Applying Log4J mitigation to Elasticsearch within the Signals Data Factory

1. From the Admin machine for a deployment, run the following command:

   `kubectl get statefulset`

2. Look for "elasticsearch-master" or "es-master" and copy the full name. In most cases it should be `sdf-core-infrastructure-elasticsearch-master`

   Example output:

   ```
   > kubectl get statefulset
   NAME                                          READY   AGE
   sdf-core-infrastructure-elasticsearch-master   1/1     16h
   ```

3. Run the following command:

   `kubectl get statefulset {{name from step #2}} -o=jsonpath="{.spec.template.spec.containers[0].env}"`

   Where {{name from step #2}} is the name you copied from above, for example,

   `kubectl get statefulset sdf-core-infrastructure-elasticsearch-master  -o=jsonpath="{.spec.template.spec.containers[0].env}"`

   The output looks similar to the following:

   ```
   [{"name":"node.name","valueFrom":{"fieldRef":{"apiVersion":"v1","fieldPath":"metadata.name"}}},{"name":"cluster.initial_master_nodes","value":"sdf-core-infrastructure-elasticsearch-master-0,"},{"name":"discovery.seed_hosts","value":"sdf-core-infrastructure-elasticsearch-master-headless"},{"name":"cluster.name","value":"sdf-core-infrastructure-elasticsearch"},{"name":"network.host","value":"0.0.0.0"},{"name":"ES_JAVA_OPTS","value":"-Xmx11g -Xms11g"},{"name":"node.data","value":"true"},{"name":"node.ingest","value":"true"},{"name":"node.master","value":"true"},{"name":"node.ml","value":"false"},{"name":"node.remote_cluster_client","value":"false"}]>
   ```

   Look for "ES_JAVA_OPTS" and make note of the value that's next to it. In this example, it is `" -Xmx11g -Xms11g "`

4. Run the following command:

```
kubectl set env statefulset {{name from step #2}} "ES_JAVA_OPTS={{value from step #3}}
-Dlog4j2.formatMsgNoLookups=true"
```

From the output above, the example command to run would be:

```
kubectl set env statefulset sdf-core-infrastructure-elasticsearch-master
"ES_JAVA_OPTS= -Xmx11g -Xms11g  -Dlog4j2.formatMsgNoLookups=true"
```

5. Rerun the command from step #3 and look at ES_JAVA_OPTS to see that it took effect. For example, the output should now be:


```
[{"name":"node.name","valueFrom":{"fieldRef":{"apiVersion":"v1","fieldPath":"metadat
a.name"}}},{"name":"cluster.initial_master_nodes","value":"sdf-core-infrastructure-el
asticsearch-master-0,"},{"name":"discovery.seed_hosts","value":"sdf-core-infrastructu
re-elasticsearch-master-headless"},{"name":"cluster.name","value":"sdf-core-infrastru
cture-elasticsearch"},{"name":"network.host","value":"0.0.0.0"},{"name":"ES_JAVA_OPT
S","value":"-Xmx11g -Xms11g -Dlog4j2.formatMsgNoLookups=true"},{"name":"node.data","v
alue":"true"},{"name":"node.ingest","value":"true"},{"name":"node.master","value":"tr
ue"},{"name":"node.ml","value":"false"},{"name":"node.remote_cluster_client","valu
e":"false"}]>
```


Note that ES_JAVA_OPTS now contains the log4j mitigation:

```
-Dlog4j2.formatMsgNoLookups=true
```