

Contenido

[Introducción](#)

[Antecedentes](#)

[La adyacencia OSPF no viene Online](#)

[Vecino OSPF pegado en el estado de la inicialización \(INIT\)](#)

[Vecino OSPF pegado en un estado de dos vías](#)

[Vecino OSPF pegado en el exstart/intercambio](#)

[Vecino OSPF pegado en un estado de carga](#)

Introducción

Este documento discute varios escenarios frecuentes encontrados donde no viene el vecino del Open Shortest Path First (OSPF) en línea como se esperaba. Para evitar este tipo de conducta inesperada en el Switches Cisco Nexus de la serie 7000, asegúrese seguir las guías de consulta y las restricciones sobre la capa 3 (L3) y el canal del puerto virtual (vPC).

Antecedentes

Antes de que usted resuelva problemas los problemas, asegúrese de que las guías de consulta y las restricciones estén resueltas. Refiera al [diseño y a la guía de configuración: Mejores prácticas para los canales del puerto virtual \(vPC\) en el Switches Cisco Nexus de la serie 7000](#) para más información sobre el L3 y el vPC.

El procedimiento usado para resolver problemas los problemas de la adyacencia OSPF en el nexa 7000 es similar a los procedimientos para el [®] del Cisco IOS, pero el nexa 7000 tiene herramientas más incorporadas y filtros para identificar fácilmente el problema.

La adyacencia OSPF no viene Online

Hay las épocas en que no viene la adyacencia OSPF en línea. El comando `show ip ospf neighbor` no muestra al vecino.

```
R3#show ip ospf neighbor
```

```
R3#
```

Este problema se pudo causar por:

- Problema de conectividad L2/L3
- OSPF no habilitado en la interfaz
- La interfaz se define como voz pasiva
- Máscara de subred unida mal
- Unido mal intervalo muerto/de saludo
- Clave de autenticación unida mal

- Discrepancia de ID de área
- Unido mal transite/stub/opción de la área no exclusiva de rutas internas (NSSA)

Para investigar el problema, marque la Conectividad, la configuración, y la lista de control de acceso (ACL)/las Políticas del plano de control (CoPP).

Marque la Conectividad L2/L3

1. Marque la Conectividad del unicast con el ping.

Si hay un problema de conectividad, descubra si es debido al Proveedor de servicios de Internet (ISP) L2, a un puerto físico, a un Convertidor de la interfaz de Gigabit (GBIC), o a un cable.

Nota: Asuma que no hay ACL/CoPP que bloquean el tráfico. Si el problema es debido al hardware defectuoso o a un cable, sustitúyalo o muévelo otro puerto para resolver problemas.

2. Marque la Conectividad del Multicast con el ping.

```
N7K1-RP# ping multicast 224.0.0.5 interface vlan 5
PING 224.0.0.5 (224.0.0.5): 56 data bytes
64 bytes from 5.5.5.2: icmp_seq=0 ttl=254 time=1.739 ms
64 bytes from 5.5.5.2: icmp_seq=1 ttl=254 time=1.253 ms
64 bytes from 5.5.5.2: icmp_seq=2 ttl=254 time=0.866 ms
64 bytes from 5.5.5.2: icmp_seq=3 ttl=254 time=1.045 ms
64 bytes from 5.5.5.2: icmp_seq=4 ttl=254 time=1.89 ms

--- 224.0.0.5 ping multicast statistics ---
5 packets transmitted,
From member 5.5.5.2: 5 packets received, 0.00% packet loss
--- in total, 1 group member responded ---
N7K1-RP#
```

Verifique que la interfaz sea limpia y que no hay descensos o errores con el comando 1/1 de los Ethernetes de la demostración internacional.

```
N7K1-RP# show int ethernet 1/20 | section RX|TX
RX
 340213 unicast packets  368092 multicast packets  2 broadcast packets
 708307 input packets  233094927 bytes
 0 jumbo packets  0 storm suppression packets
 0 runs  0 giants  0 CRC/FCS  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
 1374131 unicast packets  324752 multicast packets  3 broadcast packets
 1698886 output packets  196282264 bytes
 0 jumbo packets
 0 output error  0 collision  0 deferred  0 late collision
 0 lost carrier  0 no carrier  0 babble  0 output discard
 0 Tx pause
N7K1-RP#
```

3. Determine si estas características caen los paquetes en el linecard, la interfaz, o el CPU entrante.

- ACL - Entrante/saliente de la interfaz
- Calidad de Servicio (QoS) - En la interfaz

• CoPP

QoS

```
N7K1-RP# show policy-map interface ethernet 1/20
```

```
Global statistics status :    enabled
```

```
Ethernet1/20
```

```
Service-policy (queuing) input:    default-in-policy  
SNMP Policy Index:    301989913
```

```
Class-map (queuing):    in-q1 (match-any)  
queue-limit percent 50  
bandwidth percent 80  
queue dropped pkts : 0
```

```
Class-map (queuing):    in-q-default (match-any)  
queue-limit percent 50  
bandwidth percent 20  
queue dropped pkts : 0
```

```
Service-policy (queuing) output:    default-out-policy  
SNMP Policy Index:    301989922
```

```
Class-map (queuing):    out-pq1 (match-any)  
priority level 1  
queue-limit percent 16  
queue dropped pkts : 0
```

```
Class-map (queuing):    out-q2 (match-any)  
queue-limit percent 1  
queue dropped pkts : 0
```

```
Class-map (queuing):    out-q3 (match-any)  
queue-limit percent 1  
queue dropped pkts : 0
```

```
Class-map (queuing):    out-q-default (match-any)  
queue-limit percent 82  
bandwidth remaining percent 25  
queue dropped pkts : 0
```

CoPP

```
show policy-map interface control-plane class test1-copp-class-critical
```

```
Control Plane
```

```
service-policy input test1-copp-policy-lenient
```

```
class-map test1-copp-class-critical (match-any)  
match access-group name test1-copp-acl-bgp  
match access-group name test1-copp-acl-pim  
match access-group name test1-copp-acl-rip  
match access-group name test1-copp-acl-vpc  
match access-group name test1-copp-acl-bgp6  
match access-group name test1-copp-acl-igmp  
match access-group name test1-copp-acl-lisp  
match access-group name test1-copp-acl-msdp  
match access-group name test1-copp-acl-ospf
```

```

match access-group name test1-copp-acl-pim6
match access-group name test1-copp-acl-rip6
match access-group name test1-copp-acl-rise
match access-group name test1-copp-acl-eigrp
match access-group name test1-copp-acl-lisp6
match access-group name test1-copp-acl-ospf6
match access-group name test1-copp-acl-rise6
match access-group name test1-copp-acl-eigrp6
match access-group name test1-copp-acl-otv-as
match access-group name test1-copp-acl-mac-l2pt
match access-group name test1-copp-acl-mpls-ldp
match access-group name test1-copp-acl-mpls-oam
match access-group name test1-copp-acl-mpls-rsvp
match access-group name test1-copp-acl-mac-l3-isis
match access-group name test1-copp-acl-mac-otv-isis
match access-group name test1-copp-acl-mac-fabricpath-isis
match protocol mpls router-alert
match protocol mpls exp 6
set cos 7
police cir 39600 kbps bc 375 ms
  conform action: transmit
  violate action: drop
module 1:
  conformed 539964945 bytes,
    5-min offered rate 5093 bytes/sec
    peak rate 5213 bytes/sec
  violated 0 bytes,
    5-min violate rate 0 bytes/sec
module 2:
  conformed 784228080 bytes,
    5-min offered rate 5848 bytes/sec
    peak rate 7692 bytes/sec
  violated 0 bytes,
    5-min violate rate 0 bytes/sec
module 3:
  conformed 5114206 bytes,
    5-min offered rate 41 bytes/sec
    peak rate 6656 bytes/sec
  violated 0 bytes,
    5-min violate rate 0 bytes/sec

