ERICSSON

ENM Monitoring Troubleshooting Guide

Check List

1/159 01-AOM 901 151-2 Uen BK

Contents

1 ENM Monitoring Troubleshooting Guide

This document describes the troubleshooting tasks for the ENM Monitoring applications.

Target Group

System Administrators

2 Connect to a Service

2.1 Connect to a Virtual Machine on a Physical ENM Deployment

Prerequisites

A command window is open and you have `superuser` privileges.

Steps

1. Log on to the ENM MS as **`litp-admin`** user and switch to the **`root`** user.
2. List the contents of the host file to view all connected VMs within the deployment.

```
[root@ms-1 ~]# cat /etc/hosts
...
192.168.99.20 svc-1-pmserv # Created by LITP. Please do not edit
192.168.99.26 svc-1-netex # Created by LITP. Please do not edit
192.168.99.16 svc-1-ebc # Created by LITP. Please do not edit
192.168.99.36 svc-1-mspm # Created by LITP. Please do not edit
192.168.99.28 svc-1-uiserv # Created by LITP. Please do not edit
192.168.99.14 svc-1-supervc # Created by LITP. Please do not edit
..
192.168.99.32 svc-1-mscm # Created by LITP. Please do not edit
..
192.168.99.50 svc-1-jms # Created by LITP. Please do not edit
..
192.168.99.3 logstash # Created by LITP. Please do not edit
..
192.168.99.2 httpd # Created by LITP. Please do not edit
192.168.99.40 sso # Created by LITP. Please do not edit
..
192.168.99.12 svc-1-medrout # Created by LITP. Please do not edit
192.168.99.22 svc-1-cmserv # Created by LITP. Please do not edit
192.168.99.52 svc-1-sec # Created by LITP. Please do not edit
192.168.99.8 openidm # Created by LITP. Please do not edit
```

   The aliases for the parallel VMs take the form of `<SVC host>-<service>`.

   For example: `svc-1-cmserv, svc-2-cmserv`.

   The active-passive VMs take the form of `<service>`.

   For example: `httpd, sso, openidm`.

3. To access the VM, copy the private key of the cloud-user from its secure location to the `MS` or `SVC` node.

```
[root@ms-1 ~]# /root/.ssh/vm_private_key
```

Refer to *VM Security Tasks* in the *ENM System Administrator Guide* to learn more about the `vm_private_key`.

4. Connect by `SSH` to the VM you want.

   To access the `VM`, use the `cloud-user` user ID and include the path to the `VM` private key. For example:

```
[root@ms-1 ~]# ssh -i /root/.ssh/vm_private_key cloud-user@svc-1-cmserv
Last login: Thu Feb 26 10:14:43 2015 from 192.110.0.59
[cloud-user@svc-1-cmserv ~]# sudo su - root
[root@svc-1-cmserv ~]#
```

2.1.1 Connect to each ENM Physical Node

Prerequisites

— The root password was changed during the installation process and must be known by the system administrator. This must be repeated on all newly deployed ENM nodes.

— A command window is open.

Steps

1. Log on to each physical node from the MS

```
[root@ms-1 ~]$ ssh litp-admin@<node_hostname>
litp-admin@<node_hostname>'s password:
Last login: Mon Feb 23 11:25:13 2015 from ms-1
[litp-admin@<node_hostname> ~]$ su - root
Password:
[root@<node_hostname> ~]#
```

Note: Once connected, after the initial deployment, the passwords for both the `litp-admin` and `root` users must be changed.

2.2 Connect to a Virtual Machine on an ENM on Cloud Deployment

Prerequisites

— A command window is open and you have `superuser` privileges.

— You have access to the private key file for authentication, contact your OpenStack administrator

Steps

1. List the virtual machine aliases from the consul service:

   Using the private key for authentication, copy the key to the EMP server. Log on to EMP server and list the consul members to view all connected VMs within the deployment:

```
> scp -i <cloud-user private key> <cloud-user private key> cloud-user@<EMP I →
P Address>:/var/tmp/vm_private_key
> ssh -i <cloud-user private key> cloud-user@<EMP IP Address>
[cloud-user@ostk003-emp-0 ~]$ chmod 700 /var/tmp/vm_private_key
[cloud-user@ostk003-emp-0 ~]$ sudo su -
[root@ostk003-emp-0 ~]# consul members
Node                      Address          Status  Type    Build  Protocol  →
 DC
haproxy                   10.3.2.31:8301   alive   client  0.8.1  2         →
 dc1
opendj-1                  10.3.2.83:8301   alive   client  0.8.1  2         →
 dc1
opendj-2                  10.3.2.84:8301   alive   client  0.8.1  2         →
 dc1
openidm                   10.3.2.85:8301   alive   client  0.8.1  2         →
 dc1
ostk003-accesscontrol-0   10.3.1.251:8301  alive   client  0.8.1  2         →
 dc1
ostk003-accesscontrol-1   10.3.1.252:8301  alive   client  0.8.1  2         →
 dc1
ostk003-elasticsearch-0   10.3.2.15:8301   alive   client  0.8.1  2         →
 dc1
...
ostk003-neo4j-2           10.3.2.77:8301   alive   client  0.8.1  2         →
 dc1
ostk003-nfscommon-0       10.3.0.81:8301   alive   client  0.8.1  2         →
 dc1
ostk003-nfsnrbk-0         10.3.0.83:8301   alive   client  0.8.1  2         →
 dc1
ostk003-nfspm-0           10.3.0.85:8301   alive   client  0.8.1  2         →
 dc1
ostk003-nfspm-1           10.3.0.82:8301   alive   client  0.8.1  2         →
 dc1
...
ostk003-secserv-1         10.3.2.98:8301   alive   client  0.8.1  2         →
 dc1
ostk003-serviceregistry-0 10.3.2.100:8301  alive   server  0.8.1  2         →
 dc1
ostk003-serviceregistry-1 10.3.2.101:8301  alive   server  0.8.1  2         →
 dc1
ostk003-serviceregistry-2 10.3.2.102:8301  alive   server  0.8.1  2         →
 dc1
ostk003-uiserv-0          10.3.2.116:8301  alive   client  0.8.1  2         →
 dc1
ostk003-uiserv-1          10.3.2.117:8301  alive   client  0.8.1  2         →
 dc1
ostk003-vnflaf-services   10.3.1.249:8301  alive   client  0.8.1  2         →
 dc1
...
svc-2-httpd               10.3.2.35:8301   alive   client  0.8.1  2         →
 dc1
svc-2-sps                 10.3.2.111:8301  alive   client  0.8.1  2         →
 dc1
svc-2-sso                 10.3.2.113:8301  alive   client  0.8.1  2         →
 dc1
```

2. SSH to the VM you want.

   To access the VM, use the cloud-user user ID and include the path to the VM private key. The VM can be accessed using either the node identifier or it's IP address. For example:

```
[cloud-user@ostk003-emp-0 ~]$ ssh -i /var/tmp/vm_private_key cloud-user@10.3 →
.2.31
The authenticity of host 'haproxy (10.3.2.31)' can't be established.
RSA key fingerprint is b9:4f:ca:4f:bc:55:00:de:a8:77:e5:08:56:7c:db:98.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added 'haproxy,10.3.2.31' (RSA) to the list of known ho →
sts.
[cloud-user@haproxy ~]$
```

2.3 View Log Files and Dump Locations on a Virtual Machine

The following are details of log files available within each service in ENM.

Logs

All logs are configured to be forwarded to the Central Log Service. As such they are visible in Log Viewer using the ENM Launcher.

JBOSS Logs

All JBOSS logs are stored locally in `/ericsson/3pp/jboss/standalone/log`

3PP & System Logs

As standard, most 3PP and system logs are available locally in `/var/log`

Dumps

All application memory and core dump files are located in `/ericsson/enm/dumps`

3 Restarting a Service

3.1 Restart a Service on a Physical ENM Deployment

Prerequisites

— Root access to MS.

Steps

1. Establish the service instances installed on the ENM deployment using `grep` for a particular service instance:

```
[root@<MS> ~]# /opt/ericsson/enminst/bin/vcs.bsh --groups | grep <service_na →
me>
```

Example

```
[root@ieatlms4352 ~]#  /opt/ericsson/enminst/bin/vcs.bsh --groups | grep msp →
m

svc_cluster Grp_CS_svc_cluster_mspm ieatrcxb2539-1 parallel vm ONLINE OK -
svc_cluster Grp_CS_svc_cluster_mspm ieatrcxb4373    parallel vm ONLINE OK -
svc_cluster Grp_CS_svc_cluster_mspm ieatrcxb4374    parallel vm ONLINE OK -
```

2. Restart the VCS service group:

```
/opt/ericsson/enminst/bin/vcs.bsh --restart -g <service_group> -s <system>
```

Note: The `-s` command restarts only one service at a time. To restart multiple services, repeat the command and modify the system name.

It is not recommended (unless specifically instructed) to restart more than one instance of a service at the same time. Restarting more than one instance of a service at the same time impacts the service availability and also results in some application specific consequences.

Example

```
/opt/ericsson/enminst/bin/vcs.bsh --restart -g Grp_CS_svc_cluster_mspm -s ie →
atrcxb4373

[root@ms-1 bin]# bash vcs.bsh --restart -g Grp_CS_svc_cluster_mspm -s ieatrc →
xb4373
2020-07-23 12:02:04.481 INFO hagrp_offline : Offlining 1 group(s)
2020-07-23 12:02:04.515 INFO hagrp_offline : Offlining Grp_CS_svc_cluster_ms →
pm on ieatrcxb4373
2020-07-23 12:02:04.807 INFO wait_vcs_state : Waiting for Grp_CS_svc_cluster →
_mspm to go OFFLINE on ieatrcxb4373 (timeout=1800)
2020-07-23 12:05:43.185 INFO wait_vcs_state : Group Grp_CS_svc_cluster_mspm  →
now OFFLINE on ieatrcxb4373 (3m:39s)
```

```
2020-07-23 12:05:43.817 INFO hagrp_online : Onlining 1 group(s)
2020-07-23 12:05:43.822 INFO online_services : Onlining Grp_CS_svc_cluster_m →
spm on ieatrcxb4373
2020-07-23 12:05:44.057 INFO wait_vcs_state : Waiting for Grp_CS_svc_cluster →
_mspm to go ONLINE on ieatrcxb4373 (timeout=4500)
2020-07-23 12:09:03.400 INFO wait_vcs_state : Group Grp_CS_svc_cluster_mspm  →
now ONLINE on ieatrcxb4373 (3m:19s)
[root@ms-1 bin]#
```

3. Verify if the service instance is `ONLINE`:

```
/opt/ericsson/enminst/bin/vcs.bsh --groups | grep mspm
```

Example

```
[root@ieatlms4352 ~]#  /opt/ericsson/enminst/bin/vcs.bsh --groups | grep msp →
m

svc_cluster Grp_CS_svc_cluster_mspm ieatrcxb2539-1 parallel vm ONLINE OK -
svc_cluster Grp_CS_svc_cluster_mspm ieatrcxb4373    parallel vm ONLINE OK -
svc_cluster Grp_CS_svc_cluster_mspm ieatrcxb4374    parallel vm ONLINE OK -
```

4. After the service restarted in *Step 2* is `ONLINE`, you can repeat *Step 2* and *Step 3* to restart further instances of the service as per your requirement.

3.2 Restart a Service on an ENM on Cloud Deployment

Prerequisites

— User connected to EMP server.

Steps

1. Establish the service instances installed on the vENM deployment using grep for a particular service instance.

```
#consul members | grep <service name>
```

Example

```
#consul members | grep mscm
```

2. Connect to the VM of the service group by following *section 3.2* and trigger a healthcheck failure of the VM by killing consul.

```
#pkill consul
```

3. Verify if the service instance is ONLINE.

4. After the restarted service is ONLINE, repeat the preceding two steps to restart further instances of the service as per your requirement.

4 Configuring PIB Parameters

> To configure a Platform Integration Bridge (PIB) parameter, it is necessary to determine what environment you are working on and follow the task relevant to your environment.

4.1 Configuring PIB Parameters on a Physical ENM Deployment

To access the PIB (Platform Integration Bridge) script to update parameters for ENM applications, it is necessary to determine what environment you are working on and follow the task relevant to your environment. This task outlines the steps to read and configure PIB parameters on a physical ENM Deployment.

Prerequisites

— A command window is open and you have super user privileges.

— You are connected to the ENM MS as per the Connect to a Virtual Machine on a Physical ENM Deployment on page 2.

Steps

1. Find the hostname for the service instance:

```
grep <service_name> /etc/hosts
```

2. Choose one of the returned hostnames for the next steps.
3. Navigate to the following directory:

```
[root @ms-1 ~]# cd /ericsson/pib-scripts/etc/
```

4. Check a configuration parameter on sample VM:

```
./config.py read --app_server_address=<service VM hostname>:8080 --service_i →
dentifier=<service_identifier_name> --name=<parameter_name>
```

Note: `--service_identifier=<service_identifier_name>` is optional for this command.

Example

To check value of the `SMRS_ERBS_NoOf_BACKUP_FILES` parameter:

```
./config.py read --app_server_address=svc-1-smrsserv:8080 --name=SMRS_ERBS_N
oOf_BACKUP_FILES
```

5. Update a configuration parameter on a deployed VM:

```
./config.py update --app_server_address=<service VM hostname>:8080 --service
_identifier=<service_identifier_name> --name=<parameter_name> --value=<new_v
alue>
```

Note: `--service_identifier=<service_identifier_name>` is optional for this command.

Example

To update the `SMRS_ERBS_NoOf_BACKUP_FILES` value to 4:

```
./config.py update --app_server_address=svc-1-smrsserv:8080 --name=SMRS_ERBS
_NoOf_BACKUP_FILES --value=4
```

Results

You have updated an application parameter using the PIB script.

4.2 Configuring PIB Parameters on ENM on Cloud Deployment

To access the PIB (Platform Integration Bridge) script to update parameters for ENM applications, it is necessary to determine what environment you are working on and follow the task relevant to your environment. This task outlines the steps to read and configure PIB parameters on an ENM on Cloud Deployment.

Note: ENM concepts are explained in the *ENM Product Description*.

Prerequisites

— A command window is open and you have super user privileges.

— You are connected to an EMP VM using Connect to a Virtual Machine on an ENM on Cloud Deployment on page 3.

Steps

1. As cloud-user change to root:

```
[cloud-user@emp ~]$ sudo su -
[root@emp ~]#
```

2. Find the hostname for the service instance:

```
consul members|grep <service_name>
```

3. Choose one of the returned hostnames for the next steps.
4. Change directory to where the config.py script is located:

```
[root@emp ~]# cd /ericsson/pib-scripts/etc/
[root@gat-emp-0 etc]#
```

5. Read the current parameter value:

```
./config.py read --app_server_address=<service VM hostname>:8080 --service_i →
dentifier=<service_identifier_name> --name=<parameter_name>
```

6. Set the parameter to the required value:

```
./config.py update --app_server_address=<service VM hostname>:8080 --service →
_identifier=<service_identifier_name> --name=<parameter_name> --value=<new_v →
alue>
```

Results

You have updated an application parameter using the PIB script.

5 FM CORBA NBI Troubleshooting

This section provides the troubleshooting steps recommended to diagnose, and fix common problems in the FM NBI.

5.1 Connection to ENM Is Not Available from Network Management System

This section describes how to troubleshoot if the Network Management System (NMS) loses connection to ENM.

Prerequisites

— Basic knowledge of Linux and JBoss

— On physical: Access to the Management Server (MS) and Virtual Machines (VMs)

— On Cloud:

- Access to the private key file for authentication (contact your Openstack administrator)
- Access to the Ericsson Management Portal (EMP) VM in the ENM on Cloud deployment

Steps

1. Verify that `visinamingnb` VM is online.

    a. Log on to the ENM Management Server or ENM Management Portal (EMP) Server on Cloud environment.

    On physical:

    ```
    [root@ms-1 ~]# ssh litp-admin@cloud-svc-1
    [litp-admin@cloud-svc-1 ~]$su root
    Password:
    [root@cloud-svc-1 litp-admin]#
    ```

    b. Execute the following command to check whether the `visinamingnb` VM is online:

    ```
    $ hastatus -sum | grep <VMname>
    ```

Example

```
[root@cloud-svc-1 litp-admin]#  hastatus -sum | grep visinamingnb
B  Grp_CS_svc_cluster_visinamingnb cloud-svc-1          Y
N               ONLINE
B  Grp_CS_svc_cluster_visinamingnb cloud-svc-2          Y
N               OFFLINE
```

Note: Only one `visinamingnb` VM instance is online at any time.

If all the VMs are offline, the issue is with the deployment of VMs. Online VMs by executing the following command, and try connecting to Northbound Interface (NBI) from NMS.

```
$ hagrp -online <groupName> -sys <MasterServerName>
```

Example

```
[root@cloud-svc-1 litp-admin]# hagrp -online rp_CS_svc_cluster_visi
namingnb -sys cloud-svc-2
```

If `visinamingnb` VM is online, check the hostnames of VMs.

On physical use `cat /etc/hostname | grep <VMname>`:

```
[root@cloud-svc-1 litp-admin]#   cat /etc/hosts | grep nbalarmirp
<ipAddress>   svc-3-nbalarmirp      # Created by LITP. Please do no
t edit
<ipAddress>   svc-4-nbalarmirp      # Created by LITP. Please do no
t edit
```

On Cloud use `consul members | grep <VMname>` as the root user on EMP.

```
[root@gat-emp-0 data]# consul members | grep nbalarmirp
gat-nbalarmirp-0 <ipAddress>:8301 alive client 0.8.1 2 dc1
gat-nbalarmirp-1 <ipAddress>:8301 alive client 0.8.1 2 dc1
```

To connect to the `visinamingnb` VM instance, see Connect to a Service on page 2.

c. Execute the following command to check if the `visinaming` service is running:

   If the status of `visinaming` is `'not running'`, execute the following command and try connecting to NBI from NMS.

   ```
   $ service visinaming start
   ```

   If the status of `visinaming` is `'running'`, continue to step 2.

2. Verify that the `nbalarmirp` VMs are online.

   a. Execute the following command to check if JBoss is running:

```
$ service <serviceName> monitor
```

b. Execute the following command to check whether VMs are online:

```
 $ hastatus -sum |grep <VMname>
```

Example

```
[root@cloud-svc-1 litp-admin]#  hastatus -sum | grep nbalarmirp
B  Grp_CS_svc_cluster_nbalarmirp cloud-svc-3          Y          N   →
               ONLINE
B  Grp_CS_svc_cluster_nbalarmirp cloud-svc-4          Y          N   →
               ONLINE
```

If VMs are not online, the issue is with the deployment of VMs.

Execute the following command to online VMs:

```
$ hagrp -online <groupName> -sys <MasterServerName>

Example:
[root@svc-1 litp-admin]# hagrp -online Grp_CS_svc_cluster_nbalarmir  →
p -sys cloud-svc-3
```

If VMs are online, check hostnames of VMs:

On physical use `cat /etc/hostname | grep <VMname>`:

```
[root@cloud-svc-1 litp-admin]#   cat /etc/hosts | grep nbalarmirp
<ipAddress>   svc-3-nbalarmirp      # Created by LITP. Please do no  →
t edit
<ipAddress>   svc-4-nbalarmirp      # Created by LITP. Please do no  →
t edit
```

On Cloud use `consul members | grep <VMname>` as the root user on EMP.

```
[root@gat-emp-0 data]# consul members | grep nbalarmirp
gat-nbalarmirp-0 <ipAddress>:8301 alive client 0.8.1 2 dc1
gat-nbalarmirp-1 <ipAddress>:8301 alive client 0.8.1 2 dc1
```

To connect to the `nbalarmirp` VM instance, see Connect to a Service on page 2.

c. If JBoss is not running, start JBoss by executing the following command:

```
$ service JBoss start
```

d. Verify any active subscriptions exist by using the following command.

```
[root@svc-1-nbalarmirp test_client]# ./testclient.sh subscriptionDa  →
ta
```

Results

You can troubleshoot the issue if NMS loses connection to ENM.

5.2 User Is Not Able to Receive Alarms to NMS Using CORBA-Based NBI

> This section describes how to troubleshoot when the Network Management System (NMS) is unable to receive alarms through CORBA Northbound Interface (NBI).

Prerequisites

— Basic knowledge of Linux and JBoss.

— On physical: Access to the Management Server (MS) and Virtual Machines (VMs).

— On Cloud:

- Access to the private key file for authentication (contact your OpenStack administrator).
- Access to the Ericsson Management Portal (EMP) VM in the ENM on Cloud deployment.

— Access to log viewer and ENM UI

— Active alarm load

Steps

1. Verify if `nbalarmirp` VMs are online.

    a. Log on to the ENM Management Server or ENM Management Portal (EMP) Server on Cloud environment.

    On physical:

    ```
    [root@ms-1 ~]# ssh litp-admin@cloud-svc-1
    [litp-admin@cloud-svc-1 ~]$su root
    Password:
    [root@cloud-svc-1 litp-admin]#
    ```

    b. Execute the following command to check whether VM is online.

    ```
    $ hastatus -sum |grep <VMname>
    ```

Example

```
[root@cloud-svc-1 litp-admin]#  hastatus -sum | grep nbalarmirp
B  Grp_CS_svc_cluster_nbalarmirp cloud-svc-3          Y          N   →
               ONLINE
B  Grp_CS_svc_cluster_nbalarmirp cloud-svc-4          Y          N   →
               ONLINE
```

c. Execute the following command to bring offline VMs online.

```
$ hagrp -online <groupName> -sys <MasterServerName>
```

Example

```
[root@svc-1 litp-admin]# hagrp -online Grp_CS_svc_cluster_nbalarmir →
p -sys cloud-svc-3
```

d. Check the hostname of VMs if they are online.

On physical use `use cat /etc/hostname | grep <VMname>`:

```
root@cloud-svc-1 litp-admin]#   cat /etc/hosts | grep nbalarmirp
<ipAddress>   svc-3-nbalarmirp      # Created by LITP. Please do no →
t edit
<ipAddress>   svc-4-nbalarmirp      # Created by LITP. Please do no →
t edit
```

On Cloud use `consul members | grep <VMname>` as the root user on EMP.

```
[root@gat-emp-0 data]# consul members | grep nbalarmirp
gat-nbalarmirp-0 <ipAddress>:8301 alive client 0.8.1 2 dc1
gat-nbalarmirp-1 <ipAddress>:8301 alive client 0.8.1 2 dc1
```

To connect to `nbalarmirp` VM instance, refer Connect to a Virtual Machine.

e. Execute the following command to check if JBoss and `visinotify` services are running:

```
$ service <serviceName> monitor
```

f. If they are not running, start them by executing the following command:

```
$ service <serviceName> start
```

g. If JBoss and `visinotify` services are running properly, verify that active subscriptions exist by using the following command:

```
[root@svc-3-nbalarmirp cloud-user]# cd /opt/ericsson/com.ericsson.o →
ss.nbi.fm/test_client/
[root@svc-3-nbalarmirp test_client]# ./testclient.sh subscriptionDa →
ta
```

If active subscriptions exist, continue to step 2.

2. Perform the same procedure for `fmalarmprocessing` VMs except verification of `visinotify` status by connecting to `fmalarmprocessing` VMs.

3. If the issue still exists after performing the above two steps, verify the issue using the log viewer, as described below:

    a. Record the host names of the different `nbalarmirp` instances. To do this, SSH into ENM MS and run command: `grep nbalarmirp /etc/hosts`.

    Example

    ```
    [root@ieatlms4899~]# grep nbalarmirp /etc/hosts
    10.247.246.89svc-1-nbalarmirp        nbalarmirp-1-internal   # Crea →
    ted by LITP. Please donot edit
    10.247.246.90svc-2-nbalarmirp        nbalarmirp-2-internal   # Crea →
    ted by LITP. Please donot edit
    ```

    b. Using a web browser such as Firefox, open the log viewer: `https://<<FQDN>>/#logviewer`.

    To retrieve the FQDN, SSH into ENM MS and run command: `grep httpd /etc/hosts`.

    Example

    ```
    [root@ieatlms4899~]# grep httpd /etc/hosts
    10.247.246.102httpd-instance-2iorfile2.ieatENM5325-1.athtem.eei.eri →
    csson.se   # Created by LITP. Please do not edit
    10.247.246.102svc-2-httpd     httpd-2-internal        # Created by   →
    LITP. Please do not edit
    10.247.246.101svc-1-httpd     httpd-1-internal        # Created by   →
    LITP. Please do not edit
    10.247.246.101httpd-instance-1iorfile1.ieatENM5325-1.athtem.eei.eri →
    csson.se   # Created by LITP. Please do not edit
    [root@ieatlms4899~]#
    ```

    Use the hostname of the `httpd-instance` but drop the `iorfile[NUM]`.

    Example

    ```
    ieatenm5325-75.athtem.eei.ericsson.se
    ```

    Use this to form the Fully Qualified Domain Name (FQDN):

    ```
    https://ieatenm5325-75.athtem.eei.ericsson.se/#logviewer
    ```

    c. Enter the following details:

    — Enter **To** and **From** dates as required.

— For widget **Message**, select **contains** and then enter **NotificationService**.

— Select severity type **Error**.

— Add additional criteria and choose **host**, **contains**, followed by one of the `nbalarmirp` instances, for example, `svc-1-nbalarmirp`.

d. Perform search.

e. Verify errors containing `Failed to load object - no data found for:` are present.

Example of log:

```
ERROR [root] (VBJ ThreadPool Worker id=2se=iiop_tp scm=iiop_tp orb=
13b2f057) Failed to load object - no data found for: 41c67461-0a70-
11e8-b9fc-9abaf6787b96
```

f. Repeat the log viewer search for all `nbalarmirp` instances and note the problematic instances reporting the above error.

All `nbalarmirp` instances can experience the same issue and report the log shown above.

4. Perform steps to recover.

a. Follow the steps to verify issues using the log viewer as described in the previous step to identify all problematic `nbalarmirp` instances.

b. Log on to the ENM MS as `root` user.

c. SSH into the problematic `nblarmirp` instance:

```
ssh -i /root/.ssh/vm_private_key cloud-user@<<nbalarmirp-instances>
>
```

Example

```
[root@ieatlms4899~]# ssh -i /root/.ssh/vm_private_key cloud-user@sv
c-1-nbalarmirp
The authenticity of host 'svc-1-nbalarmirp (10.247.246.89)'can't be
 established.
RSA key fingerprint is 8f:61:56:19:16:92:20:55:1a:3e:ae:80:22:22:b0
:c6.
Are you sure you want to continueconnecting (yes/no)? yes
Warning: Permanently added 'svc-1-nbalarmirp,10.247.246.89'(RSA) to
 the list of known hosts.
[cloud-user@svc-1-nbalarmirp ~]$
```

d. Swap into `root` user and move the `notificationservice` healthcheck script. If present, also move the `application_healthceck.sh` script:

```
[cloud-user@svc-1-nbalarmirp ~]$ sudo su
[root@svc-1-nbalarmirp cloud-user]#
[root@svc-1-nbalarmirp cloud-user]# ll /usr/lib/ocf/resource.d
total 24
-rwxr-xr-x. 1root root 983Mar 2313:36application_healthcheck.sh
-rwxr-xr-x. 1root root 971Jan 2916:03cloud_init_ocf.sh
-rwxr-xr-x. 1root root 843Mar  916:41jboss_healthcheck.sh
-rwxr-xr-x. 1root root 892Jan 2916:03nfs_mount_check_ocf.bsh
-rwxr-xr-x. 1root root  94Mar 2313:36notificationservice-lsb-monito →
r.bsh
-rwxr-xr-x. 1root root 880Jan 2916:03oom_check_ocf.bsh


[root@svc-1-nbalarmirp cloud-user]# mv /usr/lib/ocf/resource.d/noti →
ficationservice-lsb-monitor.bsh /tmp
[root@svc-1-nbalarmirp cloud-user]# mv /usr/lib/ocf/resource.d/appl →
ication_healthcheck.sh /tmp
```

e. Stop the `notificationservice` using command: `service notificationservice stop`

   Example

```
[root@svc-1-nbalarmirp cloud-user]# service notificationservice sto →
p
0[main] WARN root  - Stopping service: NotificationService
102[main] DEBUG root  - Server NotificationService: Stopped
NotificationService: INFORMATION (): Stopped notificationservice pr →
ocess with pid : 6597
[root@svc-1-nbalarmirp cloud-user]#
```

f. Verify the `notificationservice` has stopped using command: `service notificationservice status`

   Example

```
[root@svc-1-nbalarmirp cloud-user]# service notificationservice sta →
tus
is stopped
```

g. Start the `notificationservice` using command: `service notificationservice start`

   Example

```
[root@svc-1-nbalarmirp cloud-user]# service notificationservice sta →
rt
...
1537
```

h. Verify the `notificationservice` has successfully started using command: `service notificationservice status`

   Example

```
[root@svc-1-nbalarmirp cloud-user]# service notificationservice sta →
tus
(pid  1537) is running...
```

i. Move the healthcheck scripts back into `/usr/lib/ocf/resource.d`:

```
[root@svc-1-nbalarmirp cloud-user]# mv /tmp/notificationservice-lsb
-monitor.bsh /usr/lib/ocf/resource.d
[root@svc-1-nbalarmirp cloud-user]# mv /tmp/application_healthcheck
.sh /usr/lib/ocf/resource.d/
```

j. Exit from the `nbalarmirp` instance and return to the ENM MS:

```
[root@svc-1-nbalarmirp test_client]# exit
[cloud-user@svc-1-nbalarmirp test_client]# exit
logout
Connection to svc-1-nbalarmirp closed.
[root@ms~]#
```

k. Repeat above steps **a** to **j** for all remaining problematic `nbalarmirp` instances.

l. After steps **a** to **j** have been completed for all problematic `nbalarmirp` instances, verify that FM NBI is forwarding alarms. To do this, use the `testclient`, as described in the following steps.

m. From the ENM MS, SSH into any `nbalarmirp` instance using command: `ssh -i /root/.ssh/vm_private_key cloud-user@<<nbalarmirp-instances>>`

```
[root@ms~]# ssh -i /root/.ssh/vm_private_key cloud-user@svc-1-nbala
rmirp
The authenticity of host 'svc-1-nbalarmirp (10.247.246.89)'can't be
 established.
RSA key fingerprint is 8f:61:56:19:16:92:20:55:1a:3e:ae:80:22:22:b0
:c6.
Are you sure you want to continueconnecting (yes/no)? yes
Warning: Permanently added 'svc-1-nbalarmirp,10.247.246.89'(RSA) to
 the list of known hosts.
[cloud-user@svc-1-nbalarmirp ~]$
```

n. To run the test client, change to directory `/opt/ericsson/com.ericsson.oss.nbi.fm/test_client` using command: `cd /opt/ericsson/com.ericsson.oss.nbi.fm/test_client`. From here, launch the test client using command: `./testclient.sh subscribe category 1f1`

   Example

```
[root@svc-1-nbalarmirp cloud-user]# cd /opt/ericsson/com.ericsson.o
ss.nbi.fm/test_client
[root@svc-1-nbalarmirp test_client]# ./testclient.sh subscribe cate
gory 1f1
<<Subscription Details>>
<<Alarms>>
....
```

   Ensure that there is an active alarm load to verify that alarms are being forwarded by FM NBI.

o. Verify that alarms are being received by the test client you just launched. This shows that alarms are now being forwarded by FM NBI.

