排除SMF CNDP "network-receive-error"；故障在eno6/bd0接口上

目录

简介

问题

- 确定警报的来源

验证节点、Pod、端口状态

- 来自主要VIP的节点和Pod验证
- K8s主VIP的端口验证
- 从SMI集群部署器进行端口验证

识别UCS服务器

- 从SMI集群部署器验证UCS服务器
- 映射主要VIP端口和UCS网络接口

识别枝叶交换机

解决方案

简介

本文档介绍如何识别特定会话管理功能(SMF)云本地部署平台(CNDP)的计算和枝叶交换机，并解决通用执行环境(CEE)中报告的“network-receive-error”警报。

问题

CEE Opcenter Rack2上会报告“network-receive-error”警报。

```
[lab0200-smf/labceed22] cee# show alerts active summary
NAME UID SEVERITY STARTS AT SOURCE SUMMARY
-------------------------------------------------------------------------------------------------------
network-receive-error 998c77d6a6a0 major 10-26T00:10:31 lab0200-smf-mas Network interface "bd0" showing
network-receive-error ea4217bf9d9e major 10-26T00:10:31 lab0200-smf-mas Network interface "bd0" showing
network-receive-error 97fad40d2a58 major 10-26T00:10:31 lab0200-smf-mas Network interface "eno6" showin
network-receive-error b79540eb4e78 major 10-26T00:10:31 lab0200-smf-mas Network interface "eno6" showin
network-receive-error e3d163ff4012 major 10-26T00:10:01 lab0200-smf-mas Network interface "bd0" showing
network-receive-error 12a7b5a5c5d5 major 10-26T00:10:01 lab0200-smf-mas Network interface "eno6" showin
```

有关警报的说明，请参阅Ultra Cloud Core Subscriber Microservices Infrastructure Operations Guide。

<#root>

Alert:

network-receive-errors

```
 Annotations:
  Type: Communications Alarm
  Summary: Network interface "{{ $labels.device }}" showing receive errors on hostname {{ $labels.hostn
  Expression:
    |
   rate(node_network_receive_errs_total{device!~"veth.+"}[2m]) > 0
For: 2m
Labels:
 Severity: major
```

确定警报的来源

登录CEE labceed22，检查bd0和eno6接口上报告的“network-receive-error”警报详细信息，以识别节点和pod。

```
[lab0200-smf/labceed22] cee# show alerts active summary
NAME                   UID           SEVERITY  STARTS AT       SOURCE                  SUMMARY
-------------------------------------------------------------------------------------------------
network-receive-error  3b6a0a7ce1a8  major     10-26T21:17:01  lab0200-smf-mas  Network interface "bd0"
network-receive-error  15abab75c8fc  major     10-26T21:17:01  lab0200-smf-mas  Network interface "eno6
```

执行show alerts active detail network-receive-error <UID>以提取警报的详细信息。

在示例中，两个警报的来源都是node lab0200-smf-primary-1 pod node-exporter-47xmm。

<#root>

[lab0200-smf/labceed22] cee#

show alerts active detail network-receive-error 3b6a0a7ce1a8

```
alerts active detail network-receive-error 3b6a0a7ce1a8
 severity    major
 type        "Communications Alarm"
 startsAt    2021-10-26T21:17:01.913Z
 source      lab0200-smf-primary-1
 summary     "Network interface \"bd0\" showing receive errors on hostname lab0200-smf-primary-1\""
 labels      [ "alertname: network-receive-errors" "cluster: lab0200-smf_cee-labceed22" "component: nod
```

hostname: lab0200-smf-primary-1

" "instance: 10.192.1.42:9100" "job: kubernetes-pods" "monitor: prometheus" "namespace: cee-labceed22"

pod: node-exporter-47xmm

" "pod_template_generation: 1" "replica: lab0200-smf_cee-labceed22" "severity: major"]

```
 annotations [ "summary: Network interface \"bd0\" showing receive errors on hostname lab0200-smf-prima

[lab0200-smf/labceed22] cee#
```

show alerts active detail network-receive-error 15abab75c8fc

```
alerts active detail network-receive-error 15abab75c8fc
 severity    major
 type        "Communications Alarm"
 startsAt    2021-10-26T21:17:01.913Z
 source      lab0200-smf-primary-1
 summary     "Network interface \"eno6\" showing receive errors on hostname lab0200-smf-primary-1\""
 labels      [ "alertname: network-receive-errors" "cluster: lab0200-smf_cee-labceed22" "component: nod
```

hostname: lab0200-smf-primary-1

```
" "instance: 10.192.1.42:9100" "job: kubernetes-pods" "monitor: prometheus" "namespace: cee-labceed22"
```

pod: node-exporter-47xmm

```
" "pod_template_generation: 1" "replica: lab0200-smf_cee-labceed22" "severity: major" ]
 annotations [ "summary: Network interface \"eno6\" showing receive errors on hostname lab0200-smf-prim
```

