DCAP CENTRAL

jSONAR Simplifying Security

tech note 1801

jSonar Cassandra/DataStax Auditing

Forwarding audit data from DataStax Cassandra nodes to - a jSonar machine only requires two types of operations:

- **Editing text configuration files**
- **Restarting services**

These steps could be automated for deployment across a large number of nodes.

Other features that can be confgured are:

- **Certificate-based TLS encryption**
- **Nodes forward to a jump server from which data is then forwarded to jSonar** (in cases where advanced network segmentation is needed)

We show configuration steps for two forwarding methods:

- **Direct Forwarding** via syslog
- **Monitoring** a logfile on the node

The two methods could be mixed and matched on different nodes if necessary. Each node will require configuration.

PER NODE: ENABLE DSE AUDITING

Both Direct Forwarding and Monitoring require enabling audit_logging_options in the dse.yaml file.

Edit audit_logging_options
$ sudo vi /etc/dse/dse.yaml

find the audit_logging_options section and change enabled to true, will look like this:

```
....
....
system_key_directory: /etc/dse/conf
solr_index_stats_options:
  enabled: false
audit_logging_options:
  enabled: true
  cassandra_audit_writer_options:
        mode: sync
        write_consistency: QUORUM
        dropped_event_log: /var/log/cassandra/
dropped_audit_events.log
        day_partition_millis: 3600000
cql_slow_log_options:
  enabled: true
....
....
```

We will restart the dse service after the configuration of either Direct Forwarding or Monitoring.

PER NODE: DIRECT FORWARDING

This method forwards audit data directly to jSonar. We also disable writing to the audit.log file under the assumption that the client does cannot persist the audit data for security or storage reasons.
Copy the following XML into /etc/dse/cassandra/logback.xml.Look for the <!--audit log--> section in the file and put it below there.

Replace the syslogHost element (highlighted) with the IP address of the jSonar machine.

```
$ sudo vi /etc/dse/cassandra/logback.xml
  <!--audit log-->
....
....

  <appender name="SYSLOG" class="ch.qos.logback.
classic.net.SyslogAppender">
        <syslogHost>18.216.236.92</syslogHost>
        <port>10551</port>
        <facility>AUDIT</facility>
        <throwableExcluded>true</throwableExcluded>
        <suffixPattern>%5level [%thread] %date{ISO8601}
%X{service} %F:%L - %msg%n</suffixPattern>
  </appender>
  <logger name="SLF4JAuditWriter" level="INFO"
additivity="false">
        <appender-ref ref="SYSLOG"/>
  </logger>
....
....
```

Note the port number 10551. The jSonar machine will be listening on this port for audit data from this (and other) nodes.

Disable DSE audit to file
Comment out or delete the following lines from logback.xml. In the example below, we comment them out. The comment syntax to add is highlighted in green.

```
$ sudo vi /etc/dse/cassandra/logback.xml
<!--audit log-->
....
....
 <!--
 <logger name="SLF4JAuditWriter" level="INFO"
additivity="false">
        <appender-ref ref="SYSLOG"/>
 </logger>
 -->
....
....
```

Restart the DSE service
$ sudo service dse restart
--> restarting dse can take a few minutes

PER NODE: MONITORING

This method monitors the default Cassandra audit.log file that is used when audit_logging_options is enabled. There are options to control the size of the audit.log file, rolling policy, etc. These can be configured via the /etc/dse/cassandra/logback.xml file. See the DataStax documentation for how to do this.

Instead of adding a SyslogAppender to the logback.xml as we did for direct forwarding, we use the native rsyslog application (already installed on the node) to monitor and forward the contents of the audit.log file.

Create a cassandra.conf rsyslog configuration file
Copy the following contents into /etc/rsyslog.d/cassandra.conf.
Replace the ***target*** field ***(highlighted)*** with the IP address of the jSonar machine.

```
$ sudo vi /etc/rsyslog.d/cassandra.conf
---------- start snip here ---------
module(load="imfile")
input(type="imfile"
        file="/var/log/cassandra/audit/audit.log"
        tag="cassandra_audit"
        ruleset="10550_cassandra_audit")
ruleset(name="10550_cassandra_audit") {
        action(type="omfwd"
        target="18.217.214.119"
        port="10550"
        protocol="tcp"
        ziplevel="9"
          compression.mode="single"
        keepalive="on")
        stop
}
---------- end snip here ---------
```

Note the port number 10552. The jSonar machine will be listening on this port for audit data from this (and other) nodes.

