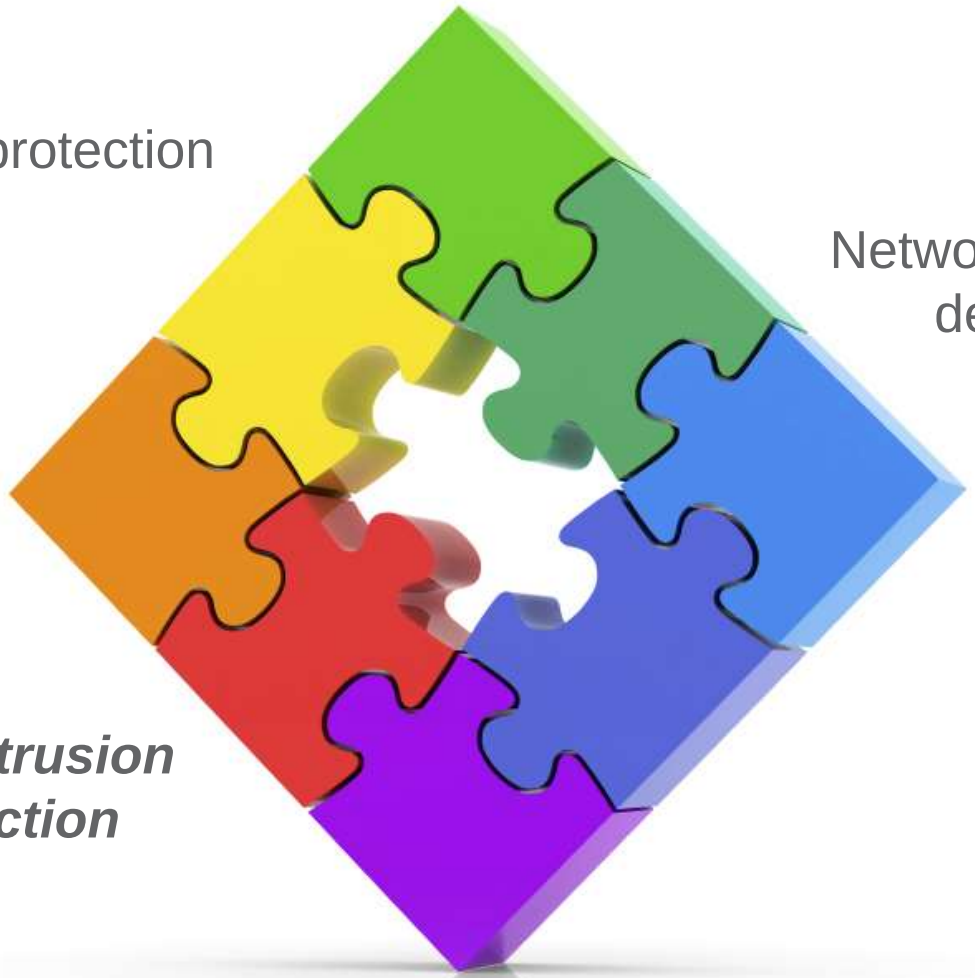

Host intrusion detection

OSSEC

Why Use Host Intrusion Detection?