验证节点、Pod、端口状态

来自主要VIP的节点和Pod验证

登录Rack2的K8s主VIP以验证源节点和Pod的状态。

在本例中，两者都处于良好状态：准备就绪，正在运行。

<#root>

```
cloud-user@lab0200-smf-primary-1:~$
```

kubectl get nodes

```
NAME                          STATUS    ROLES                  AGE    VERSION
```

lab0200-smf-primary-1

Ready

```
    control-plane          105d   v1.21.0
lab0200-smf-primary-2   Ready    control-plane          105d   v1.21.0
lab0200-smf-primary-3   Ready    control-plane          105d   v1.21.0
lab0200-smf-worker-1    Ready    <none>                 105d   v1.21.0
lab0200-smf-worker-2    Ready    <none>                 105d   v1.21.0
lab0200-smf-worker-3    Ready    <none>                 105d   v1.21.0
lab0200-smf-worker-4    Ready    <none>                 105d   v1.21.0
lab0200-smf-worker-5    Ready    <none>                 105d   v1.21.0
```

<#root>

cloud-user@lab0200-smf-primary-1:~$

kubectl get pods -A -o wide

 | grep node-exporter--47xmm
cee-labceed22   node-exporter-47xmm   1/1

Running

 0   18d   10.192.1.44   lab0200-smf-primary-1   <none>   <none>

K8s主VIP的端口验证

验证bd0和eno6接口是否为UP状态，并使用ip addr | grep eno6 and ip addr | grep bd0.

> 注意：当过滤器应用于bd0时，eno6显示在输出中。原因是eno5和eno6配置为bd0下的绑定接口，可以在SMI集群部署器中验证。

<#root>

cloud-user@lab0200-smf-primary-1:~$

ip addr | grep eno6

3:

eno6

: <BROADCAST,MULTICAST,SECONDARY,UP,LOWER_UP> mtu 1500 qdisc mq primary bd0

state UP

 group default qlen 1000

cloud-user@lab0200-smf-primary-1:~$

ip addr | grep bd0

2: eno5: <BROADCAST,MULTICAST,SECONDARY,UP,LOWER_UP> mtu 1500 qdisc mq primary bd0 state UP group defau
3: eno6: <BROADCAST,MULTICAST,SECONDARY,UP,LOWER_UP> mtu 1500 qdisc mq primary bd0 state UP group defau
12:

bd0

: <BROADCAST,MULTICAST,PRIMARY,UP,LOWER_UP> mtu 1500 qdisc noqueue

state UP

 group default qlen 1000
13: vlan111@bd0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP group default qlen 1
14: vlan112@bd0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP group default qlen 1
182: cali7a166bd093d@if4: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1440 qdisc noqueue state UP group defau

从SMI集群部署器进行端口验证

登录Cluster Manager VIP，然后ssh访问运营(Ops)中心ops-center-smi-cluster-deployer。

<#root>

```
cloud-user@lab-deployer-cm-primary:~$
```

kubectl get svc -n smi-cm

```
NAME                                          TYPE        CLUSTER-IP       EXTERNAL-IP     PORT(S)
cluster-files-offline-smi-cluster-deployer    ClusterIP   10.102.53.184    <none>          8080/TCP
iso-host-cluster-files-smi-cluster-deployer   ClusterIP   10.102.38.70     172.16.1.102    80/TCP
iso-host-ops-center-smi-cluster-deployer      ClusterIP   10.102.83.54     172.16.1.102    3001/TCP
netconf-ops-center-smi-cluster-deployer       ClusterIP   10.102.196.125   10.241.206.65   3022/TCP,22/
```

ops-center-smi-cluster-deployer

```
              ClusterIP   10.102.12.170    <none>          8008/TCP,2024/TCP,2022/TCP,7681/TCP,3000/TCP
squid-proxy-node-port                     NodePort    10.102.72.168    <none>          3128:32572/T
```