Results

User is able to troubleshoot the issue if NMS is not able to receive alarms through CORBA NBI.

6 FM SNMP NBI Troubleshooting

This section provides the troubleshooting steps recommended to diagnose and fix common problems in the FM SNMP NBI application.

6.1 Network Management System Cannot Reach ENM Using SNMP

Diagnose why Network Management System (NMS) cannot reach ENM using SNMP on the UDP port.

Prerequisites

- Basic knowledge of Linux, OpenStack, JBoss, and SNMP protocol
- FM VIP IPv4/IPv6 external address of the ENM deployment is known.
- Access to the shell terminal on NMS host with a MIB browser or an SNMP utility like `net-snmp-utils`.

Verify UDP Connectivity between NMS and ENM

1. Log on to `nbfmsnmp` VM by following the instructions in Connect to a Service on page 2
2. Check `nbfmsnmp` respond to the keepalive request of `lvsrouter` on TCP 35161:

```
[cloud-user@svc-8-nbfmsnmp ~]$ sudo su
[root@svc-8-nbfmsnmp cloud-user]# tcpdump tcp -i any dst port 35161
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on any, link-type LINUX_SLL (Linux cooked), capture size 65535 byt →
es
13:58:21.754861 IP svc-2-lvsrouter.58854 > svc-8-nbfmsnmp.35161: Flags [.],  →
ack 4284863371, win 229, options [nop,nop,TS val 9673199 ecr 9581021], lengt →
h 0
13:58:21.754925 IP svc-2-lvsrouter.58854 > svc-8-nbfmsnmp.35161: Flags [P.],  →
 seq 0:1, ack 1, win 229, options [nop,nop,TS val 9673199 ecr 9581021], leng →
th 1
13:58:21.754954 IP svc-2-lvsrouter.58854 > svc-8-nbfmsnmp.35161: Flags [F.],  →
 seq 1, ack 1, win 229, options [nop,nop,TS val 9673199 ecr 9581021], length →
 0
13:58:21.755167 IP svc-2-lvsrouter.58854 > svc-8-nbfmsnmp.35161: Flags [.],  →
ack 9, win 229, options [nop,nop,TS val 9673200 ecr 9581022], length 0
13:58:21.755187 IP svc-2-lvsrouter.58854 > svc-8-nbfmsnmp.35161: Flags [.],  →
ack 10, win 229, options [nop,nop,TS val 9673200 ecr 9581022], length 0
13:58:22.576109 IP svc-1-lvsrouter.25096 > svc-8-nbfmsnmp.35161: Flags [S],  →
seq 407185252, win 29200, options [mss 1460,sackOK,TS val 9674678 ecr 0,nop, →
wscale 7], length 0
13:58:22.576280 IP svc-1-lvsrouter.25096 > svc-8-nbfmsnmp.35161: Flags [.],  →
ack 1404881823, win 229, options [nop,nop,TS val 9674678 ecr 9581843], lengt →
```

```
h 0
13:58:22.576819 IP svc-1-lvsrouter.25096 > svc-8-nbfmsnmp.35161: Flags [P.], →
 seq 0:1, ack 1, win 229, options [nop,nop,TS val 9674678 ecr 9581843], leng →
th 1
13:58:22.576823 IP svc-1-lvsrouter.25096 > svc-8-nbfmsnmp.35161: Flags [F.], →
 seq 1, ack 1, win 229, options [nop,nop,TS val 9674678 ecr 9581843], length →
 0
13:58:22.577008 IP svc-1-lvsrouter.25096 > svc-8-nbfmsnmp.35161: Flags [.], →
ack 9, win 229, options [nop,nop,TS val 9674679 ecr 9581844], length 0
13:58:22.577042 IP svc-1-lvsrouter.25096 > svc-8-nbfmsnmp.35161: Flags [.], →
ack 10, win 229, options [nop,nop,TS val 9674679 ecr 9581844], length 0
 ^C 11 packets captured 28 packets received by filter 0 packets dropped by k →
ernel
[root@svc-8-nbfmsnmp cloud-user]#
```

If you do not see any output from `tcpdump` connect to the other instance and repeat the command to check that you are connected to the master instance.

3. Run `tcpdump` again, filtering the `udp` packets:

```
[root@svc-8-nbfmsnmp cloud-user]# tcpdump udp -i any dst port 35161
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on any, link-type LINUX_SLL (Linux cooked), capture size 65535 byt →
es
```

`tcpdump` is now listening for incoming requests on the UDP 35161, which is the SNMP agent listening port.

4. Run `snmpget` on NMS host to verify the request is forwarded properly to the `nbfmsnmp` master instance.

   You can use a MIB browser or any SNMP utility like `net-snmp-utils`, which is available for many OSs and Linux distributions. Open a new terminal on NMS host perform a `snmpwalk` toward the FM_VIP on UDP agent listening port.

   Here is an example using `snmpwalk` from `net-snmp-utils`:

```
snmpwalk -c <community> <ip address>:<port>
```

   You can use any string, at this point you do not need to know which community string is configured. Verify that the UDP packets are captured to confirm connectivity. The IP address is the external IPv4 or IPv6 FM_VIP address.

   Example:

```
netsim@ieatnetsimv7009-53:~/.snmp> snmpwalk -v2c -c anystring 131.160.168.12 →
6:35161
Timeout: No Response from 131.160.168.126:35161
netsim@ieatnetsimv7009-53:~/.snmp
```

   On the other terminal on `nbfmsnmp` VM where you have `tcpdump` listening, see the captured packages:

```
[root@svc-8-nbfmsnmp cloud-user]# tcpdump udp -i any dst port 35161
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on any, link-type LINUX_SLL (Linux cooked), capture size 65535 byt →
```

```
es
14:32:30.437311 IP 10.149.20.98.53604 > svc-8-nbfmsnmp.35161: UDP, length 43
14:32:31.438344 IP 10.149.20.98.53604 > svc-8-nbfmsnmp.35161: UDP, length 43
14:32:32.443449 IP 10.149.20.98.53604 > svc-8-nbfmsnmp.35161: UDP, length 43
14:32:33.440410 IP 10.149.20.98.53604 > svc-8-nbfmsnmp.35161: UDP, length 43
14:32:34.441494 IP 10.149.20.98.53604 > svc-8-nbfmsnmp.35161: UDP, length 43
```

If you have a result like in the example, the connectivity between NMS and ENM is good and you can proceed with other diagnostic steps.

6.2 Network Management System Cannot Query Alarm Table, Alert Tables, and Scalar Objects

> Diagnose why MS cannot query Alarm Table, Alert Table, and Scalar Objects.

Prerequisites

- Access to the Management Server and Virtual Machines (VMs)
- FM VIP IPv4/IPv6 external address of the ENM deployment is known
- Access to shell terminal on NMS host and ability to run a MIB browser or an SNMP utility like `net-snmp-utils`
- Basic knowledge of Linux, JBoss, and SNMP protocol
- Completed *Verify UDP Connectivity between NMS and ENM* in Network Management System Cannot Reach ENM Using SNMP on page 21.

Verify the SNMP Agent Is Running in Unlocked Mode and the Subscription Is Present and Correctly Configured.

- From ENM web CLI app, get information about the SNMP agent:

```
fmsnmp get agent
```

For example:

```
fmsnmp get agent
Administrative state : locked
```

If the SNMP agent is locked, unlock it:

```
fmsnmp set agent -as unlocked
```

Verify The Subscription Is Present and Correctly Configured.

1. From ENM web CLI app, get a subscription list:

   ```
   fmsnmp get nmslist
   ```

   For example:

   ```
   fmsnmp get nmslist
   sub01
   1 instance(s)
   ```

2. Check that the subscription is unlocked:

   ```
   fmsnmp get target -n sub01
   ```

   For example:

   ```
   fmsnmp get target -n sub01
   name : sub01
   ip address : 10.149.21.248
   port : 10162
   type : v2c
   community : com01
   administrative state : locked
   ```

   If the subscription is locked, unlock it:

   ```
   fmsnmp set target -n <subsname> -as unlocked
   ```

   For example:

   ```
   fmsnmp set target -n sub01 -as unlocke
   ```

Check the Subscription Is Configured as Expected.

1. Check for SNMPv2c that the type parameter is set to v2c and the community parameter is the same use from NMS.

   If these mandatory parameters are wrong:

   a. Delete the subscription:

      ```
      fmsnmp delete -n <subsname>
      ```

      For example:

      ```
      fmsnmp delete -n sub01
      ```

   b. Re-create it with correct parameters.

For example:

```
fmsnmp create v2c -n myTarget --ipaddress 10.149.21.248 --community pub
lic
create  command executed
```

2. Check for SNMPv3 that the type parameter is set to v3 and the authentication and privacy parameters are correct.

   For example:

```
fmsnmp get target -n myNMS
fmsnmp get target -n myNMS
name : myNMS
ip address : 10.144.37.46
port : 30162
type : v3
user : snmpuser2
security level : authpriv
authentication protocol : SHA
privacy protocol : DES
administrative state : unlocked
```

   Note: The `authkey` and `privkey` are not shown for security reasons. If you do not know the configured values, delete and configure them again with the correct ones:

```
»fmsnmp create v3 authpriv --name myNMS --ipaddress 10.144.37.46 --
user snmpuser2 --authkey <password> --privkey <password> --authprot
ocol SHA --privprotocol DES --port 30162
fmsnmp create v3 authpriv --name myNMS --ipaddress 10.144.37.46 --u
ser snmpuser2 --authkey <password> --privkey <password> --authproto
col SHA --privprotocol DES --port 30162
fmsnmp create command executed
```

Verify the SNMP Agent Is Alive and Running

1. From the `nbfmsnmp` VM Linux shell, run the test script to query all the supported scalars to verify that the agent is storing alarms, managing summary counters, and processing synchronous requests (SNMP GET/ WALK).

   For example:

```
[cloud-user@svc-9-nbfmsnmp test_client]$ cd
[cloud-user@svc-9-nbfmsnmp ~]$ cd /opt/ericsson/ERICnbfmsnmpservice_CXP90348
20/test_client/
[cloud-user@svc-9-nbfmsnmp test_client]$
```

2. Check the agent on both VM instances:

```
./snmpnbi-get-summary.sh 127.0.0.1 <community>
```

   The agent IP address is the internal VLAN IPv4 address binding the eth0 interface, or external VLAN IPv6 address binding the eth2 interface. The community string is one of the communities for the configured subscription in ENM (`com01` in the example).

For example:

```
[cloud-user@svc-8-nbfmsnmp test_client]$ ./snmp_test_scalar.sh 10.247.246.24
5 com01
================================================================================
=
START GETTING SCALARS USING FM_VIP: [10.247.246.244] community [com01]
================================================================================
=
--->> ALARMS SUMMARY
eriAlarmSumIndeterminate.0 = 14
eriAlarmSumCritical.0 = 92
eriAlarmSumMajor.0 = 54
eriAlarmSumMinor.0 = 0
eriAlarmSumWarning.0 = 111
________________________________________________________________________________
_
--->> ALARMS COUNTERS
eriAlarmActiveNumber.0 = 98977
eriAlarmActiveLastChanged.0 = 2017-12-6,16:24:23.7,+0:0
eriAlarmActiveLastSequenceNo.0 = 0
________________________________________________________________________________
_
--->> ALERT COUNTERS
eriAlarmAlertNumber.0 = 640
eriAlarmAlertLastChanged.0 = 2017-12-6,14:3:0.3,+0:0
eriAlarmAlertLastSequenceNo.0 = 0
________________________________________________________________________________
_
--->> SNMP-FRAMEWORK-MIBSNMP-FRAMEWORK-MIB::snmpEngineID.0 = Hex-STRING: 80
00 13 70 04 45 4E 4D 2D 4E 42 2D 49 6E 74 65
72 66 61 63 65
SNMP-FRAMEWORK-MIB::snmpEngineBoots.0 3
SNMP-FRAMEWORK-MIB::snmpEngineTime.0 19392 seconds
========================================================================
[cloud-user@svc-9-nbfmsnmp test_client]$
```

Note: This script is based on Linux utility `net-snmp-utils`. You can also perform the check as follows:

```
snmpget -v2c -c <community> <agent_address>:<listening_port> SNMP-F
RAMEWORK-MIB::snmpEngineTime.0 -Oq
```

For example:

```
[cloud-user@svc-8-nbfmsnmp test_client]$ snmpget -v2c -c com01 10.2
47.246.244:35161 SNMP-FRAMEWORK-MIB::snmpEngineTime.0 -Oq
SNMP-FRAMEWORK-MIB::snmpEngineTime.0 4363 seconds
you can also use the OID in this way:
[cloud-user@svc-8-nbfmsnmp test_client]$ snmpget -v2c -c com01 10.2
47.246.244:35161 .1.3.6.1.6.3.10.2.1.3.0
SNMP-FRAMEWORK-MIB::snmpEngineTime.0 = INTEGER: 4807 seconds
```

Verify SNMP GET/WALK on NMS Host Using a Test Tool

Use a test client running on the NMS host and check the SNMP answer to SNMP GET/WALK requests.

You can use any SNMP tool like a MIB browser or the the Linux SNMP utility `net-snmp-utils`.

1. Configure your tool and insert the external FM_VIP IPv4 or IPv6 address:
    a. Insert UDP 35161 as SNMP agent listening port.

b. Insert one of communities for the configured subscription in ENM.

c. Select the `snmpEngineTime` OID .1.3.6.1.6.3.10.2.1.3 to check the agent time.

d. If the NMS host is Linux-based, use `net-snmp-utils` or equivalent tool. Copy the SNMP-FRAMEWORK-MIB, ERICSSON-ALARM-MIB, ERICSSON-ALARM-PC-MIB, ERICSSON-ALARM-TC-MIB, ERICSSON-ALARM-XPATH-MIB, ERICSSON-TC-MIB, ERICSSON-TOP-MIB MIB files in `$HOME/.snmp/mibs`.

2. Check that the SNMP agent is alive:

```
[root@test-nms ~]# snmpget -v2c -c com01 10.247.246.244:35161  SNMP-FRAMEWOR →
K-MIB::snmpEngineTime.0 -Oq
SNMP-FRAMEWORK-MIB::snmpEngineTime.0 8997 seconds
```

You can also use OID instead:

```
[root@test-nms ~]# snmpget -v2c -c com01 10.247.246.244:35161 .1.3.6.1.6.3.1 →
0.2.1.3.0
SNMP-FRAMEWORK-MIB::snmpEngineTime.0 = INTEGER: 8977 seconds
```

3. Check if alarms can be retrieved:

```
snmpwalk -v2c -c <community> <FM_VIP address>:<listening_port> .1.3.6.1.4.1. →
193.183.4.1.3.5 -Oq.0 -Oq
```

For example:

```
[cloud-user@svc-2-nbfmsnmp test_client]$ snmpwalk -v2c -c public  131.160.16 →
8.126:35161 .1.3.6.1.4.1.193.183.4.1.3.5 -Pde -Os
enterprises.193.183.4.1.3.5.1.2.1514981 = Gauge32: 193
enterprises.193.183.4.1.3.5.1.3.1514981 = Gauge32: 1514981
enterprises.193.183.4.1.3.5.1.4.1514981 = STRING: "test"
enterprises.193.183.4.1.3.5.1.5.1514981 = STRING: "SubNetwork=Europa_Test,Su →
bNetwork=PMIC,MeContext=LTE06dg2ERBS00005,PowerDistribution - LTE06dg2ERBS00 →
005 -1758"
enterprises.193.183.4.1.3.5.1.6.1514981 = INTEGER: 6
enterprises.193.183.4.1.3.5.1.7.1514981 = Hex-STRING: 07 E1 0C 07 0A 1F 2D 0 →
0 2B 00 00
enterprises.193.183.4.1.3.5.1.8.1514981 = Hex-STRING: 07 E1 0C 07 0A 1F 2D 0 →
0 2B 00 00
enterprises.193.183.4.1.3.5.1.9.1514981 = INTEGER: 315 enterprises.193.183.4 →
.1.3.5.1.10.1514981 = INTEGER: 4
enterprises.193.183.4.1.3.5.1.11.1514981 = INTEGER: 4 enterprises.193.183.4. →
1.3.5.1.11.1514986 = INTEGER: 4
enterprises.193.183.4.1.3.5.1.12.1514981 = STRING: "My comment" enterprises. →
193.183.4.1.3.5.1.13.1514981 = STRING: "My comment"
enterprises.193.183.4.1.3.5.1.14.1514981 = OID: zeroDotZero enterprises.193. →
183.4.1.3.5.1.15.1514981 = ""
```

6.3 Network Management System Cannot Receive Alarm or Alert Notifications Using SNMP

Diagnose why NMS cannot receive SNMP traps.

Prerequisites

— Basic knowledge of Linux, OpenStack, SNMP protocol.

— Access to ENM web CLI.

— Access to shell terminal on NMS host and ability to run a packet capturing tool (`tcpdump`).

Verify the Configured Target Destination SNMP Version Is of the Same Type Where the NMS SNMP Manager Is Listening

1. From the ENM web CLI app, get subscription information:

   ```
   fmsnmp get target -n <target_name>
   ```

   For example:

   ```
   fmsnmp get target -n sub01
   name : sub01
   ip address : 10.149.21.248
   port : 10162
   type : v2c
   community : com01
   administrative state : unlocked
   ```

2. Check the subscription version.

   If it is not as expected, delete the subscription and configure it again.

Verify the Configured Target Destination Port Is the Same Where the NMS SNMP Manager Is Listening

1. From ENM web CLI app, get subscription information:

   ```
   fmsnmp get target -n <target_name>
   ```

2. Check that the port configured is the same as the one configured on the NMS SNMP Manager.

   If the port is wrong, delete the subscription and configure it again.

Verify the Configured Target Destination IP Address Is the Same Where the NMS SNMP Manager Is Listening

1. From ENM web CLI app, get subscription information:

   ```
   fmsnmp get target -n <target_name>
   ```

2. Check the IP address configured is the same on which NMS SNMP Manager is listening for traps receiving.

   If the IP is wrong, delete the subscription and configure it again.

Verify the SNMP Agent Is Running in Unlocked Mode and the Subscription Is Present and Correctly Configured

— From ENM web CLI app, get information about the SNMP agent:

```
fmsnmp get agent
```

For example:

```
fmsnmp get agent
Administrative state : unlocked
```

If the SNMP agent is locked, unlock it:

```
fmsnmp set agent -as unlocked
```

Verify the Subscription Is Present and Correctly Configured

1. From ENM Web CLI app, get a subscription list:

   ```
   fmsnmp get nmslist
   ```

   For example:

   ```
   fmsnmp get nmslist
   sub01
   1 instance(s)
   ```

2. Check that the subscription is unlocked:

   ```
   fmsnmp get target -n sub01
   ```

   For example:

   ```
   fmsnmp get target -n sub01
   name : sub01
   ip address : 10.149.21.248
   port : 10162
   ```

```
type : v2c
community : com01
administrative state : unlocked
```

If the subscription is locked, unlock it:

```
fmsnmp set target -n <subsname> -as unlocked
```

For example:

```
fmsnmp set target -n sub01 -as unlocked
```

3. Check that the subscription is configured as expected.

Verify the Configured Target USM (for SNMPv3) Is the Same Where the NMS SNMP Manager Is Listening

1. From ENM web CLI app, get subscription information:

```
fmsnmp get target -n <target_name>
```

For example:

```
fmsnmp get target -n myNMS
name : myNMS
ip address : 10.144.37.46
port : 30162
type : v3
user : snmpuser2
security level : authpriv
authentication protocol : SHA
privacy protocol : DES
administrative state : unlocked
```

2. Check that the user you want to use as the SNMPv3 user is configured in ENM as ENM user with the `NbiFmSnmpManager` role:

   If the user is not present, configure a user in ENM with the `NbiFmSnmpManager` role and use it as the user parameter in the SNMPv3 subscription.

3. Check that all authentication parameters are configured the same at NMS SNMP Manager.

Note: The `authkey` and `privkey` are not shown for security reasons. If you do not know the configured values, delete and configure them again with the correct ones:

```
»fmsnmp create v3 authpriv --name myNMS --ipaddress 10.144.37.46 --
user snmpuser2 --authkey <password> --privkey <password> --authprot
ocol SHA --privprotocol DES --port 30162
fmsnmp create v3 authpriv --name myNMS --ipaddress 10.144.37.46 --u
ser snmpuser2 --authkey <password> --privkey <password> --authproto
col SHA --privprotocol DES --port 30162
fmsnmp create command executed
```

The user snmpuser2 shown in this example have been also configured as ENM user with `NbiFmSnmpManager` role.

Verify the Configured Target USM (for SNMPv3) Is the Same Where the NMS SNMP Manager Is Listening

Use a tool to capture UDP packets like `tcpdump`, running on the NMS host and check the traps are received.

1. If the NMS host is Linux-based, open a shell terminal as superuser and perform a capture running `tpcpdump` on the trap port configured in the subscription on ENM.

   For example, assuming in ENM you have configured the subscription on the default trap port 30162:

   ```
   [root@test-nms ~]#  tcpdump udp -i any dst port 30162
   tcpdump: verbose output suppressed, use -v or -vv for full protocol decode l
   istening on any, link-type LINUX_SLL (Linux cooked), capture size 96 bytes
   ```

2. If the subscription is configured and is unlocked, generate an alarm on ENM and check the packets are getting received and captured in the `tcpdump` at the NMS host.

   Connectivity is good if the packages are captured. You can proceed to checking the settings in NMS application.

7 Fault Management Troubleshooting

This section provides the troubleshooting steps recommended to diagnose and fix common problems in ENM Fault Management.

7.1 After Selecting an Alarm, Alarm Details Pages Are Not Displayed on the GUI

This section describes how to troubleshoot the following issue:

— An alarm is selected, but the relevant Alarm Details GUI is not displayed.

Prerequisites

— Basic knowledge of Linux and JBoss

— On physical: Access to the Management Server (MS) and Virtual Machines (VMs)

— On Cloud:

- Access to the private key file for authentication (contact your Openstack administrator)
- Access to the Ericsson Management Portal (EMP) VM in the ENM on Cloud deployment

Steps

1. Access the cliapp from the following address:

```
https://<apache_host_name>/#cliapp
```

2. From the cliapp, execute the following command to view alarms on the Command Line Interface (CLI):

```
alarm get * --objectOfReference MeContext=LTE02ERBS00007 --list
```

The following command and system output response are given as an example:

```
alarm get * --objectOfReference MeContext=LTE02ERBS00007 --list

eventPoId                  objectOfReference
281474978194244           MeContext=LTE02ERBS00007,ManagedElement=1
```

```
Total number of alarms fetched for the given query is :1
```

Result: If no alarms are displayed on the CLI, it means that no alarms are present in the database and the procedure ends.

If one or more alarms are present then, from the Alarm Details GUI, compare the values of the eventPoId displayed before and after execution of the command. If the two values are the same then skip to step 4 to verify the current fmserv VM status is online. Otherwise, proceed to the next step.

3. Log on to the ENM Management Server and connect to any SVC node or to ENM Management Portal (EMP) server on cloud environment.

```
[root@ms-1 ~]# ssh litp-admin@cloud-svc-1
```

   a. Switch to `root` user:

```
[litp-admin@cloud-svc-1 ~]$su root
```

   b. Type the required password at the root prompt:

```
Password:
[root@cloud-svc-1 litp-admin]#
```

4. Check the status of the fmserv VMs:

```
 $ hastatus -sum |grep <VMname>
```

See the following command and system output response as an example (where the status of the fmserv VMs is `online`):

```
[root@cloud-svc-1 litp-admin]# hastatus -sum |grep fmserv
B  Grp_CS_svc_cluster_fmserv cloud-svc-3          Y          N               →
 ONLINE
B  Grp_CS_svc_cluster_fmserv cloud-svc-4          Y          N               →
 ONLINE
```

   a. If the VMs are not online, the issue is with the VM deployment. In that case proceed to the next step (5). If the VMs are online, skip to step 6.

5. Switch to online VMs and the procedure is ended:

```
$ hagrp -online <groupName> -sys <MasterServerName>
```

Example

```
[root@svc-1 litp-admin]# hagrp -online Grp_CS_svc_cluster_fmserv -sys svc-3
```

6. If the VMs are online, check the hostname of the VMs:

On physical use `cat /etc/hostname | grep <VMname>`:

```
[root@cloud-svc-1 litp-admin]# cat /etc/hosts | grep fmserv
<ipAddress>  svc-3-fmserv    # Created by LITP. Please do not edit
<ipAddress>  svc-4-fmserv    # Created by LITP. Please do not edit
```

On Cloud use `consul members | grep <VMname>` as the root user on EMP.

```
[root@gat-emp-0 data]# consul members | grep fmserv
gat-fmserv-0 10.5.1.82:8301 alive client 0.8.1 2 dc1
gat-fmserv-1 10.5.1.83:8301 alive client 0.8.1 2 dc1
[root@gat-emp-0 data]#
```

7. To connect to the `fmserv` VM instance, see Connect to a Service on page 2.

```
[root@ms-1 ~]# ssh -i /root/.ssh/vm_private_key cloud-user@svc-1-fmserv
```

8. Verify that the JBoss service is running:

```
$ service jboss monitor
```

Result: If the JBoss service is not running then go to the next step (9) to start it.

If the JBoss is running, skip to step 10.

9. Start JBoss:

```
$ service jboss start
```

10. If the JBoss instances are running properly, verify that both the `alarmcontroldisplayservice` ear and the `AlarmQueryService` ear are deployed properly. If not, deploy them again.

11. Check the fmserv VM DEBUG logs.

    Example

    ```
    [root@svc-1-fmserv~]# cd /ericsson/3pp/jboss/standalone/log
    [root@svc-1-fmservlog]# vi server.log
    ```

    a. Verify that the request reaches AlarmQueryService:

    ```
    Request hits AlarmQueryService with QueryData
    ```

    b. Check the number of AlarmRecords fetched for a given Query or Request:

    ```
    Total Alarms found with given PO Ids
    ```

    c. Verify that the Comments History is fetched successfully:

```
Number of comments fetched from db are
```

12. If the fmserv VM DEBUG logs are properly displayed, exit from the fmserv VM and go to the next step.
13. Repeat steps 1 to 13 for the fmhistory VMs.
    a. Verify the `historical-query-service` ear is deployed properly, otherwise deploy it again.
    b. If `historical-query-service` ear is deployed properly, exit from fmhistory VM and go to the next step.
14. Repeat the steps 1 to 13 for the solr VMs.

    Note: For information on collecting DEBUG logs refer to the *ENM Data Collection Guidelines*.

Results

The Alarm Details GUI is displayed if an alarm is selected.

7.2 Alarm List Table Not Responding in Alarm Monitor GUI

This section describes how to troubleshoot the following issue:

— The Alarm List table not responding in the Alarm Monitor GUI.

Prerequisites

— Basic knowledge of Linux and JBoss

— On physical: Access to the Management Server (MS) and Virtual Machines (VMs)

— On Cloud:

- Access to the private key file for authentication (contact your OpenStack administrator)
- Access to the Ericsson Management Portal (EMP) VM in the ENM on Cloud deployment

Steps

1. Log on to the ENM Management Server and connect to any SVC node or to ENM Management Portal (EMP) server on Cloud environment.

```
[root@ms-1 ~]# ssh litp-admin@cloud-svc-1
```

a. Switch to `root` user:

```
[litp-admin@cloud-svc-1 ~]$su root
```

b. Type the password at the root prompt:

```
Password:
[root@cloud-svc-1 litp-admin]#
```

2. Check the status of the `fmserv` VMs:

```
 $ hastatus -sum |grep <VMname>
```

See the following command and system output response as an example (where the `fmserv` VM status is online):

```
[root@cloud-svc-1 litp-admin]# hastatus -sum |grep fmserv
B  Grp_CS_svc_cluster_fmserv cloud-svc-3          Y          N               →
 ONLINE
B  Grp_CS_svc_cluster_fmserv cloud-svc-4          Y          N               →
 ONLINE
```

a. If the VMs are not online, the issue is with the deployment of VMs. In that case, go to step 3 to bring the VMs online.

b. If the VMs are online, skip to step 4.

3. Bring the VMs online:

```
$ hagrp -online <groupName> -sys <MasterServerName>
```

Example

```
[root@svc-1 litp-admin]# hagrp -online Grp_CS_svc_cluster_fmserv -sys svc-3
```

Result: When you have brought the VMs online, this procedure ends. The remaining steps pertain to situations in which VMs were found to be online in step 2.

4. If the VMs are online, check the hostname of the VMs:

On physical use `cat /etc/hostname | grep <VMname>`:

```
[root@cloud-svc-1 litp-admin]# cat /etc/hosts | grep fmserv
<ipAddress>  svc-3-fmserv    # Created by LITP. Please do not edit
<ipAddress>  svc-4-fmserv    # Created by LITP. Please do not edit
```

On Cloud use `consul members | grep <VMname>` as the root user on EMP.

```
[root@gat-emp-0 data]# consul members | grep fmserv
gat-fmserv-0            10.5.1.82:8301   alive   client  0.8.1  2        d →
c1
gat-fmserv-1            10.5.1.83:8301   alive   client  0.8.1  2        d →
```

```
c1
[root@gat-emp-0 data]#
```

5. To connect to the `fmserv` VM instance, see Connect to a Service on page 2.

```
[root@ms-1 ~]# ssh -i /root/.ssh/vm_private_key cloud-user@svc-3-fmserv
```

6. Verify the JBoss service is running:

```
$ service jboss monitor
```

If the JBoss service is not running, go to the next step. If JBoss is already running, skip to step 8.

7. Start JBoss:

```
$ service jboss start
```

8. Access the cliapp from the following address:

```
https://<apache_host_name>/#cliapp
```

9. Execute in the cliapp the following command to view alarms on CLI:

```
alarm get <nodeName>
```

Example

```
alarm get LTE02ERBS00007
```

If no Alarms are displayed on the CLI that means no issue with Alarm Query Service.

10. Check the `fmserv` VM DEBUG logs.

    a. Verify that request reaches postAllEventPoIdList:

    ```
    Request for postAllEventPoIdList from UI is
    ```

    If exceptional messages are displayed, request does not reach the control display service.

    b. If the request reaches postAllEventPoIdList, check the response message:

    ```
    Response for postAllEventPoIdList is
    ```

If the log is not displayed, there is an issue in data conversion in the control display service. If both the logs are displayed, the HTTP request does not reach the UI from the control display service.

