

Implementing HSRP Over LANE

Document ID: 10446

Introduction

Prerequisites

- Requirements
- Components Used
- Conventions

Background Information

Case Studies

- 1) Native HSRP Over LANE
- 2) HSRP Over Routers Behind LANE
- 3) Mixed Environment

Conclusion

Related Information

Introduction

The purpose of this document is to outline the issues that may be encountered when implementing Hot Standby Router Protocol (HSRP) in a LAN Emulation (LANE) environment. It describes many of the specifics of HSRP over LANE and provides troubleshooting tips for various scenarios.

Prerequisites

Requirements

There are no specific requirements for this document.

Components Used

This document is not restricted to specific software and hardware versions.

Conventions

For more information on document conventions, see the Cisco Technical Tips Conventions.


Background Information

In summary, the purpose of HSRP is to allow hosts in a subnet to use a single "virtual" router as the default gateway multiple routers participate in the HSRP protocol in order to elect the active router, which assumes the role of default gateway and a backup router in case the active one fails. The result is that the default gateway will always appear to be up even if the physical first hop router changes. A complete description of HSRP can be found in RFC 2281 .


HSRP was designed for use over multi-access, multicast, or broadcast capable LANs (typically Ethernet, Token Ring, or Fiber Distributed Data Interface [FDDI]). Therefore, HSRP should work well over ATM LANE.

Several situations involving HSRP and LANE interaction may arise:


1. Since Cisco IOS® Software Release 11.2, HSRP can run "natively" over LANE . In this case, the **standby** commands are configured directly on the ATM subinterfaces where LAN Emulation Clients (LECs) reside. See the following illustration.



2. There also is an instance where HSRP is configured on LAN interfaces, but part of the subnet spans a LANE cloud. This is accomplished by the intermediate of a LAN switch with an ATM interface (such as a Cisco Catalyst 5000 with a LANE module). See the following illustration.



3. Finally, there is a "hybrid" situation where some HSRP routers are LANE-attached and others are on a LAN behind a LAN switch.



Case Studies

1) Native HSRP Over LANE

Routers participating in HSRP send "hello" packets over the broadcast medium in order to learn about each other and elect the active and standby routers. These packets are sent to multicast address 224.0.0.2 with a Time to Live (TTL) of 1 and a multicast destination MAC address of 0100 5E00 0002.

LANE introduces no new issues here so the details described in RFC 2281 still apply through the exchange of hello, coup, and resign packets, the active and standby routers are elected.

The hello packets are sent over the broadcast and unknown server (BUS) and the following is what a **debug atm packet** (on the Multicast Forward virtual circuit [VC]) and a **debug standby** would reveal:

```
Medina#show run
[snip]interface ATM3/0.1 multipoint
ip address 1.1.1.3 255.255.255.0
no ip redirects
no ip directed-broadcast
lane client ethernet HSRP
standby 1 ip 1.1.1.1
[snip]

Medina#show lane client
LE Client ATM3/0.1 ELAN name: HSRP Admin:
up State: operational
Client ID: 2
LEC up for 14 minutes 34 seconds
ELAN ID: 0
Join Attempt: 7
Last Fail Reason: Config VC being released
HW Address: 0050.a219.5c54 Type: ethernet
Max Frame Size: 1516
ATM Address: 47.00918100000000604799FD01.0050A2195C54.01
VCD rxFrames txFrames Type ATM Address
0 0 0 configure 47.00918100000000604799FD01.00604799FD05.00
12 1 3 direct 47.00918100000000604799FD01.00604799FD03.01
13 2 0 distribute 47.00918100000000604799FD01.00604799FD03.01
14 0 439 send 47.00918100000000604799FD01.00604799FD04.01
15 453 0 forward 47.00918100000000604799FD01.00604799FD04.01
```

```

Medina#show atm vc 15
ATM3/0.1: VCD: 15, VPI: 0, VCI: 40
UBR, PeakRate: 149760
LANE-LEC, etype:0xE, Flags: 0x16C7, VCmode: 0x0
OAM frequency: 0 second(s)
InARP DISABLED
Transmit priority 4
InPkts: 601, OutPkts: 0, InBytes: 48212, OutBytes: 0
InProc: 0, OutProc: 0, Broadcasts: 0
InFast: 0, OutFast: 0, InAS: 0, OutAS: 0
InPktDrops: 0, OutPktDrops: 0
CrcErrors: 0, SarTimeOuts: 0, OverSizedSDUs: 0
OAM cells received: 0
OAM cells sent: 0
Status: UP
TTL: 0
interface = ATM3/0.1, call remotely initiated,
call reference = 8388610
vcnum = 15, vpi = 0, vci = 46, state = Active(U10)
, multipoint call
Retry count: Current = 0
timer currently inactive, timer value = 00:00:00
Root Atm Nsap address: 47.00918100000000604799FD01.00604799FD04.01
, VC owner: ATM_OWNER_UNKNOWN