<#root>

```
cloud-user@lab-deployer-cm-primary:~$
```

ssh -p 2024 admin@10.102.12.170

```
admin@10.102.12.170's password:
      Welcome to the Cisco SMI Cluster Deployer on lab-deployer-cm-primary
      Copyright © 2016-2020, Cisco Systems, Inc.
      All rights reserved.
admin connected from 172.16.1.100 using ssh on ops-center-smi-cluster-deployer-5cdc5f94db-bnxqt
[lab-deployer-cm-primary]
```

SMI Cluster Deployer#

验证节点的集群、节点默认值、接口和参数模式。在本例中， lab0200-smf。

<#root>

```
[lab-deployer-cm-primary] SMI Cluster Deployer#  show running-config clusters
```

clusters lab0200-smf

```
 environment lab0200-smf-deployer_1
…
 node-defaults initial-boot netplan ethernets eno5
  dhcp4 false
  dhcp6 false
 exit
```

```
 node-defaults initial-boot netplan ethernets eno6
  dhcp4 false
  dhcp6 false
 exit
 node-defaults initial-boot netplan ethernets enp216s0f0
  dhcp4 false
  dhcp6 false
 exit
 node-defaults initial-boot netplan ethernets enp216s0f1
  dhcp4 false
  dhcp6 false
 exit
 node-defaults initial-boot netplan ethernets enp94s0f0
  dhcp4 false
  dhcp6 false
 exit
 node-defaults initial-boot netplan ethernets enp94s0f1
  dhcp4 false
  dhcp6 false
 exit
```

node-defaults initial-boot netplan bonds bd0

```
  dhcp4      false
  dhcp6      false
  optional   true
```

interfaces [eno5 eno6]

parameters mode   active-backup

```
  parameters mii-monitor-interval 100
  parameters fail-over-mac-policy active
 exit
```

在主VIP中，验证接口bd0和eno6上的错误和/或丢弃。

当两个接口都发生丢弃时，必须检查UCS或枝叶交换机硬件是否存在任何硬件问题。

<#root>

```
cloud-user@lab0200-smf-primary-1:~$
```

ifconfig bd0

```
bd0: flags=5187<UP,BROADCAST,RUNNING,PRIMARY,MULTICAST>  mtu 1500
        inet6 fe80::8e94:1fff:fef6:53cd  prefixlen 64  scopeid 0x20<link>
        ether 8c:94:1f:f6:53:cd  txqueuelen 1000  (Ethernet)
        RX packets 47035763777  bytes 19038286946282 (19.0 TB)
        RX errors 49541
```

dropped 845484

```
  overruns 0  frame 49541
        TX packets 53797663096  bytes 32320571418654 (32.3 TB)
        TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0

cloud-user@lab0200-smf-primary-1:~$
```

ifconfig eno6

```
eno6: flags=6211<UP,BROADCAST,RUNNING,SECONDARY,MULTICAST>  mtu 1500
        ether 8c:94:1f:f6:53:cd  txqueuelen 1000  (Ethernet)
        RX packets 47035402290  bytes 19038274391478 (19.0 TB)
        RX errors 49541
```

dropped 845484

```
  overruns 0  frame 49541
        TX packets 53797735337  bytes 32320609021235 (32.3 TB)
        TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0
```

识别UCS服务器

从SMI集群部署器验证UCS服务器

在SMI集群部署器中运行show running-config clusters <cluster name> nodes <node name> 以查找UCS服务器的CIMC IP地址。

<#root>

```
[lab-deployer-cm-primary] SMI Cluster Deployer#
```

show running-config clusters lab0200-smf nodes primary-1

```
clusters lab0200-smf

nodes primary-1

  maintenance  false
  host-profile cp-data-r2-sysctl
  k8s node-type       primary
  k8s ssh-ip          10.192.1.42
  k8s sshd-bind-to-ssh-ip true
  k8s node-ip         10.192.1.42
  k8s node-labels smi.cisco.com/node-type oam
  exit
  k8s node-labels smi.cisco.com/node-type-1 proto
  exit
  ucs-server cimc user admin
...
```

ucs-server cimc ip-address 172.16.1.62

```
...
  exit
```

通过Active CM通过SSH连接到CIMC IP地址172.16.1.62，并验证服务器名称。

在本例中，服务器名称为LAB0200-Server8-02。

<#root>

```
cloud-user@lab-deployer-cm-primary:~$
```

ssh admin@172.16.1.62

```
Warning: Permanently added '172.16.1.62' (RSA) to the list of known hosts.
admin@172.16.1.62's password:
```

