3.1. Introduction

Redundant topologies provide multiple paths for traffic to flow, so that the network remains operational even if a single component fails. This is especially important for real-time voice and video services, where even brief interruptions can be disruptive.

Physical redundancy is the simplest way to eliminate single points of failure. For example, building a reliable switched network requires additional switches and redundant links between devices. When a switch or link fails, the redundant one takes over, providing unnoticeable network interruptions.

Benefits of redundant topologies:

- Increased reliability: Redundant networks are less likely to experience outages due to single component failures.
- Improved performance: Redundant networks can distribute traffic across multiple paths, which can improve performance and reduce congestion.
- Increased scalability: Redundant networks can be more easily scaled to meet growing demand.
- Reduced downtime: Redundant networks can be maintained and upgraded without disrupting service.

Problems that can result from a redundant Layer 2 network:

- Broadcast storms: Switches forward broadcast frames to all ports except the port on which the frame was received. If there is a loop in the network, broadcast frames can circulate indefinitely, consuming bandwidth and causing network outages.
- MAC database instability: Switches maintain MAC address tables to learn which MAC addresses are reachable on which ports. If there is a loop in the network, switches may receive conflicting MAC address information, causing the MAC address tables to become unstable.
- Duplicate unicast frames: If there is a loop in the network, unicast frames may be forwarded multiple times, causing duplicate frames to be received by destination hosts.

A network of roads is a good example of a redundant topology. If one road is closed for repair, there is likely an alternate route to the destination. Similarly, a community with two bridges across a river is more resilient than a community with only one bridge. If one bridge is damaged, the other bridge can still be used.

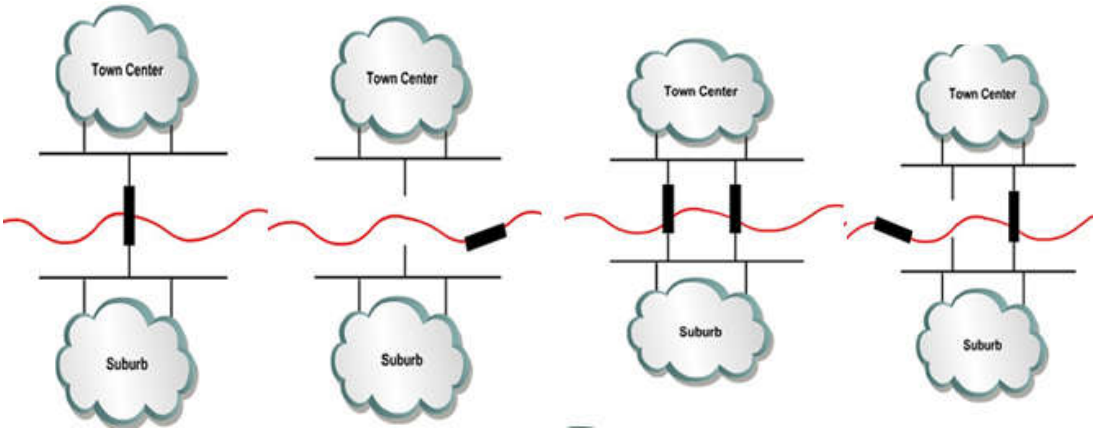

3.2. Hierarchical Network Design: Fault Domains

In the field of networking, hierarchical design divides the network into distinct layers. Each layer, or level, of the hierarchy offers specific functions that define its role in the network. This allows the network designer and architect to select the appropriate network hardware and software, as well as the features needed for the roles of that network layer. Hierarchical models apply to both LAN and WAN designs

To optimize bandwidth in an enterprise network, the network must be organized so that traffic is maintained at the local level and does not spread unnecessarily to other parts of the network. The use of the three-layer hierarchical design model helps to organize the network.

In following figure, the functionality of the network is divided into three different layers.