Results

The Alarm List table is properly loaded in the Alarm Monitor GUI.

7.3 Alarms Are Not Received in FM Mediation from the NE

This section describes how to troubleshoot the following issue:

— Fault Management (FM) Mediation cannot receive alarms from the NE.

Prerequisites

— User requires basic knowledge of Linux and JBoss.

— On physical: Access to the Management Server (MS) and Virtual Machines (VMs).

— On Cloud:

- Access to the private key file for authentication (contact your OpenStack administrator)
- Access to the Ericsson Management Portal (EMP) VM in the ENM on Cloud deployment

Steps

1. Verify that supervision is enabled on the NE from the Alarm Monitoring GUI of the ENM.

   For further details about how to enable supervision, refer to the *Alarm Monitor* tutorial in the ENM online help at the following path:

   ```
   https:/<apache_host_name>/#help/app/alarmviewer/concept/tutorials/NodeOperat →
   ions
   ```

2. Remove any filter criteria, if applied.

3. Check the supervision state of the Network Element from CliApp using the following command:

   Note: CliApp can be accessed through `https://<apache_host_name>/#cliapp`

   $ alarm status `<nodename>`

Example

```
&187;alarm status LTE02ERBS00003 -l
Subscription Status : true
Heartbeat Failure : false
Synchronization Ongoing : false
Type : ERBS
Alarm Suppressed : false
Node Suspended : false
Out Of Synchronization : false
Supervision Status : on
Technician Present : false
NodeName : LTE02ERBS00003

Total number of node status data fetched for the given query is :1
```

4. Check the status of active attribute.

    a. If it is `false`, enable the supervision and send an alarm from Network Element. If the problem persists, then follow the previous steps.

    b. If it is `true`, then check for node name in all `msfm` VM instances with the following command:

    ```
    $ grep "Supervision Started on node  FDN" server.log | grep <nodena →
    me>
    ```

    Example

    ```
    [root@cloud-svc-3-msfm log]# grep "Supervision Started on node FDN" →
     server.log | grep LTE02ERBS00003
    2015-10-07 07:08:48,497 INFO [com.ericsson.oss.mediation.alarm.ra.c →
    ommon.CorbaSupervision] (default-threads - 2) Supervision Started o →
    n node FDN:           NetworkElement=LTE02ERBS00003 with Subscripti →
    on ID: 22
    ```

    Note: Take the latest log of all VM instances by checking the time-stamp.

    c. If Subscription ID is '`null`' in the log, then check if HeartBeat alarm is present on the Network Element in the alarm viewer.

    d. If HeartBeat alarm is present, then collect snoop or tcpdump from `msfm` where the Network Element is being supervised with the following command:

    ```
    $ tcpdump -i any -B 16384 -Uw <filename.pcap> '(host<nodeIP>)'
    ```

    Example

    ```
    [root@cloud-svc-3-msfm ]# tcpdump -i any -B 16384 -Uw /var/tmp/tcpd →
    ump.pcap '(host <ipAddress>)'
    ```

    Note: File Extension must be .pcap.

    e. Collect the logs for 5 minutes and provide the log files of `msfm` and tcpdump.

f. If Subscription ID is proper, then execute the tcpdump command from `msfm`. Send an alarm from the Network Element and provide the logs from `msfm`, `fmalarmprocessing`, and `tcpdump`.

Note: Refer as `mssnmpfm` instead of `msfm` for SNMP-based nodes.

5. Check the security settings of `SnmpTargetV3=ENMFM` MO in ENM CLI are consistent with the corresponding ENM credentials in ENM CLI.

a. Get the security settings of `SnmpTargetV3=ENMFM` MO in ENM CLI:

```
cmedit get SubNetwork=<subNetworkName>,ManagedElement=<nodeName>,Sy →
stemFunctions=1,SysM=1,Snmp=1,SnmpTargetV3=ENMFM
```

Example

```
cmedit get SubNetwork=LTE01dg2ERBS00003,ManagedElement=LTE01dg2ERBS →
00003,SystemFunctions=1,SysM=1,Snmp=1,SnmpTargetV3=ENMFM
```

b. Get ENM credentials in ENM CLI:

```
secadm snmp get --plaintext show --nodelist <nodeName>
```

Example

```
secadm snmp get --plaintext show --nodelist LTE01dg2ERBS00003
```

c. If the values returned in steps (a) and (b) do not match, manually update the security settings of `SnmpTargetV3=ENMFM` MO to be consistent with the corresponding ENM credentials in ENM CLI.

Note: The value of "`user`" is always set as "`DummyUser`" in `SnmpTargetV3=ENMFM` MO when FM Supervision is enabled. This is regardless of the value set for "`user`" in ENM CLI.

Results

FM Mediation can receive alarms from the NE.

7.3.1 Alarms Are not Received in FM Mediation from the NE for Router6672

In Router6672 nodes check the following steps along with the aforementioned points in Alarms Are Not Received in FM Mediation from the NE on page 38:

Note: The steps are also applicable to Router6675, Router6274, Router6273, and Router6x71 nodes.

Steps

1. Check the address context on the Network Element from which the notifications and traps are routed to external network. The 'local' context is the default configured address-context available on the node.

   — The node can also have any other address-context other than 'local' (for example 'mgmt'), from which the traps, notifications, and events are routed to external network (SNMP Manager) from the node (SNMP Agent).

   — While configuring the SNMP target management station (ENM) receives and processes the SNMP notifications, the address-context information must be provided to the node added in the ENM.

   — In ENM, the address context value is defined in the Er6000ConnectivityInformation model through 'fmSendercontextname' attribute. It accepts string value.

   — The default value of the attribute '`fmSendercontextname`' is 'local'.

   — The context value can be provided either from Discovery or Add node commands through CM CLI in Er6000ConnectivityInformation model '`fmSendercontextname`' attribute.

   — After FM Supervision is enabled, node has the specified address-context and target destination configuration as below:

     For example if ENM IP Address is '131.160.219.162' and 'mgmt' is the address-context.

     Case SNMPV2:

     snmp target ENM 131.160.219.162 address-context "mgmt" security-name enm-public version 2c view ENM-View

     Case SNMPV3:

     snmp notify-target ENM 131.160.219.162 address-context "mgmt" tag Tag-Inform parameters Param2

2. Check the Firewall status.

   — Check from node end that the packets are routed to the ENM correctly.

   — Verify that there is no firewall blocking on real node to ENM or any network issues for routing the packets to ENM from the node.

   — If there are any issues related to firewall, then ENM internal alarm “Heartbeat Alarm” is raised.

7.4 Proper Subscription State Is Not Set on Network Element

This section describes how to troubleshoot the following issue:

— Network Element (NE) subscription state is in ENABLING or DISABLING state, which means that the alarm supervision on the NE is not properly ENABLED or DISABLED.

Prerequisites

— Access to Management Server (MS)

— Access to the involved Virtual Machines (VM)

Steps

1. Check the subscription state of the NE from CliApp using the following command:

   Note: CliApp can be accessed at `https://<apache_host_name>/#cliapp`

   ```
   $cmedit get NetworkElement=<nodename>,FmFunction=1
   ```

   Example

   ```
   &187;cmedit get NetworkElement=LTE02ERBS00015,FmFunction=1
   FDN : NetworkElement=LTE02ERBS00015,FmFunction=1
   FmFunctionId : 1
   alarmSuppressedState : false
   currentServiceState : IDLE
   lastUpdated : 1489146835692
   lastUpdatedTimeStamp : Fri Mar 10 11:53:55 GMT 2017
   subscriptionState : ENABLING
   technicianPresentState : false

   1 instance(s)


   cmedit get NetworkElement=LTE02ERBS00011,FmFunction=1
   FDN : NetworkElement=LTE02ERBS00011,FmFunction=1
   FmFunctionId : 1
   alarmSuppressedState : false
   currentServiceState : IN_SERVICE
   lastUpdated : 1489116318168
   lastUpdatedTimeStamp : Fri Mar 10 03:25:18 GMT 2017
   subscriptionState : DISABLING
   technicianPresentState : false

   1 instance(s)
   ```

2. Check the status of subscription state attribute.

   — If it is ENABLING, alarm supervision on NE is not properly enabled.

   — If it is DISABLING, alarm supervision on NE is not properly disabled.

The state of the NE is automatically recovered to ENABLED or DISABLED within a maximum duration of 34 minutes.

Results

Subscription state of a NE is set to ENABLED or DISABLED value.

7.5 Live Severity Count of Alarms Is Not Visible on Alarm Monitor GUI

Refer to this section if the live severity count of alarms is not updated in the Alarm Monitor GUI.

Prerequisites

— User requires basic knowledge of Linux and JBoss.

— On physical: Access to the Management Server (MS) and Virtual Machines (VMs).

— `wpserv` and `fmserv` are online.

— On Cloud:

- Access to the private key file for authentication (contact your OpenStack administrator).
- Access to the Ericsson Management Portal (EMP) VM in the ENM on Cloud deployment.

Steps

1. Check if Webpush channels are subscribed successfully in DEBUG logs of `wpserv`. Connect to the `wpserv` service as described in Connect to a Service on page 2 and view the `server.log` file.

   Example:

   ```
   [root@svc-1-wpserv~]# cd /ericsson/3pp/jboss/standalone/log
   [root@svc-1-wpserv log]# vi server.log
   ```

   If channels are not subscribed properly, exception messages are displayed in server log.

   The following code block shows server log that displays subscribed channels:

   ```
   2015-10-08 16:20:19,321 INFO [com.ericsson.oss.presentation.server.webpush.r →
   est.api.WebPushResource] (ajp-executor-threads - 2484) Web Push Rest - subsc →
   ribe push channels: provided subscriber ID: a9c65343-2d52-49b5-87a7-0624f9f8 →
   ```

```
f91b is registered. Subscription Timestamp: 1444317619318

2015-10-08 16:20:19,365 INFO [com.ericsson.oss.presentation.server.webpush.r →
est.api.WebPushResource] (ajp-executor-threads - 2484) Web Push Rest - subsc →
ribe push channels: subscribing 5 channels

2015-10-08 16:20:19,430 INFO [com.ericsson.oss.presentation.server.webpush.r →
est.api.WebPushResource] (ajp-executor-threads - 2484) Web Push Rest - subsc →
ribe push channels: 5 channels successfully subscribed
```

If Webpush channels are subscribed successfully, check if the events successfully reached `wpserv` VM.

Also check whether events reached the client.

If the events successfully reach `wpserv` VM, it is displayed as in the following example:

```
[root@svc-1-wpserv~]# /ericsson/3pp/jboss/bin/jboss-cli.sh -c

[standalone @localhost : 9999 /] /subsystem=logging/logger=com.ericsson.oss: →
change-log-level(level=DEBUG)

{ "outcome" =>  "success" }

HZ Received events (onMessage): 12

------------------------------------------------------------

Incoming events for Subscriber 6cd67dcf-9836-4f93-9606-82a470c0af8d : 9

Pushed events for Subscriber 6cd67dcf-9836-4f93-9606-82a470c0af8d : 9

Average socket writing time for Subscriber 6cd67dcf-9836-4f93-9606-82a470c0a →
f8d : 1.22

------------------------------------------------------------

Incoming events for Subscriber 0b164c1b-8d16-4c50-9467-7cc1b19c6c3c : 12

Pushed events for Subscriber 0b164c1b-8d16-4c50-9467-7cc1b19c6c3c : 12

Average socket writing time for Subscriber 0b164c1b-8d16-4c50-9467-7cc1b19c6 →
c3c : .25

------------------------------------------------------------

Total pushed events: 21

********** END WEB PUSH HEALTH MONITOR ****************
```

Else, exception messages are displayed.

If the events do not successfully reach `wpserv`, exit from `wpserv` VM, and go to the next step..

2. Connect to the `fmserv` service as described in Connect to a Service on page 2 and view the `server.log` file. Check if the channels are published properly by verifying DEBUG logs of `fmserv` VM. If the channels are not published, exception messages are displayed.

   The following example shows server log which displaysg published channels:

```
[root@svc-1-fmserv~]# cd /ericsson/3pp/jboss/standalone/log
```

```
[root@svc-1-fmservlog]# vi server.log


2015-10-13 05:33:57,606 INFO [com.ericsson.oss.services.alarmcontroldisplays →
ervice.webpush.broadcast.ChannelPublisher] (ServerService Thread Pool -- 85) →
 Published the nodeSummaryChannel

2015-10-13 05:33:57,672 INFO [com.ericsson.oss.services.alarmcontroldisplays →
ervice.webpush.broadcast.ChannelPublisher] (ServerService Thread Pool -- 85) →
 Published the alarmOperationChannel

2015-10-13 05:33:57,745 INFO [com.ericsson.oss.services.alarmcontroldisplays →
ervice.webpush.broadcast.ChannelPublisher] (ServerService Thread Pool -- 85) →
 Published the newAlarmChannel

2015-10-13 05:33:57,797 INFO [com.ericsson.oss.services.alarmcontroldisplays →
ervice.webpush.broadcast.ChannelPublisher] (ServerService Thread Pool -- 85) →
 Published the severityChannel

2015-10-13 05:33:57,853 INFO [com.ericsson.oss.services.alarmcontroldisplays →
ervice.webpush.broadcast.ChannelPublisher] (ServerService Thread Pool -- 85) →
 Published the alarmPolicyChannel
```

If channels are published properly, check if a Webpush event is created.

If a Webpush event is not created, exception messages are displayed.

The following example shows a server log which displays the successful creation of a Webpush event:

```
[root@svc-1-fmserv~]# cd /ericsson/3pp/jboss/standalone/log

[root@svc-1-fmserv log]# vi server.log


2015-10-13 05:33:57,606 INFO [com.ericsson.oss.services.alarmcontroldisplays →
ervice.webpush.broadcast.ChannelPublisher] (ServerService Thread Pool -- 85) →
 Published the nodeSummaryChannel

2015-10-13 05:33:57,672 INFO [com.ericsson.oss.services.alarmcontroldisplays →
ervice.webpush.broadcast.ChannelPublisher] (ServerService Thread Pool -- 85) →
 Published the alarmOperationChannel

2015-10-13 05:33:57,745 INFO [com.ericsson.oss.services.alarmcontroldisplays →
ervice.webpush.broadcast.ChannelPublisher] (ServerService Thread Pool -- 85) →
 Published the newAlarmChannel

2015-10-13 05:33:57,797 INFO [com.ericsson.oss.services.alarmcontroldisplays →
ervice.webpush.broadcast.ChannelPublisher] (ServerService Thread Pool -- 85) →
 Published the severityChannel

2015-10-13 05:33:57,853 INFO [com.ericsson.oss.services.alarmcontroldisplays →
ervice.webpush.broadcast.ChannelPublisher] (ServerService Thread Pool -- 85) →
 Published the alarmPolicyChannel

Results
```

Results

User is able to troubleshoot the issue when live severity count is not updated in Alarm Monitor GUI.

7.6 Unable to Perform Alarm Actions through the Fault Management UI

Refer to this section when the alarm actions cannot be performed through the Fault Management (FM) GUI.

Prerequisites

— User requires basic knowledge of Linux and JBoss.

— On physical: Access to the Management Server (MS) and Virtual Machines (VMs).

— On Cloud:

- Access to the private key file for authentication (contact your OpenStack administrator).
- Access to the Ericsson Management Portal (EMP) VM in the ENM on Cloud deployment.

Steps

1. Verify if fmserv VMs are online.

    a. Log on to the ENM Management Server and connect to any SVC node or to ENM Management Portal (EMP) server on Cloud environment.

    ```
    [root@ms-1 ~]# ssh litp-admin@cloud-svc-1
    [litp-admin@cloud-svc-1 ~]$su root
    Password:
    [root@cloud-svc-1 litp-admin]#
    ```

    b. Execute the following command to check whether VMs are online.

    ```
     $ hastatus -sum |grep <VMname>
    ```

    Example

    ```
    [root@cloud-svc-1 litp-admin]# hastatus -sum |grep fmserv
    B  Grp_CS_svc_cluster_fmserv cloud-svc-3          Y          N       →
             ONLINE
    B  Grp_CS_svc_cluster_fmserv cloud-svc-4          Y          N       →
             ONLINE
    ```

    If VMs are not online, the issue is with the Deployment of VM.

    Solution: Execute the following command to online VMs.

    ```
    $ hagrp -online <groupName> -sys <MasterServerName>
    ```

Example

```
[root@svc-1 litp-admin]# hagrp -online Grp_CS_svc_cluster_fmserv -s
ys svc-3
```

If VMs are online, proceed as follows:

Check hostname of VMs:

On physical, use `use cat /etc/hostname | grep <VMname>`:

```
[root@cloud-svc-1 litp-admin]# cat /etc/hosts | grep f
mserv
<ipAddress>  svc-3-fmserv    # Created by LITP. Please
 do not edit
<ipAddress>  svc-4-fmserv    # Created by LITP. Please
 do not edit
```

On Cloud, use `consul members | grep <VMname>` as the root user on EMP.

```
[root@gat-emp-0 data]# consul members | grep fmserv
gat-fmserv-0            <ipAddress>:8301   alive   cl
ient  0.8.1  2         dc1
gat-fmserv-1            <ipAddress>:8301   alive   cl
ient  0.8.1  2         dc1
[root@gat-emp-0 data]#
```

To connect to `fmserv` VM instance, see Connect to a Service on page 2.

c. Execute the following command to check if the JBoss service is running.

```
$ service jboss  monitor
```

Solution: If JBoss is not running, start JBoss by executing the following command.

```
$ service jboss start
```

If JBoss service is running properly, then proceed with the following steps.

d. Check whether the control display service ear is deployed properly.

If not, restart the JBoss service by executing the following commands:

— Stop the JBoss service:

```
$ service jboss stop
```

— Start the JBoss service:

```
$ service jboss start
```

e. Verify `fmserv` VM DEBUG logs if any exception has occurred.

If logs are proper, then exit from fmserv VM and go to step 2.

2. Perform the same procedure for `fmhistory` VMs.

a. Check whether AlarmActionService ear is deployed properly.

If not, restart the JBoss service by executing the following commands:

— Stop the JBoss service:

```
$ service jboss stop
```

— Start the JBoss service:

```
$ service jboss start
```

b. Verify `fmhistory` VM DEBUG logs.

```
[root@svc-1-fmhistory~]# cd /ericsson/3pp/jboss/standalone/log
[root@svc-1-fmhistory log]# vi server.log
```

c. When operations like Ack, Unack, Clear, or Comment are performed, then check whether the respective operation is successfully performed on alarm or not.

The following code block displays the log messages depending upon operation:

```
ACK - ACKNOWLEDGE operation is performed on the alarm with AlarmNum →
ber
UnACK - UN-ACKNOWLEDGE operation is performed on the alarm with Ala →
rmNumber
CLEAR - CLEAR operation is performed on the alarm with AlarmNumber
COMMENT - COMMENT is added on the alarm with AlarmNumber
```

d. Check for ERROR logs in the fmserv and fmhistory server logs if the alarm operation fails.

Results

User is able to troubleshoot the issue if the user is unable to perform alarm actions through GUI.

7.7 OPI Not Displayed in Alarm Monitor

This section describes how to troubleshoot the following issue:

— OPI for an alarm is not displayed in Alarm Monitor

To solve this issue, check that the correct CPI library is loaded on the ENM server, and that it is correctly configured for the relevant node type or management system.

Prerequisites

Physical Deployments

Access to the Management Server (MS)

Cloud Deployments

Access to httpd VM. See Connect to a Service on page 2.

Steps

1. Open the ENM CLI:

```
https://<apache_host_name>/#cliapp
```

   a. From the ENM CLI, check the `AlarmCpiDetails` associations:

```
fmedit get * AlarmCpiDetails
```

   The output lists the associations between the CPI library name and the node type (for example, ERBS) or the management system (for example, ENM), as follows:

```
fmedit get * AlarmCpiDetails

Target Type                Target Model Identity        Library Nam →
e                   Alarm Page Title

ENM                              lzn7030205.alx

ERBS                              en_lzn7931010.alx

RadioNode                          en_lzn7931020.alx

SGSN-MME                          en_lzn7040197.alx


Total number of Alarm OPI mapping found is : 4
```

   For the following nodes, the release attribute is mandatory:

   — HLR-FE (HLR-FE (APZ), HLR-FE (IS), HLR-FE (BSP), vHLR-FE)

— MSC (MSC-S DB, IP-STP, Compact MSC-S, Compact IP-STP, vMSC-S Compact, vIP-STP Compact, MSC-S BC (IS), MSC-S BC (BSP), vMSC-S HC)

Example:

```
fmedit get * AlarmCpiDetails.(targetType=="HLR-FE")
fmedit get * AlarmCpiDetails.(targetType=="MSC-DB-BSP")
```

The output lists:

```
Target Type         Target Model Identity         Library Name       →
    Alarms Page Title         Release
HLR-FE              4322-436-393             en_lzn7931010.alx       →
                1.9
Total number of Alarm OPI mapping found is : 1


fmedit get * AlarmCpiDetails.(targetType=="MSC-DB-BSP")
Target Type         Target Model Identity         Library Name       →
    Alarms Page Title         Release
MSC-DB-BSP          4322-436-393              en_lzn7931010.alx       →
               18A
Total number of Alarm OPI mapping found is : 1
```

If the table does not report the relevant target type, ensure that the procedure described in *Make CPI Libraries Available*, from the *ENM Launcher* section of the *ENM System Administrator Guide, 1/1543-AOM 901 151* is completely carried out. If the target type is listed, check if the associated library name is correct.

There must be only one entry per `Target Type` . If there are more entries, check that each one has a different `Target Model Identity` value. The "ENM" `Target Type` never has an associated `Target Model Identity`.

In case of multiple instances, delete the entries that refer to the older libraries using the `fmedit delete` command. Refer to the *WEB CLI online help* for more details.

2. Log on as described for Physical or Cloud and run the commands.

   — Log on to the MS (on Physical) as the `litp-admin` user and switch to the `root` user.

   — Log on to the `httpd` VM (on Cloud) as `root` user.

```
# cd /ericsson/enm/alex/libraries/
# ls -l
total 43557
-rw-r--r--. 1  502 502  1302052 Mar  9 15:28 alex_help.ahx
-rwxr-xr--. 1 root 502  3190300 May  5 14:44 en_lzn7040197.alx
-rwxr-xr--. 1 root 502 23548626 May  5 14:44 en_lzn7931010.alx
-rwxr-xr--. 1 root 502  3690155 May  5 14:44 en_lzn7931010.alx
-rwxr-xr--. 1 root 502  3954739 May  5 14:44 lzn7030205.alx
```

Check that the `.alx` files present in the directory are the same as reported in Step 1 .a, otherwise run the proper commands to correct the assignments.

3. If some alarms are found in the CPI library while others are not, check that the loaded CPI library is up to date.

    a. Connect to the ENM launcher:

    ```
    https://<apache_host_name>/#launcher
    ```

    b. To display a list of loaded libraries and version information, click the **Elex Library** link.

    Check that the versions are up-to-date. If not, obtain and install latest libraries for each type using information in the *ENM System Administrator Guide, 1/1543-AOM 901 151* and *Web CLI Online Help* for AlarmCpiDetails configuration.

    Check that the libraries associated with `AlarmCpiDetails` objects contain OPI (alarm Operating Instructions) and troubleshooting information.

4. If the corresponding OPI is not found for a specific alarm, check the relevant Node CPI library.

    a. Connect to the ENM launcher:

    ```
    https://<apache_host_name>/#launcher
    ```

    b. To display a list of loaded libraries and version information, click the **Elex Library** link.

    c. Open the relevant Node CPI library.

    d. Search for the *<Specific Problem>* of the alarm information using the Search box.

7.7.1 OPI Not Displayed in Alarm Monitor (IS-based Nodes)

This troubleshooting applies only to IS variants (MSC-BC-IS and HLR-FE-IS neTypes) of MSC and HLR-FE nodes that have multiple ELEX libraries.

Diagnostics

After multiple ELEX libraries have been configured, an error message is displayed similar to the following:

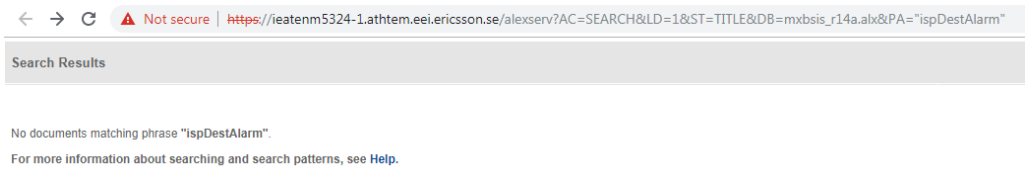

Figure 1 No documents matching phrase...

Configure the Value of the Alarms Page Title

1. Configure the value of the `Alarms Page Title` so that the OPI directly references the title of the page from where an operator can search for remedial instructions for a specific problem.

   Example

   ```
   fmedit set * AlarmCpiDetails.(targetType=="MSC-BC-IS",libraryName=="mxbsis_r
   14a.alx",componentName=="SIS") alarmsPageTitle="SIS System Notifications"
   ```

   Result: In this example, the `Alarms Page Title` is configured and OPI launches SIS System Notifications. A user can now perform a manual search for a specific problem.

Figure 2 Configuration

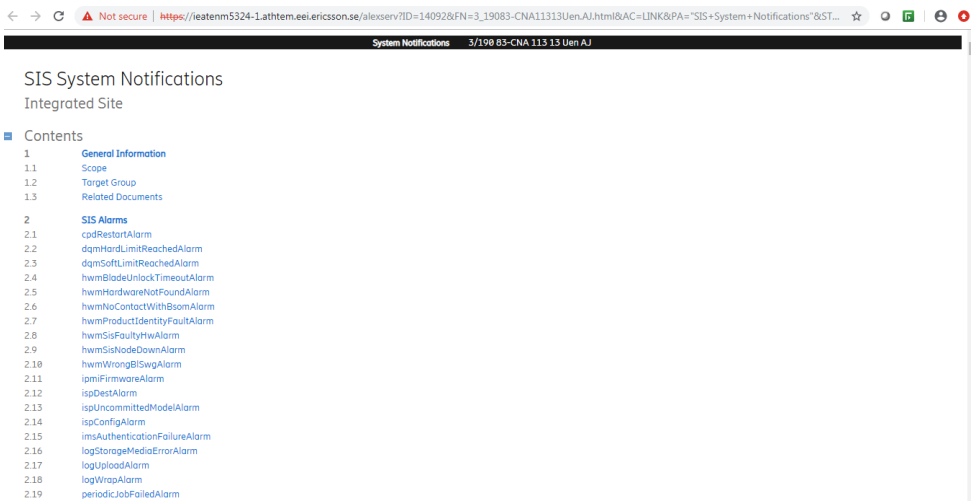

Figure 3 OPI Launched

7.8 Historical Alarms Not Storing in Solr

> FM historical alarms are not getting loaded by the DLMS Service group into the collection1 data store on the Solr Service group.

Prerequisites

— Physical deployment:

- Access to the Management Server (MS).
- Access to the Service Cluster nodes.
- Access to the ENM Command Line Interface (CLI).

— Cloud deployment:

- Access to the private key file for authentication. Contact your OpenStack administrator.
- Access to the Ericsson Management Portal (EMP) VM in the ENM on Cloud deployment.
- Access to the ENM Command Line Interface (CLI).

Cause

Alarm JSON files in the directory `/ericsson/enm/dlms/history/data` are not getting loaded into the data store collection1 in the Solr Service group. The files do not exist in the directory if they are getting successfully loaded into the data store location. This can be defined in the order of minutes.

Verify That Solr Is Running

— Connect to the Solr VM (see Connect to a Service on page 2) and check that Solr is running:

```
[root@solr ~]# service solr status
SOLR: INFORMATION (): Solr script called with the argument:status
 (pid  2187) is running...
SOLR: INFORMATION (): Current status of solr is: 0
```

Verify That collection1 Core Started Successfully

— Check in Solr that `collection1` core started successfully:

```
curl "http://localhost:8983/solr/admin/cores?action=STATUS&core=collection1&wt=json&indent=on&omitHeader=on"
```

Example response:

```
[root@solr ~]# curl "http://localhost:8983/solr/admin/cores?action=STATUS&core=collection1&wt=json&indent=on&omitHeader=on"
{
  "initFailures":{},
  "status":{
    "collection1":{
      "name":"collection1",
      "isDefaultCore":true,
      "instanceDir":"/ericsson/solr/install/solr-4.8.1/enm/solr/collection1/",
      "dataDir":"/ericsson/enm/solr/data/",
      "config":"solrconfig.xml",
      "schema":"managed-schema",
      "startTime":"2018-05-15T19:57:12.488Z",
      "uptime":577754814,
      "index":{
        "numDocs":263187685,
        "maxDoc":263187685,
        "deletedDocs":0,
        "indexHeapUsageBytes":293137167,
        "version":66519,
        "segmentCount":67,
        "current":true,
        "hasDeletions":false,
        "directory":"org.apache.lucene.store.NIOFSDirectory:NIOFSDirectory@/ericsson/enm/solr/data/index lockFactory=SimpleFSLockFactory@/ericsson/enm/solr/data/index",
```

```
        "userData":{"commitTimeMSec":"1526991907389"},
        "lastModified":"2018-05-22T12:25:07.389Z",
        "sizeInBytes":202025797191,
        "size":"188.15 GB"}}}}
[root@solr ~]#]
```

Check the Solr Logs for Errors

1. Get the Solr logs by connecting VM:

```
/ericsson/solr/install/solr-4.8.1/logs/solr* on VM
/var/log/messages on solr VM
```

2. Check the logs for errors and exceptions.

   The following exceptions can be seen when there is segment corruption in Solr:

   — `org.apache.solr.common.SolrException;null:org.eclipse.jetty.io.EofException`

   — `Caused by: java.lang.ArrayIndexOutOfBoundsException`

   To repair the corruption, follow the steps in Historical Alarms Database Issue on page 56.

Verify That JBoss Is Running

— Log on to the dlms (master instance) VM as root user (see Connect to a Service on page 2) and check that JBoss is running:

```
[root@svc-1-dlms ~]# service jboss status
jboss-as is running
```

Check the JSON File Permissions

— Check on the DLMS that the file permissions are correct for the JSON files in `/ericsson/enm/dlms/history/data`.