```

N7K1#

Marque la configuración de OSPF

Utilice estos comandos para marcar la configuración de OSPF (subred, intervalo muerto/de salud, ID de área, tipo de área, clave de autenticación (eventualmente), y no-pasivo), y asegúrese de que hace juego en los ambos lados.

1. Muestre el OSPF del funcionamiento
2. Muestre la interfaz OSPF 5 del IP
3. Muestre OSPF 5 del IP

Aquí está un ejemplo del primer comando:

N7K1-RP# **show run ospf**

```

!Command: show running-config ospf
!Time: Thu May 16 11:27:24 2013

```

```

version 6.2(2)
feature ospf

```

```
logging level ospf 7

router ospf 5
  router-id 5.5.0.1

interface Vlan5
  ip router ospf 5 area 0.0.0.0

interface loopback5
  ip router ospf 5 area 0.0.0.0
```

N7K1-RP#

Aquí está un ejemplo del comando second:

```
N7K1-RP# show ip ospf 5 interface
Vlan5 is up, line protocol is up
  IP address 5.5.5.1/24, Process ID 5 VRF default, area 0.0.0.0
  Enabled by interface configuration
  State DR, Network type BROADCAST, cost 40
  Index 2, Transmit delay 1 sec, Router Priority 1
  Designated Router ID: 5.5.0.1, address: 5.5.5.1
  Backup Designated Router ID: 5.5.0.2, address: 5.5.5.2
  1 Neighbors, flooding to 1, adjacent with 1
  Timer intervals: Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello timer due in 00:00:00
  No authentication
  Number of opaque link LSAs: 0, checksum sum 0
loopback5 is up, line protocol is up
  IP address 5.5.0.1/32, Process ID 5 VRF default, area 0.0.0.0
  Enabled by interface configuration
  State LOOPBACK, Network type LOOPBACK, cost 1
  Index 1
```

N7K1-RP#

Aquí está un ejemplo del tercer comando:

```
N7K1-RP# show ip ospf 5

Routing Process 5 with ID 5.5.0.1 VRF default
Routing Process Instance Number 3
Stateful High Availability enabled
Graceful-restart is configured
  Grace period: 60 state: Inactive
  Last graceful restart exit status: None
Supports only single TOS(TOS0) routes
Supports opaque LSA
Administrative distance 110
Reference Bandwidth is 40000 Mbps
SPF throttling delay time of 200.000 msecs,
  SPF throttling hold time of 1000.000 msecs,
  SPF throttling maximum wait time of 5000.000 msecs
LSA throttling start time of 0.000 msecs,
  LSA throttling hold interval of 5000.000 msecs,
  LSA throttling maximum wait time of 5000.000 msecs
Minimum LSA arrival 1000.000 msec
LSA group pacing timer 10 secs
Maximum paths to destination 8
Number of external LSAs 0, checksum sum 0
Number of opaque AS LSAs 0, checksum sum 0
Number of areas is 1, 1 normal, 0 stub, 0 nssa
Number of active areas is 1, 1 normal, 0 stub, 0 nssa
Install discard route for summarized external routes.
Install discard route for summarized internal routes.
```

```
Area BACKBONE(0.0.0.0)
Area has existed for 1d10h
Interfaces in this area: 2 Active interfaces: 2
Passive interfaces: 0 Loopback interfaces: 1
No authentication available
SPF calculation has run 47 times
Last SPF ran for 0.000542s
Area ranges are
Number of LSAs: 3, checksum sum 0x84d4
```

N7K1-RP#

Marque los mensajes OSPF

Ingrese el comando de la **adyacencia del historial de eventos OSPF del IP de la demostración** para verificar que los mensajes del debug son enviados y recibidos por el proceso OPSF.

Nota: Los últimos mensajes aparecen en el top.

La salida muestra todos los mensajes de la adyacencia OSPF que se intercambian entre los vecinos OSPF. Cuando se forma una adyacencia OSPF, un router pasa a través de varios cambios de estado antes de que llegue a ser completamente adyacente con su vecino. Esta salida muestra todos los cambios de estado y las negociaciones. Si hay un problema (unidad máxima de la transición (MTU), problemas de conectividad, caída de paquetes), se refleja en la salida.

N7K1-RP# **show ip ospf 5 event-history adjacency**

```
Adjacency events for OSPF Process "ospf-5"
2013 May 16 10:50:58.121128 ospf 5 [9386]: : mtu 1600, opts: 0x42, ddbits:
0, seq: 0x6f40fde4
2013 May 16 10:50:58.121124 ospf 5 [9386]: : Sent DBD with 0 entries to 5.5.5.2
on Vlan5
2013 May 16 10:50:58.121114 ospf 5 [9386]: : Sending DBD to 5.5.5.2 on Vlan5
2013 May 16 10:50:58.118030 ospf 5 [9386]: : Nbr 5.5.5.2: LOADING --> FULL,
event LDDONE
2013 May 16 10:50:58.115840 ospf 5 [9386]: : Built LS Request packet for 5.5.5.2
with 1 entries
2013 May 16 10:50:58.115835 ospf 5 [9386]: : Add 5.5.0.2(0x1)5.5.0.2
(0x8000104e)(0x7ef8) (156) to LSR
2013 May 16 10:50:58.115823 ospf 5 [9386]: : Building LS Request packet to
5.5.5.2
2013 May 16 10:50:58.112201 ospf 5 [9386]: : Nbr 5.5.5.2: EXCHANGE --> LOADING,
event EXCHDONE
2013 May 16 10:50:58.112026 ospf 5 [9386]: : seqnr 0x6f40fde4, dbdbits 0x1,
mtu 1600, options 0x42
2013 May 16 10:50:58.112022 ospf 5 [9386]: : Got DBD from 5.5.5.2 with 0 entries
2013 May 16 10:50:58.111988 ospf 5 [9386]: : seqnr 0x6f40fde4, dbdbits 0x1,
mtu 1600, options 0x42
2013 May 16 10:50:58.111984 ospf 5 [9386]: : Got DBD from 5.5.5.2 with 0 entries
2013 May 16 10:50:58.110169 ospf 5 [9386]: : mtu 1600, opts: 0x42, ddbits: 0,
seq: 0x6f40fde3
2013 May 16 10:50:58.110165 ospf 5 [9386]: : Sent DBD with 0 entries to 5.5.5.2
on Vlan5
2013 May 16 10:50:58.110155 ospf 5 [9386]: : Sending DBD to 5.5.5.2 on Vlan5
2013 May 16 10:50:58.106609 ospf 5 [9386]: : Added 1 out of 1 LSAs to
request list
2013 May 16 10:50:58.106606 ospf 5 [9386]: : Added 5.5.0.2(0x1)5.5.0.2
(0x8000104e)(0x7ef8) (156) to request list
2013 May 16 10:50:58.106586 ospf 5 [9386]: : seqnr 0x6f40fde3, dbdbits 0x3,
mtu 1600, options 0x42
```