Dilbert is a copyright of Scott Adams, Inc. – http://www.dilbert.com

To get visibility into your system's security events

OSSEC

Security Events To Look For

Root logins to nodes

Kerberos ticket granting

Failed HDFS operations

HBase REST requests

HBase logins

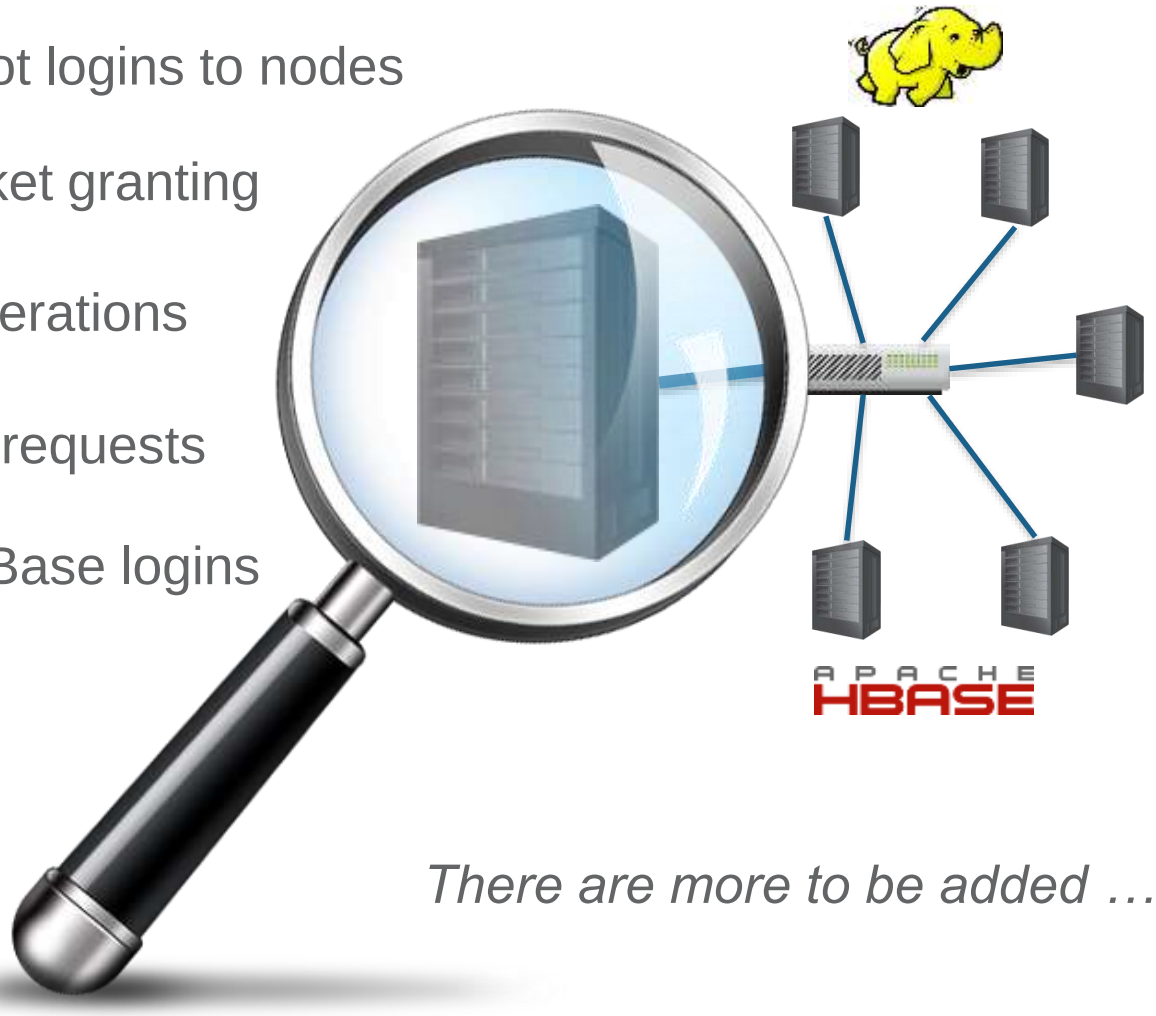

There are more to be added …

OSSEC

OSSEC in a Nutshell

What is OSSEC ?

- **O**pen **S**ource **SEC**urity
- Open Source Host-based Intrusion Detection System
- Monitors important system files and logs for signs of intrusion – file changes, log entries, etc.
- Provides protection for Windows, Linux, Mac OS, Solaris and many other *nix systems
- http://www.ossec.net
- Founded by Daniel Cid
- Managed by JB Cheng and Vic Hargrave
- Sponsored by Trend Micro

OSSEC Capabilities

- Log monitoring and analysis
- File integrity checking (Unix and Windows)
- Registry integrity checking (Windows)
- Host-based anomaly detection (for Unix – rootkit detection)
- Active Response

