

Troubleshoot ACI Fault Code F199144, F93337, F381328, F93241, F450296 : TCA

Contents

[Introduction](#)

[Background](#)

[Fault : F199144](#)

[Quick Start to Address Fault: F199144](#)

[1. Command "show platform internal hal l3 routingthresholds"](#)

[2. Command "show platform internal hal health-stats"](#)

[Next Steps Fault : F199144](#)

[Fault : F93337](#)

[Quick Start to Address Fault : F93337](#)

[1. Command "moquery -d 'comp/prov-VMware/ctrlr-\[\]- /vm-vm- '"](#)

[2. Command "moquery -c compRsHv | grep 'vm-1071'"](#)

[3. Command "moquery -c compHv -f 'comp.Hv.oid=="host-1068'"](#)

[Next Steps Fault : F93337](#)

[Fault : F93241](#)

[Quick Start to Address Fault : F93241](#)

[1. Command "moquery -d 'comp/prov-VMware/ctrlr-\[\]- /vm-vm- '"](#)

[2. Command "moquery -c compRsHv | grep 'vm-1071'"](#)

[3. Command "moquery -c compHv -f 'comp.Hv.oid=="host-1068'"](#)

[Next Steps Fault : F93241](#)

[Fault : F381328](#)

[Quick Start to Address Fault : F381328](#)

[1. Dump the highest number interfaces with CRC in fabric](#)

[2. Dump the highest number of FCS in the fabric](#)

[Next Steps Fault : F381328](#)

[Python Script for fault : F381328](#)

[Fault : F450296](#)

[Quick Start to Address Fault : F450296](#)

[1. Command "show platform internal hal health-stats asic-unit all"](#)

[Next Steps Fault : F450296](#)

Introduction

This document describes remediations steps for ACI Fault Codes: F199144, F93337, F381328, F93241, F450296

Background

If you have an Intersight connected ACI fabric, a Service Request was generated on your behalf to indicate that instance of this fault was found within your Intersight-Connected ACI fabric.

This is being actively monitored as part of [Proactive ACI Engagements](#).

This document describes next steps for remediation of the following fault:

Fault : F199144

```
"Code" : "F199144",
"Description" : "TCA: External Subnet (v4 and v6) prefix entries usage current value(eqptcapacityPrefixE
"Dn" : "topology/pod-1/node-132/sys/eqptcapacity/fault-F199144"
```

This specific fault is raised when current usage of the external subnet prefix exceeds 99%. This suggests a hardware limitation in terms of routes handled by these switches.

Quick Start to Address Fault: F199144

1. Command "show platform internal hal l3 routingthresholds"

```
module-1# show platform internal hal l3 routingthresholds
Executing Custom Handler function

OBJECT 0:
trie debug threshold           : 0
tcam debug threshold          : 3072
Supported UC lpm entries      : 14848
Supported UC lpm Tcam entries : 5632
Current v4 UC lpm Routes     : 19526
Current v6 UC lpm Routes     : 0
Current v4 UC lpm Tcam Routes : 404
Current v6 UC lpm Tcam Routes : 115
Current v6 wide UC lpm Tcam Routes : 24
Maximum HW Resources for LPM  : 20480 < ----- Maximum hardware resource
Current LPM Usage in Hardware : 20390 < -----Current usage in Hw
Number of times limit crossed  : 5198 < ----- Number of times
Last time limit crossed       : 2020-07-07 12:34:15.947 < ----- Last occurrence
```

2. Command "show platform internal hal health-stats"

```
module-1# show platform internal hal health-stats
No sandboxes exist
|Sandbox_ID: 0 Asic Bitmap: 0x0
|-----

L2 stats:
=====
bds:                : 249
...
l2_total_host_entries_norm : 4

L3 stats:
=====
l3_v4_local_ep_entries   : 40
max_l3_v4_local_ep_entries : 12288
```

```

l3_v4_local_ep_entries_norm : 0
l3_v6_local_ep_entries      : 0
max_l3_v6_local_ep_entries  : 8192
l3_v6_local_ep_entries_norm : 0
l3_v4_total_ep_entries      : 221
max_l3_v4_total_ep_entries  : 24576
l3_v4_total_ep_entries_norm : 0
l3_v6_total_ep_entries      : 0
max_l3_v6_total_ep_entries  : 12288
l3_v6_total_ep_entries_norm : 0
max_l3_v4_32_entries        : 49152
total_l3_v4_32_entries      : 6294
  l3_v4_total_ep_entries    : 221
  l3_v4_host_uc_entries     : 6073
  l3_v4_host_mc_entries     : 0
total_l3_v4_32_entries_norm : 12
max_l3_v6_128_entries       : 12288
total_l3_v6_128_entries     : 17
  l3_v6_total_ep_entries    : 0
  l3_v6_host_uc_entries     : 17
  l3_v6_host_mc_entries     : 0
total_l3_v6_128_entries_norm : 0
max_l3_lpm_entries          : 20480 < ----- Maximum
l3_lpm_entries               : 19528 < ----- Current L3 LPM entries
  l3_v4_lpm_entries          : 19528
  l3_v6_lpm_entries          : 0
l3_lpm_entries_norm         : 99
max_l3_lpm_tcam_entries     : 5632
max_l3_v6_wide_lpm_tcam_entries : 1000
l3_lpm_tcam_entries         : 864
  l3_v4_lpm_tcam_entries     : 404
  l3_v6_lpm_tcam_entries     : 460
  l3_v6_wide_lpm_tcam_entries : 24
l3_lpm_tcam_entries_norm    : 15
l3_v6_lpm_tcam_entries_norm : 2
l3_host_uc_entries          : 6090
  l3_v4_host_uc_entries     : 6073
  l3_v6_host_uc_entries     : 17
max_uc_ecmp_entries         : 32768
uc_ecmp_entries             : 250
uc_ecmp_entries_norm        : 0
max_uc_adj_entries          : 8192
uc_adj_entries              : 261
uc_adj_entries_norm         : 3
vrfs                         : 150
  infra_vrfs                 : 0
  tenant_vrfs                 : 148
rtd_ifs                     : 2
sub_ifs                      : 2
svi_ifs                      : 185

