

Troubleshooting Spanning Tree Over LANE

Document ID: 8305

Introduction

Prerequisites

Requirements

Components Used

Conventions

Introduction to STP over LANE

Topology Changes

Common Problems

Excessive LE-ARP Packets

BUS Throttling

Flush Mechanism

Conclusion

Related Information

Introduction

This document looks at the spanning tree issues you can encounter on Cisco LAN switches which are interconnected via LAN Emulation (LANE).

Prerequisites

Requirements

Cisco recommends that you have knowledge of this topic:

- Basic prior knowledge of spanning-tree is required in order to benefit from this document. Refer to Spanning Tree Protocol (STP) technology for more information.

Components Used

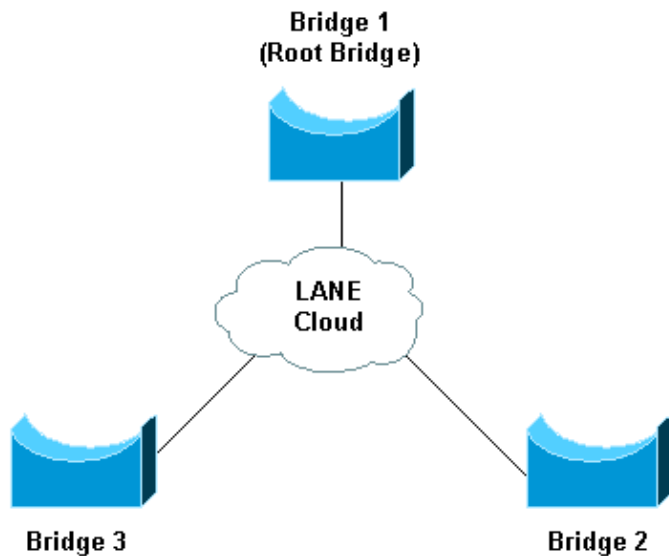
This document is not restricted to specific software and hardware versions.

Conventions

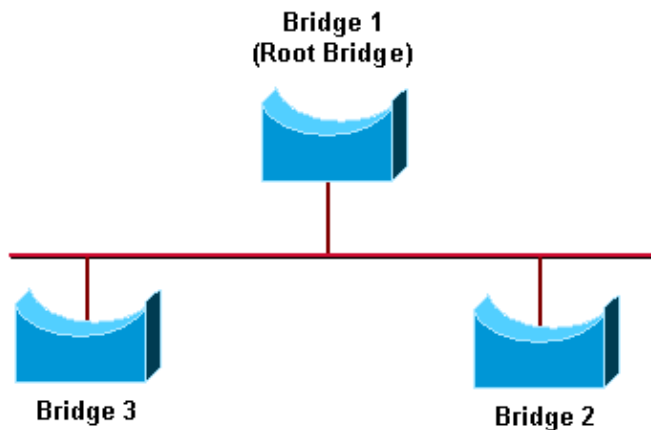
Refer to Cisco Technical Tips Conventions for more information on document conventions.

Introduction to STP over LANE

There is not a great deal in spanning tree technology that is specific to LANE. LANE is simply another method of trunking, which conveys several Virtual LANs (VLANs) over one physical medium. But unlike Ethernet trunking methods such as Inter-Switch Link (ISL) and dot1q, which are supported on point-to-point links, an Emulated LAN (ELAN) represents a shared medium. This diagram illustrates this point:



From a spanning tree point of view, this is equivalent to this diagram:



In a given ELAN, a bridge sends its Bridge Protocol Data Units (BPDUs) over the multicast send virtual circuit (VC) towards the Broadcast and Unknown Server (BUS). The BUS, in turn, then forwards this to all the LAN Emulation Clients (LECs) in the ELAN. Thus, all the others see each BPDU that a bridge sends in the ELAN, just as if they were all attached to the same Ethernet segment.

The consequence is that, provided these conditions are met, all the ATM ports forward for all ELANs:

- The bridges only have one LEC in each ELAN, only one LANE module.
- There are no other connections between these bridges such as Ethernet trunks.

Another peculiarity of the forwarding of BPDUs to the BUS is that the max age timer of a given LEC never expires without the expiration on all LECs in that ELAN. If the max age timer expires on all the LECs unexpectedly, it is possible that the BUS misbehaves. If a single LEC sees the problem, the BUS must be okay. Take care in order to see where the BPDU is dropped. This can either be dropped somewhere in the ATM cloud, by the LANE module or possibly by the bridge itself.

Topology Changes

As previously mentioned, STP over LANE is fairly straightforward. BPDUs are simply sent over the BUS.

There is one exception to this, though, as topology changes are handled in a special way. Refer to Understanding Spanning-Tree Protocol Topology Changes for more information.