2013 May 16 10:50:58.106582 ospf 5 [9386]: : Got DBD from 5.5.5.2 with 1 entries
2013 May 16 10:50:58.106537 ospf 5 [9386]: : seqnr 0x6f40fde3, dbdbits 0x3,
mtu 1600, options 0x42
2013 May 16 10:50:58.106532 ospf 5 [9386]: : Got DBD from 5.5.5.2 with 1 entries
2013 May 16 10:50:58.104462 ospf 5 [9386]: : Built reply LSU with 2 LSAs for
5.5.5.2 128 bytes
2013 May 16 10:50:58.104439 ospf 5 [9386]: : Added 5.5.5.2(0x2)5.5.0.2
(0x80000045) (0xaf32) (156)
2013 May 16 10:50:58.104431 ospf 5 [9386]: : Added 5.5.0.1(0x1)5.5.0.1
(0x80000ecf) (0xd834) (8)(0)
2013 May 16 10:50:58.104408 ospf 5 [9386]: : Building reply LSU to 5.5.5.2
2013 May 16 10:50:58.104404 ospf 5 [9386]: : 2 requests in LSR (2 left)
2013 May 16 10:50:58.104370 ospf 5 [9386]: : Answering LSR from 5.5.5.2
2013 May 16 10:50:58.100790 ospf 5 [9386]: : Recv LSR from Nbr 5.5.5.2
2013 May 16 10:50:58.099055 ospf 5 [9386]: : mtu 1600, opts: 0x42, ddbits:
0x2, seq: 0x6f40fde2
2013 May 16 10:50:58.099051 ospf 5 [9386]: : Sent DBD with 3 entries to 5.5.5.2
on Vlan5
2013 May 16 10:50:58.099038 ospf 5 [9386]: : Sending DBD to 5.5.5.2 on Vlan5
2013 May 16 10:50:58.095072 ospf 5 [9386]: : seqnr 0x6f40fde2, dbdbits 0x7,
mtu 1600, options 0x42
2013 May 16 10:50:58.095068 ospf 5 [9386]: : Got DBD from 5.5.5.2 with 0 entries
2013 May 16 10:50:58.095024 ospf 5 [9386]: : Nbr 5.5.5.2: EXSTART --> EXCHANGE,
event NEGDONE
2013 May 16 10:50:58.094895 ospf 5 [9386]: : We are SLAVE, 5.5.5.2 is master
2013 May 16 10:50:58.094890 ospf 5 [9386]: : seqnr 0x6f40fde2, dbdbits 0x7,
mtu 1600, options 0x42
2013 May 16 10:50:58.094886 ospf 5 [9386]: : Got DBD from 5.5.5.2 with 0 entries
2013 May 16 10:50:58.093037 ospf 5 [9386]: : mtu 1600, opts: 0x42, ddbits: 0x7,
seq: 0x7273409a
2013 May 16 10:50:58.093033 ospf 5 [9386]: : Sent DBD with 0 entries to 5.5.5.2
on Vlan5
2013 May 16 10:50:58.093029 ospf 5 [9386]: : Sending DBD to 5.5.5.2 on Vlan5
2013 May 16 10:50:58.092915 ospf 5 [9386]: : Nbr 5.5.5.2: INIT --> EXSTART,
event TWOWAYRCVD
2013 May 16 10:50:58.092862 ospf 5 [9386]: : Nbr 5.5.5.2: TWOWAY --> EXSTART,
event ADJOK
2013 May 16 10:50:58.092763 ospf 5 [9386]: [9446]: Interface Vlan5 ---> BDR
2013 May 16 10:50:58.092757 ospf 5 [9386]: [9446]: Elected 5.5.0.2 as DR,
5.5.0.1 as BDR
2013 May 16 10:50:58.092690 ospf 5 [9386]: [9446]: This nbr 5.5.5.2 promoted
to current dr
2013 May 16 10:50:58.092687 ospf 5 [9386]: [9446]: Walking neighbor 5.5.5.2
(0x93e3524), state TWOWAY
2013 May 16 10:50:58.092683 ospf 5 [9386]: [9446]: Neighbor not declared DR,
ignoring
2013 May 16 10:50:58.092680 ospf 5 [9386]: [9446]: Walking neighbor 5.5.5.1
(0xaclf7514), state SELF
2013 May 16 10:50:58.092676 ospf 5 [9386]: [9446]: DR election starting
2013 May 16 10:50:58.092673 ospf 5 [9386]: [9446]: This neighbor is greater
than 2way
2013 May 16 10:50:58.092670 ospf 5 [9386]: [9446]: Walking neighbor 5.5.5.2
(0x93e3524), state TWOWAY
2013 May 16 10:50:58.092666 ospf 5 [9386]: [9446]: Compare done, new current
bdr 5.5.5.1
2013 May 16 10:50:58.092663 ospf 5 [9386]: [9446]: Current BDR set to this
neighbor
2013 May 16 10:50:58.092660 ospf 5 [9386]: [9446]: This neighbor is in
consideration for bdr
2013 May 16 10:50:58.092657 ospf 5 [9386]: [9446]: This neighbor is greater
than 2way
2013 May 16 10:50:58.092654 ospf 5 [9386]: [9446]: Walking neighbor 5.5.5.1
(0xaclf7514), state SELF
2013 May 16 10:50:58.092650 ospf 5 [9386]: [9446]: BDR election starting

```

2013 May 16 10:50:58.092647 ospf 5 [9386]: [9446]: DR/BDR Status of this router
changed, new election run
2013 May 16 10:50:58.092643 ospf 5 [9386]: [9446]: This nbr 5.5.5.2 promoted
to current dr
2013 May 16 10:50:58.092639 ospf 5 [9386]: [9446]: Walking neighbor 5.5.5.2
(0x93e3524), state TWOWAY
2013 May 16 10:50:58.092635 ospf 5 [9386]: [9446]: Neighbor not declared DR,
ignoring
2013 May 16 10:50:58.092632 ospf 5 [9386]: [9446]: Walking neighbor 5.5.5.1
(0xa1cf7514), state SELF
2013 May 16 10:50:58.092628 ospf 5 [9386]: [9446]: DR election starting
2013 May 16 10:50:58.092625 ospf 5 [9386]: [9446]: This neighbor is greater
than 2way
2013 May 16 10:50:58.092622 ospf 5 [9386]: [9446]: Walking neighbor 5.5.5.2
(0x93e3524), state TWOWAY
2013 May 16 10:50:58.092618 ospf 5 [9386]: [9446]: Compare done, new current
bdr 5.5.5.1
2013 May 16 10:50:58.092613 ospf 5 [9386]: [9446]: Current BDR set to this
neighbor
2013 May 16 10:50:58.092610 ospf 5 [9386]: [9446]: This neighbor is in
consideration for bdr
2013 May 16 10:50:58.092607 ospf 5 [9386]: [9446]: This neighbor is greater
than 2way
2013 May 16 10:50:58.092604 ospf 5 [9386]: [9446]: Walking neighbor 5.5.5.1
(0xa1cf7514), state SELF
2013 May 16 10:50:58.092597 ospf 5 [9386]: [9446]: BDR election starting
2013 May 16 10:50:58.092573 ospf 5 [9386]: [9446]: Current 0.0.0.0 as DR,
0.0.0.0 as BDR
2013 May 16 10:50:58.092567 ospf 5 [9386]: [9446]: Begin OSPF DR election on
Vlan5
2013 May 16 10:50:58.092432 ospf 5 [9386]: : Nbr 5.5.5.2: DOWN --> INIT,
event HELLORCVD

