

Dépanner SMF CNDP "network-receive-error" sur les interfaces eno6/bd0

Contenu

[Introduction](#)

[Problème](#)

[Identifier la source des alertes](#)

[Valider le noeud, le pod et l'état des ports](#)

[Validation des noeuds et des pods à partir du VIP principal](#)

[Validations de ports à partir du VIP principal K8s](#)

[Validations de ports à partir de SMI Cluster Deployer](#)

[Identifier le serveur UCS](#)

[Validation du serveur UCS à partir du dépoyeur de cluster SMI](#)

[Mappage des ports VIP principaux et des interfaces réseau UCS](#)

[Identification du commutateur leaf](#)

[Solution](#)

Introduction

Ce document décrit comment identifier le commutateur de calcul et de terminal pour une plateforme de déploiement cloud native (CNDP) SMF (Session Management Function) spécifique et résoudre l'alerte « network-receive-error » signalée dans Common Execution Environment (CEE).

Problème

Les alertes « network-receive-error » sont signalées sur le rack CEE Opcenter Rack2.

```
[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 receive errors on hostname lab0200-s...  
network-receive-error ea4217bf9d9e major 10-26T00:10:31 lab0200-smf-mas Network interface "bd0"  
showing receive errors on hostname lab0200-s...  
network-receive-error 97fad40d2a58 major 10-26T00:10:31 lab0200-smf-mas Network interface "eno6"  
showing receive errors on hostname lab0200-...  
network-receive-error b79540eb4e78 major 10-26T00:10:31 lab0200-smf-mas Network interface "eno6"  
showing receive errors on hostname lab0200-...  
network-receive-error e3d163ff4012 major 10-26T00:10:01 lab0200-smf-mas Network interface "bd0"  
showing receive errors on hostname lab0200-s...  
network-receive-error 12a7b5a5c5d5 major 10-26T00:10:01 lab0200-smf-mas Network interface "eno6"  
showing receive errors on hostname lab0200-...
```

Pour obtenir la description de l'alerte, reportez-vous au [Guide d'exploitation de l'infrastructure des microservices de l'abonné](#) de base [Ultra Cloud](#).

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

Identifier la source des alertes

Connectez-vous à CEE labceed22, vérifiez les détails d'alerte « network-receive-error » signalés sur les interfaces bd0 et eno6 pour identifier le noeud et le pod.

```
[lab0200-smf/labceed22] cee# show alerts active summary
NAME                               UID                               SEVERITY  STARTS AT          SOURCE                SUMMARY
-----
network-receive-error 3b6a0a7c1a8 major      10-26T21:17:01 lab0200-smf-mas Network
interface "bd0" showing receive errors on hostname tpc...
network-receive-error 15abab75c8fc major      10-26T21:17:01 lab0200-smf-mas Network
interface "eno6" showing receive errors on hostname tp...
```

Exécutez la commande **show alerts active detail network-receive-error <UID>** pour extraire les détails de l'alerte.

Dans l'exemple, la source des deux alertes est le noeud lab0200-smf-primary-1 pod node-exporter-47xmm.

```
[lab0200-smf/labceed22] cee# show alerts active detail network-receive-error 3b6a0a7c1a8
alerts active detail network-receive-error 3b6a0a7c1a8
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: node-exporter" "controller_revision_hash: 75c4cb979f" "device: bd0" "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-primary-1\" "type: Communications Alarm" ]
```

```
[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: node-exporter" "controller_revision_hash: 75c4cb979f" "device: eno6" "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-primary-1\" \"type: Communications Alarm" ]
```

Valider le noeud, le pod et l'état des ports

Validation des noeuds et des pods à partir du VIP principal

Connectez-vous au K8 VIP principal du rack2 pour valider l'état du noeud source et du pod.

Dans l'exemple, les deux sont en bon état : Prêt et en cours.

```
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

```
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>
```

Validations de ports à partir du VIP principal K8s

Validez que les interfaces bd0 et eno6 sont actives avec `ip addr | grep eno6` et adresse ip | `grep bd0`.

Note: Lorsque le filtre est appliqué pour bd0, l'eno6 est affiché dans la sortie. La raison est que eno5 et eno6 sont configurés en tant qu'interfaces liées sous bd0, qui peuvent être validées dans le SMI Cluster Deployer.