1. Access layer
2. Distribution layer
3. Core layer

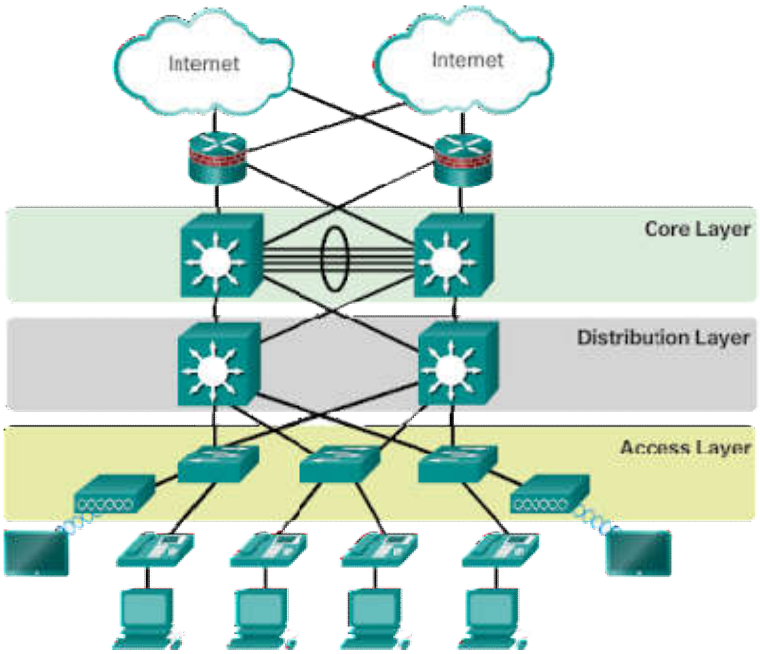

Each layer is designed to fulfill specific functions. The access layer provides connectivity to users. The distribution layer is used to send traffic from one local network to another. Finally, the core layer represents a high-speed backbone between the dispersed networks. User traffic starts at the access layer and passes through the other layers if you need to use the functionality of those layers.

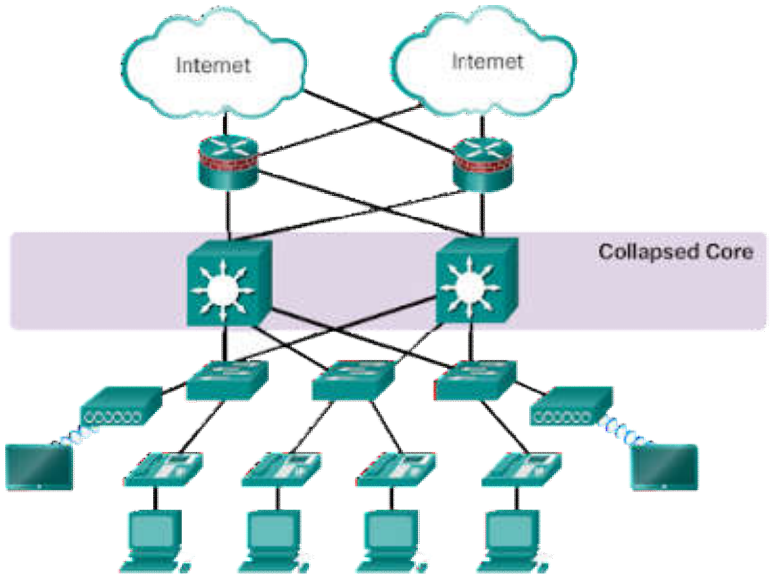

Although the hierarchical model consists of three layers, a two-tier hierarchical design may be implemented in some small business networks. As shown in the bellow figure, in a two-tier

hierarchical design , the core and distribution layers are combined into one, which reduces cost and complexity.

Fault Domains

A well-designed network not only controls traffic, but also limits the size of failure domains. A failure domain is the area of a network that is impacted when a critical device or network service experiences problems.

The function of the device that initially fails determines the impact of a failure domain. For example, a malfunctioning switch on a network segment normally affects only the hosts on that segment. However, if the router that connects this segment to others fails, the impact is much greater.

The use of redundant links and reliable enterprise-class equipment minimize the chance of disruption in a network. Smaller failure domains reduce the impact of a failure on company productivity. They also simplify the troubleshooting process, thereby, shortening the downtime for all users.

The following figure shows the associated failure domain to each network device.

