

High-Availability Networks with Spanning Tree and Rapid Spanning Tree Protocol

This Application Note relates to the following Dell PowerConnect™ products:

- PowerConnect 33xx
- PowerConnect 52xx

Abstract

As organizations and their networks expand, high availability becomes a key requirement. Even small periods of network downtime can lead to losses in productivity, revenue, and employee and customer satisfaction.

Adding redundant links to the network can increase availability, but also carries with it the danger of traffic loops. Network loops increase the processing burden on hosts and lead to generally unpredictable network performance.


The spanning tree protocol (STP) addresses the need for redundancy while providing mechanisms for loop detection and prevention. Because STP is a standards-based solution, it allows interconnection of multiple vendors' equipment. Further, STP can run at multiple link speeds simultaneously.

This application note describes STP and rapid spanning tree protocol (RSTP). This document also provides step-by-step instructions for implementing these technologies.

Applicable Network Scenarios


Spanning tree protocol is applicable for essentially any network that has two or more switches and requires redundancy.

Consider a network that connects a highly automated manufacturing facility with a control center. As shown in the diagram below, one switch in the factory connects to manufacturing machines and one switch in the control center connects to several computers. A single link interconnects the two switches.



In this example, the single switch link represents a single point of failure between the manufacturing floor and the control center. Since the manufacturing floor relies on computers in the control center for instructions, any failure of the interswitch link would bring production to a halt.


Adding a second link provides redundancy but it also creates a loop, with equally undesirable results.



In the diagram above, suppose the control center sends a “batch start” command to instruct factory machines to begin production. If either switch does not know the destination host’s address, it will flood (broadcast) traffic to all interfaces of both switches – including interfaces connected back to the control center. Note that this situation is quite common, especially since by default PowerConnect switches and most other switches remember host addresses for only 5 minutes.

Besides reducing network availability by consuming excessive bandwidth, broadcasts also degrade application performance. Every broadcast frame must be processed by all hosts on both switches, reducing the amount of processing power the hosts have available for other tasks. Further, switch interfaces at the factory will receive the batch start command multiple times.

By enabling the spanning tree protocol, we can prevent traffic loops while still providing redundancy. STP achieves this result by blocking one of the interswitch links until it is needed. Now there is only one path active between the switches.



Technology Background

As defined in IEEE standard 802.1D, the spanning tree protocol is a method to detect loops and shut down redundant links. Using control messages called bridge protocol data units (BPDUs), switches and bridges running STP elect a “root bridge” at the center, or root, of the spanning tree. Each switch and bridge in the network determines which of its interfaces is the closest to the root bridge, and designates that interface as its root interface. A switch or bridge can only have one root interface active at any one time. At any given time, each switch or bridge interface has only one path to any host on the network.

If a switch or bridge determines it has a second path available via another interface, it places that interface in blocking mode so it won’t forward traffic. However, the interface in blocking mode continues to listen for changes in network topology. If a link or interface fails, the spanning tree process begins again. The 802.1D version of STP typically takes between 30 and 60 seconds to “converge.” Convergence occurs when switches and bridges define a stable tree and traffic can pass freely around the network. Until the network has converged, all interfaces remain in a blocking state and the network is temporarily inoperable.

For many networks, convergence times of 30 to 60 seconds are simply too high. RSTP is designed to achieve faster convergence times through the use of acknowledged communications between devices rather than the passive method used by STP. RSTP also is backward compatible with STP, a helpful feature when gradually migrating networks to RSTP. However, mixed STP/RSTP networks have the slower convergence times of STP.

Proposed Solution

Overview


Implementing spanning tree is a simple two-step procedure:

1. Enable spanning tree.
2. Verify correct spanning tree operation.

In the following example, we configure a Dell PowerConnect 5224 switch and a Dell PowerConnect 3348 switch to run RSTP. This configuration might be used in the network of the example given earlier in this application note.

Typical Network Design

The diagram below again shows a manufacturing facility equipped with a Dell PowerConnect 5224 switch and a control center equipped with a Dell PowerConnect 3348 switch. We will use interfaces 17 and 18 on each device as members of the spanning tree.



Step-By-Step Instructions

1. Enable spanning tree in RSTP mode.

PowerConnect 5224:

```
Dell-5224# configure
Dell-5224 (config)# spanning-tree
Dell-5224 (config)# spanning-tree mode rstp
Dell-5224 (config-if)# exit
Dell-5224 (config)# exit
```

Now the PowerConnect 5224 switch is ready to function in a spanning tree with other RSTP devices.