OSSEC In Action

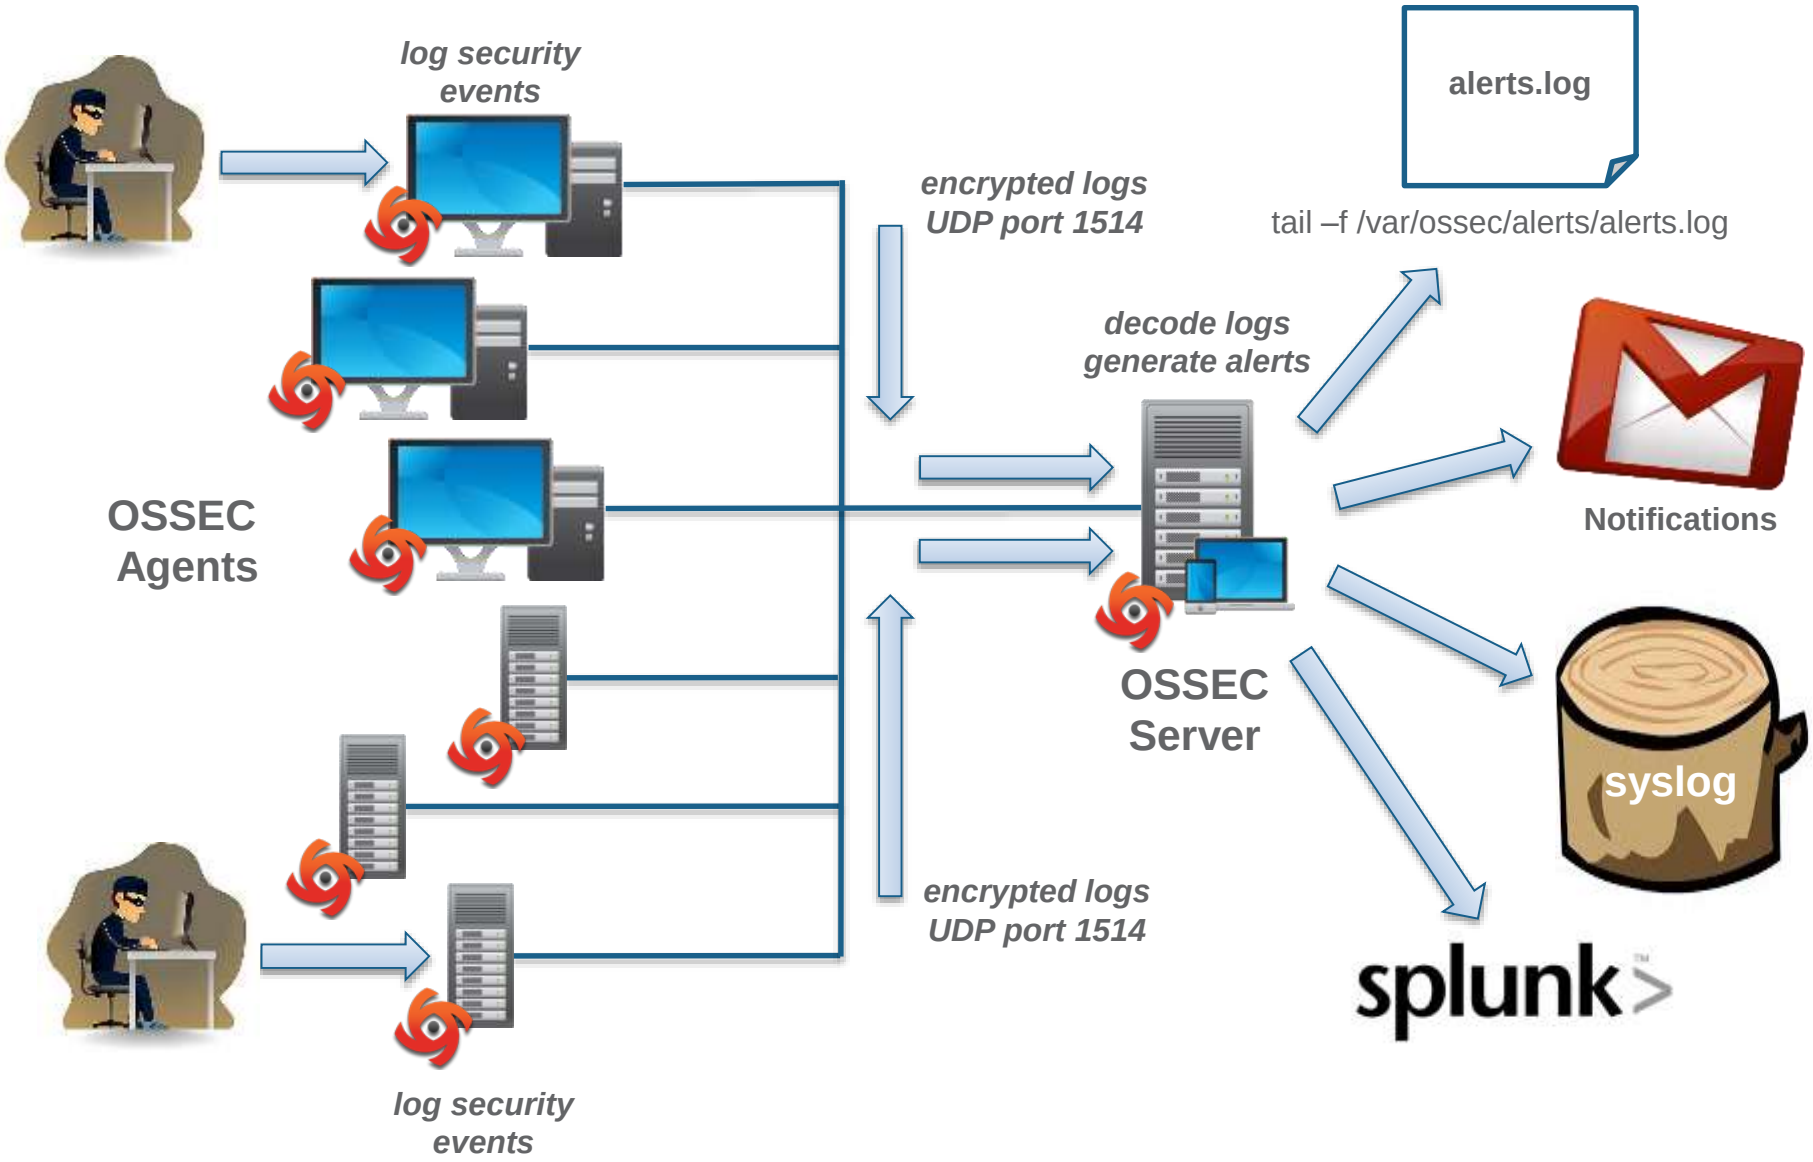

OSSEC

OSSEC Downloads

- Source packages and virtual appliance
  - http://www.ossec.net/?page_id=19
- RPM packages
  - http://www.atomicorp.com/channels/atomic/
- DEB packages
  - (Unofficial) https://launchpad.net/~nicolas-zin/+archive/ossec-ubuntu
  - Official packages coming soon…

Installing OSSEC

1. Install server and agents
   - Default location */var/ossec*
2. Register agent IPs and get agent keys* from server
   - */var/ossec/bin/manage_agents*
3. Connect each agent to server
   - Install key on agent – */var/ossec/bin/manage_agents*
   - Restart agent – */var/ossec/bin/ossec-control restart*
4. Restart server
   - */var/ossec/bin/ossec-control restart*
5. Check agents are connected
   - */var/ossec/bin/agent_control –l*

*Used for agent authentication and log transfer encryption

Configuring OSSEC

1. Configure */var/ossec/etc/ossec.conf* on the OSSEC agents
   - Add files to check for changes
   - Add logs to monitor and parse
2. Add decoders to */var/ossec/etc/local_decoders.xml* to parse logs and decode fields
3. Add rules to */var/ossec/rules/local_rules.xml* to generate alerts according to decoded fields
4. Test decoders and rules, repeating steps 1 – 4
   - */var/ossec/bin/ossec-logtest*
5. Restart agents and server

OSSEC for Hadoop

Configure File Integrity Checking

- We want to know if and when Hadoop and HBase configuration and JAR files are changed
- Out of the box OSSEC checks all directories and files in /etc, /usr/bin and /usr/sbin so you are all set to go if your Hadoop and HBase config files are there
- Otherwise add the config locations to ossec.conf on each node:

```
<syscheck>
  ...
  <directories check_all="yes">/etc,/usr/bin,/usr/sbin</directories>
  <directories check_all="yes">/hadoop/conf</directories>
  ...
</syscheck>
```

Configure Real-Time Integrity Checking

```
<syscheck>
  ...
  <directories realtime="yes" check_all="yes">/etc,/usr/bin,/usr/sbin
  </directories>
  <directories realtime="yes" check_all="yes">/hadoop/conf
  </directories>

  <alert_new_files>yes</alert_new_files>
  ...
</syscheck>
```