  Permissions needed by `jboss_user`:

```
[root@svc-1-dlms cloud-user]# ls -lart /ericsson/enm/dlms/history/data/
total 2
-rw-r--r--. 1 jboss_user jboss 1291 Jun 21 10:26 changes_create_svc →
-2-fmhistory_21-06-2018_10-26-53.json
drwxrwx---. 2 root fmhistoryusers 96 Jun 19 04:32 lost+found
[root@svc-5-dlms data]#
```

Check the JBoss Log File for Errors

1. Enable the DEBUG log level:

```
[root@svc-1-dlms cloud-user]# /ericsson/3pp/jboss/bin/jboss-cli.sh -c
[standalone@localhost:9999 /] /subsystem=logging/logger=com.ericsson.oss.ser →
 →
vices.fm.alarmhistory:add(level=DEBUG)
{"outcome" => "success"}
```

2. Check the JBoss log file `/ericsson/3pp/jboss/standalone/log/server.log` for any errors or exceptions.

7.9 Historical Alarms Database Issue

> Check for and repair any issues with the historical alarms stored in Solr Collection1.

Prerequisites

— Physical deployment:

- Access to the Management Server (MS)
- Access to the Service Cluster nodes
- Access to the ENM Command-Line Interface (CLI)

— Cloud deployment:

- Access to the private key file for authentication

  Contact your OpenStack administrator.
- Access to the Ericsson Management Portal (EMP) VM in the ENM on Cloud deployment
- Access to the ENM Command-Line Interface (CLI)

Check for Index Corruptions

1. Connect to the `solr` VM instance.

   For details, see Connect to a Service on page 2.
2. Go to the directory `/ericsson/solr/install/`.
3. Run the script `SolrIndexUtility.sh` with the `checkIndex` parameter.

```
[root@solr install]# ./SolrIndexUtility.sh checkIndex /ericsson/enm/solr/dat →
a/index
```

4. Check the script report for any corrupted indexes.

   — If there are corrupted indexes (`CorruptedIndexException` in the report), go to step 5.

   — If not, go to step 6.

   Example

   Without any problems:

```
[root@solr install]# ./SolrIndexUtility.sh checkIndex /ericsson/enm/solr/dat →
a/index

Script checks for corruption of Solr indexes and may possibly repair as well →
.
Please note that it is always recommended to take a backup of existing index →
 and run this script to repair with repairIndex option passed as argument.
checkIndex called with args: /ericsson/enm/solr/data/index
Opening index @ /ericsson/enm/solr/data/index
Segments file=segments_a numSegments=9 version=4.8 format= userData={commitT →
imeMSec=1529559901108}
  1 of 9: name=_0 docCount=2
    codec=Lucene46
    compound=false
    numFiles=11
    size (MB)=0.011
    diagnostics = {timestamp=1529434801449, os=Linux, os.version=2.6.32-504. →
64.1.el6.x86_64, source=flush, lucene.version=4.8.1 1594670 - rmuir - 2014-0 →
5-14 19:22:52, os.arch=amd64, java.version=1.7.0_131, java.vendor=Oracle Cor →
poration}
    no deletions
    test: open reader.........OK
    test: check integrity.....OK
    test: check live docs.....OK
    test: fields..............OK [35 fields]
    test: field norms.........OK [3 fields]
    test: terms, freq, prox...OK [40 terms; 58 terms/docs pairs; 0 tokens]
    test: stored fields.......OK [70 total field count; avg 35 fields per do →
c]
    test: term vectors........OK [0 total vector count; avg 0 term/freq vect →
or fields per doc]
    test: docvalues...........OK [27 docvalues fields; 0 BINARY; 8 NUMERIC;  →
19 SORTED; 0 SORTED_SET]
  2 of 9: name=_1 docCount=18
    codec=Lucene46
    compound=false
    numFiles=11
    size (MB)=0.015
    diagnostics = {timestamp=1529435101854, os=Linux, os.version=2.6.32-504. →
64.1.el6.x86_64, source=flush, lucene.version=4.8.1 1594670 - rmuir - 2014-0 →
5-14 19:22:52, os.arch=amd64, java.version=1.7.0_131, java.vendor=Oracle Cor →
poration}
    no deletions
    test: open reader.........OK
    test: check integrity.....OK
    test: check live docs.....OK
    test: fields..............OK [37 fields]
    test: field norms.........OK [5 fields]
    test: terms, freq, prox...OK [118 terms; 534 terms/docs pairs; 0 tokens]
    test: stored fields.......OK [630 total field count; avg 35 fields per d →
oc]
    test: term vectors........OK [0 total vector count; avg 0 term/freq vect →
or fields per doc]
    test: docvalues...........OK [29 docvalues fields; 0 BINARY; 10 NUMERIC; →
 19 SORTED; 0 SORTED_SET]

  3 of 9: name=_2 docCount=36
    codec=Lucene46
```

```
    compound=false
    numFiles=11
    size (MB)=0.019
    diagnostics = {timestamp=1529435401392, os=Linux, os.version=2.6.32-504. →
64.1.el6.x86_64, source=flush, lucene.version=4.8.1 1594670 - rmuir - 2014-0 →
5-14 19:22:52, os.arch=amd64, java.version=1.7.0_131, java.vendor=Oracle Cor →
poration}
    no deletions
    test: open reader.........OK
    test: check integrity.....OK
    test: check live docs.....OK
    test: fields..............OK [37 fields]
    test: field norms.........OK [5 fields]
    test: terms, freq, prox...OK [181 terms; 1068 terms/docs pairs; 0 tokens →
]
    test: stored fields.......OK [1260 total field count; avg 35 fields per  →
doc]
    test: term vectors........OK [0 total vector count; avg 0 term/freq vect →
or fields per doc]
    test: docvalues...........OK [29 docvalues fields; 0 BINARY; 10 NUMERIC; →
 19 SORTED; 0 SORTED_SET]

  4 of 9: name=_3 docCount=1
    codec=Lucene46
    compound=false
    numFiles=11
    size (MB)=0.009
    diagnostics = {timestamp=1529436601895, os=Linux, os.version=2.6.32-504. →
64.1.el6.x86_64, source=flush, lucene.version=4.8.1 1594670 - rmuir - 2014-0 →
5-14 19:22:52, os.arch=amd64, java.version=1.7.0_131, java.vendor=Oracle Cor →
poration}
    no deletions
    test: open reader.........OK
    test: check integrity.....OK
    test: check live docs.....OK
    test: fields..............OK [33 fields]
    test: field norms.........OK [3 fields]
    test: terms, freq, prox...OK [28 terms; 28 terms/docs pairs; 0 tokens]
    test: stored fields.......OK [33 total field count; avg 33 fields per do →
c]
    test: term vectors........OK [0 total vector count; avg 0 term/freq vect →
or fields per doc]
    test: docvalues...........OK [25 docvalues fields; 0 BINARY; 8 NUMERIC;  →
17 SORTED; 0 SORTED_SET]

  5 of 9: name=_4 docCount=1
    codec=Lucene46
    compound=false
    numFiles=9
    size (MB)=0.001
    diagnostics = {timestamp=1529558164181, os=Linux, os.version=2.6.32-504. →
64.1.el6.x86_64, source=flush, lucene.version=4.8.1 1594670 - rmuir - 2014-0 →
5-14 19:22:52, os.arch=amd64, java.version=1.7.0_131, java.vendor=Oracle Cor →
poration}
    no deletions
    test: open reader.........OK
    test: check integrity.....OK
    test: check live docs.....OK
    test: fields..............OK [3 fields]
    test: field norms.........OK [1 fields]
    test: terms, freq, prox...OK [3 terms; 3 terms/docs pairs; 0 tokens]
    test: stored fields.......OK [3 total field count; avg 3 fields per doc]
    test: term vectors........OK [0 total vector count; avg 0 term/freq vect →
or fields per doc]
    test: docvalues...........OK [0 docvalues fields; 0 BINARY; 0 NUMERIC; 0 →
 SORTED; 0 SORTED_SET]

  6 of 9: name=_5 docCount=36
    codec=Lucene46
    compound=false
    numFiles=11
    size (MB)=0.02
    diagnostics = {timestamp=1529558701589, os=Linux, os.version=2.6.32-504. →
64.1.el6.x86_64, source=flush, lucene.version=4.8.1 1594670 - rmuir - 2014-0 →
5-14 19:22:52, os.arch=amd64, java.version=1.7.0_131, java.vendor=Oracle Cor →
poration}
    no deletions
    test: open reader.........OK
    test: check integrity.....OK
    test: check live docs.....OK
```

```
    test: fields.............OK [36 fields]
    test: field norms.........OK [4 fields]
    test: terms, freq, prox...OK [213 terms; 1057 terms/docs pairs; 0 tokens →
]
    test: stored fields.......OK [1258 total field count; avg 34.944 fields  →
per doc]
    test: term vectors........OK [0 total vector count; avg 0 term/freq vect →
or fields per doc]
    test: docvalues...........OK [28 docvalues fields; 0 BINARY; 9 NUMERIC;  →
19 SORTED; 0 SORTED_SET]

  7 of 9: name=_6 docCount=3
    codec=Lucene46
    compound=false
    numFiles=11
    size (MB)=0.01
    diagnostics = {timestamp=1529559001556, os=Linux, os.version=2.6.32-504. →
64.1.el6.x86_64, source=flush, lucene.version=4.8.1 1594670 - rmuir - 2014-0 →
5-14 19:22:52, os.arch=amd64, java.version=1.7.0_131, java.vendor=Oracle Cor →
poration}
    no deletions
    test: open reader.........OK
    test: check integrity.....OK
    test: check live docs.....OK
    test: fields.............OK [35 fields]
    test: field norms.........OK [4 fields]
    test: terms, freq, prox...OK [47 terms; 90 terms/docs pairs; 0 tokens]
    test: stored fields.......OK [105 total field count; avg 35 fields per d →
oc]
    test: term vectors........OK [0 total vector count; avg 0 term/freq vect →
or fields per doc]
    test: docvalues...........OK [27 docvalues fields; 0 BINARY; 9 NUMERIC;  →
18 SORTED; 0 SORTED_SET]

  8 of 9: name=_7 docCount=2
    codec=Lucene46
    compound=false
    numFiles=11
    size (MB)=0.01
    diagnostics = {timestamp=1529559301508, os=Linux, os.version=2.6.32-504. →
64.1.el6.x86_64, source=flush, lucene.version=4.8.1 1594670 - rmuir - 2014-0 →
5-14 19:22:52, os.arch=amd64, java.version=1.7.0_131, java.vendor=Oracle Cor →
poration}
    no deletions
    test: open reader.........OK
    test: check integrity.....OK
    test: check live docs.....OK
    test: fields.............OK [35 fields]
    test: field norms.........OK [4 fields]
    test: terms, freq, prox...OK [37 terms; 60 terms/docs pairs; 0 tokens]
    test: stored fields.......OK [70 total field count; avg 35 fields per do →
c]
    test: term vectors........OK [0 total vector count; avg 0 term/freq vect →
or fields per doc]
    test: docvalues...........OK [27 docvalues fields; 0 BINARY; 9 NUMERIC;  →
18 SORTED; 0 SORTED_SET]

  9 of 9: name=_8 docCount=2
    codec=Lucene46
    compound=false
    numFiles=11
    size (MB)=0.01
    diagnostics = {timestamp=1529559901311, os=Linux, os.version=2.6.32-504. →
64.1.el6.x86_64, source=flush, lucene.version=4.8.1 1594670 - rmuir - 2014-0 →
5-14 19:22:52, os.arch=amd64, java.version=1.7.0_131, java.vendor=Oracle Cor →
poration}
    no deletions
    test: open reader.........OK
    test: check integrity.....OK
    test: check live docs.....OK
    test: fields.............OK [35 fields]
    test: field norms.........OK [4 fields]
    test: terms, freq, prox...OK [37 terms; 60 terms/docs pairs; 0 tokens]
    test: stored fields.......OK [70 total field count; avg 35 fields per do →
c]
    test: term vectors........OK [0 total vector count; avg 0 term/freq vect →
or fields per doc]
    test: docvalues...........OK [27 docvalues fields; 0 BINARY; 9 NUMERIC;  →
18 SORTED; 0 SORTED_SET]
```

```
No problems were detected with this index.
```

Example

With broken segments:

```
[root@solr index]# /ericsson/solr/install/SolrIndexUtility.sh checkIndex /er →
icsson/enm/solr/data/index/
Script checks for corruption of Solr indexes and may possibly repair as well →
.
Please note that it is always recommended to take a backup of existing index →
 and run this script to repair with repairIndex option passed as argument.
checkIndex called with args: /ericsson/enm/solr/data/index/

Opening index @ /ericsson/enm/solr/data/index/

Segments file=segments_30tv numSegments=56 version=4.8 format= userData={com →
mitTimeMSec=1554355374487}

1 of 9: name=_33ka docCount=5645513
codec=Lucene46
compound=false
numFiles=12
size (MB)=5,068.893
diagnostics = {os=Linux, java.vendor=Oracle Corporation, java.version=1.8.0_ →
172, lucene.version=4.8.1 1594670 - rmuir - 2014-05-14 19:22:52, mergeMaxNum →
Segments=-1, os.arch=amd64, source=merge, mergeFactor=8, os.version=2.6.32-7 →
54.9.1.el6.x86_64, timestamp=1552510838482}
has deletions [delGen=6]
test: open reader.........OK
test: check integrity.....OK
test: check live docs.....OK [206 deleted docs]
test: fields..............OK [48 fields]
test: field norms.........OK [5 fields]
test: terms, freq, prox...OK [16710371 terms; 180447783 terms/docs pairs; 0  →
tokens]
test (ignoring deletes): terms, freq, prox...OK [16711506 terms; 180454137 t →
erms/docs pairs; 0 tokens]
test: stored fields.......OK [212457867 total field count; avg 37.634 fields →
 per doc]
test: term vectors........OK [0 total vector count; avg 0 term/freq vector f →
ields per doc]
test: docvalues...........OK [34 docvalues fields; 0 BINARY; 10 NUMERIC; 24  →
SORTED; 0 SORTED_SET]

2 of 9: name=_3417 docCount=5209113
codec=Lucene46
compound=false
numFiles=12
size (MB)=4,969.003
diagnostics = {os=Linux, java.vendor=Oracle Corporation, java.version=1.8.0_ →
172, lucene.version=4.8.1 1594670 - rmuir - 2014-05-14 19:22:52, mergeMaxNum →
Segments=-1, os.arch=amd64, source=merge, mergeFactor=10, os.version=2.6.32- →
754.9.1.el6.x86_64, timestamp=1552593604463}
has deletions [delGen=4]
test: open reader.........OK
test: check integrity.....FAILED
WARNING: fixIndex() would remove reference to this segment; full exception:
org.apache.lucene.index.CorruptIndexException: checksum failed (hardware pro →
blem?) : expected=ffdbce3 actual=afbe03ac (resource=BufferedChecksumIndexInp →
ut(MMapIndexInput(path="/ericsson/enm/solr/data/index/_3417.fdt")))
at org.apache.lucene.codecs.CodecUtil.checkFooter(CodecUtil.java:211)
at org.apache.lucene.codecs.CodecUtil.checksumEntireFile(CodecUtil.java:268)
at org.apache.lucene.codecs.compressing.CompressingStoredFieldsReader.checkI →
ntegrity(CompressingStoredFieldsReader.java:539)
at org.apache.lucene.index.SegmentReader.checkIntegrity(SegmentReader.java:5 →
77)
at org.apache.lucene.index.CheckIndex.checkIndex(CheckIndex.java:561)
at org.apache.lucene.index.CheckIndex.main(CheckIndex.java:1967)

3 of 9: name=_34l8 docCount=5441870
codec=Lucene46
compound=false
numFiles=12
size (MB)=5,041.22
diagnostics = {os=Linux, java.vendor=Oracle Corporation, java.version=1.8.0_ →
```

```
172, lucene.version=4.8.1 1594670 - rmuir - 2014-05-14 19:22:52, mergeMaxNum →
Segments=-1, os.arch=amd64, source=merge, mergeFactor=10, os.version=2.6.32- →
754.9.1.el6.x86_64, timestamp=1552691428620}
has deletions [delGen=3]
test: open reader.........OK
test: check integrity.....OK
test: check live docs.....OK [139 deleted docs]
test: fields..............OK [47 fields]
test: field norms.........OK [5 fields]
test: terms, freq, prox...OK [16159483 terms; 173439737 terms/docs pairs; 0  →
tokens]
test (ignoring deletes): terms, freq, prox...OK [16160238 terms; 173443892 t →
erms/docs pairs; 0 tokens]
test: stored fields.......OK [204182761 total field count; avg 37.522 fields →
 per doc]
test: term vectors........OK [0 total vector count; avg 0 term/freq vector f →
ields per doc]
test: docvalues...........OK [34 docvalues fields; 0 BINARY; 10 NUMERIC; 24  →
SORTED; 0 SORTED_SET]

4 of 9: name=_3545 docCount=5103807
codec=Lucene46
compound=false
numFiles=12
size (MB)=5,043.2
diagnostics = {os=Linux, java.vendor=Oracle Corporation, java.version=1.8.0_ →
172, lucene.version=4.8.1 1594670 - rmuir - 2014-05-14 19:22:52, mergeMaxNum →
Segments=-1, os.arch=amd64, source=merge, mergeFactor=10, os.version=2.6.32- →
754.9.1.el6.x86_64, timestamp=1552784433448}
has deletions [delGen=3]
test: open reader.........OK
test: check integrity.....OK
test: check live docs.....OK [128 deleted docs]
test: fields..............OK [47 fields]
test: field norms.........OK [5 fields]
test: terms, freq, prox...OK [15250150 terms; 162754894 terms/docs pairs; 0  →
tokens]
test (ignoring deletes): terms, freq, prox...OK [15250860 terms; 162758718 t →
erms/docs pairs; 0 tokens]
test: stored fields.......OK [191434492 total field count; avg 37.509 fields →
 per doc]
test: term vectors........OK [0 total vector count; avg 0 term/freq vector f →
ields per doc]
test: docvalues...........OK [34 docvalues fields; 0 BINARY; 10 NUMERIC; 24  →
SORTED; 0 SORTED_SET]

5 of 9: name=_35og docCount=5747124
codec=Lucene46
compound=false
numFiles=12
size (MB)=5,057.196
diagnostics = {os=Linux, java.vendor=Oracle Corporation, java.version=1.8.0_ →
172, lucene.version=4.8.1 1594670 - rmuir - 2014-05-14 19:22:52, mergeMaxNum →
Segments=-1, os.arch=amd64, source=merge, mergeFactor=9, os.version=2.6.32-7 →
54.9.1.el6.x86_64, timestamp=1552883705898}
has deletions [delGen=3]
test: open reader.........OK
test: check integrity.....FAILED
WARNING: fixIndex() would remove reference to this segment; full exception:
org.apache.lucene.index.CorruptIndexException: checksum failed (hardware pro →
blem?) : expected=f222bb93 actual=c576d64d (resource=BufferedChecksumIndexIn →
put(MMapIndexInput(path="/ericsson/enm/solr/data/index/_35og.fdt")))
at org.apache.lucene.codecs.CodecUtil.checkFooter(CodecUtil.java:211)
at org.apache.lucene.codecs.CodecUtil.checksumEntireFile(CodecUtil.java:268)
at org.apache.lucene.codecs.compressing.CompressingStoredFieldsReader.checkI →
ntegrity(CompressingStoredFieldsReader.java:539)
at org.apache.lucene.index.SegmentReader.checkIntegrity(SegmentReader.java:5 →
77)
at org.apache.lucene.index.CheckIndex.checkIndex(CheckIndex.java:561)
at org.apache.lucene.index.CheckIndex.main(CheckIndex.java:1967)

6 of 9: name=_3654 docCount=4902552
codec=Lucene46
compound=false
numFiles=12
size (MB)=5,034.75
diagnostics = {os=Linux, java.vendor=Oracle Corporation, java.version=1.8.0_ →
172, lucene.version=4.8.1 1594670 - rmuir - 2014-05-14 19:22:52, mergeMaxNum →
Segments=-1, os.arch=amd64, source=merge, mergeFactor=10, os.version=2.6.32- →
754.9.1.el6.x86_64, timestamp=1552964709134}
```

```
has deletions [delGen=3]
test: open reader.........OK
test: check integrity.....OK
test: check live docs.....OK [118 deleted docs]
test: fields..............OK [47 fields]
test: field norms.........OK [5 fields]
test: terms, freq, prox...OK [14915002 terms; 156903507 terms/docs pairs; 0 →
tokens]
test (ignoring deletes): terms, freq, prox...OK [14915663 terms; 156907032 t →
erms/docs pairs; 0 tokens]
test: stored fields.......OK [184537063 total field count; avg 37.642 fields →
 per doc]
test: term vectors........OK [0 total vector count; avg 0 term/freq vector f →
ields per doc]
test: docvalues...........OK [34 docvalues fields; 0 BINARY; 10 NUMERIC; 24 →
SORTED; 0 SORTED_SET]

7 of 9: name=_36ob docCount=5722406
codec=Lucene46
compound=false
numFiles=12
size (MB)=5,059.659
diagnostics = {os=Linux, java.vendor=Oracle Corporation, java.version=1.8.0_ →
172, lucene.version=4.8.1 1594670 - rmuir - 2014-05-14 19:22:52, mergeMaxNum →
Segments=-1, os.arch=amd64, source=merge, mergeFactor=10, os.version=2.6.32- →
754.9.1.el6.x86_64, timestamp=1553058929673}
has deletions [delGen=3]
test: open reader.........OK
test: check integrity.....OK
test: check live docs.....OK [133 deleted docs]
test: fields..............OK [47 fields]
test: field norms.........OK [5 fields]
test: terms, freq, prox...OK [16672261 terms; 182633808 terms/docs pairs; 0 →
tokens]
test (ignoring deletes): terms, freq, prox...OK [16672997 terms; 182637783 t →
erms/docs pairs; 0 tokens]
test: stored fields.......OK [214687501 total field count; avg 37.518 fields →
 per doc]
test: term vectors........OK [0 total vector count; avg 0 term/freq vector f →
ields per doc]
test: docvalues...........OK [34 docvalues fields; 0 BINARY; 10 NUMERIC; 24 →
SORTED; 0 SORTED_SET]

8 of 9: name=_376e docCount=5583006
codec=Lucene46
compound=false
numFiles=12
size (MB)=5,072.371
diagnostics = {os=Linux, java.vendor=Oracle Corporation, java.version=1.8.0_ →
172, lucene.version=4.8.1 1594670 - rmuir - 2014-05-14 19:22:52, mergeMaxNum →
Segments=-1, os.arch=amd64, source=merge, mergeFactor=9, os.version=2.6.32-7 →
54.9.1.el6.x86_64, timestamp=1553147705330}
has deletions [delGen=4]
test: open reader.........OK
test: check integrity.....FAILED
WARNING: fixIndex() would remove reference to this segment; full exception:
org.apache.lucene.index.CorruptIndexException: checksum failed (hardware pro →
blem?) : expected=ebe966d3 actual=1f59a4b7 (resource=BufferedChecksumIndexIn →
put(MMapIndexInput(path="/ericsson/enm/solr/data/index/_376e.fdt")))
at org.apache.lucene.codecs.CodecUtil.checkFooter(CodecUtil.java:211)
at org.apache.lucene.codecs.CodecUtil.checksumEntireFile(CodecUtil.java:268)
at org.apache.lucene.codecs.compressing.CompressingStoredFieldsReader.checkI →
ntegrity(CompressingStoredFieldsReader.java:539)
at org.apache.lucene.index.SegmentReader.checkIntegrity(SegmentReader.java:5 →
77)
at org.apache.lucene.index.CheckIndex.checkIndex(CheckIndex.java:561)
at org.apache.lucene.index.CheckIndex.main(CheckIndex.java:1967)

9 of 9: name=_37zj docCount=5473340
codec=Lucene46
compound=false
numFiles=11
size (MB)=5,044.086
diagnostics = {os=Linux, java.vendor=Oracle Corporation, java.version=1.8.0_ →
172, lucene.version=4.8.1 1594670 - rmuir - 2014-05-14 19:22:52, mergeMaxNum →
Segments=-1, os.arch=amd64, source=merge, mergeFactor=10, os.version=2.6.32- →
754.9.1.el6.x86_64, timestamp=1553290524377}
no deletions
test: open reader.........OK
test: check integrity.....OK
```

```
test: check live docs.....OK
test: fields.............OK [47 fields]
test: field norms.........OK [5 fields]
test: terms, freq, prox...OK [16153886 terms; 174697629 terms/docs pairs; 0
tokens]
test: stored fields.......OK [205462058 total field count; avg 37.539 fields
 per doc]
test: term vectors........OK [0 total vector count; avg 0 term/freq vector f
ields per doc]
test: docvalues...........OK [34 docvalues fields; 0 BINARY; 10 NUMERIC; 24
SORTED; 0 SORTED_SET]

WARNING: 3 broken segments (containing 16539243 documents) detected
WARNING: would write new segments file, and 16539243 documents would be lost
, if -fix were specified
```

5. Run the script `SolrIndexUtility.sh` with the `repairIndex` parameter:

```
[root@solr install]# ./SolrIndexUtility.sh repairIndex /ericsson/enm/solr/da
ta/index
```