```

Of importance is looking at what the LAN Emulation Client (LEC) receives over the BUS (for example, by way of the multicast forward):

```

Medina#debug atm packet
interface atm 3/0.1 vcd 15
ATM packets debugging is on
Displaying packets on interface ATM3/0.2 VPI 0, VCI 46 only
Medina#debug standby
Hot standby protocol debugging is on
*Feb 18 06:36:05.443: SB1:ATM3/0.1 Hello in 1.1.1.2
Active pri 110 hel 3 hol 10 ip 1.1.1.1
*Feb 18 06:36:08.007: SB1:ATM3/0.1 Hello out 1.1.1.3
Standby pri 100 hel 3 hol 10 ip 1.1.1.1
*Feb 18 06:36:08.439: ATM3/0.1(I):
VCD:0xF VPI:0x0 VCI:0x40 Type:0xE, LANE, ETYPE:0x000E
LECID:0x0004 Length:0x4A
*Feb 18 06:36:08.439: 0004 0100 5E00 0002 0000 0C07
AC01 0800 45C0 0030 0000 0000 0111 D6F8 0101
*Feb 18 06:36:08.443: 0102 E000 0002 07C1 07C1 001C
AAEE 0000 1003 0A6E 0100 6369 7363 6F00 0000
*Feb 18 06:36:08.443: 0101 0101 0001 0001 000C

```

This hex-dump translates to the following:

```

VCD:0xF VPI:0x0 VCI:0x28: VCD number 15, VPI=0 and VCI=400
004: LECID from the sender of the packet
0100 5E00 0002: Destination MAC address for HSRP hellos
0000 0C07 AC01: Virtual MAC address of HSRP (the last octet is actually the standby group
0800: Type = IP
45C0 0030 0000 0000 0111 D6F8: IP header - UDP packet
0101 0102: Source IP = 1.1.1.2
E000 0002: Destination IP = 224.0.0.2
07C1 07C1 001C AAEE: UDP header - Source & Destination ports = 1985
00: HSRP version 0
00: Hello packet (type 0)
10: State (of the sender) is Active (16)
03: Hello time (3 sec)
0A: Hold time (10 sec)
6E: Priority = 110
01: Group

```

```
00: Reserved
6369 7363 6F00 0000: Authentication Data
0101 0101: Virtual IP address = 1.1.1.1
```

What is noteworthy is that the hello packets are sourced by the active router with the Virtual MAC address (VMAC) as source MAC address this is desirable because learning bridges (switches) that forward these packets will update their content-addressable memory (CAM) table with the appropriate location of the VMAC.

The key to HSRP lies within the mapping between an IP address and a MAC address.