Select Hadoop Logs to Monitor

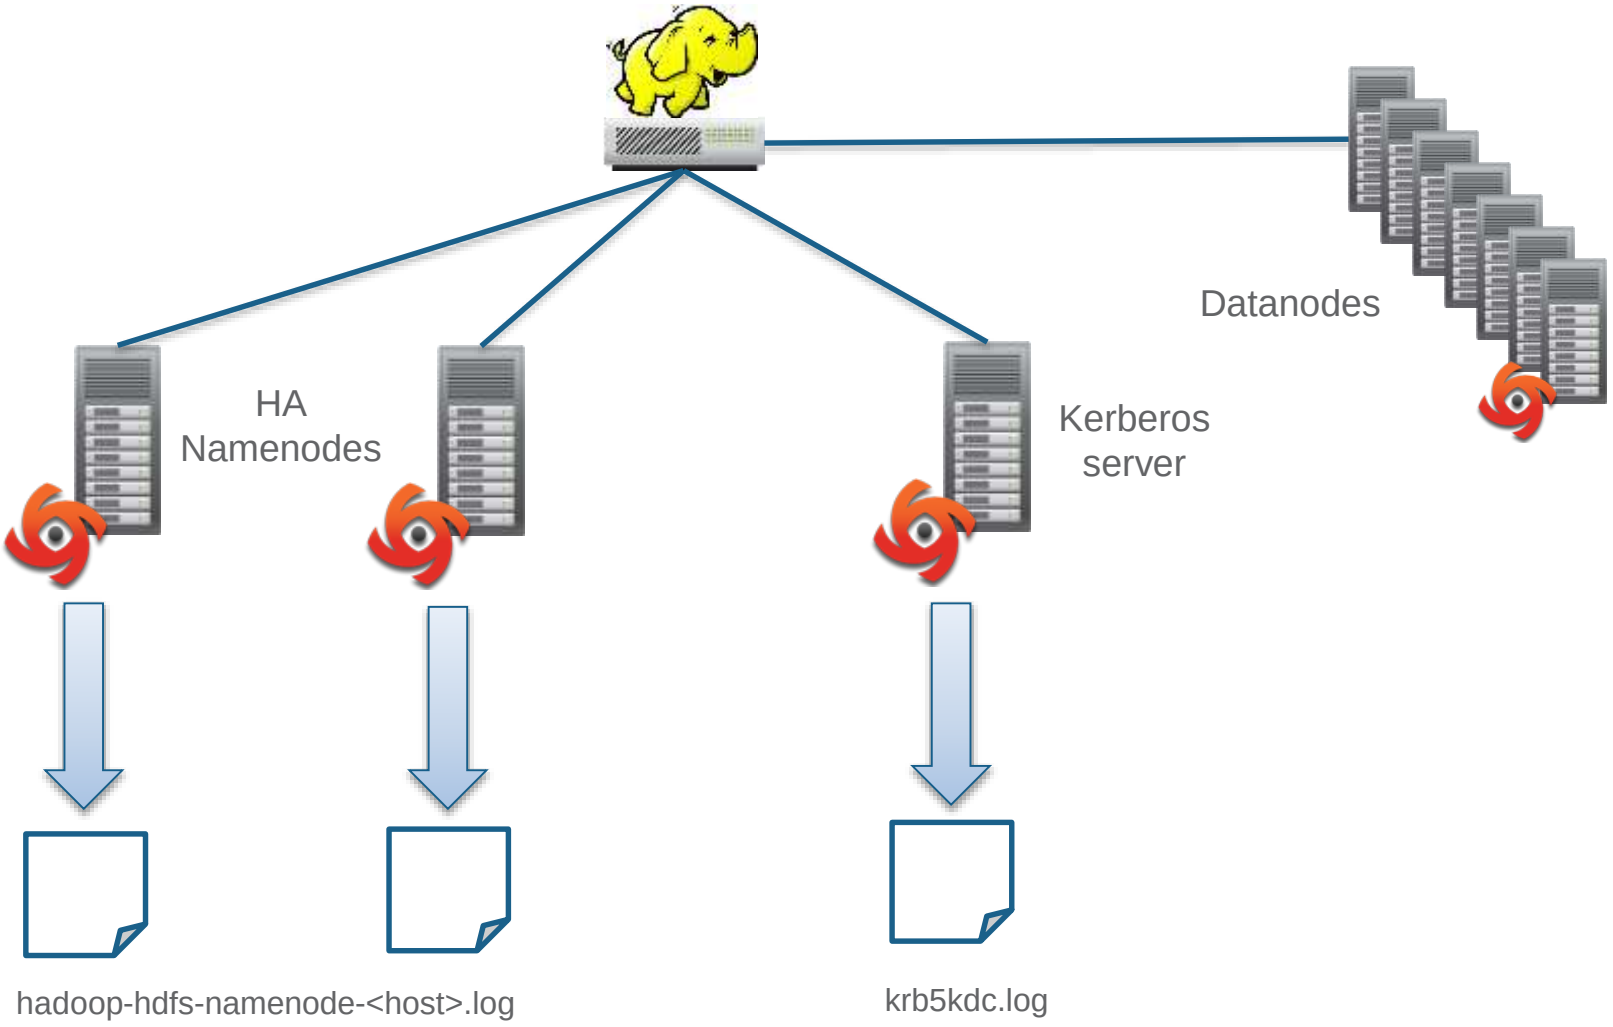

OSSEC

Add Namenode Log to ossec.conf

```
<ossec_config>
  ...
  <localfile>
    <log_format>syslog</log_format>
    <location>/var/log/hadoop-hdfs/hadoop-hdfs-namenode-HOST.log
    </location>
  </localfile>
  ...
</ossec_config>
```

Add Namenode Decoder to local_decoder.xml

```
<decoder name="hadoop">
  <prematch>org.apache.hadoop</prematch>
</decoder>

<decoder name="hdfs-auth-fail">
  <parent>hadoop</parent>
  <prematch>org.apache.hadoop.security.UserGroupInformation: </prematch>
  <regex>\S+ (\S+) as:(\S+) \S+ \S+ (\S+ \w+): \.+</regex>
  <order>extra_data,user,status</order>
</decoder>
```

Add Namenode Rule to local_rules.xml

```
<group name="hadoop">
  ...
  <rule id="700000" level="0">
    <decoded_as>hadoop</decoded_as>
    <description>Hadoop alert rules</description>
  </rule>

  <rule id="700002" level="10">
    <if_sid>700000</if_sid>
    <match>PriviledgedActionException</match>
    <description>HDFS user permission denied</description>
  </rule>
  ...
</group>
```

Alert Looks Like This

```
$ tail –f /var/log/alerts/alerts.log

** Alert 1379976873.84037: mail  - hadoop
2013 Sep 23 15:54:33 (HOST) XX->/var/log/hadoop-hdfs/hadoop-hdfs-namenode-HOST.log
Rule: 700002 (level 10) -> 'HDFS user permission denied'
User: vic@XX.ORG
2013-09-23 15:54:31,825 ERROR org.apache.hadoop.security.UserGroupInformation:
PriviledgedActionException as:vic@XX.ORG (auth:KERBEROS)
cause:org.apache.hadoop.security.AccessControlException: Permission denied: user=vic, access=WRITE,
inode="/user/user":user:supergroup:drwxr-xr-x
```