LAB0200-Server8-02

#

注意：验证客户信息调查问卷(CIQ)中服务器的名称（如果CIQ可用）。

映射主要VIP端口和UCS网络接口

在主VIP上，使用ls -la /sys/class/net 命令检查eno6的物理接口名称。在本示例中，使用lscpi 标识eno6设备时，必须使用端口1d:00.1标识eno6。

<#root>

```
cloud-user@lab0200-smf-primary-1:~$ ls -la /sys/class/net
total 0
drwxr-xr-x  2 root root    0 Oct 12 06:18 .
drwxr-xr-x 87 root root    0 Oct 12 06:18 ..
lrwxrwxrwx  1 root root    0 Oct 12 06:18 bd0 -> ../../devices/virtual/net/bd0
lrwxrwxrwx  1 root root    0 Oct 12 06:18 bd1 -> ../../devices/virtual/net/bd1
…
lrwxrwxrwx  1 root root    0 Oct 12 06:18 eno5 -> ../../devices/pci0000:17/0000:17:00.0/0000:18:00.0/00
lrwxrwxrwx  1 root root    0 Oct 12 06:18 eno6 -> ../../devices/pci0000:17/0000:17:00.0/0000:18:00.0/00
```

1d:00.1/net/eno6

注意：lspci显示有关UCS服务器上所有设备的信息，如MLOM、SLOM、PCI等。设备信息可用于映射ls -la /sys/class/net 命令输出中的接口名称。

在本例中，端口1d:00.1属于MLOM和eno6接口。eno5是1d:00.0 MLOM端口。

<#root>

```
cloud-user@lab0200-smf-primary-1:~$ lspci
……
1d:00.0 Ethernet controller: Cisco Systems Inc VIC Ethernet NIC (rev a2)
```

1d:00.1

```
 Ethernet controller: Cisco Systems Inc VIC Ethernet NIC (rev a2)
3b:00.0 Ethernet controller: Intel Corporation Ethernet Controller 10G X550T (rev 01)
3b:00.1 Ethernet controller: Intel Corporation Ethernet Controller 10G X550T (rev 01)
5e:00.0 Ethernet controller: Intel Corporation Ethernet Controller XL710 for 40GbE QSFP+ (rev 02)
5e:00.1 Ethernet controller: Intel Corporation Ethernet Controller XL710 for 40GbE QSFP+ (rev 02)
d8:00.0 Ethernet controller: Intel Corporation Ethernet Controller XL710 for 40GbE QSFP+ (rev 02)
d8:00.1 Ethernet controller: Intel Corporation Ethernet Controller XL710 for 40GbE QSFP+ (rev 02)
```

在CIMC GUI中，匹配主VIP的ifconfig输出中看到的MLOM MAC地址。

<#root>

```
cloud-user@lab0200-smf-primary-1:~$  ifconfig bd0
bd0: flags=5187<UP,BROADCAST,RUNNING,PRIMARY,MULTICAST>  mtu 1500
        inet6 fe80::8e94:1fff:fef6:53cd  prefixlen 64  scopeid 0x20<link>
        ether
```

8c:94:1f:f6:53:cd

```
  txqueuelen 1000  (Ethernet)
        RX packets 47035763777  bytes 19038286946282 (19.0 TB)
        RX errors 49541  dropped 845484  overruns 0  frame 49541
        TX packets 53797663096  bytes 32320571418654 (32.3 TB)
        TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0


cloud-user@lab0200-smf-primary-1:~$  ifconfig eno6
eno6: flags=6211<UP,BROADCAST,RUNNING,SECONDARY,MULTICAST>  mtu 1500
        ether
```