Example

```
[root@solr install]# ./SolrIndexUtility.sh repairIndex /ericsson/enm/solr/da
ta/index
Script checks for corruption of Solr indexes and may possibly repair as well
.
Please note that it is always recommended to take a backup of existing index
 and run this script to repair with repairIndex option passed as argument.
repairIndex called with args: /ericsson/enm/solr/data/index
********************************************************************************
********************************************************************************
*******
Please be noted that repair index results in loss of some documents and ensu
re backup of index is taken.If not sure of the consequences please exit imme
diately.
********************************************************************************
********************************************************************************
*******
Sleep for 10 seconds to allow user take his/her decision to proceed further
or not.
Performing repairIndex now..

Opening index @ /ericsson/enm/solr/data/index

Segments file=segments_a numSegments=9 version=4.8 format= userData={commitT
imeMSec=1529559901108}


1 of 9: name=_33ka docCount=5645513
codec=Lucene46
compound=false
numFiles=12
size (MB)=5,068.893
diagnostics = {os=Linux, java.vendor=Oracle Corporation, java.version=1.8.0_
172, lucene.version=4.8.1 1594670 - rmuir - 2014-05-14 19:22:52, mergeMaxNum
Segments=-1, os.arch=amd64, source=merge, mergeFactor=8, os.version=2.6.32-7
54.9.1.el6.x86_64, timestamp=1552510838482}
has deletions [delGen=6]
test: open reader.........OK
test: check integrity.....OK
test: check live docs.....OK [206 deleted docs]
test: fields.............OK [48 fields]
test: field norms.........OK [5 fields]
test: terms, freq, prox...OK [16710371 terms; 180447783 terms/docs pairs; 0
tokens]
test (ignoring deletes): terms, freq, prox...OK [16711506 terms; 180454137 t
erms/docs pairs; 0 tokens]
test: stored fields.......OK [212457867 total field count; avg 37.634 fields
 per doc]
test: term vectors........OK [0 total vector count; avg 0 term/freq vector f
ields per doc]
test: docvalues...........OK [34 docvalues fields; 0 BINARY; 10 NUMERIC; 24
SORTED; 0 SORTED_SET]
```

```
2 of 9: name=_3417 docCount=5209113
codec=Lucene46
compound=false
numFiles=12
size (MB)=4,969.003
diagnostics = {os=Linux, java.vendor=Oracle Corporation, java.version=1.8.0_ →
172, lucene.version=4.8.1 1594670 - rmuir - 2014-05-14 19:22:52, mergeMaxNum →
Segments=-1, os.arch=amd64, source=merge, mergeFactor=10, os.version=2.6.32- →
754.9.1.el6.x86_64, timestamp=1552593604463}
has deletions [delGen=4]
test: open reader.........OK
test: check integrity.....FAILED
WARNING: fixIndex() would remove reference to this segment; full exception:
org.apache.lucene.index.CorruptIndexException: checksum failed (hardware pro →
blem?) : expected=ffdbce3 actual=afbe03ac (resource=BufferedChecksumIndexInp →
ut(MMapIndexInput(path="/ericsson/enm/solr/data/index/_3417.fdt")))
at org.apache.lucene.codecs.CodecUtil.checkFooter(CodecUtil.java:211)
at org.apache.lucene.codecs.CodecUtil.checksumEntireFile(CodecUtil.java:268)
at org.apache.lucene.codecs.compressing.CompressingStoredFieldsReader.checkI →
ntegrity(CompressingStoredFieldsReader.java:539)
at org.apache.lucene.index.SegmentReader.checkIntegrity(SegmentReader.java:5 →
77)
at org.apache.lucene.index.CheckIndex.checkIndex(CheckIndex.java:561)
at org.apache.lucene.index.CheckIndex.main(CheckIndex.java:1967)

3 of 9: name=_3418 docCount=5441870
codec=Lucene46
compound=false
numFiles=12
size (MB)=5,041.22
diagnostics = {os=Linux, java.vendor=Oracle Corporation, java.version=1.8.0_ →
172, lucene.version=4.8.1 1594670 - rmuir - 2014-05-14 19:22:52, mergeMaxNum →
Segments=-1, os.arch=amd64, source=merge, mergeFactor=10, os.version=2.6.32- →
754.9.1.el6.x86_64, timestamp=1552691428620}
has deletions [delGen=3]
test: open reader.........OK
test: check integrity.....OK
test: check live docs.....OK [139 deleted docs]
test: fields..............OK [47 fields]
test: field norms.........OK [5 fields]
test: terms, freq, prox...OK [16159483 terms; 173439737 terms/docs pairs; 0  →
tokens]
test (ignoring deletes): terms, freq, prox...OK [16160238 terms; 173443892 t →
erms/docs pairs; 0 tokens]
test: stored fields.......OK [204182761 total field count; avg 37.522 fields →
 per doc]
test: term vectors........OK [0 total vector count; avg 0 term/freq vector f →
ields per doc]
test: docvalues...........OK [34 docvalues fields; 0 BINARY; 10 NUMERIC; 24  →
SORTED; 0 SORTED_SET]

4 of 9: name=_3545 docCount=5103807
codec=Lucene46
compound=false
numFiles=12
size (MB)=5,043.2
diagnostics = {os=Linux, java.vendor=Oracle Corporation, java.version=1.8.0_ →
172, lucene.version=4.8.1 1594670 - rmuir - 2014-05-14 19:22:52, mergeMaxNum →
Segments=-1, os.arch=amd64, source=merge, mergeFactor=10, os.version=2.6.32- →
754.9.1.el6.x86_64, timestamp=1552784433448}
has deletions [delGen=3]
test: open reader.........OK
test: check integrity.....OK
test: check live docs.....OK [128 deleted docs]
test: fields..............OK [47 fields]
test: field norms.........OK [5 fields]
test: terms, freq, prox...OK [15250150 terms; 162754894 terms/docs pairs; 0  →
tokens]
test (ignoring deletes): terms, freq, prox...OK [15250860 terms; 162758718 t →
erms/docs pairs; 0 tokens]
test: stored fields.......OK [191434492 total field count; avg 37.509 fields →
 per doc]
test: term vectors........OK [0 total vector count; avg 0 term/freq vector f →
ields per doc]
test: docvalues...........OK [34 docvalues fields; 0 BINARY; 10 NUMERIC; 24  →
SORTED; 0 SORTED_SET]

5 of 9: name=_35og docCount=5747124
codec=Lucene46
```

```
compound=false
numFiles=12
size (MB)=5,057.196
diagnostics = {os=Linux, java.vendor=Oracle Corporation, java.version=1.8.0_ →
172, lucene.version=4.8.1 1594670 - rmuir - 2014-05-14 19:22:52, mergeMaxNum →
Segments=-1, os.arch=amd64, source=merge, mergeFactor=9, os.version=2.6.32-7 →
54.9.1.el6.x86_64, timestamp=1552883705898}
has deletions [delGen=3]
test: open reader.........OK
test: check integrity.....FAILED
WARNING: fixIndex() would remove reference to this segment; full exception:
org.apache.lucene.index.CorruptIndexException: checksum failed (hardware pro →
blem?) : expected=f222bb93 actual=c576d64d (resource=BufferedChecksumIndexIn →
put(MMapIndexInput(path="/ericsson/enm/solr/data/index/_35og.fdt")))
at org.apache.lucene.codecs.CodecUtil.checkFooter(CodecUtil.java:211)
at org.apache.lucene.codecs.CodecUtil.checksumEntireFile(CodecUtil.java:268)
at org.apache.lucene.codecs.compressing.CompressingStoredFieldsReader.checkI →
ntegrity(CompressingStoredFieldsReader.java:539)
at org.apache.lucene.index.SegmentReader.checkIntegrity(SegmentReader.java:5 →
77)
at org.apache.lucene.index.CheckIndex.checkIndex(CheckIndex.java:561)
at org.apache.lucene.index.CheckIndex.main(CheckIndex.java:1967)

6 of 9: name=_3654 docCount=4902552
codec=Lucene46
compound=false
numFiles=12
size (MB)=5,034.75
diagnostics = {os=Linux, java.vendor=Oracle Corporation, java.version=1.8.0_ →
172, lucene.version=4.8.1 1594670 - rmuir - 2014-05-14 19:22:52, mergeMaxNum →
Segments=-1, os.arch=amd64, source=merge, mergeFactor=10, os.version=2.6.32- →
754.9.1.el6.x86_64, timestamp=1552964709134}
has deletions [delGen=3]
test: open reader.........OK
test: check integrity.....OK
test: check live docs.....OK [118 deleted docs]
test: fields..............OK [47 fields]
test: field norms.........OK [5 fields]
test: terms, freq, prox...OK [14915002 terms; 156903507 terms/docs pairs; 0  →
tokens]
test (ignoring deletes): terms, freq, prox...OK [14915663 terms; 156907032 t →
erms/docs pairs; 0 tokens]
test: stored fields.......OK [184537063 total field count; avg 37.642 fields →
 per doc]
test: term vectors........OK [0 total vector count; avg 0 term/freq vector f →
ields per doc]
test: docvalues...........OK [34 docvalues fields; 0 BINARY; 10 NUMERIC; 24  →
SORTED; 0 SORTED_SET]

7 of 9: name=_36ob docCount=5722406
codec=Lucene46
compound=false
numFiles=12
size (MB)=5,059.659
diagnostics = {os=Linux, java.vendor=Oracle Corporation, java.version=1.8.0_ →
172, lucene.version=4.8.1 1594670 - rmuir - 2014-05-14 19:22:52, mergeMaxNum →
Segments=-1, os.arch=amd64, source=merge, mergeFactor=10, os.version=2.6.32- →
754.9.1.el6.x86_64, timestamp=1553058929673}
has deletions [delGen=3]
test: open reader.........OK
test: check integrity.....OK
test: check live docs.....OK [133 deleted docs]
test: fields..............OK [47 fields]
test: field norms.........OK [5 fields]
test: terms, freq, prox...OK [16672261 terms; 182633808 terms/docs pairs; 0  →
tokens]
test (ignoring deletes): terms, freq, prox...OK [16672997 terms; 182637783 t →
erms/docs pairs; 0 tokens]
test: stored fields.......OK [214687501 total field count; avg 37.518 fields →
 per doc]
test: term vectors........OK [0 total vector count; avg 0 term/freq vector f →
ields per doc]
test: docvalues...........OK [34 docvalues fields; 0 BINARY; 10 NUMERIC; 24  →
SORTED; 0 SORTED_SET]

8 of 9: name=_376e docCount=5583006
codec=Lucene46
compound=false
numFiles=12
size (MB)=5,072.371
```

```
diagnostics = {os=Linux, java.vendor=Oracle Corporation, java.version=1.8.0_ →
172, lucene.version=4.8.1 1594670 - rmuir - 2014-05-14 19:22:52, mergeMaxNum →
Segments=-1, os.arch=amd64, source=merge, mergeFactor=9, os.version=2.6.32-7 →
54.9.1.el6.x86_64, timestamp=1553147705330}
has deletions [delGen=4]
test: open reader.........OK
test: check integrity.....FAILED
WARNING: fixIndex() would remove reference to this segment; full exception:
org.apache.lucene.index.CorruptIndexException: checksum failed (hardware pro →
blem?) : expected=ebe966d3 actual=1f59a4b7 (resource=BufferedChecksumIndexIn →
put(MMapIndexInput(path="/ericsson/enm/solr/data/index/_376e.fdt")))
at org.apache.lucene.codecs.CodecUtil.checkFooter(CodecUtil.java:211)
at org.apache.lucene.codecs.CodecUtil.checksumEntireFile(CodecUtil.java:268)
at org.apache.lucene.codecs.compressing.CompressingStoredFieldsReader.checkI →
ntegrity(CompressingStoredFieldsReader.java:539)
at org.apache.lucene.index.SegmentReader.checkIntegrity(SegmentReader.java:5 →
77)
at org.apache.lucene.index.CheckIndex.checkIndex(CheckIndex.java:561)
at org.apache.lucene.index.CheckIndex.main(CheckIndex.java:1967)

9 of 9: name=_37zj docCount=5473340
codec=Lucene46
compound=false
numFiles=11
size (MB)=5,044.086
diagnostics = {os=Linux, java.vendor=Oracle Corporation, java.version=1.8.0_ →
172, lucene.version=4.8.1 1594670 - rmuir - 2014-05-14 19:22:52, mergeMaxNum →
Segments=-1, os.arch=amd64, source=merge, mergeFactor=10, os.version=2.6.32- →
754.9.1.el6.x86_64, timestamp=1553290524377}
no deletions
test: open reader.........OK
test: check integrity.....OK
test: check live docs.....OK
test: fields..............OK [47 fields]
test: field norms.........OK [5 fields]
test: terms, freq, prox...OK [16153886 terms; 174697629 terms/docs pairs; 0  →
tokens]
test: stored fields.......OK [205462058 total field count; avg 37.539 fields →
 per doc]
test: term vectors........OK [0 total vector count; avg 0 term/freq vector f →
ields per doc]
test: docvalues...........OK [34 docvalues fields; 0 BINARY; 10 NUMERIC; 24  →
SORTED; 0 SORTED_SET]

WARNING: 3 broken segments (containing 16539243 documents) detected
WARNING: 16539243 documents will be lost

NOTE: will write new segments file in 5 seconds; this will remove 16539243 d →
ocs from the index. THIS IS YOUR LAST CHANCE TO CTRL+C!
  5...
  4...
  3...
  2...
  1...
Writing...
OK
Wrote new segments file "segments_b"
```

6. Clean up the Solr database manually following the procedure described in the *ENM Monitoring System Administrator Guide*, section **Fault Management Administration Tasks** > **Configure the Alarm History** > **Clean Manually Historical Alarms**.

7.10 Solr Database Fills Fast When Automatic Alarm Purging Does Not Suffice Its Requirement

Historical alarms are purged automatically based on two conditions:

Size — If the index size exceeds 70% of the configured disk space.

Time If data is stored for more than 90 days.

If, for any reason, the Solr database fills up, historical alarms can be purged from the database manually, using a script.

Prerequisites

Access to the Management Server (MS) and Virtual Machines (VMs)

Steps

1. Connect to the `solr` VM instance.

   For details, see Connect to a Service on page 2.

2. Clean up the Solr database manually, following the procedure described in the *ENM Monitoring System Administrator Guide*, section **Fault Management Administration Tasks** > **Configure the Alarm History** > **Clean Manually Historical Alarms**.

Results

Historical alarms are removed from the Solr database.

7.11 Alarm Supervision Recovery Failed after ENM Reboot or comecimpolicy VMs Reboot

Fault Management (FM) generates an internal alarm whenever:

— Alarm supervision of nodes cannot be resumed after ENM reboot.

— `comecimpolicy` VMs reboot due to unavailability of dependent resources like, among others, jms, or database.

This internal alarm has the following alarm attributes:

Alarm Attribute	Value
ObjectOfReference	FMMediation
SpecificProblem	Alarm Supervision Error
ProbableCause	System Resource Unavailable
EventType	ET_PROCESSING_ERROR_ALARM
RecordType	ERROR_MESSAGE
AdditionalText	Internal alarm raised by FM Mediation when automatic node supervision cannot be resumed because of unavailability of dependencies like JMS/Mediation Services/Database.
FDN	'ManagementSystem=ENM'

This internal alarm is mapped to the ENM Management System. The user needs to import the Management System from the Topology Browser to see this alarm in Alarm Monitor application.

Prerequisites

— User requires access to the Management Server (MS) and Virtual Machines (VMs).

— User requires basic knowledge of Linux and JBoss.

— Access to the private key file for authentication (contact your OpenStack administrator).

— Access to the Ericsson Management Portal (EMP) VM in the ENM on Cloud deployment.

Steps

1. To log on to `comecimpolicy` VM instances, see Connect to a Service on page 2.
2. Search JBoss server logs for lines with the text `Critical Error in handling FM Supervision` in all `comecimpolicy` VMs to confirm that this alarm is generated because of unavailability of dependent resources like database and jms.

   Server logs reside at the following path of the corresponding VMs: `/ericsson/3pp/jboss/standalone/log/server.log`
3. Toggle alarm supervision for nodes from CLI using alarm commands:

```
alarm disable *
```

```
alarm enable *
```

Results

You can troubleshoot the issue if Alarm Supervision recovery has failed after an ENM reboot or `comecimpolicy` VM reboot.

Once the steps are executed and Alarm Supervision is restored, internal alarms can be acknowledged from the Command Line Interface (CLI) or Alarm Monitor to remove them from the system.

7.12 ENIQ-S Events Are Not Seen in FM GUI

This section describes how to troubleshoot the following issue:

— ENIQ-S events are not seen in FM GUI.

Prerequisites

— The node to which the ENIQ-S event is sent is under supervision.

— User requires basic knowledge of Linux and JBoss.

— On physical: Access to the Management Server (MS) and Virtual Machines (VMs).

— On Cloud:

- Access to the private key file for authentication (contact your OpenStack administrator).
- Access to the Ericsson Management Portal (EMP) VM in the ENM on Cloud deployment.

Steps

1. Log on to the ENM MS as the `litp-admin` user and switch to the `root` user, or log on to the EMP server on Cloud environment as `root` user and go to following directory: `/ericsson/pib-scripts/etc/`
2. Read the configurable parameter 'enableEniqs'.

```
[root@cloud-ms-1 etc]# python config.py read --app_server_address=<fmserv_ho →
st_name>:8080 --name=enableEniqs

where <fmserv_host_name> is the fmserv hostname.
```

3. Check the value of the parameter 'enableEniqs'.
4. If the value is `false`, change the value to `true` by using following command. If the value is `true`, go to step 5.

```
python config.py update --app_server_address=<fmserv_host_name>:8080 --name= →
enableEniqs --value=truewhere <fmserv_host_name> is the fmserv hostname.
```

5. Check the supervision state of the Network Element from CliApp using the following command:

```
alarm status <nodename>
```

Example

```
&187;alarm status LTE02ERBS00003 -l
Subscription Status : true
Heartbeat Failure : false
Synchronization Ongoing : false
Type : ERBS
Alarm Suppressed : false
Node Suspended : false
Out Of Synchronization : false
Supervision Status : on
Technician Present : false
NodeName : LTE02ERBS00003
```

```
Total number of node status data fetched for the given query is :1
```

Check the status of active attribute. If it is false, enable the supervision and send an ENIQ-S event again. If it is true, go to step 6.

CliApp can be accessed at `https://<apache_host_name>/#cliapp`

6. Connect to one of the SVC from MS and check if `fmserv` VMs are online.

```
$ hastatus -sum |grep <VMname>
```

See the following command and system output response as an example (where the `fmserv` VMs status is online):

```
[root@cloud-svc-1 litp-admin]# hastatus -sum |grep fmserv
B  Grp_CS_svc_cluster_fmserv cloud-svc-3          Y          N               →
 ONLINE
B  Grp_CS_svc_cluster_fmserv cloud-svc-4          Y          N               →
 ONLINE
```

If the VMs are not online, the issue is with the Deployment of VM. Go to step 7.

7. Make the VMs online and the procedure is ended.

```
$ hagrp -online <groupName> -sys <MasterServerName>
```

Example

```
[root@svc-1 litp-admin]# hagrp -online Grp_CS_svc_cluster_fmserv -sys svc-3
```

8. If the VMs are online, then check the hostname of the VMs:

On physical, use `cat /etc/hostname | grep <VMname>`:

```
[root@cloud-svc-1 litp-admin]# cat /etc/hosts | grep fmserv
<ipAddress>  svc-3-fmserv    # Created by LITP. Please do not edit
<ipAddress>  svc-4-fmserv    # Created by LITP. Please do not edit
```

On Cloud, use `consul members | grep <VMname>` as the root user on EMP.

```
[root@gat-emp-0 data]# consul members | grep fmserv
gat-fmserv-0             <ipAddress>:8301   alive   client  0 →
.8.1  2         dc1
gat-fmserv-1             <ipAddress>:8301   alive   client  0 →
.8.1  2         dc1
[root@gat-emp-0 data]#
```

9. To connect to `fmserv` VM instance, see Connect to a Service on page 2.

```
[root@ms-1 ~]# ssh -i /root/.ssh/vm_private_key cloud-user@svc-3-fmserv
```

10. Verify the JBoss service is running.

```
$ service jboss monitor
```

If the JBoss service is not running, then go to step-11.

11. Start the JBoss.

```
$ service jboss start
```

12. Check the `fmalarmprocessing` and `wpserv` VMs in the same way from steps 6–11.

13. If all specified VMs and JBoss are running, then check the `fmserv` logs as specified in the next step.

14. To connect to `fmserv` VM instance, see Connect to a Service on page 2.

    a. Connect to JBoss CLI by executing the following command from the `fmserv` VM: `sh /ericsson/3pp/jboss/bin/jboss-cli.sh -c`

    b. Execute the following command in JBoss CLI to enable DEBUG level logging:

```
[standalone@localhost:9999 /] /subsystem=logging/logge →
r=com.ericsson.oss.services.fm.erreventsservice.erreve →
nts:add(level=DEBUG)
{"outcome" => "success"}
[standalone@localhost:9999 /] exit
```

15. Go to the following directory and open `server.log` from `fmserv` VM: `/ericsson/3pp/jboss/standalone/log/server.log`

16. Check for the following log in the `server.log` file. Verify that the received ENIQ-S event is translated to event notification and sent to APS for processing.

```
Sending the error event:
```

17. Remove the DEBUG logging level by connecting to JBoss CLI from `fmserv` VM.

```
[standalone@localhost:9999 /] /subsystem=logging/logger=com.ericsson.oss.ser →
vices.fm.erreventsservice.errevents:remove()
{"outcome" => "success"}
[standalone@localhost:9999 /] exit
```

18. If the log mentioned in step 16 is present in `server.log`, then debug in the APS and FM UI.

Results

ENIQ-S events are in the FM GUI.

7.13 Email routeType Not Visible in FM GUI

> This section is referred when the email routes are not created and updated through Fault Management (FM) GUI.

Prerequisites

— User requires access to the Management Server (MS) and Virtual Machines (VMs).

— Access to the private key file for authentication (contact your OpenStack administrator).

— Access to the Ericsson Management Portal (EMP) VM in the ENM on Cloud deployment.

Steps

1. Verify it from CLI by creating the route with routeType as email.
2. If the route is created successfully through the CLI, then refresh the GUI to reflect in GUI.
3. If the email route type is not visible after refreshing the GUI, then:
    a. Log on to the ENM Management Server or to the ENM Management Portal (EMP) server on Cloud environment.
    b. In ENM Management Server, go to the directory `/ericsson/pib-scripts/etc/`.

    ```
    [root@ms~]#cd /ericsson/pib-scripts/etc/
    [root@ms etc]#
    ```

    c. Read the configurable parameter 'ENABLE_OUTBOUND_EMAILS' in each instance of `fmhistory`

    Execute the following command to read the value of parameter ENABLE_OUTBOUND_EMAILS:

    ```
    python config.py read --app_server_address=<fmhistory_host_name>:80 →
    80 --name=ENABLE_OUTBOUND_EMAILS
    ```

    where `<fmhistory_host_name>` is the `fmhistory` hostname.

Example

```
python config.py read --app_server_address=svc-2-fmhistory:8080 --n
ame=ENABLE_OUTBOUND_EMAILS
```

4. Check the value of the parameter 'ENABLE_OUTBOUND_EMAILS'.
5. If the value is false, change the value to true by using the following command.

```
python config.py update --app_server_address=<fmhistory_host_name>:8080 --na
me=ENABLE_OUTBOUND_EMAILS --value=true
```

where <fmhistory_host_name> is the fmhistory hostname.

Example

```
python config.py update --app_server_address=svc-2-fmhistory:
8080 --name=ENABLE_OUTBOUND_EMAILS  --value=true
```

Result:

```
true
```

6. If the value is true, check for error logs in the alarmcontroldisplayservice logs in fmserv.

VM Name	Logger Level
fmserv	com.ericsson.oss.services.fm.alarmcontroldisplayservice

Results

Email routes are created and updated from UI.

7.14 Recover the Possible Alarm Loss Due to Blade Failure or FM Alarm Processing VM Crash

Alarm loss during the blade failure or fmalarmprocessing/fmhistory VM crash (abrupt shutdown).

Prerequisites

— Access to the AlarmMonitor/AlarmSupervisionStatus applications

— Access to the Management Server (MS) and Virtual Machines (VMs)

— Basic knowledge of Linux and JBoss

Steps

1. Access the Cliapp from the following address:

   ```
   https://<apache_host_name>/#cliapp
   ```

2. From the CliApp execute the following command to initiate the manual alarm synchronization on a node from Command Line Interface (CLI):

   ```
   alarm sync LTE09ERBS00001
   ```

   LTE09ERBS00001 is the node with suspected alarm loss during the scenario.

3. Alternatively, access to the AlarmMonitor/AlarmSupervisionStatus applications and initiate the manual alarm synchronization on a selected node.

Results

There should be no alarm loss during the blade failure or crash of fmalarmprocessing/fmhistory VMs.

To perform the workaround to recover the alarm loss during this scenario, take the following steps to initiate the manual alarm synchronisation.

7.15 Correlation of Heartbeat Alarm during Neo4j Unavailability

> This section describes how to troubleshoot when the Heartbeat clear alarm arrives for a Network Element which is in Heartbeat state during Neo4j unavailability.

Prerequisites

— User with FM_Administrator role.

Steps

1. Import and apply the Network Element in the Alarm Monitor UI.

   ```
   https://enmapache.athtem.eei.ericsson.se/#help/app/alarmviewer/topic/ui
   ```

2. Perform acknowledge and clear operations on the Heartbeat alarm.

   ```
   https://enmapache.athtem.eei.ericsson.se/#help/app/alarmviewer/topic/tutorials/AlarmOperations
   ```

3. Disable and enable the alarm supervision on the Network Element.

```
https://enmapache.athtem.eei.ericsson.se/#help/app/alarmviewer/topic/tutorials/NodeOperations
```

Results

Heartbeat alarm is cleared.

7.16 Recovery of FM Alarm Processing by Clearing the Messages Stuck in the ClusteredFMMediationChannel

This section describes how to troubleshoot the following issue:

— Alarms or events that are stuck in the ClusteredFMMediationChannel are not processed anymore.

Prerequisites

— User requires access to the Management Server (MS) and Virtual Machines (VMs).

— User requires basic knowledge of Linux and JBOSS.

Steps

1. Access the Cliapp from the following address:

```
https://<apache_host_name>/#cliapp
```

2. Disable alarm supervision for all network elements.

```
https:/<apache_host_name>/#help/app/cliapp/topic/syntax_alarm/disable
```

3. Log on to svc-1 as litp-admin user and change to root user.

```
[root@ms-1~]# ssh litp-admin@svc-1
[litp-admin@svc-1~]$su
[root@cloud-svc-1litp-admin]#
```

4. Offline all fmalarmprocessing VMs in the system using the following command:

   To offline fmalarmprocessing VM:

```
hagrp -offline Grp_CS_svc_cluster_fmalarmprocessing -any
```

```
[root@cloud-svc-1litp-admin]# hagrp -offline Grp_CS_svc_cluster_fmalarmprocessing -any
```

```
[root@cloud-svc-1litp-admin]# hastatus -sum | grep fmalarmprocessing
B Grp_CS_svc_cluster_fmalarmprocessing SVC-1 Y N OFFLINE
B Grp_CS_svc_cluster_fmalarmprocessing SVC-2 Y N OFFLINE
B Grp_CS_svc_cluster_fmalarmprocessing SVC-3 Y N OFFLINE
B Grp_CS_svc_cluster_fmalarmprocessing SVC-4 Y N OFFLINE
```

5. Connect to the database in which JMS is running and remove the messages accumulated in ClusteredFMMediationChannel queue.

   To remove the messages from the ClusteredFMMediationChannel queue:

```
/subsystem=messaging/hornetq-server=default/jms-queue=ClusteredFMMediationCh →
annel:remove-messages
```

```
[root@ms-1 ~]# ssh litp-admin@db-1
litp-admin@db-1's password:
[litp-admin@ieatrcxb4739-1 ~]$ su
Password:
[root@db-1 litp-admin]# hastatus -sum | grep jms
B Grp_CS_db_cluster_jms_clustered_service db-1 Y N ONLINE
B Grp_CS_db_cluster_jms_clustered_service db-2 Y N OFFLINE
[root@db-1 litp-admin]#

[root@db-1 litp-admin]# /ericsson/3pp/jboss/bin/jboss-cli.sh -c
[standalone@localhost:9999 /] /subsystem=messaging/hornetq-server=default/jm →
s-queue=ClusteredFMMediationChannel:remove-messages
{
"outcome" => "success",
"result" =>44664L
}
```

6. Check for messages count in ClusteredFMMediationChannel queue, if the messages are not successfully removed, perform JMS failover.

   To failover JMS:

```
hagrp -switch Grp_CS_db_cluster_jms_clustered_service -to db-2
```

```
[standalone@localhost:9999 /] /subsystem=messaging/hornetq-server=default/jm →
s-queue=ClusteredFMMediationChannel:read-attribute(name=message-count)
{
"outcome" => "success",
"result" =>44664L

}
[standalone@localhost:9999 /] exit
[root@db-1 litp-admin]# hagrp -switch Grp_CS_db_cluster_jms_clustered_servic →
e -to db-2
```

7. If messages are still not removed, perform step 3 and step 4.

8. Online all fmalarmprocessing VMs in the system using the following command:

   To online fmalarmprocessing VM:

```
hagrp -online Grp_CS_svc_cluster_fmalarmprocessing -any
```

```
[root@cloud-svc-1litp-admin]# hagrp -online Grp_CS_svc_cluster_fmalarmproces →
sing -any

[root@cloud-svc-1litp-admin]# hastatus -sum | grep fmalarmprocessing
B Grp_CS_svc_cluster_fmalarmprocessing SVC-1 Y N ONLINE
B Grp_CS_svc_cluster_fmalarmprocessing SVC-2 Y N ONLINE
B Grp_CS_svc_cluster_fmalarmprocessing SVC-3 Y N ONLINE
B Grp_CS_svc_cluster_fmalarmprocessing SVC-4 Y N ONLINE
```

9. Enable alarm supervision for all the network elements.

```
https:/<apache_host_name>/#help/app/cliapp/topic/syntax_alarm/enable
```

Results

Alarms or events are processed successfully by FM and the queue is cleared from the stuck messages.

7.17 Alarm Synchronization Issues for AXE Network Elements

> This section describes how to diagnose and resolve AXE Alarm Synchronization Failures.

Cause

The security MO credentials are changed after that the alarm supervision is activated

Solution

1. Disable and then enable alarm supervision.

7.18 ObjectOfReference of a BSC/MSC/HLR-FE Alarm Is Not Shown on the Correct BTS

> This section describes how to diagnose and resolve the issue where the ObjectOfReference of a BSC/MSC/HLR-FE Alarm is not shown on the correct BTS.

Prerequisites:

— Access to the Management Server (MS) and Virtual Machines (VMs).

— Basic knowledge of Linux and JBOSS.

Cause

CM Synchronization of the NE is not enabled.

Solution

1. Launch CliApp: `https://<apache_host_name>/#cliapp`
2. Run the following command: `$cmedit get NetworkElement=MSCV18,CmFunction=1`

   Example:

   ```
   »cmedit get NetworkElement=MSCV18,CmFunction=1
   FDN : NetworkElement=MSCV18,CmFunction=1
   CmFunctionId : 1
   failedSyncsCount : 0
   iposCopyConfigFrequency: 64800
   lastFailedSync : null
   lostSynchronization : null
   syncStatus : UNSYNCHRONIZED
   1 instance(s)
   ```

3. Check the value of `syncStatus` attribute.

   If it is UNSYNCHRONIZED, NE is not properly CM synchronized. Run the following command to perform CM synchronization.

   ```
   cmedit set NetworkElement=MSCV18,CmNodeHeartbeatSupervision=1 active=true
   cmedit action NetworkElement=MSCV18,CmFunction=1 sync
   ```

7.19 ObjectOfReference of an NWI-E Switch Alarm Is Not Shown on the Correct BTS

> Check the CM synchronization status of the NE and perform the appropriate resolution action.

Prerequisites

— Access to the Management Server (MS) and Virtual Machines (VMs).

— Basic knowledge of Linux and JBOSS.

Solution

1. Access the CliApp:

   ```
   https://<apache_host_name>/#cliapp
   ```

2. Get the CM Synchronization status of the NE:

   ```
   $cmedit get NetworkElement=GSM01BSC01,CmFunction=1
   ```

   Example

   ```
   »cmedit get NetworkElement=GSM01BSC01,CmFunction=1
   FDN : NetworkElement=GSM01BSC01,CmFunction=1
   CmFunctionId : 1
   failedSyncsCount : 0
   iposCopyConfigFrequency: 64800
   lastFailedSync : null
   lostSynchronization : null
   syncStatus : UNSYNCHRONIZED

   1 instance(s)
   ```

3. Check the value of the `syncStatus` attribute. If it is UNSYNCHRONIZED, NE is not properly CM synchronized. Run the following command to perform CM synchronization:

   ```
   cmedit set NetworkElement=MSCV18,CmNodeHeartbeatSupervision=1 active=true

   cmedit action NetworkElement=MSCV18,CmFunction=1 sync
   ```

7.20 Spontaneous Alarms Are Not Received from the BSC/MSC/HLR-FE

> This section describes how to diagnose and resolve the issue where spontaneous alarms are not received from the BSC/MSC/HLR-FE.

Prerequisites:

— Access to the Management Server (MS) and Virtual Machines (VMs).

— Basic knowledge of Linux and JBOSS.

Cause 1

Subscription State of the NE is not enabled.

Solution

1. Launch CliApp: `https://<apache_host_name>/#cliap`
2. Run the following command: `$cmedit get NetworkElement=<nodename>,FmFunction=1`

   Example:

```
»cmedit get NetworkElement=MSCV18,FmFunction= →
1
FDN : NetworkElement=MSCV18,FmFunction=1
FmFunctionId : 1
alarmSuppressedState : false
currentServiceState : IDLE
lastUpdated : 1489146835692
lastUpdatedTimeStamp : Fri Mar 10 11:53:55 GMT 2017
subscriptionState : ENABLING
technicianPresentState : false
1 instance(s)
```

3. Check that the `subscriptionState` attribute is ENABLED.

   If the status is ENABLING, the alarm supervision on NE is not properly enabled.

Cause 2

VMs are not online.

Solution

1. Log on to SVC as `litp-admin` user and change to root user.

```
[root@ms-1 ~]# ssh litp-admin@cloud-svc-1
[litp-admin@cloud-svc-1 ~]$su root
Password:
[root@cloud-svc-1 litp-admin]
```

2. Run the following command:

```
$ hastatus -sum |grep <VMname>
[root@cloud-svc-1 litp-admin]# hastatus -sum | grep msapgfm
B  Grp_CS_svc_cluster_msapgfm cloud-svc-3         Y          N               →
 ONLINE
B  Grp_CS_svc_cluster_msapgfm cloud-svc-4         Y          N               →
 ONLINE
```

   If VMs are not online, then the issue is with the VM deployment. Run the following command to bring the VMs online:

```
$ hagrp -online <groupName> -sys <MasterServerName>
Example:
```

```
[root@cloud-svc-1 litp-admin]# hagrp -online Grp_CS_svc_cluster_msapgfm -sys
 cloud-svc-3
```

Cause 3

JBOSS service is not running.

Solution

1. Check the hostname of VMs.

```
[root@cloud-svc-1 litp-admin]#  cat /etc/hosts | grep <VMname>
Example:
[root@cloud-svc-1 litp-admin]#  cat /etc/hosts | grep msapgfm
<ipAddress> svc-4-msapgfm   msapgfm-2-internal      # Created by LITP. Pleas
e do not edit
<ipAddress> svc-1-msapgfm   msapgfm-1-internal      # Created by LITP. Pleas
e do not edit
```

2. Log on to any of the VM instances listed in MS.

   Refer to Connect to a Virtual Machine on a Physical ENM Deployment on page 2.

3. Run the following command: `$ service jboss monitor`

   — If JBOSS is not running, start JBOSS with the following command: `$ service jboss start`

   — If JBoss instances are running properly, verify if any exceptions are present in `msapgfm` logs:

```
[root@svc-4-msapgfm ~]# cd /ericsson/3pp/jboss/standalone/log
[root@svc-4-msapgfm log]# vi server.log
```

Cause 4

WinFIOL exceptions

Solution

1. Connect to `WinFIOL` group.

   SSH to the VM to access it.

2. Use the cloud-user ID and include the path to the VM private key:

```
[root@ms-1 ~]# ssh -i /root/.ssh/vm_private_key cloud-user@scp-1-winfiol
```

3. Run the following command:

```
[root@scp-1-winfiol ~]# cd /ericsson/3pp/jboss/standalone/log
[root@scp-1-winfiol log]# vi server.log
```

7.21 Additional IP Address Change for Already Supervised BSC/MSC/HLR-FE Nodes

> This section describes how to diagnose and resolve issues when an additional IP address changes for already supervised BSC (NWI-E)/MSC/HLR-FE nodes, as these nodes support multiple IP addresses.