XX – anonymous IP or domain name

Select HBase Logs to Monitor

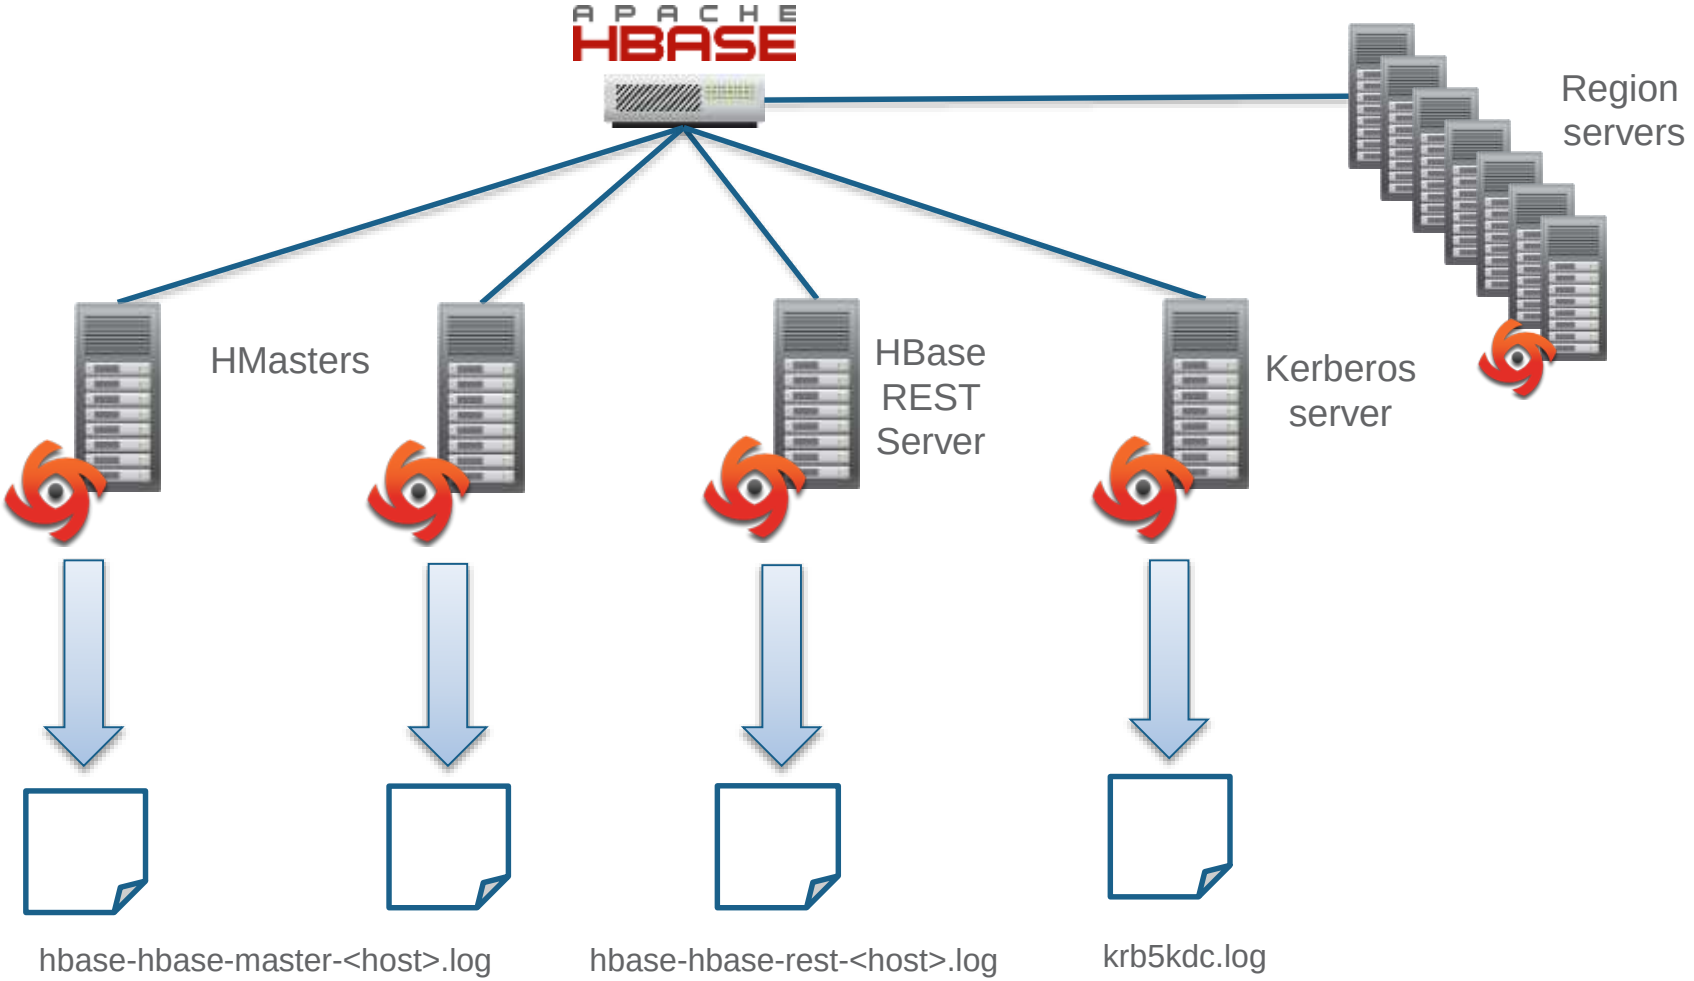

OSSEC

Add HMaster Log to ossec.conf

```
<ossec_config>
  ...
  <localfile>
    <log_format>syslog</log_format>
    <location>/var/log/hbase/hbase-hbase-master-HOST.log</location>
  </localfile>
  ...
</ossec_config>
```