In LANE, the LEC must maintain a MAC address to NSAP address table called the LAN-Emulation ARP (LE-ARP) table. This table indicates which network service access point (NSAP), and therefore which data-direct VC, if it is up, to use in order to reach a certain MAC address. Just as the bridge table, also referred to as the CAM table, must be aged out more aggressively when there are topology changes, so must the LE-ARP table. This is easily done by bridges who can possibly reduce the aging-time of their LE-ARP entries at the same time as that of the CAM entries when they receive a Topology Change Notification (TCN).

The problem is that all LECs maintain the LE-ARP table, even those that are not on bridges and therefore do not run spanning tree. For example, due to a topology change, a non-bridge LEC now has to reach a certain MAC address through a different bridge than before. If it is not informed via a TCN, the LEC has to wait for up to five minutes, default LE-ARP aging timer, in order to reverify its LE-ARP and learn the correct binding.

The ATM Forum considered this in its LANE specification and the problem is solved in this explanation:

1. For each BPDU a bridge sends to the BUS, it must also send an LE-topology-request to the LANE server (LES).
2. This, in turn, sends it to all the LECs. The LE-topology-request contains a topology change flag just like a BPDU.

The rule is simple:

- ◆ The value of this flag must always reflect that of the BPDU. In this way, even non-bridge LECs aggressively age out, 15 seconds forward delay, their LE-ARP entries in the event of a topology change.

Common Problems

While the implementation of spanning-tree over LANE is relatively simple, and the only real exception lies in the way topology changes are handled, there are some common issues to watch.

Excessive LE-ARP Packets

One of the most common problems that can occur in a large LANE environment is excessive LE-ARPing. In a large network, a typical LE-ARP table can have hundreds of entries, each of which must be reverified at least every five minutes. As a consequence, the LANE modules in the network spend a lot of CPU cycles that send and process LE-ARP packets. This also puts quite a burden on the LAN Emulation Server (LES) which must forward all these packets. Under normal circumstances, this LE-ARP activity can already represent quite a high load on the network.

If there are many topology changes that occur and the reverification timers drop to 15 seconds, the load can become excessive. The best way to see this is to use **show process cpu** and look for the LANE client process:

```
ATM#show process cpu
CPU utilization for five seconds: 51%/0% one minute: 31% five minutes: 24%
  PID  Runtime(ms)  Invoked  uSecs      5Sec   1Min   5Min  TTY Process
    1      124          424    292     0.00%  0.00%  0.00%  0 Load Meter
    2     1644          422   3895     0.08%  0.07%  0.08%  0 Subagent Reconne
    3    55436         2789  19876     0.00%  3.38%  2.82%  0 Check heaps
    4         0            1      0     0.00%  0.00%  0.00%  0 Pool Manager
    5         0            2      0     0.00%  0.00%  0.00%  0 Timers
[snip]
```

53	209532	50414	4156	33.24%	23.50%	17.38%	0 LANE Client
54	548	4424	123	0.00%	0.04%	0.02%	0 Cat5K ATM LED
55	0	2	0	0.00%	0.00%	0.00%	0 LECS Finder
56	19168	2575	7443	11.55%	5.06%	3.75%	1 Virtual Exec

As the load increases, it is possible that certain LE-ARPs are dropped and therefore the corresponding bindings are released. A symptom of this is to see many LE-ARP entries that point to the BUS:

```
ATM#show lane le-arp
Active le-arp entries: 1023      Hardware Addr   ATM Address      VCD   Interface
0000.cc00.0404  00.00000000000000000000000000000000.000000000000.00    11*  ATM0.1
0000.cc00.0303  00.00000000000000000000000000000000.000000000000.00    11*  ATM0.1
0000.cc00.0202  00.00000000000000000000000000000000.000000000000.00    11*  ATM0.1
0000.cc00.0101  00.00000000000000000000000000000000.000000000000.00    11*  ATM0.1
[snip]
```

Note: The asterisk after the VCD number indicates that the traffic is sent to the BUS. The VCD number 11 is therefore the multicast send VC in this example.

The Catalyst 2900XL is especially susceptible to this sort of problem because it is a software-based switch. The **show controller** on the Catalyst 2900XL gives a noteworthy counter. Look at the Threshold LANE ctrl drops counter in bold in this output:

```
ATM#show controllers
Catalyst 2900 XL ATM card Statistics
=====ATM to Ethernet Relay: Control & OAM Frames:
rx Completions                5076   rx not In VcControlTable      0
rx Buffer Starvation           0     Relay Frames:
rx Completions                185035 rx not In VcControlTable      0
rx BufferStarvation            0     EthTxCompletions              140657
Threshold LANE ctrl drops    6826 Small Frame (<64bytes)        0
[snip]
```

As soon as this counter increments on a Catalyst 2900XL ATM module, it means that LANE control packets, the LE-ARPs and flush packets, are dropped.