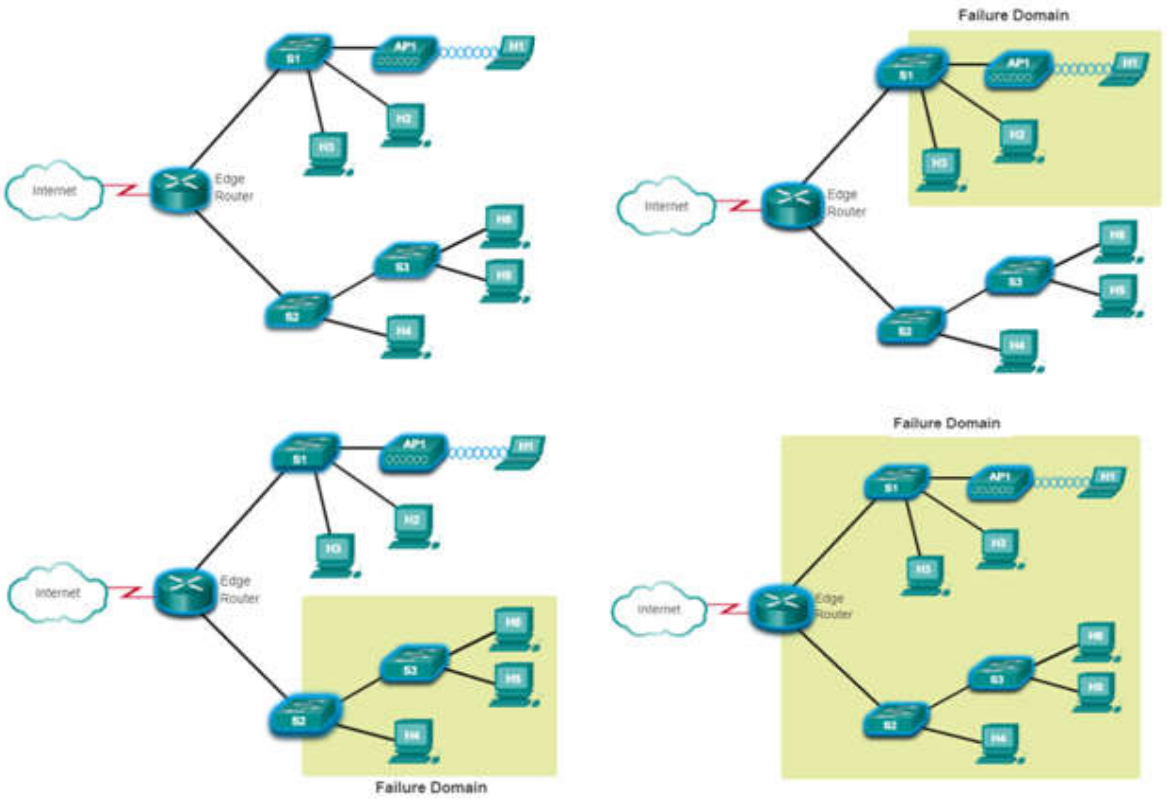

Limiting the Size of Failure Domains

Because a failure at the core layer of a network can have a potentially large impact, the network designer often concentrates on efforts to prevent failures. These efforts can greatly increase the cost of implementing the network. In the hierarchical design model, it is easiest and usually least expensive to control the size of a failure domain in the distribution layer. In the distribution layer, network errors can be contained to a smaller area; thus, affecting fewer users. When using Layer 3 devices at the distribution layer, every router functions as a gateway for a limited number of access layer users.

Switch Block Deployment

Routers, or multilayer switches, are usually deployed in pairs, with access layer switches evenly divided between them. This configuration is referred to as a building, or departmental, switch block. Each switch block acts independently of the others. As a result, the failure of a single device does not cause the network to go down. Even the failure of an entire switch block does not affect a significant number of end users.

3.3. Spanning Tree Protocol (STP)

Redundancy increases the availability of the network topology by protecting the network from a single point of failure, such as a failed network cable or switch. When physical redundancy is introduced into a design, loops and duplicate frames occur. Loops and duplicate frames

have severe consequences for a switched network. The Spanning Tree Protocol (STP) was developed to address these issues.

STP, the IEEE 802.1D bridge protocol, is a Layer 2 link-management protocol that provides path redundancy while preventing undesirable loops in the network. For a Layer 2 Ethernet network to function properly, only one active path can exist between any two stations. STP operation is transparent to end stations, which cannot detect whether they are connected to a single LAN segment or a switched LAN of multiple segments.