```

Troubleshooting

Si la investigación de la Conectividad L2/3, de la configuración, y del tráfico OSPF-permitido no destapó el problema y mostró al vecino en la lista, abra un caso del Centro de Asistencia Técnica de Cisco (TAC). Ingrese estos comandos y la fuente TAC con la información de ambas salidas de los vecinos:

- 'show run'
- Muestre el OSPF del tecnología-soporte

Vecino OSPF pegado en el estado de la inicialización (INIT)

Hay épocas cuando pegan al vecino en el estado de Init, que indica que el nexo 7000 considera los paquetes de saludo del vecino, pero no ve su router-ID en el paquete de saludo para moverse al estado siguiente de bidireccional.

```
router2#show ip ospf neighbor
```

```

Neighbor ID      Pri   State           Dead Time   Address          Interface
170.170.5.1      1     INIT/-          00:00:34   170.170.1.1     ethernet 1/1
router-2#

```

Este problema se pudo causar por:

- Uno lateral está bloqueando el paquete de saludo con el ACL.
- Un lado está traduciendo, con el Network Address Translation (NAT), los saludos OSPF.

- La capacidad de multidifusión de un lado está quebrada (L2).

Complete estos pasos para investigar el problema:

1. Verifique que el avión del control del Multicast IP funcione.

```
N7K4# ping multicast 224.0.0.5 interface Ethernet 1/1
```

Si usted no ve la dirección IP de los vecinos en los resultados del ping, después hay un problema. Compruébelo los ambos lados.

2. Verifique que los paquetes de saludo estén recibidos del vecino.

```
N7K4# show ip ospf 5 event-history adjacency
```

Nota: Los paquetes de saludo salientes no se visualizan.

3. Habilite el debug de la adyacencia OSPF en cada interfaz, y verifiquelo que los paquetes de saludo están enviados.

```
N7K4# debug logfile debug-ospf size 10000
```

```
N7K4# debug-filter ip ospf 5 interface Ethernet 1/1
```

```
N7K4# debug ip ospf 5 adjacency detail
```

Nota: No olvide inhabilitar los debugs.

```
N7K4# undebug all
```

```
N7K4# no debug-filter all
```

```
N7K4# clear debug logfile debug-ospf
```

4. Verifique que los paquetes sean enviados por el OSPF a 224.0.0.5.

```
N7K4# debug logfile ospf_vj
```

```
N7K4# debug-filter ip mpacket interface e1/5
```

```
N7K4# debug-filter ip mpacket direction outbound
```

```
N7K4# debug-filter ip mpacket dest 224.0.0.5
```

```
N7K4# debug ip ospf 5 hello
```

```
N7K4# show debug logfile ospf_vj
```

```
N7K1-RP# show debug logfile ospf_vj
```

```
2013 May 16 11:18:55.202270 ospf: 5 [9386] (default) LAN hello in, ivl 10/40,
options 0x02, mask /24, prio 1, dr 5.5.5.1, bdr 5.5.5
.2 on Vlan5 from 5.5.5.2
```

```
2013 May 16 11:19:00.527640 ospf: 5 [9386] (default) LAN hello out, ivl 10/40,
options 0x02, mask /24, prio 1, dr 5.5.5.1, bdr 5.5.
5.2 nbrs 1 on Vlan5 (area 0.0.0.0)
```

```
2013 May 16 11:19:03.500785 ospf: 5 [9386] (default) LAN hello in, ivl 10/40,
options 0x02, mask /24, prio 1, dr 5.5.5.1, bdr 5.5.5
.2 on Vlan5 from 5.5.5.2
```

```
2013 May 16 11:19:09.515150 ospf: 5 [9386] (default) LAN hello out, ivl 10/40,
options 0x02, mask /24, prio 1, dr 5.5.5.1, bdr 5.5.
5.2 nbrs 1 on Vlan5 (area 0.0.0.0)
```

```
2013 May 16 11:19:10.406800 ospf: 5 [9386] (default) LAN hello in, ivl 10/40,
options 0x02, mask /24, prio 1, dr 0.0.0.0, bdr 0.0.0
.0 on Vlan5 from 5.5.5.2
```

```
2013 May 16 11:19:10.417602 ospf: 5 [9386] (default) LAN hello in, ivl 10/40,
options 0x02, mask /24, prio 1, dr 0.0.0.0, bdr 0.0.0
.0 on Vlan5 from 5.5.5.2
```

```
N7K1-RP#
```

Nota: No olvide inhabilitar los debugs.

```
N7K4# clear debug logfile ospf_vj
```

```
N7K4# undebug all
```

```
N7K4# no debug-fil all
```

```
no debug-filter ip mpacket interface Ethernet1/5
```

```
no debug-filter ip mpacket direction outbound
```

```
no debug-filter ip mpacket dest 224.0.0.5
```

N7K4#

5. Verifique que el paquete esté en el ethalyzer.

```
N7K4# ethalyzer local interface inband capture-filter "ip proto \ospf"
```

```
N7K1# ethalyzer local interface inband capture-filter "ip proto \ospf"
```

```
Capturing on inband
```

```
50 packets captured
```

```
2013-05-16 11:06:34.387196      5.5.5.2 -> 224.0.0.5      OSPF Hello Packet
2013-05-16 11:06:34.397553      5.5.5.2 -> 224.0.0.5      OSPF Hello Packet
2013-05-16 11:06:38.895343      5.5.5.1 -> 224.0.0.5      OSPF Hello Packet
```

Nota: Debe haber uno en el contexto predeterminado del dispositivo virtual (VDC).