In the simplest expression, the virtual IP address is permanently bound to a virtual MAC address and the only aspect to worry about is that the switches always know where this virtual MAC address is located. This is ensured because the hellos are sourced by the VMAC.

```
Medina#show standby
ATM3/0.1 - Group 1
  Local state is Standby, priority 100
  Hellotime 3 holdtime 10
  Next hello sent in 00:00:00.006
  Hot standby IP address is 1.1.1.1 configured
  Active router is 1.1.1.2 expires in 00:00:08
  Standby router is local
  Standby virtual mac address is 0000.0c07.ac01
```

Another option is that the routers use their burned-in (**standby use-bia**) addresses mapped to the virtual IP address. In this case, the mapping between virtual IP and MAC address changes over time the newly active router sends out a Address Resolution Protocol (ARP) in order to announce the new virtual IP-to-MAC address mapping. An ARP is simply an unsolicited ARP response.–

Note: Certain (older) IP stacks may not understand ARPs.

```
Medina#show standby
ATM3/0.1 - Group 1
  Local state is Standby, priority 100, use bia
  Hellotime 3 holdtime 10
  Next hello sent in 00:00:02.130
  Hot standby IP address is 1.1.1.1 configured
  Active router is 1.1.1.2 expires in 00:00:09
  Standby router is local
  Standby virtual mac address is 0050.a219.5c54
```

Note: To introduce LANE, the key is that on top of the Virtual IP-to-MAC address mapping, there must be accounting for the VMAC-to-Network-Service-Access-Point (NSAP) address mapping. This mapping is simply resolved through the LAN Emulation-Address Resolution Protocol (LE-ARP) process: a LEC wishing to send traffic to the active gateway will use LE-ARP for the VMAC (or physical MAC if using the burned-in MAC address [BIA]).

Now consider what happens when a new router becomes active: in order for the LECs to be informed of the active gateway's new location (new VMAC-to-NSAP mapping), the LE-ARP table must be modified. By default, LE-ARP entries time out every five minutes but, in most cases, relying on this timeout is unacceptable convergence must be faster. The solution depends on whether the LEC assuming the new Active status is running LANE version 1 or version 2 (see ATM Forum.com for the LANE specifications):

- **LANE version 1**

When a router becomes active, in addition to the steps described in RFC 2281 , it sends out an LE-NARP in order to make the new VMAC-to-NSAP address binding known. According to the

LANE specifications, upon reception of a LE-NARP, a LEC may choose to clear or update the LE-ARP entry corresponding to the MAC address. The tendency within Cisco is to adopt the more conservative approach and choose to clear the LE-ARP entry this will cause the LEC to immediately re-LE-ARP without having to wait for the five-minute timeout.

Note: This solution may cause the compatibility issue described below.

- **LANE version 2**

In LANE version 2, certain shortcomings of LANE version 1 were alleviated: the LE-NARP has been superseded by the targetless LE-ARP and the no-source LE-NARP. The targetless LE-ARP may be seen as a vehicle to advertise new bindings whereas the no-source LE-NARP's purpose is to render obsolete an existing MAC-to-NSAP address binding. The way this is implemented is that if a router changes from Standby to Active, it sends out a targetless LE-ARP (this is used to advertise a MAC-to-NSAP mapping) and if it changes from Active to Standby, it sends out a no-source LE-NARP (this is used to render a MAC-to-NSAP binding obsolete).

Problem – Interoperability

There is a problem that arises often enough to deserve a more in-depth examination. The LANE version 1 specifications state that the LE-NARP must specify the "old binding," which is being made obsolete by specifying the (old) Target NSAP (T-NSAP) address. Typically, routers participating in HSRP do not maintain data directly between each other.

Therefore, the newly active router does not know this information and it will choose not to complete this field since it does not know better. This is a mild violation of the specifications and some vendors will ignore these packets if the T-NSAP address field is all zeros. Unfortunately, there is no workaround for this if the LE-NARP is ignored, rely on the LE-ARP timeout (typically five minutes) before the correct binding is learned.

When an LE-ARP or LE-NARP is sent with a T-NSAP address field of all zeros, it is called "targetless." As seen above, with the advent of LANE version 2 (and Multiprotocol over ATM [MPOA]), this has become standard and the problem ceases to exist.