Add HMaster Decoder to local_decoder.xml

```
<!-- NOTE: Uses the same parent rule "hadoop" as specified before -->

<decoder name="hbase-auth-success">
  <parent>hadoop</parent>
  <prematch>org.apache.hadoop.hbase.ipc.SecureServer: </prematch>
  <regex>\.+: (\.+) org\.+-(\S+) \.+</regex>
  <order>status,user</order>
</decoder>
```

Add HMaster Rule to local_rules.xml

```
<group name="hadoop">
  ...
  <rule id="700001" level="3">
    <if_sid>700000</if_sid>
    <match>Successfully authorized</match>
    <description>HBase user authorized</description>
  </rule>
  ...
</group>
```

Alert Looks Like This

$ tail –f /var/log/alerts/alerts.log

** Alert 1379976182.79643: - hadoop
2013 Sep 23 15:43:02 (HOST) XX->/var/log/hbase/hbase-hbase-master-HOST.log
Rule: 700001 (level 3) -> 'HBase user authorized'
User: vic@XX.ORG
2013-09-23 15:43:02,059 DEBUG org.apache.hadoop.hbase.ipc.SecureServer: Successfully authorized org.apache.hadoop.hbase.ipc.HMasterInterface-vic@XX.ORG (auth:SIMPLE)

XX – anonymous IP or domain name

Security Event Analysis

Security Event Capture

- *tail –f /var/ossec/logs/alerts/alerts.log*
  - OK for a look-see but not practical for most systems
- Send alerts to ancillary system via syslog
- Splunk for OSSEC
  - Free and open source application built on top of Splunk
  - Provides agent management, alert search and security trend dashboards
  - Receives OSSEC alerts via syslog or reads the alerts.log file directly when Splunk and OSSEC server are deployed on the same system

Configure Syslog Output

```
<ossec_config>
  ...

  <!-- Send syslog output to remote (syslog) server -->
  <syslog_output>
    <server>10.0.0.1</server>
    <port>9000</port>
    <format>default</format>
  </syslog_output>

  <!-- Send syslog output to local syslog server -->
  <syslog_output>
    <server>127.0.0.1</server>
    <port>514</port>
    <format>default</format>
  </syslog_output>

  ...
</ossec_config>
```

Enable Syslog Output

```
# /var/ossec/bin/ossec-control enable client-syslog
# /var/ossec/bin/ossec-control start
```