Prerequisites

— Access to the Management Server (MS) and Virtual Machines (VMs)

— Basic knowledge of Linux and JBOSS

Cause

Additional IP address (any other additional IP address attribute other than normal **IP Address** field) changed in connectivity information for BSC (NWI-E) / MSC / HLR-FE nodes.

1. Launch the **Alarm Supervision Status** application at `https://<apache_host_name>/#alarmoverview/alarm-management`.
2. Select the node for which an additional IP address has changed and choose **Disable Supervision** to disable FM supervision for this node.

   Ensure FM supervision is disabled completely for the node.
3. Select the node again (for which an additional IP address has changed and for which you have disabled FM supervision) and choose **Enable Supervision** to enable FM supervision.

7.22 Inconsistent Alarm Build-up in Neo4j

> A race condition can occur when two alarm actions are performed on the same alarm in parallel, leading to alarm data inconsistency.

Prerequisites

Access to ENM web CLI.

Solution

1. Check for a build-up of inconsistent alarms with severity CLEARED and state ACTIVE_ACKNOWLEDGED:

   ```
   alarm get * --ack --cleared --count
   ```

   A count of `0` indicates that there are no inconsistent alarms.

   Tip: Remove the `--count` parameter to view a list of inconsistent alarms.

   Example

   ```
   alarm get * --ack --cleared --count
   Total number of alarms fetched for the given query is : 1033
   ```

2. Clear all the inconsistent alarms:

   ```
   alarm clear * --cleared --ack
   ```

   Example

   ```
   alarm clear * --cleared --ack
   Operation Successful for 1033 alarm(s).
   ```

3. Verify that the inconsistent alarms were cleared (count value is `0`):

   ```
   alarm get * --ack --cleared --count
   ```

7.23 Alarms Are Not Written to the File for Save To File Alarm Routing

> Check that the enablePolicy state is true, VMs are online, and JBoss service is running.

Prerequisites

— Basic knowledge of Linux and JBoss.

— On physical: Access to the Management Server (MS) and Virtual Machines (VMs).

— On cloud:

  — Access to the Ericsson Management Portal (EMP) VM in the ENM on Cloud deployment.Access to the private key file for authentication (contact your OpenStack administrator).

Cause 1

Alarm Route is in disabled state.

Solution

1. Launch the cliapp:

```
https://%3Capache_host_name%3E/#cliapp
```

2. Get the details of the route policy:

```
Access to the private key file for authentication (contactfmedit get * Alarm →
RoutePolicy.(name==<NameOfTheRoute>) AlarmRoutePolicy.*
```

   Example

```
»fmedit get * AlarmRoutePolicy.(name==TestRoute) AlarmRoutePolicy.*
FDN : null
fileType : TXT
enablePolicy : false
fdn : null
fileHeaders : [repeatCount, insertTime, eventTime, specificProblem, probable →
Cause, eventType, commentText, alarmState, fdn, presentSeverity, alarmingObj →
ect]
beginTime : 00:00:00
endTime : 23:59:59
routeID : 281474981842356
subordinateType : All_SUBORDINATES
routeType : FILE
days : MONDAY,TUESDAY,WEDNESDAY,THURSDAY,FRIDAY,SATURDAY,SUNDAY
description : Save To File Test Route
name : TestRoute
perceivedSeverity : INDETERMINATE,CRITICAL,MAJOR,MINOR,WARNING
fileName : Save_File
outputType : file
```

3. Check if the `enablePolicy` parameter of the Alarm Route is `true` (enabled).

   If it is `false` (disabled), enable it:

```
fmedit set * AlarmRoutePolicy.(name==TestRoute) enablePolicy="true"
```

   Example:

```
fmedit set * AlarmRoutePolicy.(name==<NameOfTheRoute>) enablePolicy="true"
route [TestRoute] updated successfully
```

Cause 2

VMs are not online.

Solution

1. Log on to SVC as `litp-admin` user and change to root user:

```
[root@ms-1 ~]# ssh litp-admin@cloud-svc-1
[litp-admin@cloud-svc-1 ~]$su root
Password:
[root@cloud-svc-1 litp-admin]
```

2. Check that the VMs (`fmhistory`, `dlms`, `fmalarmprocessing`, `fmserv`, `msfm`) are online:
   - On Physical:

     ```
     $ hastatus -sum |grep <VMname>
     ```

     Example:

     ```
     [root@cloud-svc-1 litp-admin]# hastatus -sum | grep fmhistory
     B  Grp_CS_svc_cluster_fmhistory cloud-svc-3          Y          N
              ONLINE
     ```

   - As root user on Cloud:

     ```
     consul members | grep <VMname>
     ```

     Example:

     ```
     [root@gat-emp-0 data]# consul members | grep fmhistory
     gat-fmhistory-0 :8301 alive client 0.8.1 2 dc1
     gat-fmhistory-1 :8301 alive client 0.8.1 2 dc1
     [root@gat-emp-0 data]#
     ```

   If they are not online, the VM deployment is the issue. Bring the VMs online:

   ```
   $ hagrp -online <groupName> -sys <MasterServerName>
   ```

   Example:

   ```
   [root@cloud-svc-1 litp-admin]# hagrp -online Grp_CS_svc_cluster_fmhistory -s
   ys cloud-svc-3
   ```

Cause 3

JBoss service is not running.

Solution

1. Check the hostname of VMs (`fmhistory`, `dlms`, `fmalarmprocessing`, `fmserv`, `msfm`):
   - On Physical:

     ```
     cat /etc/hostname | grep <VMname>
     ```

     Example:

     ```
     [root@cloud-svc-1 litp-admin]# cat /etc/hosts | grep fmhistory
     10.247.246.63   svc-3-fmhistory fmhistory-1-internal    # Created by LI
     TP. Please do not edit
     ```

   - As root user on Cloud:

     ```
     consul members | grep <VMname>
     ```

     Example:

     ```
     [root@gat-emp-0 data]# consul members | grep fmhistory
     gat-fmhistory-0 :8301 alive client 0.8.1 2 dc1
     gat-fmhistory-1 :8301 alive client 0.8.1 2 dc1
     [root@gat-emp-0 data]#
     ```

2. Log on to any of the VM instances listed in MS.

   Refer to *Connect to a Virtual Machine*.

3. Check that JBoss is running:

   ```
   $ service jboss monitor
   ```

   - If JBoss is not running, start JBoss:

     ```
     $ service jboss start
     ```

   - If JBoss instances are running properly, verify if any exceptions are present in `fmhistory` logs:

     ```
     [root@svc-3-fmhistory ~]# cd /ericsson/3pp/jboss/standalone/log
     [root@svc-3-fmhistory log]# vi server.log
     ```

7.24 Unable to Download the Saved Alarm File from View Details of Save To File Alarm Routing

Check that the VMs are online and JBoss service is running.

Prerequisites

- Basic knowledge of Linux and JBoss.
- On physical: Access to the Management Server (MS) and Virtual Machines (VMs).
- On cloud:
  - Access to the private key file for authentication (contact your OpenStack administrator).
  - Access to the Ericsson Management Portal (EMP) VM in the ENM on Cloud deployment.

Cause 1

VMs are not online.

Solution

1. Log on to SVC as `litp-admin` user and change to root user:

   ```
   [root@ms-1 ~]# ssh litp-admin@cloud-svc-1
   [litp-admin@cloud-svc-1 ~]$su root
   Password:
   [root@cloud-svc-1 litp-admin]
   ```

2. Check that the VMs (`fmhistory`, `dlms`, `fmalarmprocessing`, `fmserv`, `msfm`) are online:
   - On Physical:

     ```
     $ hastatus -sum |grep <VMname>
     ```

     Example:

     ```
     [root@cloud-svc-1 litp-admin]# hastatus -sum | grep fmhistory
     B  Grp_CS_svc_cluster_fmhistory cloud-svc-3          Y          N
             ONLINE
     ```

   - As root user on Cloud:

```
consul members | grep <VMname>
```

Example:

```
[root@gat-emp-0 data]# consul members | grep fmhistory
gat-fmhistory-0 :8301 alive client 0.8.1 2 dc1
gat-fmhistory-1 :8301 alive client 0.8.1 2 dc1
[root@gat-emp-0 data]#
```

If they are not online, the VM deployment is the issue. Bring the VMs online and then recheck their status:

```
$ hagrp -online <groupName> -sys <MasterServerName>
```

Example:

```
[root@cloud-svc-1 litp-admin]# hagrp -online Grp_CS_svc_cluster_fmhistory -s
ys cloud-svc-3
```

Cause 2

JBoss service is not running.

Solution

1. Check the hostname of VMs (`fmhistory`, `dlms`, `fmalarmprocessing`, `fmserv`, `msfm`):
   - On Physical:

     ```
     cat /etc/hostname | grep <VMname>
     ```

     Example:

     ```
     [root@cloud-svc-1 litp-admin]# cat /etc/hosts | grep fmhistory
     10.247.246.63   svc-3-fmhistory fmhistory-1-internal    # Created by LI
     TP. Please do not edit
     ```

   - As root user on Cloud:

     ```
     consul members | grep <VMname>
     ```

     Example:

     ```
     [root@gat-emp-0 data]# consul members | grep fmhistory
     gat-fmhistory-0 :8301 alive client 0.8.1 2 dc1
     gat-fmhistory-1 :8301 alive client 0.8.1 2 dc1
     [root@gat-emp-0 data]#
     ```

2. Log on to any of the VM instances listed in MS.

   Refer to *Connect to a Virtual Machine*.

3. Check that JBoss is running:

   ```
   $ service jboss monitor
   ```

   — If JBoss is not running, start JBoss:

   ```
   $ service jboss start
   ```

   — If JBoss instances are running properly, verify if any exceptions are present in `fmhistory` logs:

   ```
   [root@svc-3-fmhistory ~]# cd /ericsson/3pp/jboss/standalone/log
   [root@svc-3-fmhistory log]# vi server.log
   ```

7.25 Alarms for SNMP-based Nodes Are Not Processed

> Troubleshoot issues after an upgrade, initial install, or a startup procedure, when many SNMP-based nodes are in IDLE state or the alarms are not processed on many SNMP-based nodes.

Prerequisites

Physical deployment prerequisites:

— Access to the Management Server (MS)

— Access to the Service Cluster nodes

Cloud deployment prerequisites:

— Access to the private key file for authentication. Contact your Openstack administrator.

— Access to the Ericsson Management Portal (EMP) VM in the ENM on Cloud deployment.

1. Log on to the ENM management server.
2. Connect to each `mssnmpfm` instance as described in Connect to a Service on page 2.
3. Search the `server.log` file for errors:

   ```
   #grep infinispan /ericsson/3pp/jboss/standalone/log/server.log | grep ERROR
   ```

4. If multiple lines are displayed, collect all the available log data and restart the faulty `mssnmp` instances.

   Note: For information on collecting DEBUG logs, refer to the ENM Data Collection Guidelines.

7.26 ENM Upgrade Fails because of nbfmsnmp Going into Faulted State

This section describes how to troubleshoot an ENM upgrade failing because of `nbfmsnmp` VM going into faulted state.

Prerequisites

— Basic knowledge of Linux and JBoss

— On physical: Access to the Management Server (MS) and Virtual Machines (VMs)

— On Cloud:

- Access to the private key file for authentication (contact your OpenStack administrator)
- Access to the Ericsson Management Portal (EMP) VM in the ENM on OpenStack Cloud deployment

Steps

1. Log on to the ENM MS as the `litp-admin` user and switch to the `root` user.

```
# ssh litp-admin@<Management Server>
# su -
```

2. Verify that the `nbfmsnmp` VM is faulted and that the VM has 4 GB of memory.

   a. Check that the VM has 4 GB of memory:

```
virsh edit nbfmsnmp
```

```
<memory unit='KiB'>4194304</memory>
<currentMemory unit='KiB'>4194304</currentMemory>
```

   b. Check the status of the `nbfmsnmp` VMs:

```
$ hastatus -sum |grep <VMname>
```

Example

```
[root@cloud-svc-1 litp-admin]# hastatus -sum | grep nbfmsnmp
B  Grp_CS_svc_cluster_nbfmsnmp cloud-svc-3          Y          N      →
            OFFLINE|FAULTED
B  Grp_CS_svc_cluster_nbfmsnmp cloud-svc-4          Y          N      →
            OFFLINE|FAULTED
```

3. To clear all the nbfmsnmp VMs from faulted state, execute the following command:

```
hagrp -clear <groupName>
```

Example

```
hagrp -clear Grp_CS_svc_cluster_nbfmsnmp
```

4. Check if the nbfmsnmp VM is now in offline state.

```
[root@cloud-svc-1 litp-admin]# hastatus -sum | grep nbfmsnmp
B  Grp_CS_svc_cluster_nbfmsnmp cloud-svc-3          Y          N              →
   OFFLINE
B  Grp_CS_svc_cluster_nbfmsnmp cloud-svc-4          Y          N              →
   OFFLINE
```

5. Update the nbfmsnmp VM memory to 5 GB.

```
virsh edit nbfmsnmp
```

   a. An XML file opens. Change the value of <memory> and <currentMemory> tag to 5 GB (5242880) and then save the file.

```
<memory unit='KiB'>5242880</memory>
<currentMemory unit='KiB'>5242880</currentMemory>
```

   b. After editing the XML file, offline and online the VM for the change to be reflected.

      Offline the VM:

```
hagrp -offline Grp_CS_svc_cluster_nbfmsnmp -any
```

      After the VM is offline, online the VM:

```
hagrp -online Grp_CS_svc_cluster_nbfmsnmp -any
```

6. Connect to the nbfmsnmp VM instance, and check the VM memory.

   a. Connect to the instance:

```
[root@ms-1 ~]# ssh -i /root/.ssh/vm_private_key cloud-user@svc-1-nb →
fmsnmp
```

b. Check if the change to the memory is reflected:

```
[cloud-user@svc-1-nbfmsnmp ~]$ less /proc/meminfo | grep MemTotal
MemTotal:        4956548 kB
```

7.27 Anomalous Heartbeat Failure Alarms Detected

> Anomalous heartbeat failure alarms can be detected on SNMP nodes due to a corruption on SnmpAddressingTrapCache.

In order to avoid this problem and then remove anomalous heartbeat alarms from SNMP mediations, execute the steps below according to your deployment.

Note: In order to perform stop and start operations in a Cloud environment, you must have access to OpenStack through a GUI or client. In the Cloud Deployment section below, commands apply to OpenStack client but the same is achievable using OpenStack GUI.

Physical Deployment

1. Log on to the ENM Management Server and connect to the SVC-1 node.

```
[root@cloud-ms-1 ~]# ssh litp-admin@svc-1
############  WARNING  #############

This system is for authorised use only. By using this system you consent to  →
monitoring and data collection.

#################################
litp-admin@svc-1's password:
Last login: Thu Oct 10 10:10:29 2019 from cloud-ms-1
[litp-admin@cloud-svc-1 ~]$
```

2. On `SVC-1`, switch to `root` user.

```
[litp-admin@cloud-svc-1 ~]$ su - root
Password:
[root@cloud-svc-1 ~]#
```

3. As `root` user on `SVC-1`, stop `msosssnmpfm` VMs on all blades.

```
hagrp -offline Grp_CS_svc_cluster_msosssnmpfm -any
```

4. As `root` user on `SVC-1`, stop `msapgfm` VMs on all blades.

```
hagrp -offline Grp_CS_svc_cluster_msapgfm  -any
```

5. As `root` user on `SVC-1`, stop `mssnmpfm` VMs on all blades.

```
hagrp -offline Grp_CS_svc_cluster_mssnmpfm  -any
```

6. Wait for the VMs to go `OFFLINE` on all blades.

```
hagrp -state | grep msosssnmpfm
Grp_CS_svc_cluster_msosssnmpfm      State                ieatrcxb4749-1 |OF →
FLINE|
Grp_CS_svc_cluster_msosssnmpfm      State                ieatrcxb4751-1 |OF →
FLINE|

hagrp -state | grep msapgfm
Grp_CS_svc_cluster_msapgfm          State                ieatrcxb4749-1 |OF →
FLINE|
Grp_CS_svc_cluster_msapgfm          State                ieatrcxb4751-1 |OF →
FLINE|

hagrp -state | grep mssnmpfm
Grp_CS_svc_cluster_mssnmpfm         State                ieatrcxb4749-1 |OF →
FLINE|
Grp_CS_svc_cluster_mssnmpfm         State                ieatrcxb4751-1 |OF →
FLINE|
```

7. As `root` user on `SVC-1`, start `msapgfm` VMs on all blades.

```
hagrp -onlineGrp_CS_svc_cluster_msapgfm  -any
```

8. As `root` user on `SVC-1`, start `mssnmpfm` VMs on all blades.

```
hagrp -online Grp_CS_svc_cluster_mssnmpfm  -any
```

9. As `root` user on `SVC-1`, start `msosssnmpfm` VMs on all blades.

```
hagrp -online Grp_CS_svc_cluster_msosssnmpfm  -any
```

10. Wait for the VMs to be `ONLINE` on all blades.

```
hagrp -state | grep msapgfm
Grp_CS_svc_cluster_msapgfm          State                ieatrcxb4749-1 |ON →
LINE|
Grp_CS_svc_cluster_msapgfm          State                ieatrcxb4751-1 |ON →
LINE|

hagrp -state | grep mssnmpfm
Grp_CS_svc_cluster_mssnmpfm         State                ieatrcxb4749-1 |ON →
LINE|
Grp_CS_svc_cluster_mssnmpfm         State                ieatrcxb4751-1 |ON →
LINE|


hagrp -state | grep msosssnmpfm
Grp_CS_svc_cluster_msosssnmpfm      State                ieatrcxb4749-1 |ON →
LINE|
Grp_CS_svc_cluster_msosssnmpfm      State                ieatrcxb4751-1 |ON →
LINE|
```

Cloud Deployment

1. From OpenStack client, stop all `msosssnmpfm` running instances.

a. Retrieve `msosssnmpfm` VM names.

```
openstack server list | grep msosssnmpfm
| 7c55780e-536f-427b-bd0e-5e12af155921 | ieatenmc5b18-msosssnmpfm-0 →
 | ACTIVE | enm_internal_network_AO_C5B18=10.10.0.213 | ERICrhel6jb →
ossimage_CXP9031560-2.70.2_CI.qcow2    | flavor_2vC4M  |             →
                                                                    →

| 39c74ef7-3635-4afc-8a1f-3c83a5a59cfc | ieatenmc5b18-msosssnmpfm-1 →
 | ACTIVE | enm_internal_network_AO_C5B18=10.10.0.215 | ERICrhel6jb →
ossimage_CXP9031560-2.70.2_CI.qcow2    | flavor_2vC4M  |
```

b. Stop `msosssnmpfm` VMs.

```
openstack server stop ieatenmc5b18-msosssnmpfm-0
openstack server stop ieatenmc5b18-msosssnmpfm-1
```

c. Wait for `msosssnmpfm` VMs to be offline (`SHUTOFF`).

```
openstack server show ieatenmc5b18-msosssnmpfm-0 | grep status
| status | SHUTOFF |
openstack server show ieatenmc5b18-msosssnmpfm-1 | grep status
| status | SHUTOFF |
```

2. From OpenStack client, stop all `msapgfm` running instances.

a. Retrieve `msapgfm` VM names.

```
openstack server list | grep msapgfm
| 59ea66d5-9c7f-4082-979d-8268a02525441 | ieatenmc5b18-msapgfm-0 |   →
ACTIVE | enm_internal_network_AO_C5B18=10.10.0.213 | ERICrhel6jboss →
image_CXP9031560-2.70.2_CI.qcow2    | flavor_2vC4M  |               →
                                                                    →

| d190b7d4-7f1f-469d-ae55-5644a23da630  | ieatenmc5b18-msapgfm-1 |   →
ACTIVE | enm_internal_network_AO_C5B18=10.10.0.215 | ERICrhel6jboss →
image_CXP9031560-2.70.2_CI.qcow2    | flavor_2vC4M  |
```

b. Stop `msapgfm` VMs.

```
openstack server stop ieatenmc5b18-msapgfm-0
openstack server stop ieatenmc5b18-msapgfm-1
```

c. Wait for `msapgfm` VMs to be offline (`SHUTOFF`).

```
openstack server show ieatenmc5b18-msapgfm-0 | grep status
| status | SHUTOFF |
openstack server show ieatenmc5b18-msapgfm-1 | grep status
| status | SHUTOFF |
```

3. From OpenStack client, stop all `mssnmpfm` running instances.

a. Retrieve `mssnmpfm` VM names.

```
openstack server list | grep mssnmpfm
| 2709de1f-2c63-41d6-8dfd-1a69a6be8258 | ieatenmc5b18-mssnmpfm-0 |   →
ACTIVE | enm_internal_network_AO_C5B18=10.10.0.213 | ERICrhel6jboss →
image_CXP9031560-2.70.2_CI.qcow2    | flavor_2vC4M  |               →
                                                                    →
```

```
| 58a24898-ef9e-4bc8-83f6-1f6d693ae2e0 | ieatenmc5b18-mssnmpfm-1 |
ACTIVE | enm_internal_network_AO_C5B18=10.10.0.215 | ERICrhel6jboss
image_CXP9031560-2.70.2_CI.qcow2    | flavor_2vC4M  |
```

b. Stop `mssnmpfm` VMs.

```
openstack server stop ieatenmc5b18-mssnmpfm-0
openstack server stop ieatenmc5b18-mssnmpfm-1
```

c. Wait for `mssnmpfm` VMs to be offline (`SHUTOFF`).

```
openstack server show ieatenmc5b18-mssnmpfm-0 | grep status
| status | SHUTOFF |
openstack server show ieatenmc5b18-mssnmpfm-1 | grep status
| status | SHUTOFF |
```

4. From OpenStack client, start all `msapgfm` running instances.

a. Start `msapgfm` VMs.

```
openstack server start ieatenmc5b18-msapgfm-0
openstack server start ieatenmc5b18-msapgfm-1
```

b. Wait for `msapgfm` VMs to be online(`ACTIVE`).

```
openstack server show ieatenmc5b18-msapgfm-0 | grep status
| status | ACTIVE |
openstack server show ieatenmc5b18-msapgfm-1 | grep status
| status | ACTIVE|
```

5. From OpenStack client, start all `mssnmpfm` running instances.

a. Start `mssnmpfm` VMs.

```
openstack server start ieatenmc5b18-mssnmpfm-0
openstack server start ieatenmc5b18-mssnmpfm-1
```

b. Wait for `mssnmpfm` VMs to be online (`ACTIVE`).

```
openstack server show ieatenmc5b18-mssnmpfm-0 | grep status
| status | ACTIVE|
openstack server show ieatenmc5b18-mssnmpfm-1 | grep status
| status | ACTIVE|
```

6. From OpenStack client, start all `msosssnmpfm` running instances.

a. Start `msosssnmpfm` VMs.

```
openstack server start ieatenmc5b18-msosssnmpfm-0
openstack server start ieatenmc5b18-msosssnmpfm-1
```

b. Wait for `msosssnmpfm` VMs to be online (`ACTIVE`).

```
openstack server show ieatenmc5b18-msosssnmpfm-0 | grep status
| status | ACTIVE |
```

```
openstack server show ieatenmc5b18-msosssnmpfm-1 | grep status
| status | ACTIVE|
```

7.28 fmalarmprocessing VM Queue Is Full

This section describes how to troubleshoot when an `fmalarmprocessing` VM queue (`<<Queue-Name>>`) is full and the following error message is seen in its `server.log`:

```
javax.ejb.EJBException: java.lang.RuntimeException: javax.jms.JMSException: HQ[N →
UM]: Address "jms.queue.<<Queue-Name>>" is full.
```

Diagnostics

To get the full stack trace of the error and troubleshoot the underlying problem, enable and collect the logs of the following packages from all the available `fmalarmprocessing` instances:

— `org.hornetq.core.client`

— `com.arjuna.ats.jta`

— `org.jboss.as.ejb3.invocation`

To proceed with the remedial steps below, you must have:

— Basic knowledge of Linux and JBoss

— On Physical:

  • Access to the Management Server (MS) and Virtual Machines (VMs)

— On Cloud:

  • Access to the private key file for authentication (contact your Openstack administrator)

  • Access to the Ericsson Management Portal (EMP) VM in the ENM on Cloud deployment

Solution

1. Connect to the `fmalarmprocessing` instance in which the logs must be collected. For more information, see Connect to a Service.
2. Switch to `root` user to be able to use the `jboss-cli`:

```
[cloud-user@svc-1-fmalarmprocessing ~]$ sudo su
```

3. Connect to the `jboss-cli`:

```
[root@svc-1-fmalarmprocessing cloud-user]# /ericsson/3pp/jboss/bin/jboss-cli →
.sh -c
```

4. Execute the following commands in `jboss-cli` to enable the log levels to `INFO`:

```
[standalone@localhost:9999 /] /subsystem=logging/logger=org.jboss.as.ejb3.in →
vocation:change-log-level(level=INFO)
{"outcome" => "success"}
[standalone@localhost:9999 /] /subsystem=logging/logger=com.arjuna.ats.jta:c →
hange-log-level(level=INFO)
{"outcome" => "success"}
[standalone@localhost:9999 /] /subsystem=logging/logger=org.hornetq.core.cli →
ent:change-log-level(level=INFO)
{"outcome" => "success"}
```

5. Repeat the previous steps for all the available `fmalarmprocessing` instances.

6. Open `server.log` using the following commands:

```
[root@svc-1-fmalarmprocessing cloud-user]$ cd /ericsson/3pp/jboss/standalone →
/log/
[root@svc-1-fmalarmprocessing log]$ less server.log
```

7. Check for the stack trace (see the following example) in the `server.log`. When the stack trace is present, proceed to the next step (wait some moments for the stack trace).

```
ERROR [org.jboss.as.ejb[NUM].invocation] (Thread-[NUM] (HornetQ-client-globa →
l-threads-[NUM])) JBAS[NUM]: EJB Invocation failed on component FmxEventSend →
er for method public abstract void com.ericsson.oss.services.fm.alarmprocess →
or.api.alarmsender.AlarmSender.sendAlarm(com.ericsson.oss.services.fm.models →
.processedevent.ProcessedAlarmEvent): javax.ejb.EJBException: java.lang.Runt →
imeException: javax.jms.JMSException: HQ[NUM]: Address "jms.queue.FmxAdaptor →
ServiceQueue" is full.
 at org.jboss.ejb.client.EJBClientInvocationContext.sendRequest(EJBClientInv →
ocationContext.java:184) [jboss-ejb-client-1.0.41.Final-redhat-1.jar:1.0.41. →
Final-redhat-1]
 at com.ericsson.oss.itpf.sdk.context.core.interceptor.ClientContextIntercep →
tor.handleInvocation(ClientContextInterceptor.java:51) [sdk-context-core-3.6 →
2.1.jar:3.62.1]
 at org.jboss.ejb.client.EJBClientInvocationContext.sendRequest(EJBClientInv →
ocationContext.java:186) [jboss-ejb-client-1.0.41.Final-redhat-1.jar:1.0.41. →
Final-redhat-1]
```

8. Collect the logs and provide them to the respective CSR or design team for further troubleshooting.

9. Repeat this step for all the available `fmalarmprocessing` instances.

   Reconnect to `jboss-cli` (see step 3) and disable the log by executing the following commands:

```
[standalone@localhost:9999 /] /subsystem=logging/logger=org.hornetq.core.cli →
ent:change-log-level(level=OFF)
{"outcome" => "success"}
[standalone@localhost:9999 /] /subsystem=logging/logger=com.arjuna.ats.jta:c →
hange-log-level(level=OFF)
```

```
{"outcome" => "success"}
[standalone@localhost:9999 /] /subsystem=logging/logger=org.jboss.as.ejb3.in →
vocation:change-log-level(level=OFF)
{"outcome" => "success"}
```

7.29 Troubleshooting SIU02 and TCU02 Nodes

The troubleshooting steps recommended to diagnose and fix common problems in FM application for the following nodes, when FM supervision is enabled but no alarms are reported in ENM.

— SIU02

— TCU02

In general, the SIU02 and TCU02 nodes support single transactions at a time. This is a node limitation. The following scenario is applicable for this limitation.

If CM supervision is enabled from ENM, the transaction is acquired by CM supervision process, and is not available for any other processes until this transaction is released. During the CM supervision enabling, if FM supervision is enabled, the supervision request is not sent to the node because transaction is not available. However, the supervision status is set to ON in ENM. If any alarms are raised on the node, these alarms are not received in ENM because the supervision is still not enabled on the node. The failure scenario No Alarms are Reported to ENM even though FM Supervision is Enabled is applicable for this limitation.

The same applies to the following scenarios:

— PM supervision and FM supervision enabled in parallel.

— CM supervision, PM supervision and FM supervision enabled in parallel.

There are other failure scenarios observed which are not due to this limitation. These failure scenarios follow:

7.29.1 Failure Scenarios

Prerequisites

— The user is logged on to ENM GUI.

— User has permissions to enable FM supervision and view Alarms.

Expected Result

The FM supervision on SIU02 or TCU02 nodes is enabled, and alarms raised on node are reported properly in ENM.

Contact Ericsson support for further assistance if alarms are not still reported in ENM after performing the following steps for mentioned failure scenarios.

7.29.1.1 No Alarms are Reported to ENM even though FM Supervision is Enabled

Steps

1. If FM supervision is enabled, disable FM supervision from ENM .
2. Enable FM supervision again from ENM.
3. Raise alarms on Node, then check whether the alarms are reported properly in ENM. In addition to the newly raised alarms, the previously raised alarms also will be reported in ENM as Synchronized alarms.
4. If the raised alarms are still not reported, then repeat the previous steps.

7.29.1.2 Heartbeat Failure Alarms Reported to ENM

Steps

1. Verify the SNMPAgent port value for the SIU02 or TCU02 node's connectivityInformation added in ENM via CM CLI or Network Explorer. (Please note that setting SNMPAgent port is optional in ENM for ConnectivityInformation. By default the port will be set to 161 if no SNMPAgent port is provided. It will be overwritten with the value provided if any) . The following CM CLI command verifies the SNMPAgent Port.

   ```
   >>cmedit get NetworkElement=<name of the NetworkElement>,STNConnectivityInfo →
   rmation=1
   ```

2. If the SNMPAgent port is other than 161, set this attribute to 161 via CM CLI using cmedit commands. The following CM CLI command sets the SNMPAgent port.