This is what is done in LANE version 1 where problems may arise:

- If the router knows the "old binding," it might as well obey the specifications. These debugs are now taken on the Control Distribute VC:

```
ATM0/0.1(I):
VCD:0xD Type:0x6, LANE, ETYPE:0x0006 LECID:0xFF00 Length:0x70
FF00 0101 0008 0000 0000 0018 0003 0000 0000 0000 0000 0000 0001 0000 0C07
AC01 4700 9181 0000 0000 101F 2D68 0100 50A2 195C 5401 0000 0000 4700 9181
0000 0000 101F 2D68 0100 102F FBA4 0101 0000 0000 0000 0000 0000 0000 0000
FF00: Marker = Control Frame
0101: ATM LANE version 10
0008: Op-code = LE_NARP_REQUEST
0000: Status
0000 0018: Transaction ID0003: Requester LECID0000: Flags
0000 0000 0000 0000: Source LAN destination
(not used for an LE-NARP)
0001 0000 0C07 AC01: Target LAN destination
(the 0001 indicates a MAC address as opposed to a route descriptor)
4700 9181 0000 0000 101F 2D68 0100 50A2 195C 5401: Source NSAP address
(new NSAP address to be bound)
0000 0000: Reserved
4700 9181 0000 0000 101F 2D68 0100 102F FBA4 0101: Target NSAP address
(old NSAP address to be rendered obsolete)
```

- If it does not know the "old binding," it does its best and at least advertises the new one:

```
ATM0/0.1(I):  
VCD:0xD Type:0x6, LANE, ETYPE:0x0006 LECID:0xFF00 Length:0x70  
FF00 0101 0008 0000 0000 0014 0003 0000 0000 0000 0000 0000 0001 0000 0C07  
AC01 4700 9181 0000 0000 101F 2D68 0100 50A2 195C 5401 0000 0000 0000 0000  
0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000 0000
```

Note: This time the T-NSAP address is blank.

Again, the behavior is completely within the specifications when using LANE version 2 clients.

Note: Software that supports MPOA also supports LANE version 2.

Troubleshooting Tips

Native HSRP over LANE should not engender too many problems other than the potential interoperability issue due to the LE-NARP devoid of the T-NSAP.

If the routers have difficulty in establishing whether they are Active or Standby, use the **debug standby** command to see if the hellos are seen on both sides. If not, then the BUS is probably not correctly forwarding the packets.

2) HSRP Over Routers Behind LANE

The situation becomes more complicated when HSRP is configured on LANE interfaces of routers located behind a LANE cloud, as illustrated in Figure 2.

Note: This figure logically depicts the fact that the router is non-ATM attached. It does not necessarily have to be in a device separate to the LAN switch (a Route Switch Module [RSM] in a Cisco Catalyst 5000 falls under this case).

Again, the difficulty arises due to the MAC-address-to-NSAP-address mapping imposed by LANE. As noted above, when the VMAC switches to a device (when a new router becomes active) that corresponds to another NSAP address, all devices attached to the LANE cloud must be informed. This is fairly easily implemented in a native HSRP over LANE environment by using the LE-NARP (or targetless LE-ARP).

The problem in this second case is that the LECs are not aware of any layer 3 information (IP), they are solely designed to bridge packets between two different mediums (the LAN and ATM).

For example, in Figure 2, if Router 2 suddenly became Active, then it would be desirable for LAN switch 2 to inform all the devices connected to the ATM (LANE) cloud about the new VMAC-to-NSAP mapping. The LEC in LAN switch 2 is said to be proxying for all the MAC addresses that are behind it. Devices across LANE wishing to send traffic to these MAC addresses must do so by way of a data-direct setup towards this LEC. Intuitively, one could think that this will not be a big problem since, as soon as Router 2 assumes the Active state, it will start sourcing hellos with the VMAC as the source MAC address. This information would then be learned by all the LAN switches and everything would converge rapidly. This is true in non-LANE environments, but LANE is special for the following reason:

In LANE, a data packet can usually be transmitted through two paths:

- The data-direct if this packet is a unicast for which the destination has been mapped to a known NSAP and if the data-direct has already been established.
- The BUS for unknown unicasts and multicasts.

