

# Network Redundancy Concepts and Configuration

## An Introduction to RSTP and MSTP

Managed Ethernet switches allow redundant connections between portions of a network, making choices based on configuration information which gives preference to certain links. In general, a knowledgeable network engineer can configure the switches to steer traffic over specific links under appropriate circumstances. This Technical Note explains network redundancy concepts and gives general direction on configuring Multiple Spanning Tree Protocol (MSTP) to make the most effective and efficient use of redundant link between switches.

Ethernet packets are hard to touch or see. But we talk of network “traffic” and a “network” of roads and highways, and automobile traffic can be a reasonable model for network traffic. The next section explains Ethernet redundancy in terms of automobile traffic in the growing hamlet of Lanford. The following section walks through specific switch configuration to achieve network redundancy similar to the traffic management on the roads and bridges of Lanford.

### Welcome to Lanford


Where the confluence of the Cerf and the Postel form the Réseau, the town of Lanford sprang up. In the very early days of the settlement, most people lived on one bank and boats carried cargo across the Réseau when needed. As traffic grew, a ferry service was established and eventually, the ferry station on the other side was joined by more and more businesses and houses until the town needed a more robust link across the river. The town undertook to build a bridge and soon a lovely two-lane span rose above the water, connecting the two parts of town.




Once the bridge was in place, Lanford became a preferred location for crossing the Réseau and traffic through town increased.



Eventually, heavy traffic on the bridge led to delays. A new, four-lane bridge was built just downriver from the original span.



Of course, nothing succeeds like success and soon traffic had grown to the capacity of the new bridge and delays began anew. Town planners were contemplating yet another bridge when someone realized that the original bridge could take some of the burden off the new bridge; it was a lot cheaper to control traffic on two bridges than to build a third. The old bridge wasn't up to modern, heavy traffic but if, say, all red vehicles could be diverted to the old bridge, the traffic on the new bridge would be lighter and everyone would be happy. Furthermore, if something happened to the new bridge, having the old bridge open would still allow at least some traffic to be diverted there quickly, even if it had to cross more slowly to account for the age and capacity of the old bridge.



When the new span opened, it could carry all the traffic with ease and the old link was retired and access roads diverted traffic to the new bridge. The old bridge was usable but it was easier to just send everyone over the new one than to manage both bridges, staff two sets of toll booths, etc. So, the old bridge sat idle but the people took comfort from knowing that if something happened to the new bridge, the old was there as a backup and they kept it in good repair just in case.



Over several generations, traffic through Lanford continued to grow. Eventually, more spans were unavoidable. In the mean time, the population had jumped the Cerf and connecting the three parts of town involved adding bridges over the Cerf and Postel as well.




Lanford remained well-known as a river crossing and local traffic control kept cars and trucks humming across the river on the various crossings with traffic sent to one or another span based on its destination.



### Making it Real

The story of Lanford is not unlike the story of LANs. As Lanford's road and bridge network grew to meet the demands of a growing population, so do Ethernet networks grow as users and applications are added. And careful management of switches and links can defer costly build out or make the most of existing infrastructure.

If East and West Lanford represent two work groups with relatively light interaction, the initial two-lane bridge in Lanford could represent a 10Mbps network link between work group switches.




There is no redundancy and no need to manage the connection. The switches can be unmanaged switches which just forward traffic across the link between them when appropriate.

As the groups and their interactions grow, more network capacity is needed between them. A 100Mbps link is added between the switches. That link is so much more capacious than the original link that it's not worth thinking about keeping both active. But the 100Mbps switches happen to be managed so the original link is kept as a hot standby. Having the original link idle most of the time doesn't require much management. RSTP is enabled and the port cost (based on speed) is enough to make the right thing happen.

**RSTP**

As long as the faster link is active, the slower link blocks so no loops are created in the network. But when the faster link fails, the slower link unblocks and, to the limit of its capacity, keeps the work groups connected. This is RSTP<sup>1</sup>.

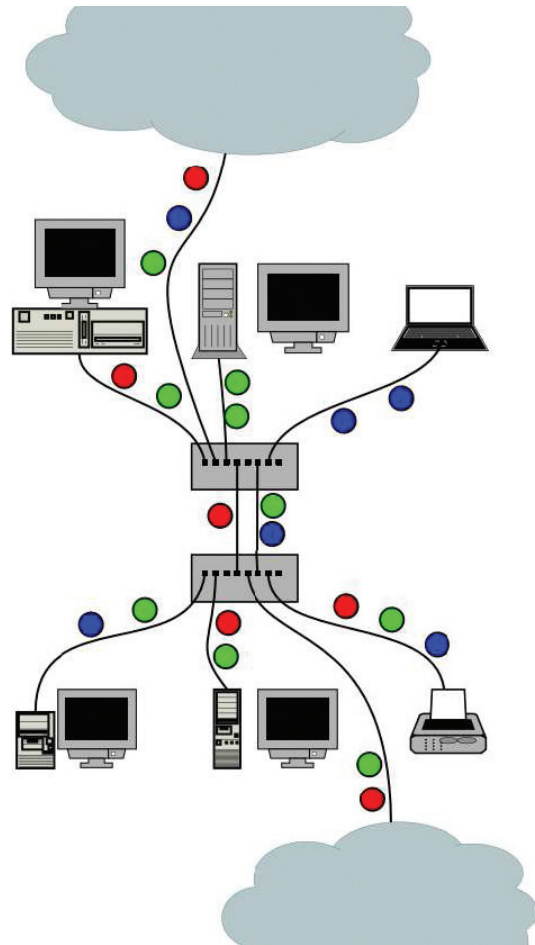


One of the distinguishing features of RSTP is that it doesn't require network-wide cooperation. Each switch acts on its own, based solely on reports from its neighbor. Only the two work group switches are involved in blocking one of the links between them. Also, just as a dispatcher in Wanville needn't know which Lanford bridge is open to send trucks there, other switches don't know or care which link between the work group switches is blocked.