```

Next Steps Fault : F199144

1. Reduce the number of routes each switch has to handle so you comply with the scalability defined for the hardware model. Please check scalability guide here

<https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/4-x/verified-scalability/Cisco-ACI-Verified-Scalability-Guide-412.html>

2. Consider changing the Forwarding Scale Profile based on the scale.

<https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/all/forwarding-scale-profiles/cisco-apic-forwarding-scale-profiles/m-overview-and-guidelines.html>

3. Removing 0.0.0.0/0 subnet in L3Out and only configure required subnets

4. If you are using Gen 1, upgrade your hardware from Gen 1 to Gen 2, as Gen 2 switches allow 20,000+ external v4 routes.

Fault : F93337

```
"Code" : "F93337",
>Description" : "TCA: memory usage current value(compHostStats15min:memUsageLast) value 100% raised above threshold",
"Dn" : "comp/prov-VMware/ctrlr-[FAB4-AVE]-vcenter/vm-vm-1071/fault-F93337"
```

This specific fault is raised when VM host is consuming memory more than the threshold. The APIC monitors these hosts via VCenter. Comp:HostStats15min is a class that represents the most current statistics for host in a 15 minute sampling interval. This class updates every 5 minutes.

Quick Start to Address Fault : F93337


1. Command "moquery -d 'comp/prov-VMware/ctrlr-[<DVS>]-<VCenter>/vm-vm-<VM id from the fault's DN>'"

This command gives information about the affected VM

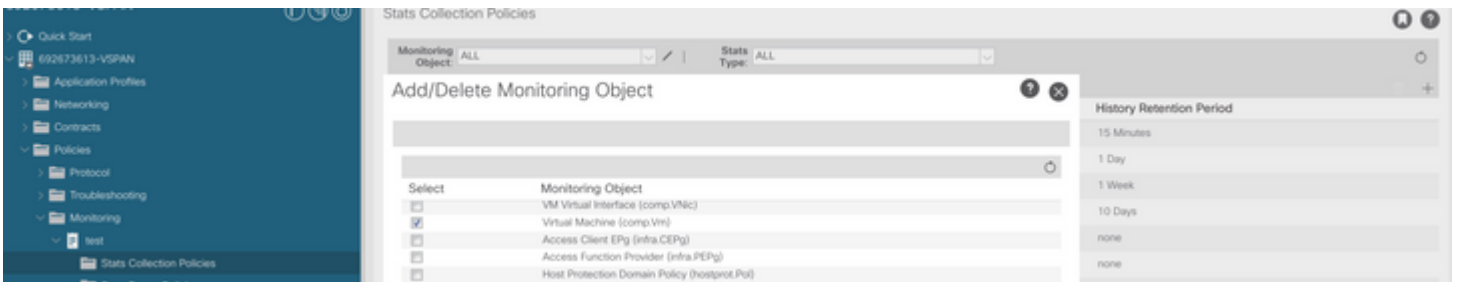
```
# comp.Vm
oid          : vm-1071
cfgd0s      : Ubuntu Linux (64-bit)
childAction  :
descr       :
dn          : comp/prov-VMware/ctrlr-[FAB4-AVE]-vcenter/vm-vm-1071
ftRole      : unset
guid        : 501030b8-028a-be5c-6794-0b7bee827557
id          : 0
issues      :
lcOwn       : local
modTs       : 2022-04-21T17:16:06.572+05:30
monPolDn    : uni/tn-692673613-VSPAN/monepg-test
name        : VM3
nameAlias   :
os          :
rn          : vm-vm-1071
state       : poweredOn
status      :
template    : no
type        : virt
uuid        : 4210b04b-32f3-b4e3-25b4-fe73cd3be0ca
```

2. Command "moquery -c compRsHv | grep 'vm-1071'"

b. Under your Monitoring policy, select stats collection policy.

