Tech Info Library

AppleTalk: Effects of Redundancy on Some Network Configurations

```
Revised:        11/11/92
Security:       Everyone
```

AppleTalk: Effects of Redundancy on Some Network Configurations

===

Article Created: 29 September 1992

TOPIC ---

We have a cisco router which provides redundant paths to several AppleTalk
networks. The internet routers now flip-flop between the two AGSs as their
next router. Is what we're doing a legal network configuration? If it is,
what causes the internet routers to swap between what they see as their
next router?

Given our number of AT devices, number of physical networks, and zones
(312), we're extremely sensitive to traffic issues. I would like to know
more about increased traffic due to redundancy and how split horizon may
help.

Also, from a desktop view, once a process is initiated from a Macintosh
computer to some service on a remote network, do all packets follow the
same path until the connection is broken? In other words, if I connect to
a server on a remote network, and the packet can be forwarded by either
CISCO 3 or CISCO 6, will all packets be forwarded through the same CISCO
for the duration of the process? Is there a difference if the connection
is to a server versus a printer or a Macintosh, and so on?

DISCUSSION --

Configuration

The network configuration you've created is perfectly legal and the
behavior you see is correct. RTMP assumes the latest route to a specific
network is the most recent information, and replaces the information
currently in the routing table, even if the route "cost" (hop count) is
equal.

DDP "Best Router" Functionality

"Best route" reduces the number of needless routers used when a DDP session

is in use. Its primary purpose is to eliminate the following situation:

```
Net 1
================================================================
        |                  |                      |
    [router A]         [router B]          [End node (Mac-A)]
        |                  |     |
        |                  |      ------to another network
Net 2   |                  ----------to yet another network
================================================================
            |
        [AFP Server]
```

Here's the order of events:

1) When end node (Mac-A) on network 1 initiates a connection to the AFP
   server on network 2, the initial packet might be sent to router B. This
   is because end nodes don't really keep any routing information. They
   rely on the routers to keep up with the routing tables, and as such they
   just sit and listen for the last RTMP packet broadcast on their cable.

2) The end nodes then use the address of that router for all routing
   functions. If router B was the last router that the end node saw
   broadcasting an RTMP packet, then the initial NBP lookUp for the AFP
   server is sent to router B, even though router A is the best choice.

3) Eventually, the packet gets to the AFP server, and the AFP server
   responds directly to the end node through router A, leaving router B
   completely out of the picture.

4) The end node then caches the address of the router that sent the packet,
   and all subsequent traffic to the server for this communications process
   uses router A. That's the basic reason for the DDP "best route"
   mechanism.

Now let's redraw the picture making router B a redundant link between
network 1 and network 2.

```
Net 1
================================================================
        |                  |                      |
    [router A]         [router B]          [End node (Mac-A)]
        |                  |     |
        |                  |      ------to another network
Net 2   |                  |
================================================================
            |
        [AFP Server]
```

In this case, either router A or router B would work as the "best route" to
network 2. In detail:

1) The end node (Mac-A) sends an NBP lookUp for the AFP server on network 2

   to either router A or router B, whichever broadcast its routing table
   last.

2) Either of the routers are capable of delivering the packet to network 2,
   but let's say router B was chosen. Router B passes the NBP packet on to
   network 2 where the AFP server picks it up and responds directly to the
   end node.

3) If the AFP server last saw an RTMP packet from router A, then router A
   would be the choice to get the packet to network 1. Since router A has
   connectivity to network 1, it passes it along and the end node (Mac-A)
   picks it up.

4) When the end node picks it up, it sees that the last router that was
   used was router A, and it then caches router A as the "best route" for
   the connection to the AFP server. The AFP server could just have easily
   last heard from router B, which would cause router B to be chosen as the
   "best route".