Troubleshooting

Si la investigación de la Conectividad L2/3, de la configuración, y del tráfico OSPF-permitido no destapa el problema y muestra que el vecino hace FULL, abra un caso TAC. Ingrese estos comandos y la fuente TAC con la información de ambas salidas de los vecinos:

- 'show run'
- Muestre el OSPF del tecnología-soporte

Vecino OSPF pegado en un estado de dos vías

Hay las épocas en que pegan al vecino OSPF en un estado de dos vías. Este escenario es normal en los tipos de red de broadcast, y reduce el periodo de inundación en el alambre. Este escenario también ocurre si configuran a todo el Routers con una prioridad igual a cero.

Nota: Solamente los routers de menor capacidad deben ser configurados con una prioridad de cero así que no participan en la elección del router designado (DR).

¿Refiera a [porqué hace el comando show ip ospf neighbor revelan a los vecinos pegados en el estado de dos vías?](#) artículo para más información.

Vecino OSPF pegado en el exstart/intercambio

Hay las épocas en que pegan al vecino OSPF en el estado Exstart/Exchange.

Si hay discordancia MTU entre las interfaces de los vecinos OSPF o no pueden hacerse ping con el tamaño de paquetes del MTU configurado debido a los medios de transmisión, después pegan al vecino OSPF en el estado Exstart/Exchange.

```
router-6# show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface	
170.170.11.7	1	EXCHANGE/	-	00:00:36	170.170.11.7	Serial2.7

```
router-6#router-7# show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface	
170.170.11.6	1	EXSTART/	-	00:00:33	170.170.11.6	Serial0.6

```
router-7#
```

Este problema se pudo causar por:

- Discordancia MTU - INTERCAMBIO en un router y EXSTART en el otro
Nota: Usted puede configurar con el **comando ip ospf mtu-ignore**.
- El ID de router de vecino (LIBRADO) es lo mismo que su vecino - EXSTART
- El unicast está quebrado - INTERCAMBIO Problema de MTU - no puede hacer ping a través con más que un paquete de la cierto-longitudACL que bloquea el unicast - después del OSPF bidireccional envía el paquete de unidifusión excepto los links de punto a punto (P2P)El NAT traduce el paquete de unidifusión

Complete estos pasos para investigar el problema:

1. Marque el ping normal con un bit del don't fragment (DF) en y con el IP MTU del máximo en la interfaz.

```
N7K4# ping 10.10.12.2 df-bit packet-size 1472
```

Nota: En el Cisco IOS, cuando usted ingresa el **comando del <size> del tamaño del ping x.x.x.x**, el tamaño refiere al tamaño del paquete IP. En Linux, usted especifica el payload del Internet Control Message Protocol (ICMP) del ping bastante que especificando el tamaño del paquete IP como en el Cisco IOS. Pues usted puede ser que sepa ya, el nexos se emplea Linux. El MTU todavía se fija a 1,500 bytes - 20 de esos bytes son el encabezado IP y otros 8 son el encabezado ICMP. $1,500 - 20 - 8 = 1,472$ bytes de carga útil. Por lo tanto, el tamaño real del IP datagram es lo mismo que en el Cisco IOS, que es 1,500 bytes.

2. Marque si los paquetes se caen en las interfaces entrantes con el **comando 1/1 de los Ethernetes de la demostración internacional**.

```
N7K1-RP# show int ethernet 1/20 | section RX|TX
```

RX

```
340213 unicast packets 368092 multicast packets 2 broadcast packets
708307 input packets 233094927 bytes
0 jumbo packets 0 storm suppression packets
0 runs 0 giants 0 CRC/FCS 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

```
1374131 unicast packets 324752 multicast packets 3 broadcast packets
1698886 output packets 196282264 bytes
0 jumbo packets
0 output error 0 collision 0 deferred 0 late collision
0 lost carrier 0 no carrier 0 babble 0 output discard
0 Tx pause
```

```
N7K1-RP#
```

3. Marque si CoPP cae el paquete OSPF con el comando de la **clase test1-copp-class-critical de la controle de plano del show policy-map interface**.

```
Control Plane
```

```
service-policy input test1-copp-policy-lenient
```

```
class-map test1-copp-class-critical (match-any)
  match access-group name test1-copp-acl-bgp
  match access-group name test1-copp-acl-pim
  match access-group name test1-copp-acl-rip
  match access-group name test1-copp-acl-vpc
  match access-group name test1-copp-acl-bgp6
  match access-group name test1-copp-acl-igmp
  match access-group name test1-copp-acl-lisp
  match access-group name test1-copp-acl-msdp
  match access-group name test1-copp-acl-ospf
```

```

match access-group name test1-copp-acl-pim6
match access-group name test1-copp-acl-rip6
match access-group name test1-copp-acl-rise
match access-group name test1-copp-acl-eigrp
match access-group name test1-copp-acl-lisp6
match access-group name test1-copp-acl-ospf6
match access-group name test1-copp-acl-rise6
match access-group name test1-copp-acl-eigrp6
match access-group name test1-copp-acl-otv-as
match access-group name test1-copp-acl-mac-l2pt
match access-group name test1-copp-acl-mps-ldp
match access-group name test1-copp-acl-mps-oam
match access-group name test1-copp-acl-mps-rsvp
match access-group name test1-copp-acl-mac-l3-isis
match access-group name test1-copp-acl-mac-otv-isis
match access-group name test1-copp-acl-mac-fabricpath-isis
match protocol mpls router-alert
match protocol mpls exp 6
set cos 7
police cir 39600 kbps bc 375 ms
  conform action: transmit
  violate action: drop
module 1:
  conformed 539964945 bytes,
    5-min offered rate 5093 bytes/sec
    peak rate 5213 bytes/sec
  violated 0 bytes,
    5-min violate rate 0 bytes/sec
module 2:
  conformed 784228080 bytes,
    5-min offered rate 5848 bytes/sec
    peak rate 7692 bytes/sec
  violated 0 bytes,
    5-min violate rate 0 bytes/sec
module 3:
  conformed 5114206 bytes,
    5-min offered rate 41 bytes/sec
    peak rate 6656 bytes/sec
  violated 0 bytes,
    5-min violate rate 0 bytes/sec

```

N7K1#

4. Marque el intercambio del descriptor de la base de datos OSPF (DBD) con el comando de la **adyacencia del historial de eventos OSPF 5 del IP de la demostración** o el comando de la **adyacencia OSPF 5 del IP del debug**.