```
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 default qlen 1000
3: eno6: <BROADCAST,MULTICAST,SECONDARY,UP,LOWER_UP> mtu 1500 qdisc mq primary bd0 state UP
group default qlen 1000
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 1000
14: vlan112@bd0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP group default
qlen 1000
182: cali7a166bd093d@if4: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1440 qdisc noqueue state UP
group default
```

Validations de ports à partir de SMI Cluster Deployer

Connectez-vous au **VIP de Cluster Manager**, puis accédez par SSH à Operations (Ops) Center `ops-center-smi-cluster-deployer`.

```
cloud-user@lab-deployer-cm-primary:~$ kubectl get svc -n smi-cm
NAME                                TYPE           CLUSTER-IP      EXTERNAL-IP
PORT(S)                             AGE
cluster-files-offline-smi-cluster-deployer ClusterIP      10.102.53.184   <none>
8080/TCP                             110d
iso-host-cluster-files-smi-cluster-deployer ClusterIP      10.102.38.70    172.16.1.102
80/TCP                                110d
iso-host-ops-center-smi-cluster-deployer ClusterIP      10.102.83.54    172.16.1.102
3001/TCP                              110d
netconf-ops-center-smi-cluster-deployer ClusterIP      10.102.196.125  10.241.206.65
3022/TCP,22/TCP                      110d
ops-center-smi-cluster-deployer      ClusterIP      10.102.12.170   <none>
8008/TCP,2024/TCP,2022/TCP,7681/TCP,3000/TCP,3001/TCP 110d
squid-proxy-node-port                NodePort       10.102.72.168   <none>
3128:32572/TCP                      110d
```

```
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#
```

Vérifiez le cluster, les valeurs par défaut du noeud, les interfaces et le mode des paramètres pour le noeud. Dans l'exemple, le **lab0200-smf**.

```
[lab-deployer-cm-primary] SMI Cluster Deployer# show running-config clusters
clusters lab0200-smf
environment lab0200-smf-deployer_1
...
node-defaults initial-boot netplan ethernet eno5
dhcp4 false
dhcp6 false
exit
node-defaults initial-boot netplan ethernet eno6
dhcp4 false
dhcp6 false
exit
node-defaults initial-boot netplan ethernet enp216s0f0
dhcp4 false
dhcp6 false
exit
node-defaults initial-boot netplan ethernet enp216s0f1
dhcp4 false
dhcp6 false
exit
node-defaults initial-boot netplan ethernet enp94s0f0
dhcp4 false
dhcp6 false
exit
node-defaults initial-boot netplan ethernet 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

```

Dans le VIP principal, validez les erreurs et/ou les abandons sur les interfaces bd0 et eno6.

Lorsque les deux interfaces sont abandonnées, le matériel du commutateur UCS ou Leaf doit être vérifié pour tout problème matériel.

```

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

```

Identifier le serveur UCS

Validation du serveur UCS à partir du déployeur de cluster SMI

Exécutez `show running-config clusters <nom du cluster> nodes <nom du noeud>` dans le SMI Cluster Deployer pour connaître l'adresse IP CIMC du serveur UCS.

```

[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
```

Envoyez SSH à l'adresse IP CIMC 172.16.1.62 via Active CM et validez le nom du serveur.

Dans l'exemple, le nom du serveur est LAB0200-Server8-02.

```
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#
```

Note: Validez le nom du serveur dans le questionnaire d'informations client (CIQ), si le CIQ est disponible.

Mappage des ports VIP principaux et des interfaces réseau UCS

Sur le VIP principal, vérifiez les noms d'interface physique pour eno6 avec la commande **ls -la /sys/class/net**. Dans l'exemple, lorsque **lspci** est utilisé pour identifier le périphérique eno6, le port **1d:0.1** doit être utilisé pour identifier **eno6**.