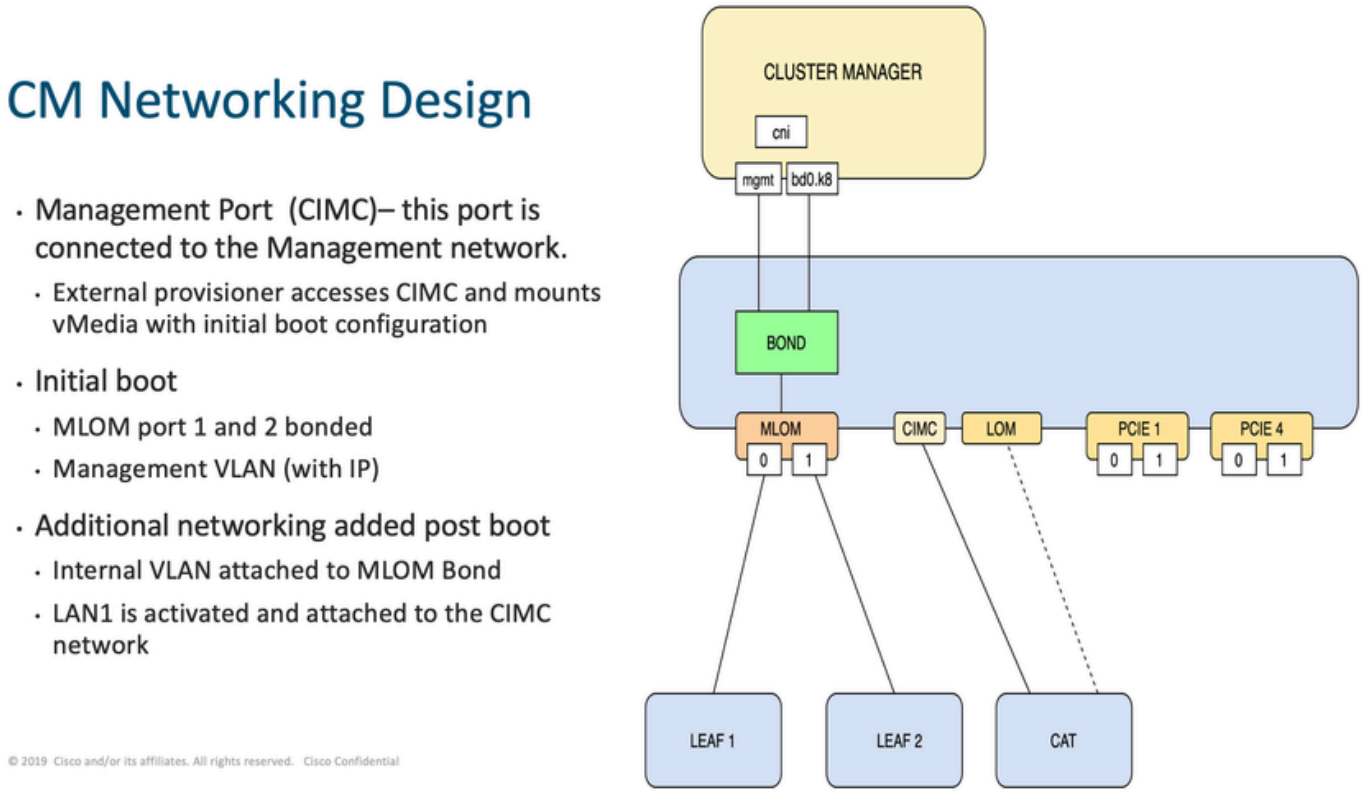
8c:94:1f:f6:53:cd

```
  txqueuelen 1000  (Ethernet)
        RX packets 47035402290  bytes 19038274391478 (19.0 TB)
        RX errors 49541  dropped 845484  overruns 0  frame 49541
        TX packets 53797735337  bytes 32320609021235 (32.3 TB)
        TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0
```

识别枝叶交换机

如图所示，在集群管理器网络中，MLOM(eno5/eno6)连接到枝叶1和2。

注意：如果CIQ可用，则验证将在CIQ中保留主机名。



同时登录到Leaves和grep服务器的名称。

在本例中，LAB0200-Server8-02 MLOM和MLOM接口连接到Leaf1和Leaf2上的接口Eth1/49。

<#root>

```
Leaf1# sh int description | inc LAB0200-Server8-02
Eth1/10       eth    40G     PCIE-01-2-LAB0200-Server8-02
Eth1/30       eth    40G     PCIE-02-2-LAB0200-Server8-02

Eth1/49       eth    40G     LAB0200-Server8-02 MLOM-P2
```

```
Leaf2# sh int description | inc LAB0200-Server8-02
Eth1/10       eth    40G     PCIE-01-1-LAB0200-Server8-02
Eth1/30       eth    40G     PCIE-02-1-LAB0200-Server8-02

Eth1/49       eth    40G     LAB0200-Server8-02 MLOM-P1
```