```

N7K1-RP# debug logfile debug-ospf size 10000
N7K1-RP# debug-filter ip ospf 5 interface Vlan 5
N7K1-RP# debug ip ospf 5 adjacency detail

```

Aquí tiene un ejemplo:

```

N7K1-RP# show debug logfile debug-ospf
2013 May 20 05:36:23.414376 ospf: 5 [8325] (default)      Nbr 5.5.5.2 FSM start:
old state FULL, event HELLORCVD
2013 May 20 05:36:23.414424 ospf: 5 [8325] (default)      Nbr 5.5.5.2: FULL -->
FULL, event HELLORCVD
2013 May 20 05:36:23.414438 ospf: 5 [8325] (default)      Nbr 5.5.5.2 FSM start:
old state FULL, event TWOWAYRCVD
2013 May 20 05:36:23.414450 ospf: 5 [8325] (default)      Nbr 5.5.5.2: FULL -->
FULL, event TWOWAYRCVD
2013 May 20 05:36:28.832638 ospf: 5 [8325] (default)      Nbr 5.5.5.2 FSM start:
old state FULL, event HELLORCVD
2013 May 20 05:36:28.832674 ospf: 5 [8325] (default)      Nbr 5.5.5.2: FULL -->

```

FULL, event HELLORCVD
2013 May 20 05:36:28.832695 ospf: 5 [8325] (default) Nbr 5.5.5.2: transitioning
to OneWay - did not find ourselves
2013 May 20 05:36:28.832709 ospf: 5 [8325] (default) Nbr 5.5.5.2 FSM start:
old state FULL, event ONEWAYRCVD
2013 May 20 05:36:28.833073 ospf: 5 [8325] (default) Nbr 5.5.5.2 FSM state
changed from FULL to INIT, event ONEWAYRCVD
2013 May 20 05:36:28.833120 ospf: 5 [8325] Begin OSPF DR election on Vlan5
2013 May 20 05:36:28.833140 ospf: 5 [8325] Current 5.5.0.1 as DR, 5.5.0.2
as BDR
2013 May 20 05:36:28.833177 ospf: 5 [8325] BDR election starting
2013 May 20 05:36:28.833196 ospf: 5 [8325] Walking neighbor 5.5.5.1
(0xaec59188), state SELF
2013 May 20 05:36:28.833211 ospf: 5 [8325] This neighbor is greater
than 2way
2013 May 20 05:36:28.833235 ospf: 5 [8325] Walking neighbor 5.5.5.2
(0x9777584), state INIT
2013 May 20 05:36:28.833249 ospf: 5 [8325] DR election starting
2013 May 20 05:36:28.833265 ospf: 5 [8325] Walking neighbor 5.5.5.1
(0xaec59188), state SELF
2013 May 20 05:36:28.833281 ospf: 5 [8325] This nbr 5.5.5.1 promoted to
current dr
2013 May 20 05:36:28.833297 ospf: 5 [8325] Walking neighbor 5.5.5.2
(0x9777584), state INIT
2013 May 20 05:36:28.833404 ospf: 5 [8325] Elected 5.5.0.1 as DR,
0.0.0.0 as BDR
2013 May 20 05:36:28.833440 ospf: 5 [8325] Interface Vlan5 ---> DR
2013 May 20 05:36:28.833456 ospf: 5 [8325] (default) Nbr 5.5.5.2 FSM start:
old state INIT, event ADJOK
2013 May 20 05:36:28.833474 ospf: 5 [8325] (default) Nbr 5.5.5.2: INIT -->
INIT, event ADJOK
2013 May 20 05:36:28.833492 ospf: 5 [8325] (default) Nbr 5.5.5.2: FULL -->
INIT, event ONEWAYRCVD
2013 May 20 05:36:28.843309 ospf: 5 [8325] (default) Nbr 5.5.5.2 FSM start:
old state INIT, event HELLORCVD
2013 May 20 05:36:28.843339 ospf: 5 [8325] (default) Nbr 5.5.5.2: INIT -->
INIT, event HELLORCVD
2013 May 20 05:36:28.843357 ospf: 5 [8325] (default) Nbr 5.5.5.2: transitioning
to OneWay - did not find ourselves
2013 May 20 05:36:28.843370 ospf: 5 [8325] (default) Nbr 5.5.5.2 FSM start:
old state INIT, event ONEWAYRCVD
2013 May 20 05:36:28.843386 ospf: 5 [8325] (default) Nbr 5.5.5.2: INIT -->
INIT, event ONEWAYRCVD
2013 May 20 05:36:34.244541 ospf: 5 [8325] (default) Got DBD from 5.5.5.2
with 0 entries
2013 May 20 05:36:34.244567 ospf: 5 [8325] (default) seqnr 0x9247f5e,
dbdbits 0x7, mtu 1600, options 0x42
2013 May 20 05:36:34.244622 ospf: 5 [8325] (default) Nbr 5.5.5.2 FSM start:
old state INIT, event TWOWAYRCVD
2013 May 20 05:36:34.244798 ospf: 5 [8325] (default) Nbr 5.5.5.2 FSM state
changed from INIT to EXSTART, event ADJOK
2013 May 20 05:36:34.244859 ospf: 5 [8325] Begin OSPF DR election on Vlan5
2013 May 20 05:36:34.244880 ospf: 5 [8325] Current 5.5.0.1 as DR, 0.0.0.0
as BDR
2013 May 20 05:36:34.244916 ospf: 5 [8325] BDR election starting
2013 May 20 05:36:34.244935 ospf: 5 [8325] Walking neighbor 5.5.5.1
(0xaec59288), state SELF
2013 May 20 05:36:34.244949 ospf: 5 [8325] This neighbor is greater
than 2way
2013 May 20 05:36:34.244965 ospf: 5 [8325] Walking neighbor 5.5.5.2
(0x9777584), state EXSTART
2013 May 20 05:36:34.244978 ospf: 5 [8325] This neighbor is greater
than 2way
2013 May 20 05:36:34.244991 ospf: 5 [8325] This neighbor is in consideration