STP ensures that there is only one logical path between all destinations on the network by intentionally blocking redundant paths that could cause a loop. A port is considered blocked when user data is prevented from entering or leaving that port. This does not include bridge protocol data unit (BPDU) frames that are used by STP to prevent loops. Blocking the redundant paths is critical to preventing loops on the network. The physical paths still exist to provide redundancy, but these paths are disabled to prevent the loops from occurring. If the path is ever needed to compensate for a network cable or switch failure, STP recalculates the paths and unblocks the necessary ports to allow the redundant path to become active. The following figure explain STP principles. All switches have STP enabled:

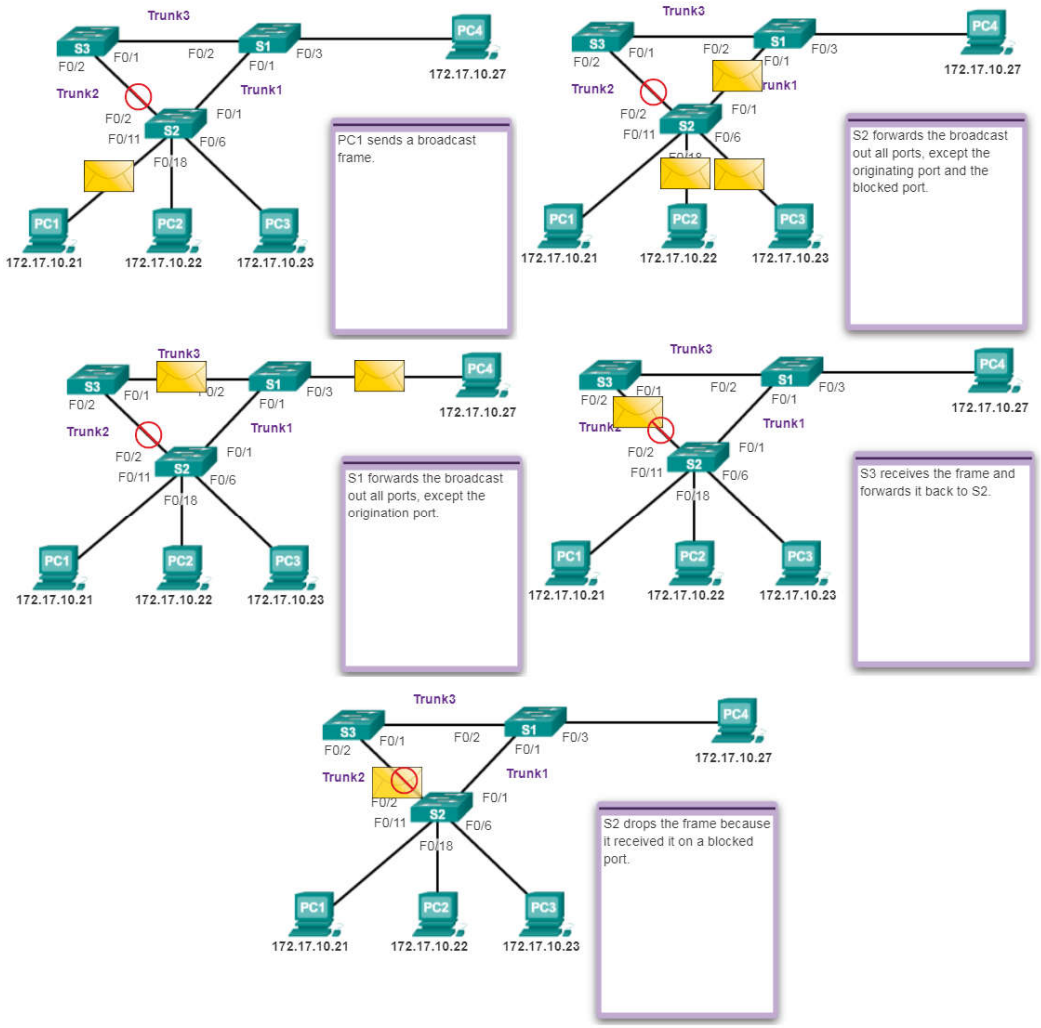

STP Operation

1. PC1 sends a broadcast out onto the network.

2. S2 is configured with STP and has set the port for Trunk2 to a blocking state. The blocking state prevents ports from being used to forward user data, thus preventing a loop from occurring. S2 forwards a broadcast frame out all switch ports, except the originating port from PC1 and the port for Trunk2.

3. S1 receives the broadcast frame and forwards it out all of its switch ports, where it reaches PC4 and S3. S3 forwards the frame out the port for Trunk2 and S2 drops the frame. The Layer 2 loop is prevented.