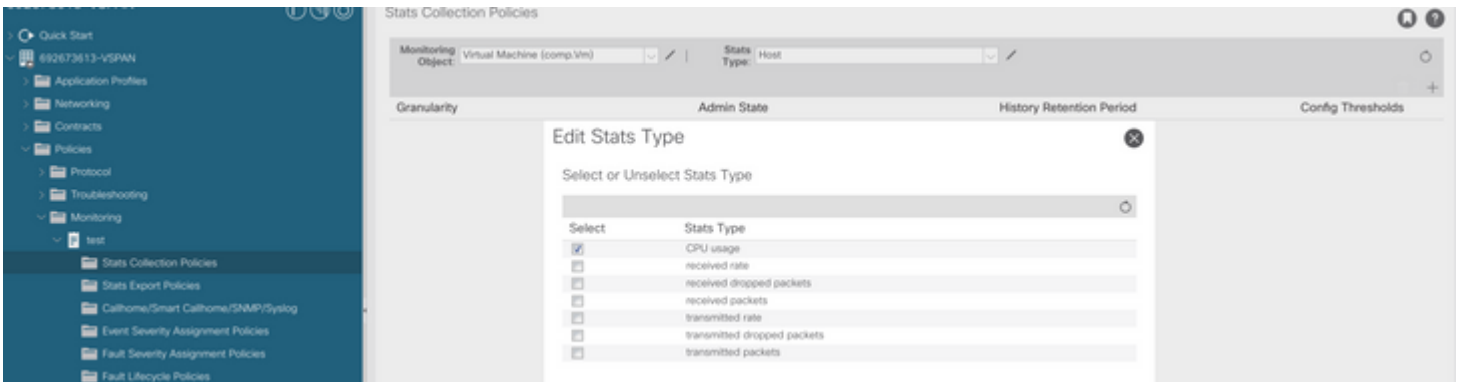


c. Click on edit icon beside Monitoring object dropdown, and check the Virtual Machine (comp.Vm) as a monitoring object. After Submitting, select the compVm object from Monitoring Object dropdown.



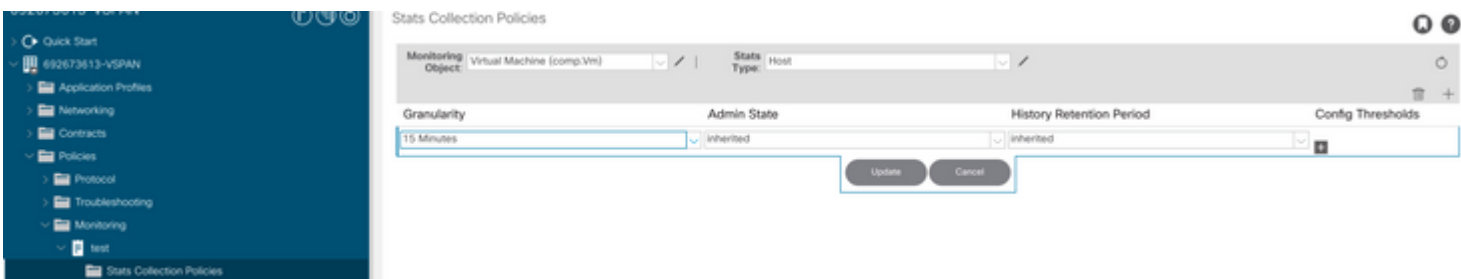
d.

d. Click on edit icon beside Stats type, then check on CPU Usage.



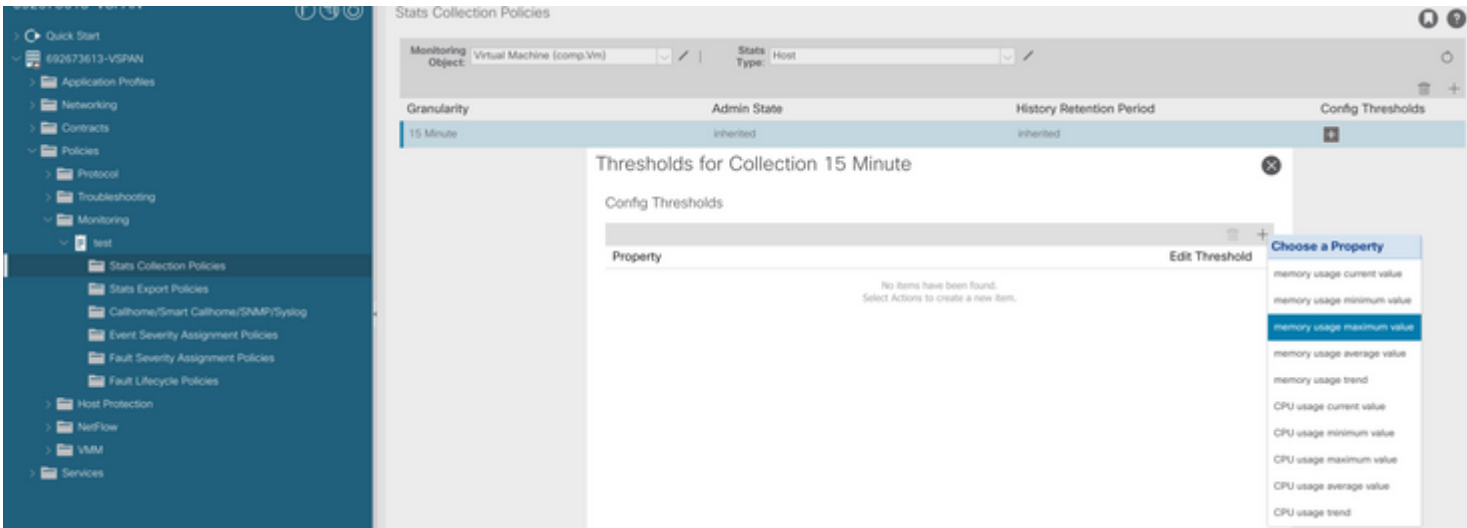
e.