   ```
   >>cmedit set NetworkElement=<name of the NetworkElement>,STNConnectivityInfo →
   rmation=1 snmpAgentPort=161
   ```

3. Disable FM supervision in ENM.
4. Enable FM supervision again from ENM.
5. Heartbeat failure alarm seen earlier should get cleared for this Network Element within 120 or 240 seconds in ENM.

8 ENM System Monitor Troubleshooting

> This section provides troubleshooting steps to diagnose common problems with ENM System Monitor (ESM).

8.1 Access ENM System Monitor When ENM Is Not Available

> This task describes the steps to access ENM System Monitor GUI when ENM goes down or when ENM launcher page is not accessible.

All the three users of ESM can access this GUI. There are cases when the users are unable to launch ESM from ENM Launcher page. This situation mostly occurs when ENM is down. In those cases, users need to follow the steps to access ESM.

Prerequisites

— The user is an Administrator or Alert Administrator.

— ESM password is available from Local Ericsson Support

Steps

1. Open the URL of ENM System Monitor in each supported browser:

   `https://<esmon ip>:7443/ - Secured login`

   (or)

   `http://<esmon ip>:7080/`

   Type the following command on MS to get the `esmon` IP:

   ```
   cat /etc/hosts | grep esmon
   ```

2. Enter the username as `esmadmin` and password supplied by local Ericsson support

Results

The user can log on to ESM application and ESM dashboard is displayed.

8.2 Cannot Access ESM GUI

> Check that the Virtual Machine partition and `/var/ericsson/esm-data` have enough disk space.

Prerequisites

The user must have SSH access to the MS and `esmon`.

Cause 1

— Virtual Machine is paused.

— System reports `100%` disk usage.

Solution

1. Verify that the Virtual Machine (VM) for ESM is running:

   ```
   [root@cloud-ms-1 ~]# virsh list | grep esmon
   ```

2. Check the amount of free disk space on the VM partition if the VM is paused:

   ```
   [root@cloud-ms-1 ~]# df -h | grep /var
   ```

   Example

   ```
   44G   39G  2.2G  95% /var
   ```

3. Free up space if the system reports `100%` disk usage:

   ```
   [root@cloud-ms-1 ~]# rm -rf /var/tmp/*
   ```

4. Confirm that space is available and then resume the VM:

   ```
   [root@cloud-ms-1 ~]# virsh resume esmon
   ```

Cause 2

— `/var/ericsson/esm-data` is full.

— `rhq-storage.log` contains an entry about reducing memtable and cache sizes:

```
Heap is 0.759394416942797 full. You may need to reduce memtable and/or cache →
 sizes. Cassandra will now flush up to the two largest memtables to free up  →
memory. Adjust flush_largest_memtables_at threshold in cassandra.yaml if you →
 don't want Cassandra to do this automatically
```

Solution

1. Log on to `esmon` using SSH.
2. Delete old data (more than 10 days old) in `/var/ericsson/esm-data/rhq-data/storage/data/rhq/raw_metrics`.
3. Log on to the MS.
4. Log on to `esmon` again.
5. Restart `rhq-storage-server` and `rhq-server`:

```
service rhq-storage-server restart
service rhq-server restart
```

8.3 Resource Is Reported as Unavailable Though Accessible with Other Means

ENM System Monitor (ESM) reports live resources as unavailable when they are accessible. All the three users of ESM are able to view the Resources. Sometimes the resources are unavailable as the agents are down. In that case, follow the next procedure to reestablish the resources as available.

Prerequisites

— The user has SSH access to all systems in the ENM network.

— Perform some alternative checks to ensure that the resource is up (among others, SSH, and ENM GUI).

Steps

1. Verify if the RHQ Agent is running on the affected platform:

```
[root@cloud-db-1 ~]# service rhq-agent status
```

   a. If the service is down, restart it:

```
[root@cloud-db-1 ~]# service rhq-agent restart
```

2. Wait for five minutes for ESM to reconnect.

Results

The resource reconnects to ESM.

8.4 Access ESM When Maintenance Mode Is Displayed (Intermittent)

Occasionally when ESM users try to access the GUI, a window is displayed on the top of the GUI stating that ESM is in maintenance mode. This occurs because storage server is either currently down or was down previously at some point. To recover, follow the recommended procedure to restart the storage server.

Prerequisites

— The user is an Administrator

Steps

1. Log on to `esmon` VM as cloud-user
2. Switch to root user
3. Restart both `rhq-storage-server` and `rhq-server` in sequence.

```
service rhq-storage-server restart
service rhq-server restart
```

Results

The user can log on to ESM application and ESM dashboard is displayed.

8.5 Agents Are Offline Because of a Missing Security Token

This section describes the steps to do when the agents are offline because of a missing security token in `prefs.xml` (`/var/ericsson/esm-data/rhq-agent/prefs/.java/.userPrefs/rhq-agent/default/prefs.xml`) file of agent. Users are able to see the agents are offline on the GUI because of a missing security token. When an agent is installed, a unique security token is created for it, which it registers with the server. All subsequent communication is done using the security token. If the security token is lost or changed, the agent is not able to connect to the server.

Note: There are two cases in which the security token is lost:

1. Restarting the agent with

```
fullcleanconfig
```

2. Unauthorized config file updates where the security token is stored.

Prerequisites

— The user has SSH access to all systems in the ENM network.

Steps

1. Log on to the management server.
2. Go to the corresponding node (SVC, MS, DB, or ESMON) where the agent is down.

   To log on the VM, see Connect to a Service on page 2.

   Example

   Log on to ESMON VM as:

```
ssh -i /root/.ssh/vm_private_key cloud-user@esmon
```

   (or)

   Log on as `litp-admin` to the nodes as,

```
ssh litp-admin@<node>
```

3. Switch to the root user account
4. Go to `/opt/rhq/rhq-agent/logs/agent.log`.
5. Check if there are any Agent registrations errors because of a missing security token are present or not.

```
2016-08-03 11:34:14,736 ERROR [RHQ Agent Registration Thread] (org.rhq.enter →
prise.agent.AgentMain)- {AgentMain.agent-registration-rejected}The server ha →
s rejected the agent registration request. Cause: [org.rhq.core.clientapi.se →
rver.core.AgentRegistrationException:The agent [ieatrcxb6222] is attempting  →
to re-register without a security token. Please consult an administrator to  →
obtain the agent's proper security token and restart the agent with the opti →
on "-Drhq.agent.security-token=<the valid security token>". An administrator →
 can find the agent's security token by navigating to the GUI page "Administ →
ration (Topology) > Agents" and drilling down to this specific agent. You wi →
ll see the long security token string there. For more information, read: htt →
ps://docs.jboss.org/author/display/RHQ/Agent+Registration]
```

6. Stop the puppet

```
service puppet stop
```

7. If the security token is lost, open the `prefs.xml` file:

```
vi /var/ericsson/esm-data/rhq-agent/prefs/.java/.userPrefs/rhq-agent/default
/prefs.xml
```

8. Add the following entry to `prefs.xml` file:

```
<entry key="rhq.agent.security-token" value="token for that corresponding ag
ent taken from Admin Tab of GUI"/>
```

9. Start `rhq-agent`:

```
service rhq-agent start
```

10. Start the puppet:

```
service puppet start
```

Results

Agents are on lined when security token is reinstated.

8.5.1 Get Security Token from Admin Tab of ESM GUI

1. Open ESM GUI.
2. Click Administration.
3. Click Agents
4. Click on the corresponding agent which is offline and copy its corresponding token

Results

The token is displayed as : tr5lesryJiqEu6Zke7+L7e9p92/H0bdhili/
iWtXI2onYNo8gb1ouLu8mwi5plhRU3o=

8.6 Issues Displaying Charts / Data in ESM

This section describes the steps to be followed when there is mismatch in the graphs/charts for two different users in ESM

Prerequisites

— Users should have credentials to log on to the ESM GUI.

Steps

1. Two users need to log on to ESM GUI. Click Inventory > Platforms. Click on one of the metric and open on of the metrics chart.
2. Both users need to refresh the page simultaneously to have the same data represented at that particular time.

   Example

   Below is the example which has the dispcrepancy in the charts for two different users.

Results

After the successful execution of the below steps, the data / charts will be aligned in such a way that there is no discrepancy.

8.7 Scenarios Where Application Common Metrics Are Not Visible

This section describes the possible scenarios where common metrics are not collected. These common metrics are visible to all the users of ESM and no specific permissions are necessary to view them. ESM collects Common Metrics

for the VMs where the JBoss is running. These metrics are collected on a per minute basis.

Scenarios:

1. For Master Server and Blades (as JBoss does not run on Master Server and Blade servers, ESM does not collect the common metrics for these servers).
2. For ESMON VM
3. In the absence of the `genjmx` XML file.

8.7.1 Check the Absence of genjmx XML File

Prerequisites

The user is logged on to ESM.

Steps

1. Log on to the server with putty.
2. Log on to the `esmon` VM.

   See *VM Security Tasks* in [2] and Connect to a Service on page 2.

   ```
   # ssh -i /root/.ssh/vm_private_key cloud-user@esmon
   ```

3. Switch to root user:

   ```
   # sudo su root
   ```

4. Go to corresponding server folder (Master Server or Blade servers):

   ```
   cd /var/ericsson/esm_ddc_data
   ls
   cd <{host_name}_TOR>
   ```

5. Check the presence of `genjmx` XML for the corresponding server:

   ```
   cd <current_date folder>/instr
   Ex: cd 260816/instr
   ls
   ```

   If no `genjmx` XML is present, ESM does not collect the common metrics for that corresponding server.

Results

Common metrics are not visible.

8.8 Platforms Duplicated in ESM GUI

ESM users are able to view the Platforms using the **Inventory** tab of ESM GUI. In some cases, when the load on the server is very high, two platforms are created for the same agent. This is not recommended and needs to be resolved.

Prerequisites

The user is logged on to `postgresql` with the password supplied by local Ericsson support.

Steps

1. Log on to the Management Server as the `litp-admin` user using SSH.
2. Log on to the `esmon` VM:

   ```
   ssh -i /root/.ssh/vm_private_key cloud-user@esmon
   ```

   See *VM Security Tasks* in [2] and Connect to a Service on page 2.
3. Log on to `postgresql` using the following command in the `esmon` VM:

   ```
   #psql -U rhqadmin -d rhq
   ```
4. Enter the password.
5. Execute the following command to get the duplicated agent information and copy the agent ID:

   ```
   select * from rhq_agent where description like '%Java capability%';
   ```
6. Find the resource information with description "Platform with 100% Java capability but not natively supported". Copy the ID.

   ```
   select id, resource_type_id, parent_resource_id  from rhq_resource where age →
   nt_id = <agent id copied from above query> and description like '%Java capab →
   ility%';
   ```
7. Find the resource information with description "Linux". Copy the parent resource ID (that is, the actual Linux platform parent resource ID).

   ```
   select id, resource_type_id, parent_resource_id  from rhq_resource where nam →
   e = 'duplicate_platform_name' and description like '%Linux%';
   ```
8. Execute the following commands to change the resource type ID to another type, such as '10016' (file system). Also set the parent resource ID. Update the row.

```
update rhq_resource set resource_type_id=10016 where id=<id of java platform →
>
update rhq_resource set parent_resource_id=<parent resource id of Linux Plat →
form> where id=< id of java platform>
```

9. Go to ESM GUI and click **Inventory** > **All Resources**.
10. Go to the resource which has the description "Platform with 100% Java capability but not natively supported". Click on it
11. Click **Uninventory** for the particular resource.
12. After few minutes, click **Refresh**.

    Result: The resource is no longer displayed.

Results

Duplicated platforms are removed

8.9 Incorrect Data Displays in Metrics Graphs in ENM System Monitor (ESM)

> The tool fails to retrieve and graph the latest metrics data from the database.

Prerequisite

User is logged on to the ESM GUI.

Cause

Metrics data is not retrieved from the database.

Solution

— Switch from the Monitoring to Summary tab.

Result: The tabs refresh and metrics data from the database loads in graphs.

8.10 Esmon Access When Underlying File System Errors (Postgres Data Corruption) Occur

This section describes the steps to be followed to access the `esmon` after fixing errors with the Postgres Database (DB corruption) and when it has not been possible to recover data.

Prerequisites

— ENM System Monitor (ESM) GUI is not launching due to Postgres database issues.

  This is caused by underlying file system errors and it is not possible to recover data (that is, all the RHQ-related data in the Postgres database is removed, but the Cassandra data remains, which is useless without Postgres data).

— The user is logged on to the ENM System - `esmon` VM.

Steps

1. Log on to the `esmon` using SSH.
2. Log on to the `esmon` VM.

   See *VM Security Tasks* in [2] and Connect to a Service on page 2.

   ```
   # ssh -i /root/.ssh/vm_private_key cloud-user@esmon
   ```

3. Switch to `root` user:

   ```
   # sudo su root
   ```

4. Go to the appropriate folder:

   ```
   cd /var/ericsson/esm_data
   ```

5. Remove all folders present at the location:

   ```
   rm -rf *
   ```

6. Reconstitute the VM

   a. `service esmon stop`

   b. `virsh undefine esmon`

   c. `service esmon start`

   Result:

   The `esmon` GUI launches with no historical data. ESM starts collecting metrics from this point in time only.

8.11 Wrong Error Message When Log In Credentials are Invalid

ESM UI displays an incorrect error message when a user tries to log in with either invalid username or password.

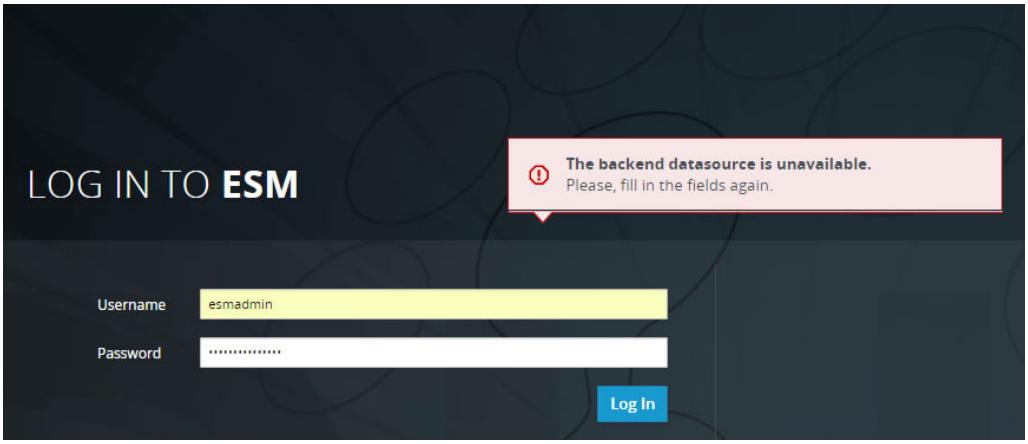

Figure 4 Invalid Username or Password Brings Up Incorrect Message

The alert message, `The backend datasource is unavailable` is incorrect. This incorrect message is due to a known issue with the combination of JDK 8.0 with RHQ 4.13.1.

For more information, refer to https://bugs.openjdk.java.net/browse/JDK-8155223.

8.12 ESM Shows Exception Report Under Timeline Tab

ESM user interface displays an exception report when the user tries to check **Summary** > **Timeline**. Exception "`Unable to compile class for JSP`" is expected .

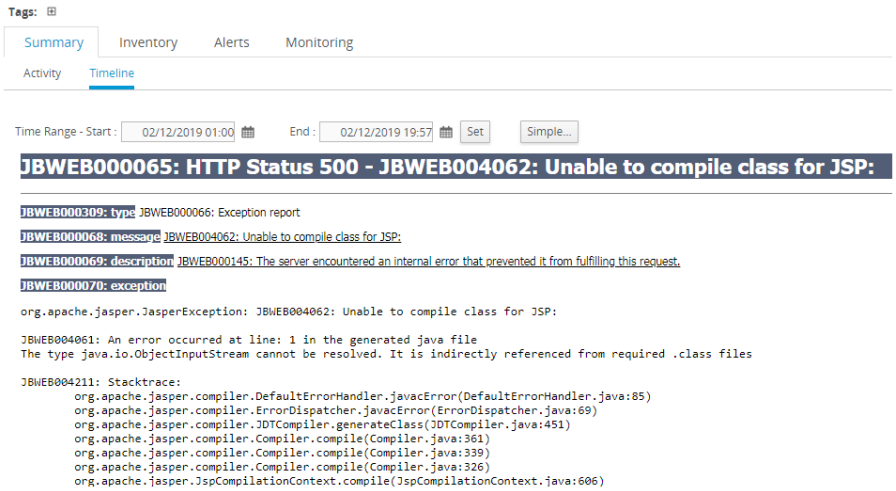

Figure 5 Timeline Error

Cause

This is a known 3PP issue with the combination of JDK8 and RHQ 4.13.1.

Solution

Please refer to https://bugs.openjdk.java.net/browse/JDK-8155223.

8.13 Expected Spikes in CPU or Memory Usage

> Metric data in ESM is aggregated to show averages.

Aggregations are done by several RHQ internal schedulers that run on the hour.

While the schedulers are running, the resource usage (CPU, memory) rises and you can see spikes in the CPU usage. This is normal and can be ignored.

For instrumentation data, there is a scheduler job that runs daily at 04:30. This job checks if any new service group has been added or if there are any new instrumentation beans exposed as EXTERNAL. If so, the job includes the related metrics in the collection and shows them in the ESM GUI. It is expected that CPU and memory usage rise while the job is running.

8.14 VM Clean Start Because of CacheBindExeption during Startup

> Service Group clean start is expected if the cache is corrupted during startup.

Diagnostics

During Service Group startup:

If the multicast port used by a cache is taken by another process, that cache fails to start and the cache manager becomes corrupted.

Solution

1. To solve this problem, the port must be freed up and the cache restarted.

   Service groups in ENM that have property - `Dsdk.cacheCorruptionExitFlag` set to `true` cause the application to exit when this state is detected by ServiceFramework, and Service Group VMs go for a restart.

   After a restart, it is expected that SG cache will start without any issues.

8.15 Unavailable Resources Can be Ignored

> Some of the plugins included in the main distribution of the RHQ software package are not suitable or not applicable for the ENM environment.

Most of the plugins are deleted when RHQ is installed in the `esmon` VM. Due to the nature of plugins in RHQ, the post-install script can disable them only when they are running in RHQ. Before the script executes, RHQ tries to find the resources and displays them in the `esmon` GUI. Some are shown as available and some as unavailable.

The resources with the following names (and names which have the following words), including their child resources, can be ignored - whether they are available or unavailable.

— EAP

— JBossAS7

— Hosts

— Samba

— Postfix

— OpenSSH

— Cobbler

— GRUB

— Aliases

— Apache

— Cron

— Sudoers

— Augeas

— Postgres

— Twitter

— hudson

— IIS

— mod_cluster

— Oracle

The postgres plugin has been disabled. This is because it opens so many connections in the DB nodes as to cause connections to be exhausted and other applications to fail. So, if the Postgres server is shown as unavailable, ignore the resource.

Some resources (showing instrumentation of the JBoss applications) are shown as unavailable. This can happen when the JBoss application is not exporting any metrics or DDC is not collecting them. It can also happen when the VM goes offline. These metrics need to stay in the `esmon` GUI to show the historical values of the instrumentation metrics. They cannot be deleted and cannot be combined with the other instances of the same service. For example, `esmon` will not show the values of `svc-1-impexpserv` in the resource `svc-2-impexpserv`. This can also be ignored as the `esmon` sends alarms to FM when all the instances of a particular service group are offline.

For more information, see *Access VM Availability Metrics* in the ENM System Monitor User Guide.

There is an option to disable the following list of plugins which are not used by ESMON and are ESMON internal instrumentation plugins.

This can be done by executingthe script disableRhqPlugins.sh in Esmon VM, plugins are disabled if they are in enabled mode.

Table 1 Plugin name and Status

Plugin Name	Status
Ant Bundle	Not Used by ESMON
Hibernate Services	Not Used by ESMON
Jboss APplication Server 5x/6x	Not Used by ESMON
Jboss Cache 2.x services	Not Used by ESMON
Jboss Cache 3.x Services	Not Used by ESMON
Performance test plugin	Not Used by ESMON
Receiver fot SNMP TRAPS	Not Used by ESMON
Script	Not Used by ESMON
Tomcat	Not Used by ESMON
Cassandra Internal RHQ instrumental	Not Used by ESMON
Jboss APplication Server 7x Internal RHQ instrumental	Not Used by ESMON
Rhq-Agent Internal RHQ instrumental	Not Used by ESMON
Rhq-Server Internal RHQ instrumental	Not Used by ESMON
Rhq Storage Internal RHQ instrumental	Not Used by ESMON

8.15.1 Steps to Execute the Script to Disable Plugins

Follow these steps to execute the script to disable the plugins:

Steps

1. Login to `Esmon VM` as root user.
2. Navigate to the following path:

   ```
   [root@ms-esmon bin]# pwd/opt/ericsson/ERICesm/bin
   ```

3. Execute the following script:

   ```
   [root@ms-esmon bin]# ./disableRhqPlugins.sh
   ```

   Result: Now the plugins are disabled.

8.15.2 Steps to Enable Plugins

Follow these steps to enable the plugins:

1. Login to `ESM GUI`.
2. Under **Administration**, go to **Configuration**.
3. Select **Agent Plugins**.
4. Select the plugin, which is to be enabled from the list of plugins.
5. Click **Enable** and then select **Yes**.

   Result: The requested plugin is enabled.

8.16 Enable or Disable Unused Platforms/Plugins from ESM GUI

8.16.1 Enable the Platform/Plugins

Follow these steps to enable the platform/plugins:

1. Login to `ESM GUI`.
2. Under **Administration**, go to **Configuration**.
3. Select **Ignored resource types**.
4. Select the resource type from the list of Platform Services or Servers, which is to be enabled from ESM GUI.
5. Click **Change** and then select **Yes**.

   Result: The requested plugin is enabled.

8.16.2 Disable the Unused Platform/Plugins

Follow these steps to disable the unused platform/plugins:

1. Login to `ESM GUI`.
2. Under **Administration**, go to **Configuration**.

3. Select **Ignored resource types**.
4. Select the resource type from the list of Platform Services or Servers, which is to be disbled from ESM GUI.
5. Click **Change** and then select **Yes**.

   Result: The requested plugin is disabled.

8.17 Irrelevant RHQ Alerts

> RHQ fires alerts when certain disk usage parameters are exceeded.

However, the predefined alerts are configured to calculate disk usage when only the Storage server is using the disk. In fact, the disk is being used by:

— Server

— Storage Server

— Agent

— Postgres

For this reason, the following RHQ alerts can be ignored:

— StorageNodeHighHeap

— StorageNodeHighDiskUsage

9 Autonomic Incident Management Troubleshooting

This section provides the troubleshooting steps recommended to diagnose and fix common problems in the ENM Autonomic Incident Management.

9.1 Install Valid License for Autonomic Incident Management

A system administrator can perform these steps when the Valid License Required error message is seen. The application cannot be used when this error is shown.

Prerequisites

— Must have a valid Autonomic Incident Management license.

— User has `Lcm_Administrator` and `Cmedit_Operator` roles.

Steps

1. On the ENM Launcher page, click **Command Line Interface (CLI)**.
2. Drag and drop the Autonomic Incident Management license file into (CLI).
3. Execute the command:

```
lcmadm install file:<Autonomic Incident Management license name>.txt
```

4. Navigate to the landing page of Autonomic Incident Management using the following address:

```
<<ENM UI IP address / fully qualified domain name>>/#incidentmanagement
```

Results

It is possible to launch Autonomic Incident Management application.

9.2 Unable to Access Autonomic Incident Management - Service Unavailable

A system administrator can perform these steps when Autonomic Incident Management GUI cannot be accessed.

Prerequisites

— User requires basic knowledge of Linux and JBOSS.

— User requires access to the Management Server (MS) and Virtual Machines (VMs).

Expected Result

User is able to troubleshoot the issue if the user is unable to access Autonomic Incident Management GUI.

Steps

1. Verify that the following VMs are online:

   Note: Number of available VMs are different depending on scale of ENM deployed.

   — `imadserv`

   — `imadserv_2`

   — `imadserv_3`

   — `imadserv_4`

   — `imfmalarmserv`

   — `imgroupingserv`

   — `imgroupingserv_2`

   — `imkbserv`

   — `imlcserv`

      a. Log on to the Management Server as the `litp-admin` user using SSH.

      b. Log on to the Automation Cluster as `litp-admin` user and change to `root` user.

```
[litp-admin@ms-1 root]$ ssh litp-admin@aut-1
############ WARNING #############


This system is for authorised use only. By using this system you co →
nsent to monitoring and data collection.


###################################
litp-admin@aut-1's password:

[litp-admin@aut-1 ~]$ su
```

```
Password:

[root@aut-1 litp-admin]#
```

Result: Successfully logged on to the Automation Cluster as `litp-admin` user and change to `root` user.

c. Execute the following command to check if all VMs are online.

```
 $hastatus -sum | grep <VMname>
```

Example

```
[root@aut-1 litp-admin]# hastatus -sum | grep 'imlcserv\|imadserv\| →
imfmalarmserv\|imkbserv\|imgroupingserv'

B Grp_CS_aut_cluster_imadserv aut-1 Y N ONLINE
B Grp_CS_aut_cluster_imadserv_2 aut-1 Y N ONLINE
B Grp_CS_aut_cluster_imadserv_3 aut-1 Y N ONLINE
B Grp_CS_aut_cluster_imadserv_4 aut-1 Y N ONLINE
B Grp_CS_aut_cluster_imfmalarmserv aut-1 Y N ONLINE
B Grp_CS_aut_cluster_imgroupingserv aut-1 Y N ONLINE
B Grp_CS_aut_cluster_imgroupingserv_2 aut-1 Y N ONLINE
B Grp_CS_aut_cluster_imkbserv aut-1 Y N ONLINE
B Grp_CS_aut_cluster_imlcserv aut-1 Y N ONLINE
```

d. If VMs are offline, bring them online:

```
$hagrp -online <Grp_CS_aut_cluster_X> -sys <MasterServerName>
```

Example

```
[root@aut-1 litp-admin]# hagrp -online Grp_CS_aut_cluster_imlcserv →
-sys aut-1

[root@aut-1 litp-admin]# hagrp -online Grp_CS_aut_cluster_imadserv →
-sys aut-1
```

2. Verify that JBoss is running.

   If VMs are online, proceed as follows:

   a. On physical, use `cat /etc/hosts | grep <VMname>`:

      Example

```
[root@ms-1 ~]# cat /etc/hosts | grep imlcserv

<IpAddress> aut-1-imlcserv imlcserv-1-internal # Created by LITP. P →
lease do not edit
```

   b. To log on to available VM instances (`imadserv`,`imadserv_2`,`imadserv_3`, `imadserv_4`, `imfmalarmserv`, `imgroupingserv`,`imgroupingserv_2`, `imkbserv`,`imlcserv`) listed from MS, see *VM Security Tasks* in [2] and Connect to a Service on page 2.

```
[root@ms-1 ~]# ssh -i /root/.ssh/vm_private_key cloud-user@aut-1-im →
lcserv
[cloud-user@aut-1-imlcserv ~]$ sudo su root
[root@aut-1-imlcserv cloud-user]#
```

c. Check if the JBoss service is running.

```
$ service jboss monitor
```

d. If JBoss is not running, start the service by executing the following command:

```
service jboss start
```

3. Analyze VM logs.

   If issues continue, investigate the logs of each VM for further analysis and identify any exceptions that are being thrown.

   a. Log on to available VM instances (`imadserv`,`imadserv_2`,`imadserv_3`, `imadserv_4`, `imfmalarmserv`, `imgroupingserv`,`imgroupingserv_2`, `imkbserv`,`imlcserv`) listed from the MS, see *VM Security Tasks* in [2] and Connect to a Service on page 2.

      Example

```
[root@ms-1 ~]# ssh -i /root/.ssh/vm_private_key cloud-user@aut-1-im →
lcserv
[cloud-user@aut-1-imlcserv ~]$ sudo less /ericsson/3pp/jboss/standa →
lone/log/server.log
```

9.3 Identify Untrained Nodes or Cells and KPIs

> A system administrator can perform these steps to identify the node and cell trained state in Autonomic Incident Management.

Prerequisites

— Autonomic Incident Management is set up with nodes and KPIs selected to monitor.

Note: A trained node or cell takes a minimum of 18 days to train which is equivalent to 1728 fifteen minute rops.

— The user can log on to `postgresql` with the password supplied by local Ericsson support.

Steps

1. Log on to the Management Server as the `litp-admin` user using SSH.
2. Log on to the `imlcserv` VM.

```
ssh -i /root/.ssh/vm_private_key cloud-user@aut-1-imlcserv
```

3. Log on to Autonomic Incident Management `postgresql` database using the following command in the `imlcserv` VM and enter the password supplied by local Ericsson support.

```
sudo /opt/rh/rh-postgresql94/root/usr/bin/psql -Upostgres -h postgresql01 -d →
 incidentpersistence
```

Result: Logged on to `incidentpersistence` database successfully.

4. Execute the following command to list cells trained.

```
\copy (SELECT count(1) as trainedKpis, k.cell_id ,c.cell_name as cell FROM k →
pi_boundary k INNER JOIN cell_topology c ON c.id = k.cell_id AND c.active='t →
' AND c.technology='LTE' WHERE k.process_count > 1728 and k.kpi_id in (selec →
t id from kpi where status = 1 and type!='DEGRADATION' and technology = 'LTE →
') GROUP BY k.cell_id,c.cell_name having count(1) = (select count(1) from kp →
i where status = 1 and type!='DEGRADATION' and technology = 'LTE') UNION SEL →
ECT count(1) as trainedKpis, k.cell_id ,c.cell_name as cell FROM kpi_boundar →
y k INNER JOIN cell_topology c ON c.id = k.cell_id  AND c.active='t' AND c.t →
echnology='WCDMA' AND c.cell_name != c.controlling_node WHERE k.process_coun →
t > 1728 and k.kpi_id in (SELECT id FROM kpi WHERE status = 1 AND type!='DEG →
RADATION' AND technology = 'WCDMA' AND name NOT IN ('CS_signaling_setup_fail →
ures','PS_signaling_setup_failures')) GROUP BY k.cell_id,c.cell_name having  →
count(1) = (SELECT count(1) FROM kpi WHERE status = 1 AND type!='DEGRADATION →
' AND technology = 'WCDMA' AND name NOT IN ('CS_signaling_setup_failures','P →
S_signaling_setup_failures')) UNION SELECT count(1) as trainedKpis, k.cell_i →
d ,c.cell_name as cell FROM kpi_boundary k INNER JOIN cell_topology c ON c.i →
d = k.cell_id  AND c.active='t' AND c.technology='WCDMA' AND c.cell_name = c →
.controlling_node WHERE k.process_count > 1728 and k.kpi_id in (SELECT id FR →
OM kpi WHERE status = 1 AND type!='DEGRADATION' AND technology = 'WCDMA' AND →
 name IN  ('CS_signaling_setup_failures','PS_signaling_setup_failures') ) GR →
OUP BY k.cell_id,c.cell_name  having count(1) = (SELECT count(1) FROM kpi WH →
ERE status = 1 AND type!='DEGRADATION' AND technology = 'WCDMA' AND name  IN →
  ('CS_signaling_setup_failures','PS_signaling_setup_failures'))) TO '/erics →
son/enm/dumps/cells_trained.csv' DELIMITER '|' CSV HEADER;
```

Result: Command output can be found in location `/ericsson/enm/dumps/cells_trained.csv`.