Once all those local decisions are made, it is possible to draw a map of the overall network. When you omit the blocked links from such a map, the overall shape of the network is what mathematicians call a "tree." Switches are branching points and unblocked links are the trunks or branches or twigs that come out of them. The overall map spans the entire network, leaving no node disconnected, and is therefore a "spanning tree."

**MSTP**


Eventually, traffic between the groups grows until even the small increment of the slow link would be useful. The network engineers already use VLANs to isolate different types of traffic and if one or two VLANs could be carried on the slower link, it would lighten the load on the faster link. Enter MSTP.



With MSTP, you create and configure multiple spanning tree instances (MSTIs) with independent settings for port cost and priority allowing different links to block in each instance. VLANs are assigned to specific MSTIs resulting in VLAN traffic being divided between multiple links.

### MSTP Region

MSTP is somewhat complex to implement and configure and must interact with existing RSTP switches. An MSTP region localizes the configuration complexity. All switches in a region must share the same region name, configuration revision number, and VLAN/MSTI assignments. Within the region, spanning tree instances allow use of multiple links between switches. From outside the region, the entire region appears to be a single RSTP switch, other parts of the network can't tell that MSTP is operating within the region.



### Configuration

Consider a switch carrying traffic on three VLANs: red, green, and blue. In the simplest case, data for all the VLANs is carried on all ports.

#### VLAN SETTINGS

VLAN mode	Standard
Core type	0x8100
Learning	Independent

Name	Type	ID	FID		Ports									
Management	Tag-based	1	0	CPU	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Red	Tag-based	10	10	CPU	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Green	Tag-based	20	20	CPU	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Blue	Tag-based	30	30	CPU	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

If we enable RSTP, a blocking port is blocked for all VLANs and remains idle, carrying only management data, no user or application data. However, if we enable MSTP, we can tune the port priorities in the spanning tree instances and assign the VLANs to the instances to split the traffic between ports under normal circumstances.

Begin by enabling MSTP and creating two MSTIs.

#### SPANNING TREE SETTINGS

Redundancy protocol	Multiple Spanning Tree Protocol
Bridge priority (0-61440)	32768
Maximum age (6-40 seconds)	20
Hello time (1-10 seconds)	2
Forward delay (4-30 seconds)	15
Transmission limit (1-10)	6
Region Name (MSTP)	Lanford
Configuration Revision (MSTP)	0
Max Hops (MSTP)	20

MST Instance	MSTID (1-4094)	Bridge priority (0-61440)
Big	1	32768
Small	2	32768

Then adjust the port priorities for the MSTIs. If the switches are connected via ports 1 and 2, and port 1 is the higher-capacity link (like the 4-lane span in Lanford), we configure the Big MSTI by giving port 1 the best priority (0) and a lower cost than port 2 and do the opposite in the Little MSTI.

**SPANNING TREE PORT SETTINGS**

Select a spanning tree instance to configure: **Big (customized)**

**BIG PORT SETTINGS**

Port	Name	Priority	Path Cost
1	port_1	0	Auto 20000
2	port_2	240	Auto 200000

**SPANNING TREE PORT SETTINGS**

Select a spanning tree instance to configure: **Small (customized)**

**SMALL PORT SETTINGS**

Port	Name	Priority	Path Cost
1	port_1	240	Auto 200000
2	port_2	0	Auto 20000

Then assign one VLAN to Small and two to Big.

**VLAN SETTINGS**

VLAN mode: Standard  
 Core type: 0x8100  
 Learning: Independent

Name	Type	ID	FID	Ports	MSTI
Management	Tag-based	1	0	CPU 1 2 3 4 5 6 7 8	
Red	Tag-based	10	10	CPU 1 2 3 4 5 6 7 8	Small
Green	Tag-based	20	20	CPU 1 2 3 4 5 6 7 8	Big
Blue	Tag-based	30	30	CPU 1 2 3 4 5 6 7 8	Big

As long as ports 1 and 2 are linked to the same redundant switch, the Small MSTI will choose to block port 1 so the Red VLAN will communicate on port 2 and the Big MSTI will choose to block port 2 so the Green and Blue VLANs will be carried on port 1. If one port fails, the other will pick up all of its traffic, up to its capacity.

<sup>1</sup>For a simple, point-to-point connection like this, port trunking could allow both ports to be unblocked and give more dynamic load balancing and link utilization with less configuration. However, this does not apply to the general case of redundant links among a set of switches and only works between two switches when the links have the same speed.

**Ultra Electronics**

NUCLEAR SENSORS & PROCESS INSTRUMENTATION

707 Jeffrey Way, PO Box 300

Round Rock, TX 78680-0300 USA

Tel: +1 512 434 2850

Fax: +1 512 434 2901

e-mail: fiberop@ultra-nspi.com

www.ultra-nspi.com

Ultra Electronics, Nuclear Sensors & Process Instrumentation is a business name of Weed Instrument Co., Inc.

Ultra Electronics reserves the right to vary these specifications without notice.

© Ultra Electronics 2010.

Printed in the USA.

03/10