e. From the stats type Dropdown click select host, click on the + sign and enter your Granularity, Admin state and History Retention Period and then click on update.

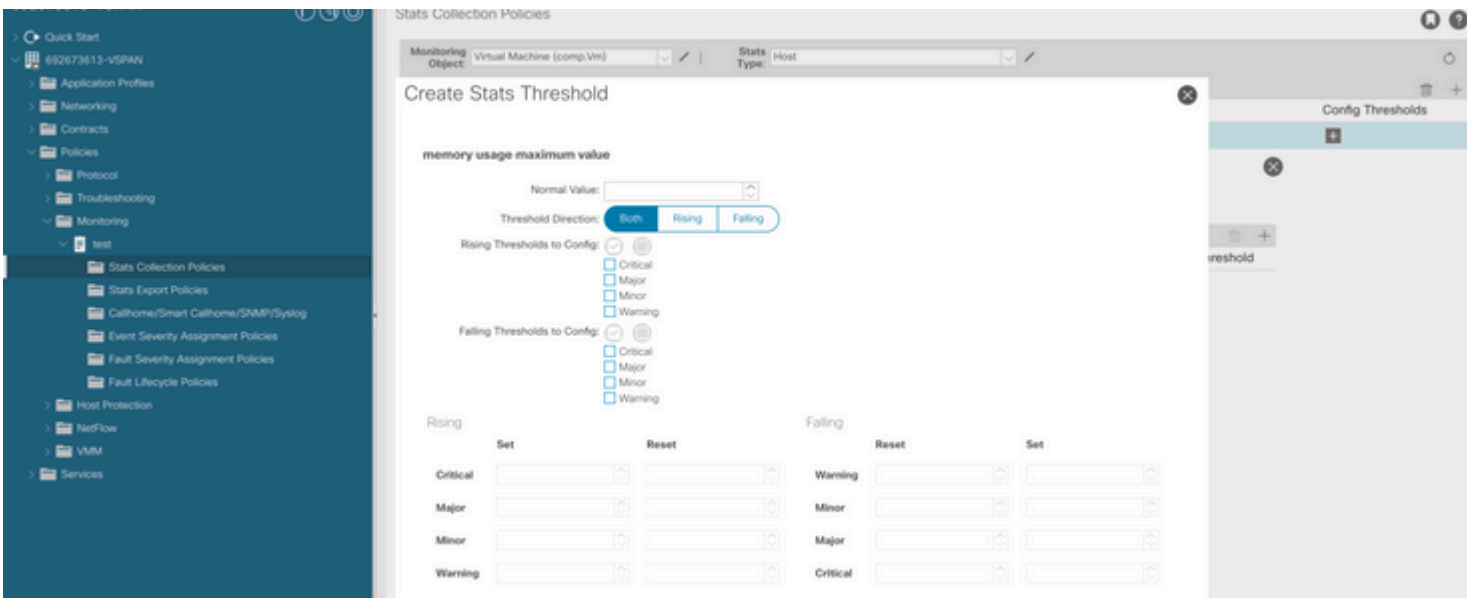


f.

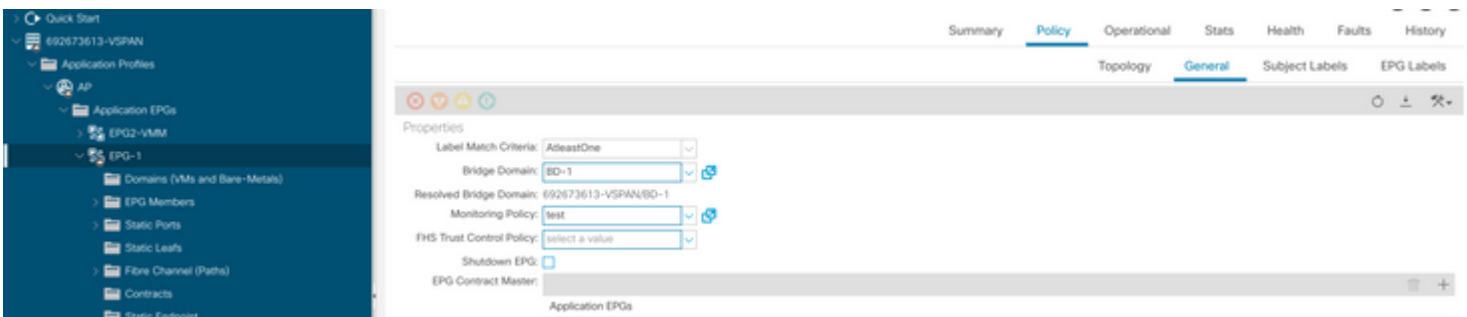
f. Click on the + Sign under config threshold and add "memory usage maximum value" as property.