PowerConnect 3348:

```
Dell-3348> en
Dell-3348# configure
Dell-3348 (config)# spanning-tree
Dell-3348 (config)# spanning-tree mode rstp
```

```
Dell-3348 (config-if)# exit
Dell-3348 (config)# exit
```

Now the PowerConnect 3348 switch is ready to function in a spanning tree with other RSTP devices.

2. Verify correct spanning tree operation.

PowerConnect 5224:

```
Dell-5224# show spanning-tree Ethernet 1/17
```

This should bring up, among other things, the current status of the spanning tree interface.

```
Admin status      :enable
Role              :root
State             :forwarding
```

Now we check the other interface on the PowerConnect 5224 switch.

```
Dell-5224# show spanning-tree Ethernet 1/18
```

```
Admin status      :enable
Role              :alternate
State             :discarding
```

In this case, the 5224 has designated interface 1/17 as root interface and put it into a forwarding state. The switch puts interface 1/18 into blocking mode.

PowerConnect 3348:

```
Dell-3348> en
Dell-3348# show spanning-tree Ethernet 1/e17
```

Interface Name	Port ID Prio	Sts	Enb	Cost	Designated Cost	Bridge Id	Port ID Prio.Nbr
1/e17	128	FRW	TRUE	19	0	8000 00:0b:db:f5:90:5d	80 11

```
Dell-3348# show spanning-tree Ethernet 1/e18
```

Interface Name	Port ID Prio	Sts	Enb	Cost	Designated Cost	Bridge Id	Port ID Prio.Nbr
1/e18	128	FRW	TRUE	19	0	8000 00:0b:db:f5:90:5d	80 12

Note: This example shows both interfaces of the PowerConnect 3348 switch in a forwarding state. This is correct behavior of RSTP. Only one side blocks the redundant link. In this case, the PowerConnect 5224 switch puts an interface into blocking state. Both interfaces on the other side of the redundant links remain in forwarding state.

As a final step, we can verify RSTP functionality using the ping utility on most computers. From one computer in the control center (192.168.1.100), we will use the ping utility to reach another system at the manufacturing facility (192.168.1.200). From a machine running Microsoft® Windows®, use the “-t” switch to ping the specified host until stopped:

```
C:\ >ping -t 192.168.0.200
```

```
Pinging 192.168.0.200 with 32 bytes of data:
```

```
Reply from 192.168.0.200: bytes=32 time<1ms TTL=64
Reply from 192.168.0.200: bytes=32 time<1ms TTL=64
Reply from 192.168.0.200: bytes=32 time<1ms TTL=64
Reply from 192.168.0.200: bytes=32 time<1ms TTL=64
Reply from 192.168.0.200: bytes=32 time<1ms TTL=64
```

Once the ping stream begins, we physically disconnect the cable from one of the spanning tree interfaces. The ping requests will fail to reach the destination host for some interval. After spanning tree reroutes traffic onto the redundant link, the ping request will again reach the destination host, which in turn will resume sending replies:

```
Reply from 192.168.0.200: bytes=32 time<1ms TTL=64
Reply from 192.168.0.200: bytes=32 time<1ms TTL=64
Reply from 192.168.0.200: bytes=32 time<1ms TTL=64
Request timed out.
Request timed out.
Reply from 192.168.0.200: bytes=32 time<1ms TTL=64
Reply from 192.168.0.200: bytes=32 time<1ms TTL=64
Reply from 192.168.0.200: bytes=32 time<1ms TTL=64
```

← Cable disconnected

← RSTP reconverges

Conclusion

We have now enabled rapid spanning tree protocol between the Dell PowerConnect 5224 switch and the Dell PowerConnect 3348 switch. The result is a network designed to have higher availability, reduced broadcast traffic, and improved application performance.

Information in this document is subject to change without notice.

© 2003 Dell Inc. All rights reserved.

This Application Note is for informational purposes only, and may contain typographical errors and technical inaccuracies. The content is provided as is, without express or implied warranties of any kind.

Dell, the DELL logo, and PowerConnect are trademarks of Dell Inc. Microsoft and Windows are registered trademarks of Microsoft Corporation. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell Inc. disclaims any proprietary interest in trademarks and trade names other than its own.