```

for bdr
2013 May 20 05:36:34.245004 ospf: 5 [8325] Current BDR set to this neighbor
2013 May 20 05:36:34.245019 ospf: 5 [8325] Compare done, new current
bdr 5.5.5.2
2013 May 20 05:36:34.245033 ospf: 5 [8325] DR election starting
2013 May 20 05:36:34.245049 ospf: 5 [8325] Walking neighbor 5.5.5.1
(Oxaec59288), state SELF
2013 May 20 05:36:34.245065 ospf: 5 [8325] This nbr 5.5.5.1 promoted to
current dr
2013 May 20 05:36:34.245080 ospf: 5 [8325] Walking neighbor 5.5.5.2
(0x9777584), state EXSTART
2013 May 20 05:36:34.245094 ospf: 5 [8325] Neighbor not declared DR,
ignoring
2013 May 20 05:36:34.245202 ospf: 5 [8325] Elected 5.5.0.1 as DR,
5.5.0.2 as BDR
2013 May 20 05:36:34.245247 ospf: 5 [8325] Interface Vlan5 ---> DR
2013 May 20 05:36:34.245262 ospf: 5 [8325] (default) Nbr 5.5.5.2 FSM start:
old state EXSTART, event ADJOK
2013 May 20 05:36:34.245299 ospf: 5 [8325] (default) Nbr 5.5.5.2:
EXSTART --> EXSTART, event ADJOK
2013 May 20 05:36:34.245318 ospf: 5 [8325] (default) Nbr 5.5.5.2:
INIT --> EXSTART, event TWOWAYRCVD
2013 May 20 05:36:34.245335 ospf: 5 [8325] (default) We are SLAVE,
5.5.5.2 is master
2013 May 20 05:36:34.245348 ospf: 5 [8325] (default) Nbr 5.5.5.2 FSM start:
old state EXSTART, event NEGDONE
2013 May 20 05:36:34.245366 ospf: 5 [8325] (default) Preparing DBD exchange
for nbr 5.5.5.2, 387/5
2013 May 20 05:36:34.245463 ospf: 5 [8325] (default) Nbr 5.5.5.2 FSM state
changed from EXSTART to EXCHANGE, event NEGDONE
2013 May 20 05:36:34.245483 ospf: 5 [8325] (default) Nbr 5.5.5.2: EXSTART -->
EXCHANGE, event NEGDONE
2013 May 20 05:36:34.245843 ospf: 5 [8325] (default) Got DBD from 5.5.5.2
with 0 entries
2013 May 20 05:36:34.245862 ospf: 5 [8325] (default) seqnr 0x9247f5e,
dbdbits 0x7, mtu 1600, options 0x42
2013 May 20 05:36:34.245997 ospf: 5 [8325] (default) Sending DBD to
5.5.5.2 on Vlan5
2013 May 20 05:36:34.246031 ospf: 5 [8325] (default) Add 5.5.0.2(0x1)5.5.0.2
(0x80000084) (0x2c26) (109) to DBD
2013 May 20 05:36:34.246062 ospf: 5 [8325] (default) Add 5.5.0.1(0x1)5.5.0.1
(0x8000007f) (0xa3c7) (5) (0) to DBD
2013 May 20 05:36:34.246078 ospf: 5 [8325] (default) Filled DBD to 5.5.5.2
with 2 entries
2013 May 20 05:36:34.246111 ospf: 5 [8325] (default) Sent DBD with 2 entries to
5.5.5.2 on Vlan5
2013 May 20 05:36:34.246128 ospf: 5 [8325] (default) mtu 1600, opts: 0x42,
ddbbits: 0x2, seq: 0x9247f5e
2013 May 20 05:36:34.258616 ospf: 5 [8325] (default) Recv LSR from Nbr 5.5.5.2
2013 May 20 05:36:34.258634 ospf: 5 [8325] (default) schedule flood
2013 May 20 05:36:34.258674 ospf: 5 [8325] (default) Answering LSR from 5.5.5.2
2013 May 20 05:36:34.258690 ospf: 5 [8325] (default) 1 requests in LSR (1 left)
2013 May 20 05:36:34.258707 ospf: 5 [8325] (default) Building reply LSU to 5.5.5.2
2013 May 20 05:36:34.258726 ospf: 5 [8325] (default) Found requested LSA
5.5.0.1(1)5.5.0.1 for 5.5.5.2
2013 May 20 05:36:34.258791 ospf: 5 [8325] (default) Added 5.5.0.1(0x1)
5.5.0.1 (0x8000007f) (0xa3c7) (5) (0)
2013 May 20 05:36:34.258872 ospf: 5 [8325] (default) Built reply LSU with 1 LSAs
for 5.5.5.2 96 bytes
2013 May 20 05:36:34.286591 ospf: 5 [8325] (default) Got DBD from 5.5.5.2
with 2 entries
2013 May 20 05:36:34.286615 ospf: 5 [8325] (default) seqnr 0x9247f5f,
dbdbits 0x3, mtu 1600, options 0x42
2013 May 20 05:36:34.286751 ospf: 5 [8325] (default) Got DBD from 5.5.5.2

```

with 2 entries
2013 May 20 05:36:34.286784 ospf: 5 [8325] (default) seqnr 0x9247f5f,
dbdbits 0x3, mtu 1600, options 0x42
2013 May 20 05:36:34.286804 ospf: 5 [8325] (default) Found 5.5.5.1(0x2)
5.5.0.1 (0x80000004) (0x46de) (111) in DBD
2013 May 20 05:36:34.286870 ospf: 5 [8325] (default) Added 5.5.5.1(0x2)
5.5.0.1 (0x80000004) (0x46de) (111)(DO) to request li
st
2013 May 20 05:36:34.286889 ospf: 5 [8325] (default) Found 5.5.0.2(0x1)
5.5.0.2 (0x800000085) (0x91d0) (5) in DBD
2013 May 20 05:36:34.286917 ospf: 5 [8325] (default) Added 5.5.0.2(0x1)
5.5.0.2 (0x800000084) (0x2c26) (109) to request list
2013 May 20 05:36:34.286932 ospf: 5 [8325] (default) Added 2 out of 2 LSAs
to request list
2013 May 20 05:36:34.287046 ospf: 5 [8325] (default) Sending DBD to
5.5.5.2 on Vlan5
2013 May 20 05:36:34.287066 ospf: 5 [8325] (default) Filled DBD to
5.5.5.2 with 0 entries
2013 May 20 05:36:34.287101 ospf: 5 [8325] (default) Sent DBD with 0 entries to
5.5.5.2 on Vlan5
2013 May 20 05:36:34.287121 ospf: 5 [8325] (default) mtu 1600, opts: 0x42,
dbdbits: 0, seq: 0x9247f5f
2013 May 20 05:36:34.291760 ospf: 5 [8325] (default) Got DBD from 5.5.5.2
with 0 entries
2013 May 20 05:36:34.291789 ospf: 5 [8325] (default) seqnr 0x9247f60,
dbdbits 0x1, mtu 1600, options 0x42
2013 May 20 05:36:34.291915 ospf: 5 [8325] (default) Got DBD from 5.5.5.2
with 0 entries
2013 May 20 05:36:34.291934 ospf: 5 [8325] (default) seqnr 0x9247f60,
dbdbits 0x1, mtu 1600, options 0x42
2013 May 20 05:36:34.291953 ospf: 5 [8325] (default) Nbr 5.5.5.2 FSM start:
old state EXCHANGE, event EXCHDONE
2013 May 20 05:36:34.292101 ospf: 5 [8325] (default) Nbr 5.5.5.2 FSM state
changed from EXCHANGE to LOADING, event EXCHDONE
2013 May 20 05:36:34.292124 ospf: 5 [8325] (default) Nbr 5.5.5.2: EXCHANGE -->
LOADING, event EXCHDONE
2013 May 20 05:36:34.293200 ospf: 5 [8325] (default) Building LS Request packet
to 5.5.5.2
2013 May 20 05:36:34.293231 ospf: 5 [8325] (default) Add 5.5.0.2(0x1)
5.5.0.2 (0x800000084) (0x2c26) (110) to LSR
2013 May 20 05:36:34.293262 ospf: 5 [8325] (default) Add 5.5.5.1(0x2)
5.5.0.1 (0x80000004) (0x46de) (111)(DO) to LSR
2013 May 20 05:36:34.293281 ospf: 5 [8325] (default) Built LS Request packet for
5.5.5.2 with 2 entries
2013 May 20 05:36:34.297954 ospf: 5 [8325] (default) Nbr 5.5.5.2 FSM start:
old state LOADING, event LDDONE
2013 May 20 05:36:34.298069 ospf: 5 [8325] (default) Nbr 5.5.5.2 FSM state
changed from LOADING to FULL, event LDDONE
2013 May 20 05:36:34.298206 ospf: 5 [8325] (default) Nbr 5.5.5.2: LOADING -->
FULL, event LDDONE
2013 May 20 05:36:34.299179 ospf: 5 [8325] (default) Sending DBD to 5.5.5.2
on Vlan5
2013 May 20 05:36:34.299199 ospf: 5 [8325] (default) Filled DBD to 5.5.5.2
with 0 entries
2013 May 20 05:36:34.299233 ospf: 5 [8325] (default) Sent DBD with 0 entries to
5.5.5.2 on Vlan5
2013 May 20 05:36:34.299253 ospf: 5 [8325] (default) mtu 1600, opts: 0x42,
dbdbits: 0, seq: 0x9247f60
2013 May 20 05:36:38.746942 ospf: 5 [8325] (default) Nbr 5.5.5.2 FSM start:
old state FULL, event HELLORCVD
2013 May 20 05:36:38.747010 ospf: 5 [8325] (default) Nbr 5.5.5.2: FULL -->
FULL, event HELLORCVD
2013 May 20 05:36:38.747024 ospf: 5 [8325] (default) Nbr 5.5.5.2 FSM start:
old state FULL, event TWOWAYRCVD