g. Change the normal value to the threshold you would prefer.



h. Apply the monitoring policy on the EPG



I. To confirm if the policy is applied on the VM run "moquery -c compVm -f 'comp.Vm.oid = "vm-<vm-id>"'"

```
apic1# moquery -c compVm -f 'comp.Vm.oid == "vm-1071"' | grep monPolDn
```

monPolDn : uni/tn-692673613-VSPAN/monepg-test <== Monitoring Policy test has been applied

â€f

Fault : F93241

```
"Code" : "F93241",
>Description" : "TCA: CPU usage average value(compHostStats15min:cpuUsageAvg) value 100% raised above th
"Dn" : "comp/prov-VMware/ctrlr-[FAB4-AVE]-vcenter/vm-vm-1071/fault-F93241"
```

This specific fault is raised when VM host is consuming CPU more than the threshold. The APIC monitors these hosts via VCenter. Comp:HostStats15min is a class that represents the most current statistics for host in a 15 minute sampling interval. This class updates every 5 minutes.

Quick Start to Address Fault : F93241

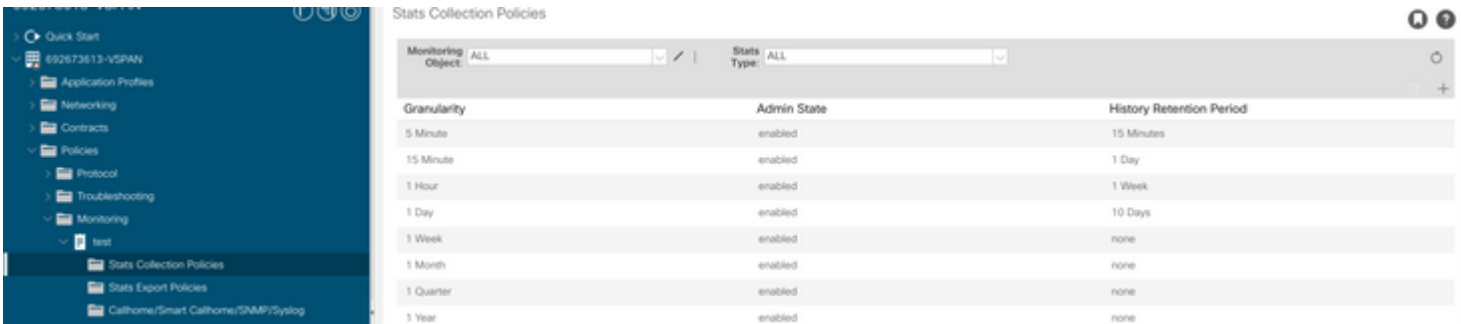
1. Command "moquery -d 'comp/prov-VMware/ctrlr-[<DVS>]-<VCenter>/vm-vm-<VM id from the fault's DN>'"

This command gives information about the affected VM


```
# comp.Vm
oid : vm-1071
cfgdOs : Ubuntu Linux (64-bit)
childAction :
descr :
dn : comp/prov-VMware/ctrlr-[FAB4-AVE]-vcenter/vm-vm-1071
ftRole : unset
guid : 501030b8-028a-be5c-6794-0b7bee827557
id : 0
issues :
lcOwn : local
modTs : 2022-04-21T17:16:06.572+05:30
monPolDn : uni/tn-692673613-VSPAN/monepg-test
name : VM3
nameAlias :
os :
rn : vm-vm-1071
state : poweredOn
status :
template : no
type : virt
uuid : 4210b04b-32f3-b4e3-25b4-fe73cd3be0ca
```

2. Command "moquery -c compRsHv | grep 'vm-1071'"

b. Under your Monitoring policy, select stats collection policy.




c. Click on edit icon beside Monitoring object dropdown, and check the Virtual Machine (comp.Vm) as a monitoring object. After Submitting, select the compVm object from Monitoring Object dropdown.