In a 3rd scenario, let's separate the redundant path by yet another
router.

```
Net 1
===================================================================
                   |                        |
          [End node (Mac-A)]           [router C]
                                            |
Net 2                                       |
===================================================================
        |                   |
    [router A]          [router B]
        |                   |     |
        |                   |     ------to another network
Net 3   |                   |
===================================================================
                                        |
                                    [router D]
Net 4                                   |
===================================================================
            |
        [AFP Server]
```

This actually looks more like what you've installed. In this scenario, DDP
"best route" does nothing to allow better use of the redundant path. Let's
look at the example:

1) If the end node (Mac-A) wants to talk to the AFP server on network 4, it
   performs an NBP lookUp.

2) The lookUp is sent to router C, and router C uses either router A or
   router B to deliver the packet to router D and network 4.

3) Router C uses the router that has a valid path and was the last one

heard from on network 2, either A or B. This is because the RTMP tables
on router C are constantly changing state for the entry to network 4,
and since both are offering the same cost, the newest entry is used.

4) The packet eventually gets to network 4, and the AFP server responds
   to the end node.

5) The end node will cache router C as the "best route" because it was the
   last router to touch the packet.

Nothing in the DDP "best route" caching mechanism does anything to load
balance the link or to use one router over the other.

Increased NBP and RTMP Traffic and RTMP Packet Size

Increased NBP and RTMP traffic and large RTMP packets usually don't cause
problems unless your network is already heavily loaded.

Increased NBP traffic can be attributed to additional devices on the cable.
When anyone does an NBP lookUp for those particular types of devices,
you'll have increased NBP traffic.

Your RTMP tables will be the same size; there are no additional routes to
advertise, so the size of the table won't increase. Please note that on
the cables where redundant routers are connected, the number of RTMP
packets will double as both routers will be sending RTMP packets. In an
internet where redundant routers are present, the split horizon mechanism
does nothing to reduce the size of the RTMP tables over the size reduction
already realized.

Problem with Redundant AppleTalk Routers

There is at least one small problem with redundant AppleTalk routers, and
there may be others. Large enterprise networks have routes that sometimes
flap. A route that flaps is stable (marked as good), then unstable (marked
as suspect), then stable, then unstable (you get the idea). This is a
common occurrence in large AppleTalk environments. You should ultimately
find the cause of the flapping route and fix it, but sometimes it's caused
by nothing more than a router that was momentarily overloaded.

```
Net 1
==================================================================
       |
     [router A]
         |
         |
Net 2    |
==================================================================
                                       |
                                   [router B]
                                       |
Net 3                                  |
==================================================================
```

```
                                       |
                                  [router C]
                                       |
       Net 4                           |
       =================================================================
             |                 |
         [router D]        [router E]
             |                 |
             |                 |
       Net 5 |                 |
       =============================================================
```

Here's what happens in a situation with redundant routers. Let's say
router A gets overloaded and isn't able to get its RTMP packets broadcast
as often as needed to sustain the valid routes throughout the internet.
Router B is the first router to discover this, and marks the route as bad.
Router B then tells router C, and router C then tells router D and router
E.

Router D gets the message that the route is bad. Router E does the same
thing, but at the same time router E was in the process of sending out
an RTMP packet to network 5. This RTMP packet says that router E has a
valid route to network 1, and updates its tables accordingly. Router D
then advertises that on network number 4, and the message is propagated
through the network. This problem will only manifest itself for about 60
seconds, since router E will eventually time out the entry and cause the
whole chain to break down. But it does create a "black hole" for 60
seconds and can cause problems for end users in the intervening time. If
you have a rock solid network, you'll never see this problem. It's common
for large AppleTalk networks to occasionally suffer from flapping routes.
You can even see this type of anomaly under normal conditions when a router
is taken down for maintenance purposes.

This may seem to be a trivial problem, but it caused our network here at
Apple considerable grief when we experienced this phenomenon. After going
through this experience, we hesitate to recommend redundant routers in a large
AppleTalk network. Redundant AppleTalk routers are less useful than redundant
links for protocols such as TCP/IP that have routing protocols to deal with
redundancy. They are perfectly legal, but may not be very useful.

Conclusion

Redundant routers may be an excellent way to provide reliable service to
critical parts of the network, and may work great if your network is rock
solid, but as it stands today we recommend that you leave one of them turned
off until it's actually needed.

Tech Info Library Article Number:10724