APPLICATION NOTE

CISCO NONSTOP FORWARDING FOR ENHANCED INTERIOR GATEWAY ROUTING PROTOCOL

OVERVIEW

Cisco Nonstop Forwarding (NSF) with Stateful Switchover (SSO) is a Cisco innovation for routers with dual route processors. Cisco NSF with SSO allows a router, which has experienced a hardware or software failure of an active route processor, to maintain data link layer connections and continue forwarding packets during the switchover to the Standby route processor. This forwarding can continue despite the loss of routing protocol peering arrangements with other routers. Routing information is recovered dynamically, in the background, while packet forwarding proceeds uninterrupted.

Service Provider environments can benefit from the initial release of Cisco NSF with SSO in Cisco IOS® Software Release 12.0(22)S. It includes support for the Border Gateway Protocol (BGP), Open Shortest Path First (OSPF) and Intermediate System-Intermediate System (IS-IS) routing protocols. However, as Cisco NSF with SSO is ported to more traditional Enterprise platforms, such as the Cisco Catalyst 6500 Series Switch, there is a corresponding requirement to support the Enhanced Interior Gateway Routing Protocol (EIGRP).

There are two components of Cisco NSF for EIGRP:

- *NSF-capability:* re-startable EIGRP component that is run by a router and supports dual route processors. Cisco NSF for EIGRP is available in Cisco IOS Software Release 12.2(18)S for the Cisco 7500 Series Router and in Release12.2(18)SXD for the Cisco Catalyst® 6500 Series Switch. This functionality may also be available in other platform-specific software releases for other Cisco dual route processor devices.
- *NSF-awareness:* compatible components that run on the neighbors of the restarting router and help the restarting router reacquire its routing information. Available in Cisco IOS Software Release 12.2(15)T.

TECHNICAL DETAILS

Internally, a Cisco router equipped with dual route processors can maintain Layer 2 data link connections and up-to-date “next-hop” information to continue forwarding packets in the event of a route processor switchover.

However, all of these innovations would be for naught if routers that are neighbors with the router performing the switchover (hereafter, the ***neighbor routers***) did not continue to forward packets to that router. In order for a neighbor router to continue packet forwarding, several conditions must be met:

- Restarting routers and their neighbor routers must each support the appropriate EIGRP extensions.
- Neighbor routers must not prematurely declare the restarting router as unavailable.
- Neighbor routers must not communicate any state change in the restarting router to any of its own neighbors. This avoids the significant detrimental effect on network performance associated with the failure of a router.
- Restarting router must signal its neighbors that it has restarted.
- Neighbor routers must send EIGRP topology updates to help the restarting NSF router reacquire its EIGRP Topology Database.
- Neighbor routers must signal the completion of the initial routing update by sending the End-of-Table marker.
- In the interim (before the restarting router has reacquired the routing information), the neighbor routers must mark any routes associated with the restarting router as “stale”, but *continue to use those routes for packet forwarding*.

To accomplish these conditions, some enhancements were made to EIGRP. A new bit—the Restart bit—was introduced in the EIGRP UPDATE and HELLO packets. In addition, an End-of-Table (EOT) signal was introduced, so that neighbor routers could tell a restarting router when it had completed sending its updates. The EOT allows the restarting router to begin topology and route calculation as quickly as possible, and speeds convergence. Figure 1 illustrates the process that occurs when an EIGRP/NSF capable router restarts.

Figure 1. Restart of an EIGRP/NSF Capable Router

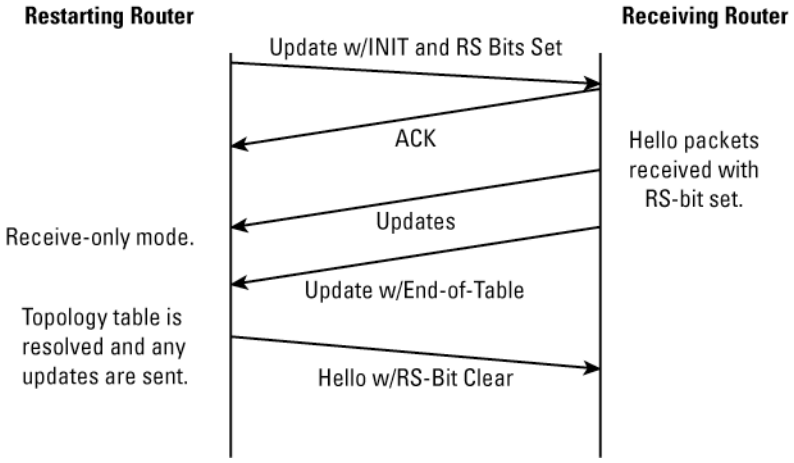

As quickly as possible after switching over to its redundant route processor, the restarting router will send out an empty update packet with both the INIT and RESTART (RS) bits set. This notifies neighbors of the restarting router that a restart has occurred, and that their assistance will be required to refresh the routing database of the restarting router.

During this process, the restarting router may also send HELLO packets, in order to maintain neighbor adjacency. These HELLO packets will also have the RS bit set.

Upon receiving the INIT and RS, the neighbors of the restarting router will acknowledge the update while realizing that the restarting router has restarted. As such, the restarting router will have no routing information, and will need to reacquire it. Each neighbor of the restarting router will begin sending updates containing routing information. In its last update, each neighbor will set an EOT signal so the restarting router knows that it has all available information and may begin the process of calculating routes.