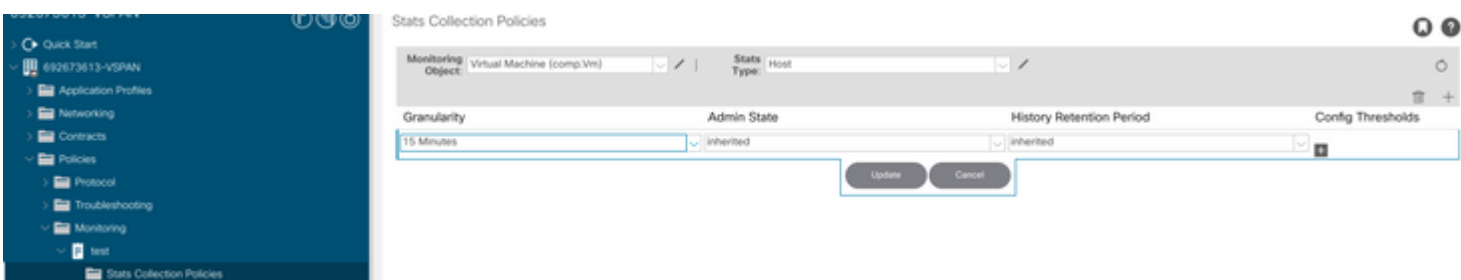
d.

d. Click on edit icon beside Stats type, then check on CPU Usage.

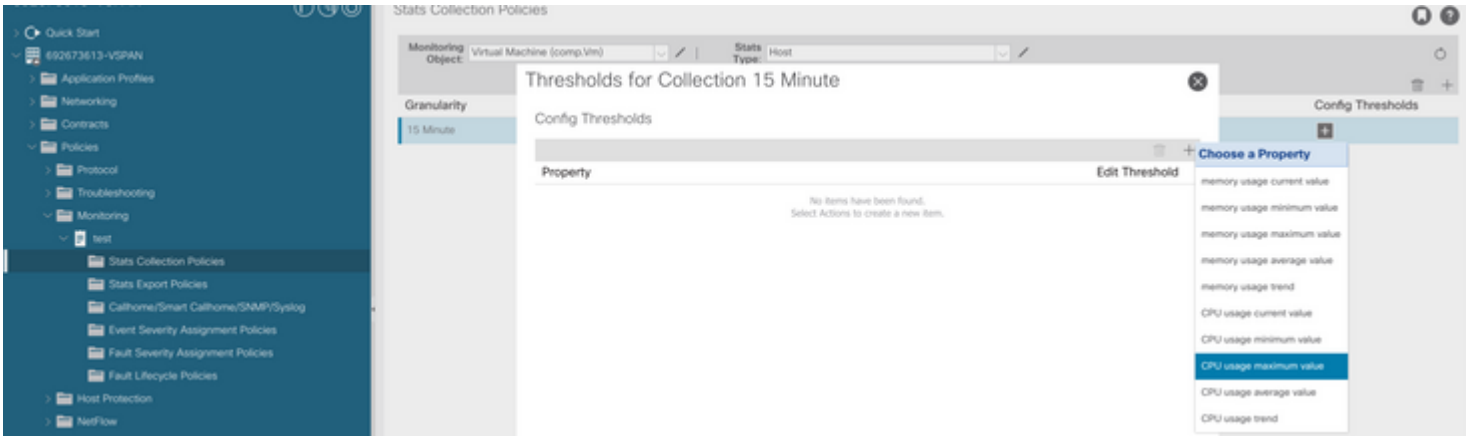


e.

e. From the stats type Dropdown click select host, click on the + sign and enter your Granularity, Admin state and History Retention Period and then click on update.