```
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/pci0000:17/0000:17:00.0/0000:18:00.0/0000:19:01.0/0000:1b:00.0/0000:1c:00.0/0000:1d:00.0/net/eno5  
lrwxrwxrwx  1 root root    0 Oct 12 06:18 eno6 ->  
../../devices/pci0000:17/0000:17:00.0/0000:18:00.0/0000:19:01.0/0000:1b:00.0/0000:1c:00.0/0000:1d:00.1/net/eno6
```

Note: La **lspci** affiche des informations sur tous les périphériques sur le serveur UCS tels que MLOM, SLOM, PCI, etc. Les informations de périphérique peuvent être utilisées pour mapper avec les noms des interfaces dans la sortie de commande **ls -la /sys/class/net**.

Dans l'exemple, le port 1d:0.1 appartient à l'interface **MLOM** et **eno6**. Le **eno5** est le port MLOM 1d:0.0.

```
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)
```

Dans l'interface graphique de CIMC, faites correspondre l'adresse MAC MLOM vue sur le résultat de `ifconfig` du VIP principal.

```
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
```


Identification du commutateur leaf

Dans le réseau Cluster Manager, comme illustré dans l'image, les **MLOM (eno5/eno6)** sont connectés aux Leafs 1 et 2.

Note: La validation laisse les noms d'hôte dans CIQ, si le CIQ est disponible.

CM Networking Design

- Management Port (CIMC)– this port is connected to the Management network.
- External provisioner accesses CIMC and mounts vMedia with initial boot configuration
- Initial boot
 - MLOM port 1 and 2 bonded
 - Management VLAN (with IP)
- Additional networking added post boot
 - Internal VLAN attached to MLOM Bond
 - LAN1 is activated and attached to the CIMC network



© 2019 Cisco and/or its affiliates. All rights reserved. Cisco Confidential

Connectez-vous aux deux leafs et saisissez le nom du serveur.

Dans l'exemple, les interfaces LAB0200-Server8-02 MLOM et MLOM sont connectées aux interfaces **Eth1/49** sur Leaf1 et Leaf2.

```
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
```

Solution

Important : Chaque problème nécessite sa propre analyse. Si aucune erreur n'est détectée du côté Nexus, vérifiez les interfaces du serveur UCS.

Dans le scénario, le problème est lié à la défaillance de liaison sur Leaf1 int **eth1/49** qui est connecté avec LAB0200-Server8-02 MLOM eno6.

Le serveur UCS est validé et aucun problème matériel n'a été détecté, MLOM et les ports étaient en bon état.

Leaf1 a montré des erreurs de sortie TX :


```

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

```

L'alerte « network-receive-error » a été résolue par le remplacement du câble sur int eth1/49 Leaf1.

La dernière défaillance de liaison d'interface a été signalée juste avant le remplacement du câble.

```

2021 Nov 17 07:36:48 TPLF0201 %BFD-5-SESSION_STATE_DOWN: BFD session 1090519112 to neighbor
10.22.101.1 on interface Vlan2201 has gone down. Reason: Control
Detection Time Expired.
2021 Nov 17 07:37:30 TPLF0201 %BFD-5-SESSION_STATE_DOWN: BFD session 1090519107 to neighbor
10.22.101.2 on interface Vlan2201 has gone down. Reason: Control
Detection Time Expired.
2021 Nov 18 05:09:12 TPLF0201 %ETHPORT-5-IF_DOWN_LINK_FAILURE: Interface Ethernet1/48 is down
(Link failure)

```

Les alertes sont effacées sur eno6/bd0 du labceed22 après le remplacement du câble.

```
[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 alerting pipeline is functional. This ale...
```

À propos de cette traduction

Cisco a traduit ce document en traduction automatisée vérifiée par une personne dans le cadre d'un service mondial permettant à nos utilisateurs d'obtenir le contenu d'assistance dans leur propre langue.

Il convient cependant de noter que même la meilleure traduction automatisée ne sera pas aussi précise que celle fournie par un traducteur professionnel.