Now the following figure shows STP recalculation when a failure occurs.

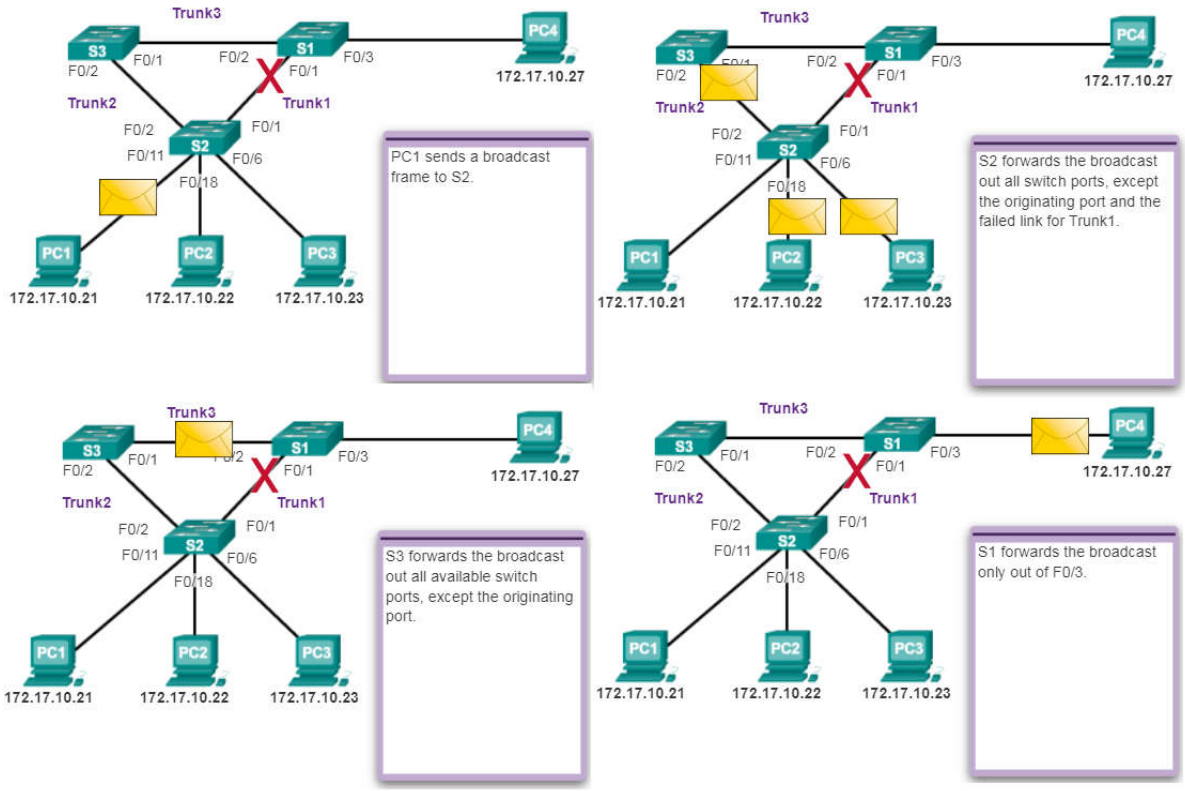

STP Compensates for Network Failure

1. PC1 sends a broadcast out onto the network.

2. The broadcast is then forwarded around the network, just as in the previous animation.

3. The trunk link between S2 and S1 fails, resulting in the previous path being disrupted.

4. S2 unblocks the previously blocked port for Trunk2 and allows the broadcast traffic to traverse the alternate path around the network, permitting communication to continue. If this link comes back up, STP reconverges and the port on S2 is again blocked.

STP prevents loops from occurring by configuring a loop-free path through the network using strategically placed "blocking-state" ports. The switches running STP are able to compensate for failures by dynamically unblocking the previously blocked ports and permitting traffic to traverse the alternate paths.

Up to now, we have used the term Spanning Tree Protocol and the acronym STP. The usage of the Spanning Tree Protocol term and the STP acronym can be misleading. Many professionals generically use these to refer to various implementations of spanning tree, such as Rapid Spanning Tree Protocol (RSTP) and Multiple Spanning Tree Protocol (MSTP). In order to communicate spanning tree concepts correctly, it is important to refer to the particular implementation or standard in context. The latest IEEE documentation on spanning tree, IEEE-802-1D-2004, says "STP has now been superseded by the Rapid Spanning Tree Protocol (RSTP)"; so one sees that the IEEE uses "STP" to refer to the original implementation of spanning tree and "RSTP" to describe the version of spanning tree specified in IEEE-802.1D-2004. In this curriculum, when the original Spanning Tree Protocol is the context of a discussion, the phrase “original 802.1D spanning tree” is used to avoid confusion.