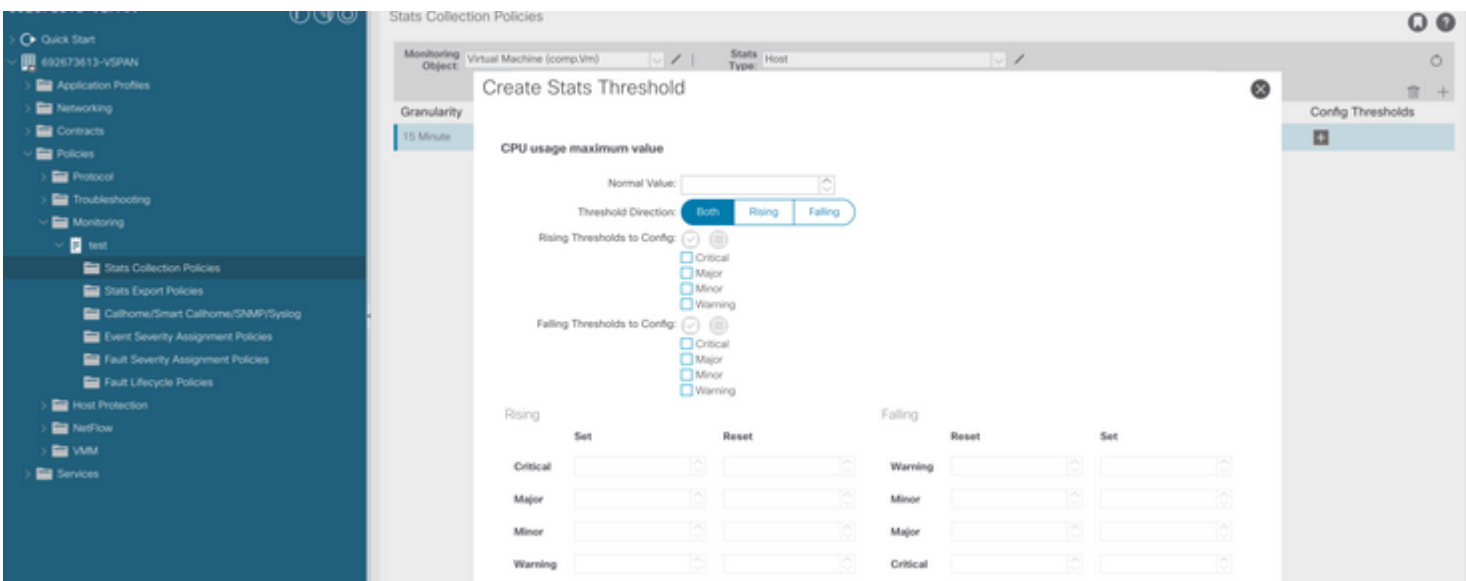


f. Click on the + Sign under config threshold and add "CPU usage maximum value" as property.

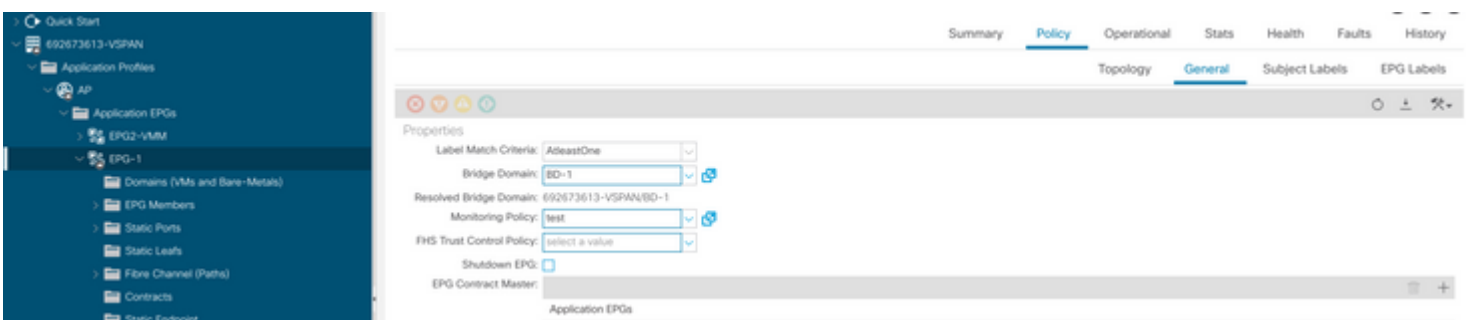


â€f

g. Change the normal value to the threshold you would prefer.



h. Apply the monitoring policy on the EPG



I. To confirm if the policy is applied on the VM run "moquery -c compVm -f 'comp.Vm.oid = "vm-<vm-id>'"

```
apic1# moquery -c compVm -f 'comp.Vm.oid == "vm-1071"' | grep monPolDn
monPolDn      : uni/tn-692673613-VSPAN/monepg-test <== Monitoring Policy test has been applied
```

Fault : F381328

```
"Code" : "F381328",  
"Description" : "TCA: CRC Align Errors current value(eqptIngrErrPkts5min:crcLast) value 50% raised above  
"Dn" : "topology/<pod>/<node>/sys/phys-<[interface]>/fault-F381328"
```

This specific fault is raised when CRC errors on an interface exceeds the threshold. There are two common types of CRC errors seen - FCS errors and CRC Stomped errors. CRC errors are propagated due to a cut-through switched path and are the result of initial FCS errors. Since ACI follows cut-through switching these frames end up traversing the ACI fabric and we see stomp CRC errors along the path, this does not mean that all the interfaces with CRC errors are faults. Recommendation is to identify the source of CRC and fix the problematic SFP/Port/Fibre.

Quick Start to Address Fault : F381328

1. Dump the highest number interfaces with CRC in fabric

```
moquery -c rmonEtherStats -f 'rmon.EtherStats.crcAlignErrors>="1"' | egrep "dn|crcAlignErrors" | egrep -  
topology/pod-1/node-103/sys/phys-[eth1/50]/dbgEtherStats 399158  
topology/pod-1/node-101/sys/phys-[eth1/51]/dbgEtherStats 399158  
topology/pod-1/node-1001/sys/phys-[eth2/24]/dbgEtherStats 399158
```

2. Dump the highest number of FCS in the fabric

```
moquery -c rmonDot3Stats -f 'rmon.Dot3Stats.fcSErrors>="1"' | egrep "dn|fcSErrors" | egrep -o "\S+$" |
```

Next Steps Fault : F381328

1. If there are FCS errors in the fabric address those errors. These errors typically indicate layer 1 issues.
2. If there are CRC stomp errors on front panel port, then check the connected device on the port and identify why stomps are coming from that device.