Therefore, a same MAC address will source packets that will be received by a LAN switch over two different paths. Multicasts and unknown unicasts will arrive by way of the BUS whereas known unicasts arrive by way of data-directs. If no particular effort had been made, a LAN switch would keep learning this MAC address either over a data-direct or over the BUS depending on the last packet received. This is undesirable because the BUS should only be used to send packets for unknown unicasts or multicasts. At this stage, nothing is learned over the BUS, but in reality, choose to do the following:

Packets received over the BUS are marked with the Conditional Learn (CL) bit set to 1 (this bit is in a control overhead specific to Cisco LAN switches). The LAN switch will only update its CAM table with this entry if it does not already have an entry for the MAC address (in this VLAN). The idea is that if a switch receives a packet from a source that it does not know about, at least it will now know that it is located somewhere across the LANE cloud. Future packets for that MAC address will be forwarded to the BUS only as opposed to being flooded in the entire VLAN.

To return to the example, it is safe to assume that all the LECs in this ELAN are already aware of the VMAC-NSAP mapping for Router 1 prior to when Router 2 becomes Active. All LAN switches also know the VMAC is behind LAN switch 1. When Router 2 becomes Active and sources the hello packets, these are forwarded to the LANE cloud over the BUS. Therefore, none of the LAN switches will update their CAM tables with this new information and all packets sent to this VMAC will be misdirected until the LAN switches "forget" about this entry (the default aging being five minutes).

Note: Overall connectivity might actually be lost for up to 10 minutes since the LE-ARP aging-timer on the LECs also are five minutes by default. Reducing the aging-timer for MAC addresses will help, but does not actually resolve the problem.

There are two solutions for this:

1. If LAN switches are non-Cisco, revert to a method described above: using the burned-in address. If the routers only use their MAC address to source the hello packets and that the virtual-IP address changes mapping whenever a switch-over occurs, there is no confusion possible as to where these MAC addresses are located.
2. If LAN switches are Cisco Catalysts, then keep using the VMAC due to the modifications provided by the Distributed Defect Tracking System (DDTS) covered in Cisco bug IDs CSCdj58719 (registered customers only) and CSCdj60431 (registered customers only).

In essence, when a router assumes the Active state, in addition to the ARP (unsolicited ARP response) that it sends in accordance with RFC 2281, the router sends a second ARP with a destination MAC address of 0100.0CCD.CDCD. When a Cisco Catalyst receives this packet it does two things:

- ◆ It clears the LE-ARP entry it has for the VMAC.
- ◆ It learns the VMAC over the BUS.

Because of this, there are no more stale LE-ARP entries in the various LECs and the new location of the VMAC is propagated to all switches (for example, beyond the LANE cloud). In order for this to work correctly, the following minimum software requirements must be met:

- Routers must have at least Cisco IOS Software Release 11.1(24), version 11.2(13), or all of version 12.0.
- LANE modules must have at least version 3.2(8). 11.3W4 versions and later are acceptable.

Cisco recommends using the latest software.

3) Mixed Environment

There is one final issue that can arise in mixed environments. Taking the scenario above and adding a directly connected LANE end-device (router or workstation), the end-device needs to be informed about a change of location of the active gateway the same way as in scenario 1. If the newly active router is connected behind a switch, the only solution is for the switch itself to send out the LE-NARP on behalf of the router and this is exactly what to do.

In addition to the steps described above, if a Cisco Catalyst picks up a packet destined to 0100 0CCD CDCD, it sends out a LE-NARP (no-source LE-NARP if running LANE version 2), which its sole purpose is to clear the LE-ARP caches for the VMAC.

Conclusion

As demonstrated, HSRP over LANE works well in principle but, under certain circumstances, users can lose connectivity for short periods of time if falling into one of the loopholes described above.

Important!: In order to ensure success with HSRP over LANE, follow at least these two recommendations:

- To be safe, upgrade to at least the latest version of Cisco IOS Software Release 12.0.
- In multi-vendor environments, it is best to use LANE version 2 or the burned-in address in order to avoid problems.

Related Information

- [ATM Technology Support Pages](#)
- [Technical Support – Cisco Systems](#)

[Contacts & Feedback](#) | [Help](#) | [Site Map](#)

© 2008 – 2009 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)

Updated: Jun 05, 2005

Document ID: 10446