At this point, the restarting router exits “receive-only mode” and performs a Diffusing Update Algorithm (DUAL) calculation to select the best loop-free routes for each destination in the topology database. Once the DUAL calculation is complete, the restarting router will send updates to each of its neighbors regarding routing destinations accessible through it. Once all updates have been sent and acknowledged, convergence is complete.

Once convergence is complete, the restarting router will clear the RS bit in its HELLO packets, and network operation will continue as normal.

Note that the restarting router and all of its neighbors have continued to forward packets throughout this entire operation.

DEPLOYMENT CONSIDERATIONS

The primary deployment scenario for EIGRP/NSF is in single-point-of-failure routers. A good example of a single-point-of-failure router is an Enterprise WAN edge router that has no redundant multi-homed path to particular destinations.

EIGRP/NSF can also be deployed in multi-homed environments or on the Enterprise Distribution or Core layers. Note one important caveat for this type of deployment: Cisco NSF maintains the capability of forwarding packets by “freezing” the Cisco Express Forwarding table at the point where the RP switchover occurs. Forwarding occurs to the last-known-good next-hop for any particular IP destination. Therefore, if a routing topology

change occurs prior to the restarting router reacquiring the most recent routing information, then transient routing loops, asymmetrical routing, or routing “black holes” may occur.

The following points mitigate this caveat:

- NSF must be compared to its alternative: the complete reset of a router, which can also result in routing loops, asymmetrical routing, or routing “black holes”.
- In most network designs, asymmetrical routing is the most common occurrence. While this is usually undesirable, it does not prohibit packets from reaching their ultimate destination.
- Depending on what topology change occurs, only a small portion of the traffic may be subject to routing loops or black holes. Other traffic continues to flow to its appropriate destinations.
- NSF is a self-correcting protocol. Routing loops or black holes disappear after convergence. In addition, EIGRP converges very quickly, so any routing problems will be shortly resolved.

FREQUENTLY ASKED QUESTIONS

Do all EIGRP Neighbors Need to be NSF-Aware?

No, all neighbors of an EIGRP NSF-capable router do not need to be NSF-aware. If a non-NSF-aware router exists, it ignores the RS bit set in the HELLO and UPDATE packets, and resets the adjacency with the restarting router. It treats the restarting router as if it had rebooted.

Mixing NSF-Aware and non-NSF aware neighbors has both advantages and limitations:

- It allows for the gradual migration to an NSF network—perhaps on a segment-by-segment basis.
- At switchover, non-NSF aware routers will have a different view of the network than NSF-Aware routers. This increases the potential for non-desirable events (ie: routing loops).

Cisco recommends that users do not mix NSF-Aware and non-NSF aware neighbors, if possible. If NSF is deployed on a segment-by-segment basis, all routers on a particular segment should be NSF-aware.

Are There any Safeguards Integrated into These Protocol Changes?

EIGRP/NSF exposes three timers that allow the user to set upper bounds on the length of EIGRP/NSF convergence:

- If the SIGNAL timer expires, the restarting router will notify the Routing Information Base (RIB) that it is converged if it has not learned of any neighbors, or has not learned of any NSF-aware neighbor, or has received all the updates from the neighbors.
- The CONVERGE timer is set on the restarting router and determines the maximum interval over which the router will wait for End-of-Table from all neighbors.
- The ROUTE-HOLD timer is set on the neighbors of the restarting router. It determines how long a neighbor will wait for complete reconvergence with the restarting router. If this timer expires, the neighbor flushes any stale routes from the restarting router and begins “normal” EIGRP recovery procedures.

The combination of these three timers provides adequate safeguards against any unforeseen problems in the NSF convergence process.

Can NSF/EIGRP Still be Used, if EIGRP HELLO/HOLD Timers Have Been Lowered to Promote Fast Convergence?

Possibly. Some customers choose to lower the EIGRP HELLO/HOLD timers in order to speed the detection of a non-responsive routing neighbor. This helps convergence, because traffic is directed to an alternate routing path more quickly.

Cisco will conduct the majority of its NSF/EIGRP testing with the EIGRP timers set at their default values. Although there may be a few test cases with lowered timer values, it may not be representative of all customer environments. Hence, if customers wish to use lowered timers with

NSF/EIGRP, it is strongly recommended that they conduct supplemental testing in their own labs, using a testbed that simulates the actual network conditions.

Can EIGRP/NSF be Used in Conjunction with the Hot Standby Router Protocol?

Yes, but the results may not be what you expect.

Because Hot Standby Router Protocol (HSRP) is not currently NSF-aware or NSF-capable, the first HSRP HELLO that is sent after an RP switchover occurs will have the STATE byte set to INITIAL. If the restarting router was the ACTIVE forwarder for the HSRP group, it will cause a failover to the STANDBY router in the group. If preemption is configured, and the restarting router has the highest priority in the group, there will be a subsequent failover back to the restarting router. Although your downtime in this scenario is no worse than a situation in which HSRP was deployed without NSF, customers may not anticipate these results.

Cisco is currently developing an NSF-capable version of HSRP.

CONCLUSION

Cisco Nonstop Forwarding for EIGRP can significantly reduce downtime on Enterprise networks during a failure on a route processor. Provided that careful attention is dedicated to the potential deployment scenarios, and that due diligence is accorded the potential caveats when incorporating NSF into the overall network design, users can achieve unprecedented levels of network availability.

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on **the Cisco Website at www.cisco.com/go/offices.**

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia • Cyprus Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland • Israel Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden • Switzerland • Taiwan Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

204170.g_ETMG_AE_2.05

Printed in the USA