There are two ways that you can reduce the LE-ARP activity in a LANE network:

1. The first is to avoid topology changes as much as possible. This is why it is important to enable port-fast on all access ports. This is desirable in a switched Ethernet environment but it is crucial in a LANE network. Refer to Understanding How PortFast Works for further information.
2. The second is to enable a feature that only exists on the Catalyst 5000 and Catalyst 6000 LANE modules:

- ◆ Local LE-ARP reverification

Thanks to the presence of a CAM table that resides on the Supervisor module and points MAC address destinations to the correct data-direct VC, we can assume that:

- ◆ If a packet with a given source MAC address is recently received over the data-direct, within the last five minutes, its entry is still present in the CAM table.
- ◆ It is therefore sufficient for the LE-ARP reverification process to check whether or not a MAC address is still present in the CAM table and associated with the same VC as in the LE-ARP table. If this is not the case, only then is a true LE-ARP request sent.

BUS Throttling

When, for any reason such as excessive LE-ARPing, a MAC address is not resolved to an NSAP address,

traffic to this destination is referred to as an unknown unicast and must be sent over the BUS. The ATM Forum LANE specification recommends that this type of traffic be throttled in order to avoid the overload of the BUS. Throttling means that the LEC sends traffic to unknown unicasts at a rate of one to ten packets per second (pps). In truth, the Catalyst 5000 and Catalyst 6000 do not throttle traffic sent to the BUS but the Catalyst 2900XL and all routers do.

Normally, BUS throttling is not a problem. In a healthy network, a LE-ARP is responded to quickly and there is not any reason for clearing it, at least as long as you can avoid the excessive LE-ARPing problem. But, as previously mentioned, the Catalyst 2900XL can experience problems. In response to this, Cisco has added a feature that allows BUS throttling to be disabled on this platform. This is filed under Cisco bug ID CSCdv44257 (registered customers only) and is integrated in Cisco IOS® Software Release 12.1(12).

Flush Mechanism

When traffic must be sent to a destination which is not yet resolved through the LE-ARP process, or when the data-direct is not yet established, this traffic is sent as unknown unicast traffic over the BUS. When the data-direct is finally established, the traffic must switch from the BUS to this data-direct. There is a potential problem here for out-of-order packets. If a packet N is sent to the BUS (and possibly throttled) and packet N+1 immediately follows over the data-direct, there is a chance that packet N+1 arrives before packet N because the data-direct represents the shortest path between the two LECs.

This is problematic for two reasons:

- This normally never happens on a LAN, for example, switched Ethernet, and since LANE is supposed to emulate the LAN completely, this must be avoided.
- Certain higher layer protocols do not react well to out-of-order packets.

The LANE specification therefore implements a flush mechanism. The principle is simple:

1. Once a LEC sees that it wishes to switch from the BUS to a newly-established data-direct, it first sends out a flush request over the BUS.
2. This arrives at the destination LEC which responds with a flush response over the LES.
3. During this time, it is possible the source LEC does not send any traffic to the destination until it receives the flush response.
4. When it does, this is the indication that it can now send traffic over the data-direct.

The problem is that traffic received during this flush mechanism is dropped. Since there are some applications which are very sensitive to loss of packets and less sensitive to out-of-order packets, and since the likelihood of out-of-order packets is still rare, it is possible to disable the flush mechanism with the configuration of **no lane client flush** at the global configuration level. This is implemented on most platforms under Cisco bug ID CSCdr06796 (registered customers only) and on the Catalyst 2900XL under Cisco bug ID CSCdv44243 (registered customers only) .

Conclusion

There are several issues to watch out for when you use spanning-tree and related issues over LANE. It is easy to see that excessive LE-ARPing can cause certain LE-ARP entries to be purged. This, in turn, causes traffic sent to that destination to be sent to the BUS and potentially throttled. When the MAC address is relearned, the flush mechanism must activate again, which leads to possible packet loss. Finally, it is possible this cycle repeats itself indefinitely until packet loss becomes noticeable.

If certain applications that run over a LANE network cannot cope with any packet loss, then it is possibly a good idea to disable the flush mechanism. But when you do this, be aware of the risk of out-of-order packets.

The key thing to remember is to avoid excessive topology changes in LANE. This is due to the inherent nature of LANE and innovations like local LE-ARP verification can only slightly alleviate them. By far, the best solution is to enable **port-fast** wherever applicable.

Related Information

- [ATM Technical Support](#)
 - [More ATM Information](#)
 - [Technical Support & Documentation – Cisco Systems](#)
-

All contents are Copyright © 2006–2007 Cisco Systems, Inc. All rights reserved. Important Notices and Privacy Statement.

Updated: Aug 07, 2007

Document ID: 8305