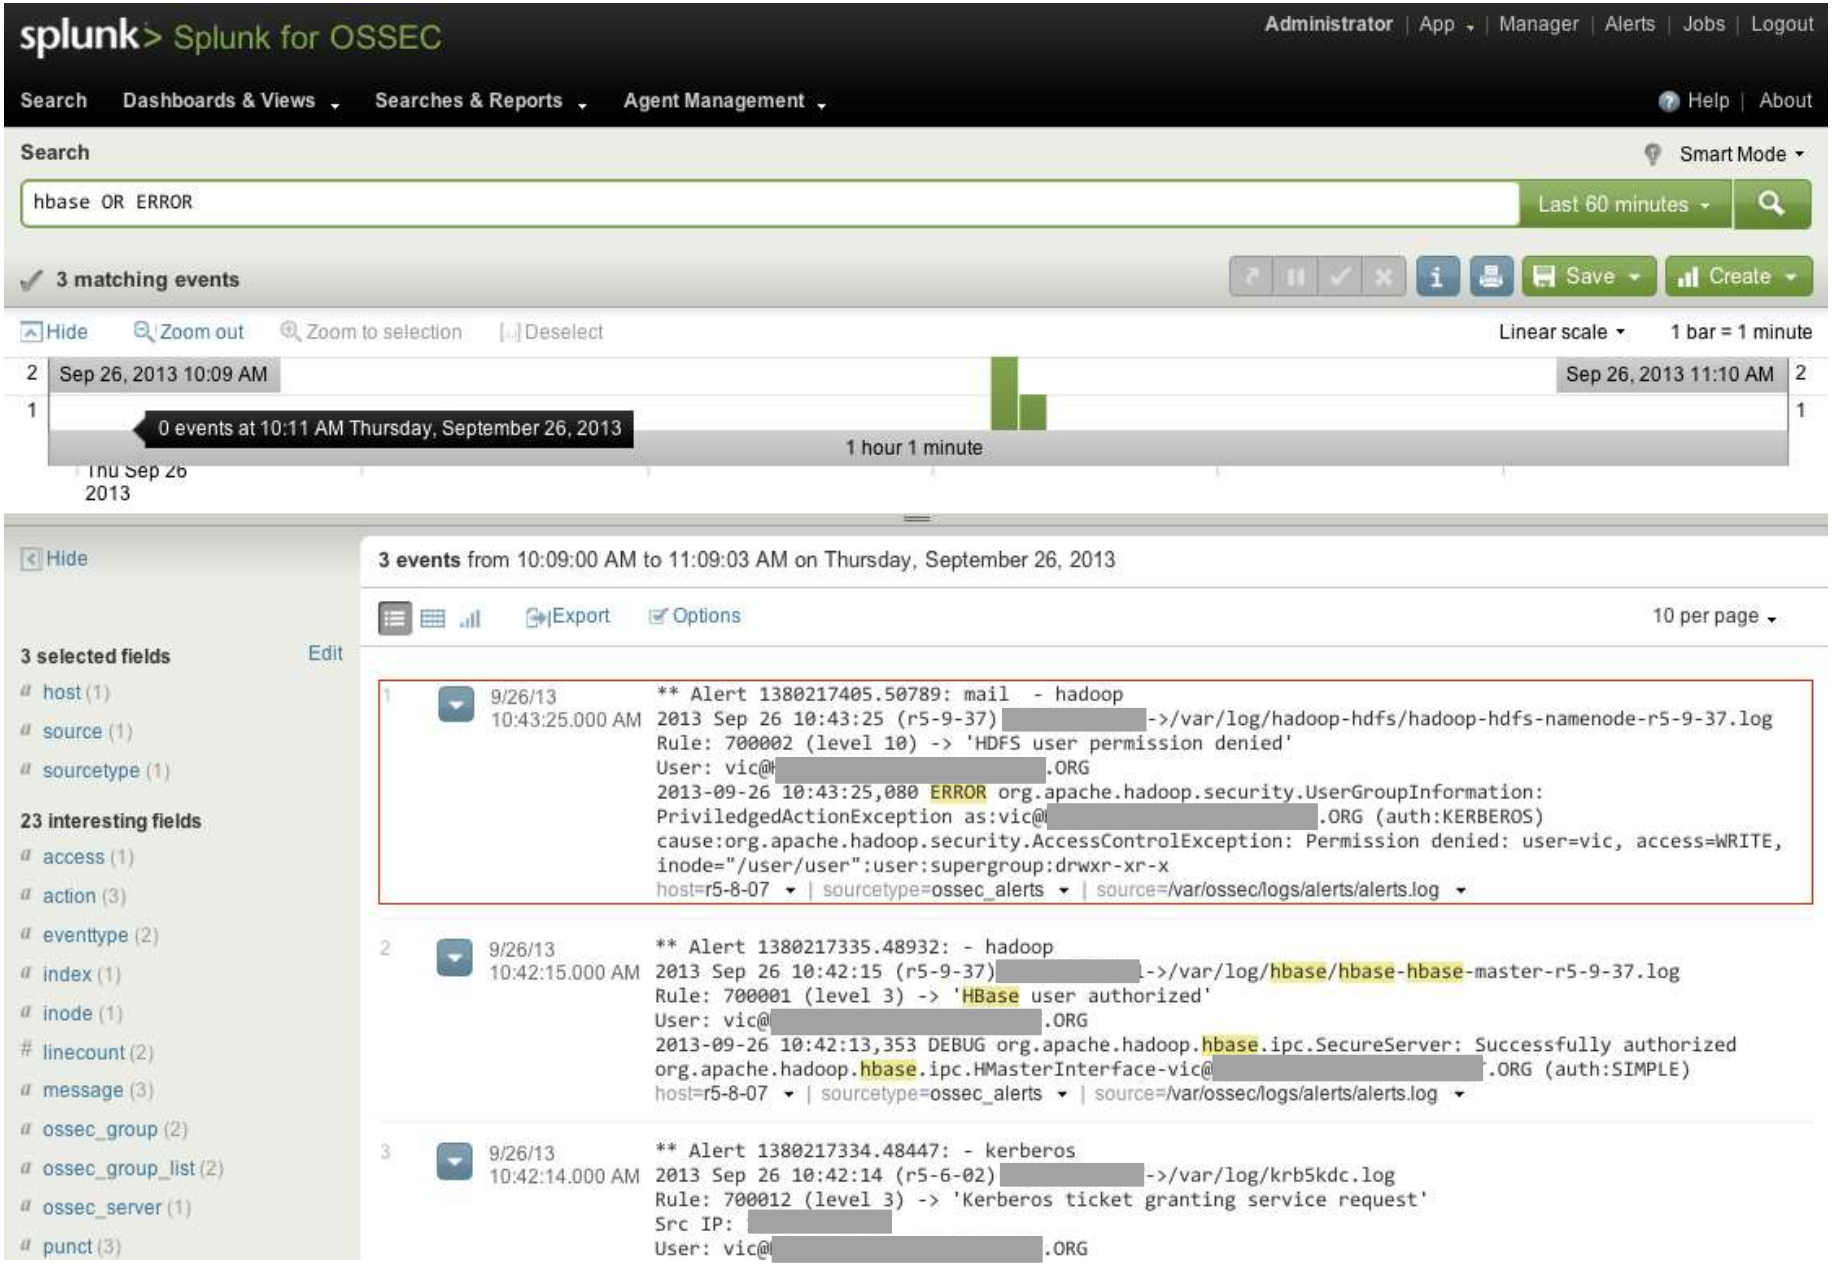
splunk> Splunk for OSSEC
Administrator | App | Manager | Alerts | Jobs | Logout
Search | Dashboards & Views | Searches & Reports | Agent Management
Help | About
Search
Smart Mode
hbase OR ERROR
Last 60 minutes
3 matching events
Save
Create
Hide
Zoom out
Zoom to selection
Deselect
Linear scale
1 bar = 1 minute
Sep 26, 2013 10:09 AM
Sep 26, 2013 11:10 AM
0 events at 10:11 AM Thursday, September 26, 2013
1 hour 1 minute
Thu Sep 26 2013
Hide
3 events from 10:09:00 AM to 11:09:03 AM on Thursday, September 26, 2013
Export
Options
10 per page
3 selected fields
Edit
host (1)
source (1)
sourcetype (1)
23 interesting fields
access (1)
action (3)
eventtype (2)
index (1)
inode (1)
linecount (2)
message (3)
ossec_group (2)
ossec_group_list (2)
ossec_server (1)
punct (3)
1
9/26/13 10:43:25.000 AM
** Alert 1380217405.50789: mail - hadoop
2013 Sep 26 10:43:25 (r5-9-37) [redacted]->/var/log/hadoop-hdfs/hadoop-hdfs-namenode-r5-9-37.log
Rule: 700002 (level 10) -> 'HDFS user permission denied'
User: vic@[redacted].ORG
2013-09-26 10:43:25,080 ERROR org.apache.hadoop.security.UserGroupInformation:
PriviledgedActionException as:vic@[redacted].ORG (auth:KERBEROS)
cause:org.apache.hadoop.security.AccessControlException: Permission denied: user=vic, access=WRITE,
inode="/user/user":user:supergroup:drwxr-xr-x
host=r5-8-07 | sourcetype=ossec_alerts | source=/var/ossec/logs/alerts/alerts.log
2
9/26/13 10:42:15.000 AM
** Alert 1380217335.48932: - hadoop
2013 Sep 26 10:42:15 (r5-9-37) [redacted]->/var/log/hbase/hbase-hbase-master-r5-9-37.log
Rule: 700001 (level 3) -> 'HBase user authorized'
User: vic@[redacted].ORG
2013-09-26 10:42:13,353 DEBUG org.apache.hadoop.hbase.ipc.SecureServer: Successfully authorized
org.apache.hadoop.hbase.ipc.HMasterInterface-vic@[redacted].ORG (auth:SIMPLE)
host=r5-8-07 | sourcetype=ossec_alerts | source=/var/ossec/logs/alerts/alerts.log
3
9/26/13 10:42:14.000 AM
** Alert 1380217334.48447: - kerberos
2013 Sep 26 10:42:14 (r5-6-02) [redacted]->/var/log/krb5kdc.log
Rule: 700012 (level 3) -> 'Kerberos ticket granting service request'
Src IP: [redacted]
User: vic@[redacted].ORG