解决方案

重要信息： 每个问题都需要有自己的分析。如果在Nexus端未发现错误，则检查UCS服务器

⚠ 接口是否存在错误。

在本场景中，此问题与与LAB0200-Server8-02 MLOM eno6连接的Leaf1 int eth1/49上的链路故障有关。

UCS服务器已验证，未发现硬件问题，MLOM和端口状态良好。

Leaf1显示TX输出错误：

<#root>

```
Leaf1# sh int Eth1/49
Ethernet1/49 is up
admin state is up, Dedicated Interface
  Hardware: 10000/40000/100000 Ethernet, address: e8eb.3437.48ca (bia e8eb.3437.48ca)
  Description: LAB0200-Server8-02 MLOM-P2
  MTU 9216 bytes, BW 40000000 Kbit , DLY 10 usec
  reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, medium is broadcast
  Port mode is trunk
  full-duplex, 40 Gb/s, media type is 40G
  Beacon is turned off
  Auto-Negotiation is turned on  FEC mode is Auto
  Input flow-control is off, output flow-control is off
  Auto-mdix is turned off
  Rate mode is dedicated
  Switchport monitor is off
  EtherType is 0x8100
  EEE (efficient-ethernet) : n/a
    admin fec state is auto, oper fec state is off
  Last link flapped 5week(s) 6day(s)
  Last clearing of "show interface" counters never
  12 interface resets
  Load-Interval #1: 30 seconds
    30 seconds input rate 162942488 bits/sec, 26648 packets/sec
    30 seconds output rate 35757024 bits/sec, 16477 packets/sec
    input rate 162.94 Mbps, 26.65 Kpps; output rate 35.76 Mbps, 16.48 Kpps
  Load-Interval #2: 5 minute (300 seconds)
    300 seconds input rate 120872496 bits/sec, 22926 packets/sec
    300 seconds output rate 54245920 bits/sec, 17880 packets/sec
    input rate 120.87 Mbps, 22.93 Kpps; output rate 54.24 Mbps, 17.88 Kpps
  RX
    85973263325 unicast packets  6318912 multicast packets  55152 broadcast packets
    85979637389 input packets  50020924423841 bytes
    230406880 jumbo packets  0 storm suppression bytes
    0 runts  0 giants  0 CRC  0 no buffer
    0 input error  0 short frame  0 overrun   0 underrun  0 ignored
    0 watchdog  0 bad etype drop  0 bad proto drop  0 if down drop
    0 input with dribble  0 input discard
    0 Rx pause
  TX


    76542979816 unicast packets  88726302 multicast packets  789768 broadcast packets
    76632574981 output packets  29932747104403 bytes
    3089287610 jumbo packets
```

79095 output error

```
 0 collision  0 deferred  0 late collision
    0 lost carrier  0 no carrier  0 babble  0 output discard
    0 Tx pause
```

int eth1/49 Leaf1上使用电缆更换解决了“network-receive-error”警报。

上一次接口链路故障报告是在更换电缆之前。

<#root>

```
2021 Nov 17 07:36:48 TPLF0201 %BFD-5-SESSION_STATE_DOWN: BFD session 1090519112 to neighbor 10.22.101.1
Detection Time Expired.
2021 Nov 17 07:37:30 TPLF0201 %BFD-5-SESSION_STATE_DOWN: BFD session 1090519107 to neighbor 10.22.101.2
Detection Time Expired.
```

2021 Nov 18 05:09:12 TPLF0201 %ETHPORT-5-IF_DOWN_LINK_FAILURE: Interface Ethernet1/48 is down (Link fail

更换电缆后，在标签22的eno6/bd0上清除警报。

<#root>

```
[lab0200-smf/labceed22] cee#
```

show alerts active summary

```
NAME UID SEVERITY STARTS AT SOURCE SUMMARY
--------------------------------------------------------------------------------------------------------
watchdog a62f59201ba8 minor 11-02T05:57:18 System This is an alert meant to ensure that the entire alert
```

关于此翻译
思科采用人工翻译与机器翻译相结合的方式将此文档翻译成不同语言，希望全球的用户都能通过各自的语言得到支持性的内容。

请注意：即使是最好的机器翻译，其准确度也不及专业翻译人员的水平。

Cisco Systems, Inc. 对于翻译的准确性不承担任何责任，并建议您总是参考英文原始文档（已提供链接）。