5. Execute the following command to list cells not trained.

```
\copy (SELECT count(1) as trainedKpis, k.cell_id ,c.cell_name as cell FROM k →
pi_boundary k INNER JOIN cell_topology c ON c.id = k.cell_id AND c.active='t →
' AND c.technology='LTE' WHERE k.process_count > 1728 and k.kpi_id in (selec →
t id from kpi where status = 1 and type!='DEGRADATION' and technology = 'LTE →
') GROUP BY k.cell_id,c.cell_name having count(1) < (select count(1) from kp →
i where status = 1 and type!='DEGRADATION' and technology = 'LTE') UNION SEL →
ECT count(1) as trainedKpis, k.cell_id ,c.cell_name as cell FROM kpi_boundar →
y k INNER JOIN cell_topology c ON c.id = k.cell_id AND c.active='t' AND c.te →
chnology='WCDMA' AND c.cell_name != c.controlling_node WHERE k.process_count →
```

```
 > 1728 and k.kpi_id in (SELECT id FROM kpi WHERE status = 1 AND type!='DEGR →
ADATION' AND technology = 'WCDMA' AND name NOT IN ('CS_signaling_setup_failu →
res','PS_signaling_setup_failures') ) GROUP BY k.cell_id,c.cell_name having  →
count(1) < (SELECT count(1) FROM kpi WHERE status = 1 AND type!='DEGRADATION →
' AND technology = 'WCDMA' AND name NOT IN ('CS_signaling_setup_failures','P →
S_signaling_setup_failures') ) UNION SELECT count(1) as trainedKpis, k.cell_ →
id ,c.cell_name as cell FROM kpi_boundary k INNER JOIN cell_topology c ON c. →
id = k.cell_id AND c.active='t' AND c.technology='WCDMA' AND c.cell_name = c →
.controlling_node WHERE k.process_count > 1728 and k.kpi_id in (SELECT id FR →
OM kpi WHERE status = 1 AND type!='DEGRADATION' AND technology = 'WCDMA' AND →
 name IN ('CS_signaling_setup_failures','PS_signaling_setup_failures') ) GRO →
UP BY k.cell_id,c.cell_name having count(1) < (SELECT count(1) FROM kpi WHER →
E status = 1 AND type!='DEGRADATION' AND technology = 'WCDMA' AND name IN (' →
CS_signaling_setup_failures','PS_signaling_setup_failures') ) UNION select 0 →
 as trainedKpis, id as cell_id , cell_name as cell from cell_topology where  →
active='t' AND id not in (select cell_id from kpi_boundary)) TO '/ericsson/e →
nm/dumps/cells_not_trained.csv' DELIMITER '|' CSV HEADER;
```

Result: Command output can be found in location `/ericsson/enm/dumps/cells_not_trained.csv`.

6. Execute the following command to list the KPIs trained.

```
\copy (SELECT id, name FROM kpi WHERE status = 1 AND type!='DEGRADATION' AND →
 id NOT IN ( SELECT distinct k.kpi_id FROM kpi_boundary k INNER JOIN cell_to →
pology c ON c.id = k.cell_id AND c.active='t' AND c.technology='LTE' INNER J →
OIN kpi kp ON kp.id = k.kpi_id WHERE k.process_count < 1728 AND k.kpi_id in  →
(select id from kpi where status = 1 and type!='DEGRADATION' and technology  →
= 'LTE') UNION SELECT id from kpi WHERE status = 1 and type!='DEGRADATION' a →
nd technology = 'LTE' AND id not in (select kpi_id from kpi_boundary) UNION  →
SELECT distinct k.kpi_id FROM kpi_boundary k INNER JOIN cell_topology c ON c →
.id = k.cell_id AND c.active='t' AND c.technology='WCDMA' AND c.cell_name != →
 c.controlling_node INNER JOIN kpi kp ON kp.id = k.kpi_id WHERE k.process_co →
unt < 1728 AND k.kpi_id in (select id from kpi where status = 1 and type!='D →
EGRADATION' and technology = 'WCDMA' AND name NOT IN ('CS_signaling_setup_fa →
ilures','PS_signaling_setup_failures') ) UNION SELECT id from kpi WHERE stat →
us = 1 and type!='DEGRADATION' and technology = 'WCDMA' AND name NOT IN ('CS →
_signaling_setup_failures','PS_signaling_setup_failures') AND id not in (sel →
ect kpi_id from kpi_boundary) UNION SELECT distinct k.kpi_id FROM kpi_bounda →
ry k INNER JOIN cell_topology c ON c.id = k.cell_id AND c.active='t' AND c.t →
echnology='WCDMA' AND c.cell_name = c.controlling_node INNER JOIN kpi kp ON  →
kp.id = k.kpi_id WHERE k.process_count < 1728 AND k.kpi_id in (select id fro →
m kpi where status = 1 and type!='DEGRADATION' and technology = 'WCDMA' AND  →
name IN ('CS_signaling_setup_failures','PS_signaling_setup_failures') ) UNIO →
N SELECT id from kpi WHERE status = 1 and type!='DEGRADATION' and technology →
 = 'WCDMA' AND name IN ('CS_signaling_setup_failures','PS_signaling_setup_fa →
ilures') AND id not in (select kpi_id from kpi_boundary))) TO '/ericsson/enm →
/dumps/kpis_trained.csv' DELIMITER '|' CSV HEADER;
```

Result: Command output can be found in location `/ericsson/enm/dumps/kpis_trained.csv`.

7. Execute the following command to list the KPIs not trained.

```
\copy (SELECT distinct k.kpi_id, kp.name FROM kpi_boundary k  →
INNER JOIN cell_topology c ON c.id = k.cell_id AND c.active=' →
t' AND c.technology='LTE' INNER JOIN kpi kp ON kp.id = k.kpi_ →
id WHERE k.process_count < 1728 AND k.kpi_id in (select id fr →
om kpi where status = 1 and type!='DEGRADATION' and technolog →
y = 'LTE') UNION SELECT id, name from kpi WHERE status = 1 an →
d type!='DEGRADATION' and technology = 'LTE' AND id not in (s →
elect kpi_id from kpi_boundary) UNION SELECT distinct k.kpi_i →
d, kp.name FROM kpi_boundary k INNER JOIN cell_topology c ON  →
c.id = k.cell_id AND c.active='t' AND c.technology='WCDMA' AN →
D c.cell_name != c.controlling_node INNER JOIN kpi kp ON kp.i →
```

```
d = k.kpi_id WHERE k.process_count < 1728 AND k.kpi_id in (se
lect id from kpi where status = 1 and type!='DEGRADATION' and
 technology = 'WCDMA' AND name NOT IN ('CS_signaling_setup_fa
ilures','PS_signaling_setup_failures') ) UNION SELECT id, nam
e from kpi WHERE status = 1 and type!='DEGRADATION' and techn
ology = 'WCDMA' AND name NOT IN ('CS_signaling_setup_failures
','PS_signaling_setup_failures') AND id not in (select kpi_id
 from kpi_boundary) UNION SELECT distinct k.kpi_id, kp.name F
ROM kpi_boundary k INNER JOIN cell_topology c ON c.id = k.cel
l_id AND c.active='t' AND c.technology='WCDMA' AND c.cell_nam
e = c.controlling_node INNER JOIN kpi kp ON kp.id = k.kpi_id
WHERE k.process_count < 1728 AND k.kpi_id in (select id from
kpi where status = 1 and type!='DEGRADATION' and technology =
 'WCDMA' AND name IN ('CS_signaling_setup_failures','PS_signa
ling_setup_failures') ) UNION SELECT id, name from kpi WHERE
status = 1 and type!='DEGRADATION' and technology = 'WCDMA' A
ND name IN ('CS_signaling_setup_failures','PS_signaling_setup
_failures') AND id not in (select kpi_id from kpi_boundary))
TO '/ericsson/enm/dumps/kpis_not_trained.csv' DELIMITER '|' C
SV HEADER;
```

Result: Command output can be found in location `/ericsson/enm/dumps/kpis_not_trained.csv`.

8. Execute the following command to list the KPI process count on all cell.

```
\copy (select c.cell_name, kpi.name, kpi_boundary.process_count from kpi,kpi
_boundary,cell_topology as c where kpi_boundary.cell_id = c.id and kpi.id =
kpi_boundary.kpi_id and kpi.name not like '%deg%') TO '/ericsson/enm/dumps/k
pi_process_count.csv' DELIMITER '|' CSV HEADER;
```

Result: Command output can be found in location `/ericsson/enm/dumps/kpis_process_count.csv`.

9. Execute the following to log out of Autonomic Incident Management `postgresql` database.

```
\quit
```

Result: Successfully logged out of `postgresql`.

10. Collect generated CSV files from `/ericsson/enm/dumps/` and analyze.

11. Delete CSV files when analysis is completed.

```
rm /ericsson/enm/dumps/cells_trained.csv
```

```
rm /ericsson/enm/dumps/cells_not_trained.csv
```

```
rm /ericsson/enm/dumps/kpis_trained.csv
```

```
rm /ericsson/enm/dumps/kpis_not_trained.csv
```

Result: CSV files deleted.

9.4 Clean Start of imlcserv VM triggered

> imlcserv VM clean starts unexpectedly..

Prerequisites

— You have access to the Log Viewer application.

Steps

1. Log on to the ENM Launcher with your user ID and password.
2. Open Log Viewer.
3. In Search criteria, ensure the following:
   — **Message** contains `"CREDENTIAL MANAGER CLI RESTART VM SCRIPT"`.
   — Under **Severity**, the **Info** checkbox is selected.
   — **Hosts** contains the `imlcserv` VM where the clean start occurred.
4. Check if log `"CREDENTIAL MANAGER CLI RESTART VM SCRIPT : file .state has been found, VM will be restarted"` is present in the search result.

   If it is present, this indicates that the clean start was triggered to install a new certificate, which is the expected behavior.

Results

JBOSS was restarted by `credentialmanagercli` for the VM to acquire a new certificate. This is expected behavior and is not a fault with the VM.

10 Network Health Monitor Troubleshooting

This section provides the tasks recommended to diagnose and resolve common problems in Network Health Monitor, Network Health Analysis, Node Monitor, and KPI Management.

10.1 N/A in the Nodes Breached per KPI Widget in the Network Health Monitor Page

A KPI displays N/A in the Nodes Breached per KPI widget if the KPI does not contain any data or is inactive.

This topic describes the tasks to perform to determine why a KPI does not contain or cannot access data. All steps are performed in NHM apps.

Prerequisites

— Nodes are synced

— PMIC is collecting PM stats files

— Access to the `KPIServ` Virtual Machine (VM). Refer to Connect to a Virtual Machine on a Physical ENM Deployment on page 2 or Connect to a Virtual Machine on an ENM on Cloud Deployment on page 3.

— Access to NHM Help

Steps

1. Verify in the Network Health Monitor app that the KPI is active.

   On the right side of Network Health Monitor, click **Settings** ⚙

   If the KPI is listed among the Active KPIs, proceed to Step 2.

   If the KPI is listed among the Inactive KPIs, activate it. Wait until the next ROP so that the relevant PM Stats file is collected. If the selected KPI in the Widget still displays N/A, proceed to Step 2.

   To activate a KPI, select**Inactive KPIs list** > **Manage KPIs** and choose the KPI. This highlights the KPI in the Key Performance Indicators table. Right-click the **KPI** and select **Activate**

   .

2. Verify that the correct selection of nodes are added to the KPI.

   The KPI displays N/A if no nodes or the wrong selection of nodes are added.

   In KPI Management, select the relevant KPI. Right-click the **KPI** and select **Edit**. In the **Edit Key Performance Indicator** page, select **Nodes** > **Add Topology Data**. Then select the relevant option to add the correct nodes to the KPI. In the top left click **Save** to add Nodes successfully.

   Wait until the next ROP so that the relevant PM Stats file is collected, then check if the Widget still displays N/A. If it does proceed to Step 3.

3. Check that the PM stats files have been collected for the latest ROP for the identified node.

   Use the FLS to check this, see "*Interwork Description for File Lookup Service (FLS), 1/155 19-CNA4033301 Uen*. If the files are collected go to step 4. If the files have not been collected, refer to *Identify the Cause of Missed File Collection for Node(s)* in this document.

4. Verify that the PM stats file contains the relevant counters that are used to calculate the KPI.

   Connect to each of the `KPIServ` VMs.

   ```
   grep <COUNTERNAME> /ericsson/pmic1/XML/<node name>/<latest ROP filname>.xml
   grep <COUNTERNAME> /ericsson/pmic2/XML/<node name>/<latest ROP filname>.xml
   ```

   Example

   To verify that a specific PM stats file contains the counter `pmA3InterFBestCellEvalReport`:

   ```
   grep pmA3InterFBestCellEvalReport /ericsson/pmic1/XML/SubNetwork=ERBS-SUBNW-
   1,MeContext=ieatnetsimv5048-01_LTE01ERBS00043/A20160222.1100+0000-1115+0000_
   SubNetwork=ERBS-SUBNW-1,MeContext=ieatnetsimv5048-01_LTE01ERBS00043_statsfil
   e.xml
   ```

   If the PM stats file does not contain the relevant counters, then add the counters to a PM subscription using PMIC.

   If the PM stats file contains the counters but N/A displayed in the widget for the KPI, then consult Ericsson local support.

Results

User is able to identify why KPI does not have any Data and is displaying N/A in the Nodes Breached per KPI Widget.

10.2 No Data Loads in Network Health Manager (NHM)

> Widgets and applications do not display results if the file parsing and KPI calculation EPS flows are not running.

NHM uses the file parsing and kpi calculation EPS flows to ensure data loading. This topic describes the steps to determine why EPS flows have failed to deploy or have stopped. These steps are performed by connecting to the `KPIServ` Virtual Machine (VM).

Prerequisites

— PMIC is collecting PM stats files.

— Access to the `kpiserv` VM. Refer to Connect to a Virtual Machine on a Physical ENM Deployment on page 2 or Connect to a Virtual Machine on an ENM on Cloud Deployment on page 3.

Steps

1. Connect to each of the `KPIServ` VMs and run the following commands to verify that both EPS flows were successfully deployed at startup:

```
grep "Successfully deployed module" /ericsson/3pp/jboss/standalone/log/serve →
r.log*

server.log.2:2016-06-23 21:11:53,218 INFO  [com.ericsson.oss.services.eps.de →
ployer.file.FileSystemListenerImpl] (ServerService Thread Pool -- 100) Succe →
ssfully deployed module from [file                              :/var/ericsson →
/eps/kpiFlow/kpi-calculation-flow.xml] - module identifier is com.ericsson.o →
ss.services.kpi_kpi-calculation-flow_1.0.0

server.log.2:2016-06-23 21:11:53,647 INFO  [com.ericsson.oss.services.eps.de →
ployer.file.FileSystemListenerImpl] (ServerService Thread Pool -- 100) Succe →
ssfully deployed module from [file                              :/var/ericsson →
/eps/kpiFlow/kpi-pm-parsing-flow.xml] - module identifier is com.ericsson.os →
s.services.kpi_kpi-pm-parsing-flow_1.0.0
```

   The example output shows that both the `kpi-calculation-flow` and `kpi-parsing-pm-flow` are successfully deployed.

2. Verify that both EPS flows were deployed at startup:

```
grep "Successfully auto-deployed" /ericsson/3pp/jboss/standalone/log/server. →
log*
/ericsson/3pp/jboss/standalone/log/server.log.2:2016-06-23 21:11:53,647 INFO →

[com.ericsson.oss.services.eps.deployer.file.FileSystemListenerImpl] (Server →
Service Thread Pool -- 100) Successfully auto-deployed 2 event-flows on star →
tup
```

   The example output shows that both EPS flows are deployed

   `Successfully auto-deployed 2 event-flows on startup`

3. Check if either of the EPS flows have gone down since startup:

```
grep -i "Destroyed component" /ericsson/3pp/jboss/standalone/log/server.log*
grep -i "Removed all components" /ericsson/3pp/jboss/standalone/log/server.l →
og*
```

These commands show if any flows were undeployed:

```
2016-06-16 15:07:47,464 INFO  [com.ericsson.oss.services.eps.pe.java.IOAdapt →
erComponentInstaller] (ServerService Thread Pool -- 92) Destroyed component  →
com.ericsson.oss.services.nhm.kpicalculationservice.eps.adapter.KpiFlowOutpu →
tAdapter@522b22e6
2016-06-16 15:07:47,464 INFO  [com.ericsson.oss.services.eps.pe.java.IOAdapt →
erComponentInstaller] (ServerService Thread Pool -- 92) Removed all componen →
ts for module [com.ericsson.oss.services.kpi_kpi-calculation-flow_1.0.0]
```

Note: EPS flows are undeployed when the `kpiserv` VM is stopped or restarted. This is correct behavior and not an error. To redeploy the EPS flows, restart the `kpiserv` VM.

Results

User is able to check if the NHM EPS flows are operational.

10.3 Clean Start of mskpirt VM Triggered

> mskpirt VM clean starts unexpectedly.

Prerequisites

You have access to the Log Viewer application.

Steps

1. Log on to ENM Launcher with your user ID and password.
2. Open Log Viewer.
3. In search criteria, ensure the following:
   - Under **Severity**, the **Info** checkbox is selected.
   - **Hosts** contains the `mskpirt` VM where the clean start occurred.
   - **Message** contains "`CREDENTIAL MANAGER CLI RESTART VM SCRIPT`".
4. Check if log "`CREDENTIAL MANAGER CLI RESTART VM SCRIPT : file .state has been found, VM will be restarted`" is present in the search result.

If it is present, this indicates that the clean start was triggered to install a new certificate, which is expected behavior.

Results

JBOSS was restarted by `credentialmanagercli` for the VM to acquire a new certificate. This is expected behavior and is not a fault with the VM.

11 Analytic Session Record Troubleshooting

This section provides the tasks recommended to diagnose and resolve common ASR issues.

Use it with the *Stream Termination and Parsing Troubleshooting* section in the *ENM Performance Management Troubleshooting Guide*.

11.1 Verify that ASR is Licensed

11.1.1 License for ASR-L

Check that the VP_Analytics_Session_record_NBI_onOffScope_LTE_5MHzSC license is present and valid.

Prerequisite

— Access to the ENM UI.

1. Log on to the ENM UI and open the CLI.
2. List the licenses:

```
lcmadm list
```

3. Ensure that an entry for the `VP_Analytics_Session_record_NBI_onOffScope_LTE_5MHzSC` license is present in the results and that it is within the stated expiration date.

11.1.2 License for ASR-N

Check that the VP_Analytics_Session_record_NBI_onOffScope_5G_5MHzSC license is present and valid.

Prerequisite

— Access to the ENM UI.

1. Log on to the ENM UI and open the CLI.
2. List the licenses:

```
lcmadm list
```

3. Ensure that an entry for the `VP_Analytics_Session_record_NBI_onOffScope_5G_5MHzSC` license is present in the results and that it is within the stated expiration date.

11.2 Verify That ASR Forwarders Are Online

> Check that each ASRL forwarder instance service is online and running.

Prerequisite

Access to the forwarder KVMs using the `cloud-user` account from the MS.

1. Get the status of each forwarder using the following command from the MS:

```
for asr_host in `egrep asr-[[:digit:]]-asrlfwd[[:digit:]] /etc/hosts | awk { →
'print $2'}`; do ssh -i /root/.ssh/vm_private_key cloud-user@$asr_host 'echo →
 $HOSTNAME;/etc/init.d/fwdr status'; done
```

2. Examine the output and ensure for each host listed that the forwarder instance is online and running.

11.3 Verify That ASR Forwarder Consumer Is Connected to the Kafka Partition

> Check that the assigned partitions have been set and are not currently revoked.

Prerequisite

Access to the forwarder KVMs using the `cloud-user` account.

Repeat these steps for all forwarder KVMs:

1. List all forwarder KVMs by running the following command from MS:

```
egrep asr-[[:digit:]]-asrlfwd[[:digit:]] /etc/hosts | awk {'print $2'} | sor →
t
```

2. SSH to one of the forwarder KVMs as cloud-user using the following command from MS:

```
ssh -i /root/.ssh/vm_private_key cloud-user@asr-<n>-asrlfwd<n>
```

3. Check that the partitions are still assigned (run the command from within the `fwdr kvm`):

```
grep 'assigned partitions' /ericsson/log/fwdr/fwdr.log
```

4. Compare the time stamps from step 3 and verify that the latest assigned partitions time stamp is later than the latest revoked partitions time stamp.

11.4 Verify That Each ASR Forwarder is Connected to the Northbound Client

> Check that the ASR forwarder is properly connected to the designated Northbound Client.

Prerequisite

Access to the forwarder KVMs using the `cloud-user` account.

Repeat these steps for all forwarder KVMs:

1. List all forwarder KVMs by running the following command from MS:

```
egrep asr-[[:digit:]]-asrlfwd[[:digit:]] /etc/hosts | awk {'print $2'} | sor →
t
```

2. SSH to a given ASR forwarder KVM as cloud-user using the following command:

```
ssh -i /root/.ssh/vm_private_key cloud-user@asr-<n>-asrlfwd<n>
```

3. Run the following command to check if Forwarder is currently connected to a northbound client, where `<n>` is the number of the current forwarder:

```
grep 'client running on' /ericsson/log/fwdr/fwdr.log | tail -1
```

4. Look for a `Connected to client running on…` message in the console output.

11.5 Validate Avro Schema Is Present In The Schema Registry

Check that the schema registry contains the ASR Avro Schema.

Prerequisites

- Fully deployed ENM system with a licensed ASR-L or ASR-N Value Pack installed.
- Valid user with the appropriate ASR NBI capabilities to access and perform requests on the ASR REST API.
- ASR-L_Schema_Operator and ENM Administrator can perform HTTP GET requests on the ASR REST API.

1. Open a web browser and navigate to the following address, where `<UI_Address>` is ENM Application GUI URL:

   For ASR-L Avro Schema:

   ```
   https://<customer-domain>/session-record/v1/record-encoding?type=ASR_L
   ```

   For ASR-N Avro Schema:

   ```
   https://<customer-domain>/session-record/v1/record-encoding?type=ASR_N
   ```

2. Log on to the ENM UI when requested.

   User type must be either Administrator or ASR User.

3. Verify that the response returns an array populated with the required schemas.

   If the schema registry is empty, a blank array will be returned, denoted by `[]`.

11.6 Verify Current Ingestion Rate into ASR-L

View metrics for current event ingestion and record production.

Prerequisites

- The ENM Deployment variant includes the ASR cluster.
- Access to the Management Server (MS-1).

1. Log on to the Management Server (MS-1).
2. Run the ASR-L metrics script:

   ```
   python /ericsson/tor/data/asrl/bin/asrl_metrics.py
   ```

   Result: The current filtered event ingestion rate into ASR-L and complete & suspect record rate out of ASR-L displays.

   Example output:

```
|    Matrix Name     |       Total        |    Per Second (in last minute)    |
-------------------------------------------------------------------------------
|  Filtered Events   |        139         |                 6                 |
|  Complete Records  |         13         |                 2                 |
|  Suspect Records   |         10         |                 0                 |
-------------------------------------------------------------------------------
```

3. View the CSV file containing the history of the metrics script runs:

   ```
   cat /ericsson/tor/data/asrl/asrl_metrics/asrl_metrics.csv
   ```

   Example output:

```
Date & Time            Filtered      Complete      Suspect      Filtered Events      Complete Records      Suspect Records
                       Events        Records       Records      Per Second           Per Second            Per Second
15-01-18 10:12:58        79            4             0            6                    1                     0
15-01-18 11:34:20        102           9             3            5                    2                     0
15-01-18 11:50:12        139           13            10           6                    2                     0
```

4. View the log for the metrics script:

   ```
   cat /ericsson/tor/data/asrl/bin/asrl_metrics.log
   ```

   Note: There are scenarios when metrics for ASR-L are unavailable. The script displays a notification message describing the issue when they occur:

   — ASR configuration was activated following a system upgrade or install. Metrics are unavailable until after midnight.

   — ASR-L Correlation job was restarted for any reason. Metrics are unavailable for up to 30 minutes as they are recalibrated.

11.7 Verify the ASR-L Topology Service

> Check the state of the ASR-L topology service KVMs and sizes of their eNodeB and Cell caches to determine why topology data is not getting enriched.

Prerequisites

— The `sessionStatus` field shows no topology enrichment for over 15 minutes after ASR-L startup. Refer to *Analytic Session Record Specifications* for information on `sessionStatus`.

— Access to the ASR-L topology service KVMs using the `cloud-user` account from the MS.

1. Check that the ASR-L Topology Service KVM is in `ONLINE` state:

   ```
   /opt/ericsson/enminst/bin/vcs.bsh --groups | grep asrltopologyservice
   ```

   If it is in `ONLINE` state proceed with step 2, else contact local Ericsson support.

2. List all ASR-L Topology Service KVMs from the MS:

   ```
   grep asrltopologyservice /etc/hosts | awk {'print $2'}
   ```

Repeat steps 3—4 for each ASR-L Topology Service KVM and take note of the outputs.

3. Connect through SSH to the ASR-L Topology Service KVM as cloud-user from the MS:

   ```
   ssh -i /root/.ssh/vm_private_key cloud-user@asr-<n>-asrltopologyservice
   ```

4. Check and take note of the eNodeB cache size and cell cache size:

   ```
   grep "METRICS:TOPOLOGY_EXPORT :Total number of Nodes found" /ericsson/3pp/jb →
   oss/standalone/log/server.log | tail -1
   ```

5. Examine all the outputs for one with the latest time stamp. If it has an eNodeB Cache size and Cell Cache size as `0` go to step 6, else contact local Ericsson support.

6. Restart the ASR-L Topology Service Group from the MS:

   ```
   /opt/ericsson/enminst/bin/vcs.bsh --restart -g Grp_CS_asr_cluster_asrltopolo →
   gyservice
   ```

   If the issue persists, contact local Ericsson support.

11.8 Overload Protection Is Enabled for an Extended Duration

> Verify that the Kafka Decoded Partitions are load balanced evenly when Overload Protection is enabled for over four hours.

Prerequisites

- The `asrMaxStageOneDurationInMilliseconds` and `asrAcceptableStageOneDurationInMilliseconds` parameters are set appropriately. Refer to *Configure Overload Protection Threshold Parameters* in the *ENM System Administrator Guide* for more information.
- Overload protection is enabled.
- Check the time when overload protection was last enabled by using the Log Viewer to search messages:
  - Contain: `overload protection`
  - Severity: `warning`
  - Tag: `sparkWorker`

  If overload protection is enabled and not disabled in the last four hours, then check that the Kafka decoded partitions are load balanced evenly.

  Refer to *Check That Kafka Decoded Partitions Are Load Balanced Evenly* in the *ENM System Administrator Guide* for more information.

Reference List

[1] *ENM System Administrator Guide*, 1/1543-AOM 901 151

[2] *ENM Network Integration Guideline*, 1/102 72-AOM 901 151.

[3] *ENM Operator's Guide*, 1/1553-AOM 901 151

[4] *ENM Security System Administrator Guide*, 2/1543-AOM 901 151

[5] *VNF-LCM Installation Instructions*, 1/1531-CNA 403 3313, available in the ENM Installation and Support CPI Library.

[6] *Typographic Conventions*, 3/1551-FCK 101 05

[7] *ENM Backup and Restore System Administrator Guide*, 3/1543-AOM 901 151

[8] *ENM Configuration System Administrator Guide*, 1/1543-AOM 901 151-1

[9] *ENM Data Collection Guideline*, available from local Ericsson Support.

[10] *ENM Interwork Description for File Lookup Service (FLS)*, 1/155 19-CNA4033301

[11] *ENM Performance Management System Administrator Guide*, 1/1543-AOM 901 151-3

[12] *ENM Installation Instructions*, Available from local Ericsson Support

[13] *Manage Security User Guide*, *18/1553-LZA 701 6014*, available from relevant WCDMA RAN CPI Library.

[14] *ENM Configuration Troubleshooting Guide*, 1/159 01-AOM 901 151-1

[15] *ENM Monitoring Troubleshooting Guide*, 1/159 01-AOM 901 151-2

[16] *ENM Performance Management Troubleshooting Guide*, 1/159 01-AOM 901 151-3

[17] *ENM Security Management Troubleshooting Guide*, 1/159 01-AOM 901 151-4

[18] *ENM on Cloud Backup and Restore Administrator Guide*, 5/1543-AOM 901 151

[19] *ENM Troubleshooting Guide*, 1/159 01-AOM 901 151

[20] *ENM Identity and Access Management System Administrator Guide*, 2/1543-AOM 901 151-1

[21] *ENM Network Security Configuration System Administrator Guide*, 2/1543-AOM 901 151-2

[22] *ENM Public Key Infrastructure System Administrator Guide*, 2/1543-AOM 901 151-3

[23] *OSS-RC Configuration for ENIQ*, 1/1546-AOM 901 076, available from Ericsson Network IQ Events CPI Library.

[24] *ENM Identity and Access Management Programmers Guide*, 19817-cna 403 3016

[25] *ENM Product Description*, 1/1551-AOM 901 151

[26] *ENM Monitoring System Administrator Guide*, 1/1543-AOM 901 151-2

[27] *LITP Node Hardening Instructions*, 2/1531-CSA 113 110

[28] *CAS Software Dropbox End-User Guide*, 6/1553-HSC 901 110, available from Ericsson Support.
[29] *ENM Neo4j Troubleshooting Guide*159 01-CNA 403 3405
[30] *ENM System Monitor User Guide*, 1/1553-cna 403 3115
[31] ENM NIG Connection Matrix
[32] *ENM System Security Configuration Programmers Guide*
1/19817-cna 403 3065 Uen
[33] *Neo4j Documentation*
https://support.neo4j.com/hc/en-us/articles/360006361794-Causal-Cluster-FAQ-for-heavy-workloads
[34] *ENM Parameter List*, 1/19059-AOM 901 151
[35] *Managed Object Model (MOM) RNC*, 155 54-CXC 173 5901/1-V1
[36] *SAPC, Measurements Ericsson Service-Aware Policy Controller User Guide*, 1/15553 - AXB 901 33/7
[37] *SAPC, Managed Object Model*, 155 54-LZN 708 0672/5-V1
[38] *ENM Privacy User Guide*, 2/1553-AOM 901 151