Python Script for fault : F381328

This entire process can also be automated using python script. Please refer <https://www.cisco.com/c/en/us/support/docs/cloud-systems-management/application-policy-infrastructure-controller-apic/217577-how-to-use-fcs-and-crc-troubleshooting-s.html>

Fault : F450296

```
"Code" : "F450296",
>Description" : "TCA: Multicast usage current value(eqptcapacityMcastEntry5min:perLast) value 91% raised
"Dn" : "sys/eqptcapacity/fault-F450296"
```

This specific fault is raised when number of multicast entries exceeds the threshold.

Quick Start to Address Fault : F450296

1. Command "show platform internal hal health-stats asic-unit all"

```
module-1# show platform internal hal health-stats asic-unit all
|Sandbox_ID: 0 Asic Bitmap: 0x0
|-----
```

L2 stats:

=====

```
bds: : 1979
max_bds: : 3500
  external_bds: : 0
  vsan_bds: : 0
  legacy_bds: : 0
  regular_bds: : 0
  control_bds: : 0
fds : 1976
max_fds : 3500
  fd_vlans : 0
  fd_vxlans : 0
vlans : 3955
max vlans : 3960
  vlan_xlates : 6739
  max vlan_xlates : 32768
  ports : 52
pcs : 47
hifs : 0
nif_pcs : 0
l2_local_host_entries : 1979
max_l2_local_host_entries : 32768
l2_local_host_entries_norm : 6
l2_total_host_entries : 1979
max_l2_total_host_entries : 65536
l2_total_host_entries_norm : 3
```

L3 stats:

=====

```
l3_v4_local_ep_entries : 3953
max_l3_v4_local_ep_entries : 32768
l3_v4_local_ep_entries_norm : 12
l3_v6_local_ep_entries : 1976
max_l3_v6_local_ep_entries : 24576
l3_v6_local_ep_entries_norm : 8
l3_v4_total_ep_entries : 3953
max_l3_v4_total_ep_entries : 65536
l3_v4_total_ep_entries_norm : 6
l3_v6_total_ep_entries : 1976
max_l3_v6_total_ep_entries : 49152
l3_v6_total_ep_entries_norm : 4
```

```

max_l3_v4_32_entries      : 98304
total_l3_v4_32_entries    : 35590
  l3_v4_total_ep_entries  : 3953
  l3_v4_host_uc_entries   : 37
  l3_v4_host_mc_entries   : 31600
total_l3_v4_32_entries_norm : 36
max_l3_v6_128_entries     : 49152
total_l3_v6_128_entries   : 3952
  l3_v6_total_ep_entries  : 1976
  l3_v6_host_uc_entries   : 1976
  l3_v6_host_mc_entries   : 0
total_l3_v6_128_entries_norm : 8
max_l3_lpm_entries       : 38912
l3_lpm_entries           : 9384
  l3_v4_lpm_entries       : 3940
  l3_v6_lpm_entries       : 5444
l3_lpm_entries_norm      : 31
max_l3_lpm_tcam_entries  : 4096
max_l3_v6_wide_lpm_tcam_entries: 1000
l3_lpm_tcam_entries      : 2689
  l3_v4_lpm_tcam_entries  : 2557
  l3_v6_lpm_tcam_entries  : 132
  l3_v6_wide_lpm_tcam_entries : 0
l3_lpm_tcam_entries_norm : 65
l3_v6_lpm_tcam_entries_norm : 0
l3_host_uc_entries       : 2013
  l3_v4_host_uc_entries   : 37
  l3_v6_host_uc_entries   : 1976
max_uc_ecmp_entries      : 32768
uc_ecmp_entries          : 1
uc_ecmp_entries_norm     : 0
max_uc_adj_entries       : 8192
uc_adj_entries           : 1033
uc_adj_entries_norm      : 12
vrfs                      : 1806
  infra_vrfs              : 0
  tenant_vrfs             : 1804
rtd_ifs                   : 2
sub_ifs                    : 2
svi_ifs                   : 1978

```

Mcast stats:

```

=====
mcast_count                : 31616 <<<<<<<
max_mcast_count            : 32768

```

Policy stats:

```

=====
policy_count                : 127116
max_policy_count            : 131072
policy_otcam_count          : 2920
max_policy_otcam_count      : 8192
policy_label_count          : 0
max_policy_label_count      : 0

```

Dci Stats:

```

=====
vlan_xlate_entries          : 0
vlan_xlate_entries_tcam    : 0
max_vlan_xlate_entries     : 0
sclass_xlate_entries       : 0
sclass_xlate_entries_tcam  : 0

```

max_sclass_xlate_entries : 0

Next Steps Fault : F450296

1. Consider moving some of multicast traffic to other Leafs.
2. Explore various forwarding scale profile to increase multicast scale. refer link <https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/all/forwarding-scale-profiles/cisco-apic-forwarding-scale-profiles/m-forwarding-scale-profiles-523.html>