Restart the DSE service
$ sudo service dse restart
--> restarting dse can take a few minutes

Restart the rsyslog service
$ sudo service rsyslog restart

JSONAR MACHINE: ENABLE THE SONARGATEWAY MAPPING

Regardless of whether you chose Direct Forwarding or Monitoring you need to configure the jSonar machine to listen for the incoming audit data.

Listen for Cassandra connections

```
$ sudo vi /etc/rsyslog.d/sonargateway.conf
# uncomment this cassandra line
....
$IncludeConfig /etc/rsyslog.d/sonar/gateway/
rulesets/cassandra.conf
....
```

Restart the rsyslog service
$ sudo systemctl restart rsyslog

Assuming you have restarted dse on all your nodes, you should see connections since the DSE OpsCenter tends to generate a steady level of audit events. We can use the unix utility lsof to confirm that we are ready to receive connections from the DataStax Cassandra nodes.

$ sudo yum install -y lsof

Direct Forwarding connections
$ sudo lsof -i udp | grep rsyslog

```
rsyslogd 11251   root   9u  IPv4 176422  0t0  UDP *:10551
rsyslogd 11251   root   10u  IPv6 176423 0t0  UDP *:10551
```

This tells us that the jSonar machine is ready to receive via Direct Forwarding from the nodes.

tech note 1801

Monitoring Connections

$ sudo lsof -i :10550

```
COMMAND     PID   USER   FD   TYPE   DEVICE  SIZE/OFF  NODE   NAME
rsyslogd  22550   root   7u   IPv4    43425       0t0   TCP    *:10550 (LISTEN)
rsyslogd  22550   root   8u   IPv6    43426       0t0   TCP    *:10550 (LISTEN)
rsyslogd  22550   root  16u   IPv4    51017       0t0   TCP
ip-172-31-43-83.us-east-2.compute.internal:10550->ec2-13-59-120-223.us-east-2.compute.amazonaws.com:38508 (ESTABLISHED)
rsyslogd  22550   root  18u   IPv4    39717       0t0   TCP
ip-172-31-43-83.us-east-2.compute.internal:10550->ec2-18-216-167-133.us-east-2.compute.amazonaws.com:21249 (ESTABLISHED)
rsyslogd  22550   root  20u   IPv4    35311       0t0   TCP
ip-172-31-43-83.us-east-2.compute.internal:10550->ec2-18-222-37-75.us-east-2.compute.amazonaws.com:42016 (ESTABLISHED)
```

Notice that **ip-172-31-43-83** is the IP address of the SonarG machine you are running the command and that (**ec2-13-59-120-223, ec2-18-216-167-133, ec2-18-222-37-75**) are the IP addresses of the cassandra nodes. In this case we were working with the default three node Datastax cluster template and connections have already been made.

Confirm that audit data is flowing

In the default DSE installation on AWS there's a constant low level amount of audit data flowing because of the OpsCenter.

Enable INFO level SonarGateway logging

$ sudo vi /etc/sonar/gateway/gateway-logging.conf

```
....
....
* WARNING:
  ENABLED          =  true
* INFO:
  ENABLED          =  true
....
....

# restart the rsyslog service
$ sudo systemctl restart rsyslog

# tail the sonargateway log file
$ sudo tail -f /var/log/sonar/gateway/sonargateway.log

### assuming steady OpsCenter activity, you should see something like:
....
....
2018-03-16 21:26:42,761 INFO Insert 1000 bsons to sonarg.instance
2018-03-16 21:26:43,761 INFO Insert 1000 bsons to sonarg.instance
2018-03-16 21:26:44,761 INFO Insert 1000 bsons to sonarg.instance
2018-03-16 21:26:45,761 INFO Insert 1000 bsons to sonarg.instance
....
....
```

At this point the audit data will be available in the **Discover** or **Analyze** browsers from the DCAP Central main page.