```

2013 May 20 05:36:38.747046 ospf: 5 [8325] (default)      Nbr 5.5.5.2: FULL -->
FULL, event TWOWAYRCVD
2013 May 20 05:36:38.747073 ospf: 5 [8325] (default)      Different BDR in hello,
invoking nbrchange
2013 May 20 05:36:38.747090 ospf: 5 [8325] (default)      Neighbor
priority/options/DR/BDR value changed
2013 May 20 05:36:38.747265 ospf: 5 [8325]      Begin OSPF DR election on Vlan5
2013 May 20 05:36:38.747288 ospf: 5 [8325]      Current 5.5.0.1 as DR,
5.5.0.2 as BDR
2013 May 20 05:36:38.747329 ospf: 5 [8325]      BDR election starting
2013 May 20 05:36:38.747348 ospf: 5 [8325]      Walking neighbor 5.5.5.1
(0xaec59478), state SELF
2013 May 20 05:36:38.747362 ospf: 5 [8325]      This neighbor is greater
than 2way
2013 May 20 05:36:38.747648 ospf: 5 [8325]      Walking neighbor 5.5.5.2
(0x9777584), state FULL
2013 May 20 05:36:38.747662 ospf: 5 [8325]      This neighbor is greater
than 2way
2013 May 20 05:36:38.747676 ospf: 5 [8325]      This neighbor is in consideration
for bdr
2013 May 20 05:36:38.747689 ospf: 5 [8325]      Current BDR set to this neighbor
2013 May 20 05:36:38.747705 ospf: 5 [8325]      Compare done, new current bdr
5.5.5.2
2013 May 20 05:36:38.747733 ospf: 5 [8325]      DR election starting
2013 May 20 05:36:38.747750 ospf: 5 [8325]      Walking neighbor 5.5.5.1
(0xaec59478), state SELF
2013 May 20 05:36:38.747766 ospf: 5 [8325]      This nbr 5.5.5.1 promoted to
current dr
2013 May 20 05:36:38.747782 ospf: 5 [8325]      Walking neighbor 5.5.5.2
(0x9777584), state FULL
2013 May 20 05:36:38.747796 ospf: 5 [8325]      Neighbor not declared DR,
ignoring
2013 May 20 05:36:38.747948 ospf: 5 [8325]      Elected 5.5.0.1 as DR,
5.5.0.2 as BDR
2013 May 20 05:36:38.748004 ospf: 5 [8325]      Interface Vlan5 ---> DR

```

Nota: No olvide inhabilitar los debugs.

```
N7K1-RP# clear debug logfile debug-ospf
```

```
N7K1-RP# undebug all
```

```
N7K1-RP# no debug-fil all
```

Extremidades de la investigación:

Look for unió mal los mensajes MTU. Siga el número de secuencia y busque una retransmisión debido al descenso del DBD. Marque para saber si hay recibo de un número de secuencia inesperado del DBD.

Troubleshooting

Si la investigación de la Conectividad L2/3, de la configuración, y del tráfico OSPF-permitido no destapó el problema y mostró que el vecino está en línea, abra un caso TAC. Ingrese estos comandos y la fuente TAC con la información de ambas salidas de los vecinos:

- 'show run'
- Muestre el OSPF del tecnología-soporte

Vecino OSPF pegado en un estado de carga

Hay las épocas en que pegan al vecino OSPF en un estado de carga.

Este problema se pudo causar por:

- Se hace una petición del estado del link (LS), y el vecino envía un mún paquete o la corrupción de la memoria existe. Ingrese comando **IP OSPF de la demostración el mún** para ver el mún anuncio del estado del link (LSA).El comando **show log** muestra el mensaje **OSPF-4-BADLSATYPE**.Se hace una petición LS, y el vecino ignora la petición.

Si un router recibe un LSA anticuado, corrompido, o perdido, el OSPF permanece en un estado de carga, y genera el mensaje de error **OSPF-4-BADLSA**.

Ingrese este comando para investigar:

```
7K1-RP# show ip ospf traffic vlan 5
OSPF Process ID 5 VRF default, Packet Counters (cleared 1d12h ago)
Interface Vlan5, Area 0.0.0.0
Total: 15214 in, 15214 out
LSU transmissions: first 88, rxmit 9(13), for req 0, nbr xmit 289801235
Flooding packets output throttled (IP/tokens): 0 (0/0)
Ignored LSAs: 0, LSAs dropped during SPF: 0
LSAs dropped during graceful restart: 0
Errors: drops in      0, drops out      0, errors in      0,
        errors out    0, hellos in      0, dbds in        0,
        lsreq in      0, lsu in        0, lsacks in      0,
        unknown in    0, unknown out  0, no ospf       0,
        bad version   0, bad crc      0, dup rid       0,
        dup src       0, invalid src  0, invalid dst   0,
        no nbr        0, passive      0, wrong area    0,
        pkt length    0, nbr changed  0,
        bad auth      0,
In:      hellos      dbds      lsreqs     lsus      acks
Out:     14957      42        13        108       94
        14957      46        14        110       87
```

N7K1-RP#

Troubleshooting

Si la investigación de la Conectividad L2/3, de la configuración, y del tráfico OSPF-permitido no destapó el problema y mostró que el vecino está en línea, abra un caso TAC. Ingrese estos comandos y la fuente TAC con la información de ambas salidas de los vecinos:

- Muestre el fichero de diario del registro
- Muestre el OSPF del tecnología-soporte