OSSEC

splunk> Splunk for OSSEC

Administrator | App | Manager | Alerts | Jobs | Logout

Search | Dashboards & Views | Searches & Reports | Agent Management | Help | About

OSSEC Dashboard | Actions

OSSEC Server

All OSSEC Servers

Last 30 days

OSSEC - Top Signatures Over Time

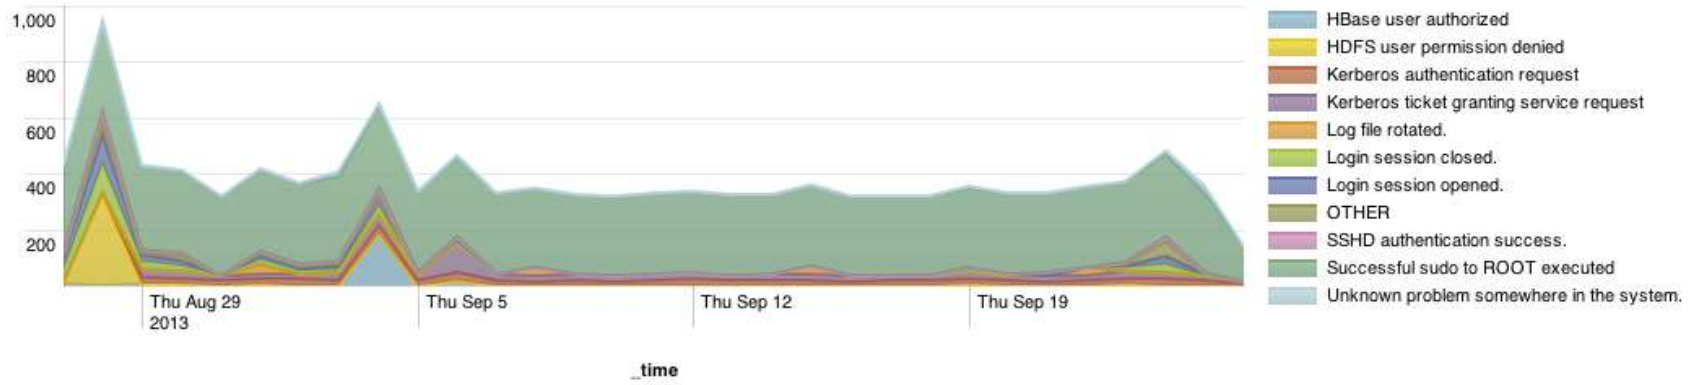

OSSEC - Top Signatures

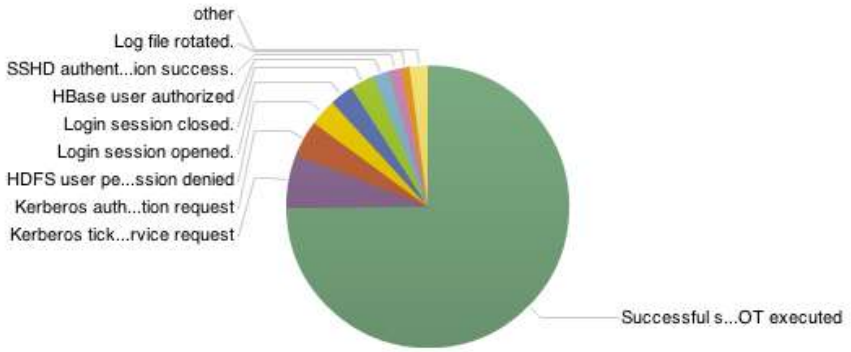

	signature	count
1	Successful sudo to ROOT executed	8888
2	Kerberos ticket granting service request	736
3	Kerberos authentication request	486
4	HDFS user permission denied	366
5	Login session opened.	329
6	Login session closed.	323
7	HBase user authorized	219
8	SSHD authentication success.	166
9	Log file rotated.	122
10	Unknown problem somewhere in the system	75

OSSEC

Summing Up

Benefits of OSSEC

- Provides visibility into your Hadoop and HBase cluster security events
- Tracks system activity – use and abuse
- Free and open source
- Easy to deploy and customize
- Large community of users and developers that share their rules, code and other findings

There's More to Be Done…

- Improve coverage of security events for Hadoop and HBase
- Additional coverage for other Hadoop facilities – MapReduce, Pig, etc.
- Add rules for new vulnerabilities as they are announced

OSSEC Resources

- OSSEC
  - OSSEC downloads – http://www.ossec.net/?page_id=19
  - OSSEC documentation – http://www.ossec.net/doc/
  - OSSEC user group – http://www.ossec.net/?page_id=21#ossec-list
- OSSEC Book
  - Instant OSSEC Host-based Intrusion Detection System by Brad Lhotsky
- Splunk for OSSEC
  - Application site – http://www.splunkbase.com/app/300/
  - Splunk + OSSEC Integration – http://www.ossec.net/?p=402
- OSSEC Log Management with Elasticsearch
  - http://vichargrave.com/ossec-log-management-with-elasticsearch/

Thanks for Attending!

Vic Hargrave | ossec@vichargrave.com | @vichargrave

Source : http://talkofthetail.wordpress.com/2011/08/25/we-have-barn-cats/

OSSEC