Note: STP is based on an algorithm invented by Radia Perlman while working for Digital Equipment Corporation, and published in the 1985 paper "An Algorithm for Distributed Computation of a Spanning Tree in an Extended LAN".

Spanning Tree Versions

Several varieties of spanning tree protocols have emerged since the original IEEE 802.1D.

The varieties of spanning tree protocols include:

- **STP** - This is the original IEEE 802.1D version (802.1D-1998 and earlier) that provides a loop-free topology in a network with redundant links. Common Spanning Tree (CST) assumes one spanning tree instance for the entire bridged network, regardless of the number of VLANs.

- **PVST+** - This is a Cisco enhancement of STP that provides a separate 802.1D spanning tree instance for each VLAN configured in the network. The separate instance supports PortFast, UplinkFast, BackboneFast, BPDU guard, BPDU filter, root guard, and loop guard.

- **802.1D-2004** - This is an updated version of the STP standard, incorporating IEEE 802.1w.

- **Rapid Spanning Tree Protocol (RSTP) or IEEE 802.1w** - This is an evolution of STP that provides faster convergence than STP.

- **Rapid PVST+** - This is a Cisco enhancement of RSTP that uses PVST+. Rapid PVST+ provides a separate instance of 802.1w per VLAN. The separate instance supports PortFast, BPDU guard, BPDU filter, root guard, and loop guard.

- **Multiple Spanning Tree Protocol (MSTP)** - This is an IEEE standard inspired by the earlier Cisco proprietary Multiple Instance STP (MISTP) implementation. MSTP maps multiple VLANs into the same spanning tree instance. The Cisco implementation of MSTP is MST, which provides up to 16 instances of RSTP and combines many VLANs with the same physical and logical topology into a common RSTP instance. Each instance supports PortFast, BPDU guard, BPDU filter, root guard, and loop guard.

STP characteristics

Protocol	Type	Description	CPU/Memory Requirements	Convergence Time	Default on Cisco Catalyst Switches
STP	IEEE standard	Creates a single spanning tree instance for the entire network	Low	Slow	No
PVST+	Cisco proprietary	Creates a separate spanning tree instance for each VLAN	Medium	Medium	Yes
RSTP	IEEE standard	Provides faster convergence than STP	Medium	Fast	No
Rapid PVST+	Cisco proprietary	Combines the features of PVST+ and RSTP	High	Very fast	No
MSTP	IEEE standard	Groups multiple VLANs with the same traffic flow requirements into the same spanning tree instance	Medium	Medium	No
MST	Cisco proprietary	Provides up to 16 instances of RSTP and combines many VLANs with the same physical and logical topology into a common RSTP instance	Medium	Medium	No

3.4. Rapid Spanning Tree Protocol

Rapid Spanning Tree Protocol (RSTP) is an improved version of Spanning Tree Protocol (STP) that enhances network performance. Like STP, RSTP creates a single spanning tree, which means it cannot consider VLANs when managing redundant paths. However, RSTP's performance enhancements make it a better choice than STP in environments without VLANs.

RSTP complies with the IEEE 802.1w standard and uses a strict set of guidelines to determine the optimal way to forward network traffic without redundancy. When enabled, RSTP automatically configures the spanning tree.

In RSTP, the topmost bridge in the network is the root bridge. The root bridge is responsible for sending network topology information to other switches in the network. This is important in case of hardware failures or topology changes, as it ensures that optimal alternative routes are established immediately. RSTP protocol defines four port roles:

- **Root ports** connect to the root bridge, which has the lowest bridge ID in the network. Root ports are always in forwarding state and forward all traffic to and from the root bridge.
- **Designated ports** connect to other switches and forward traffic to and from their segments. Designated ports are also always in forwarding state and are selected based on the lowest path cost to the root bridge.
- **Alternate ports** provide an alternative path to the root bridge in case of a failure of a root or designated port. Alternate ports are initially in discarding state and do not forward any traffic unless they become root or designated ports.
- **Backup ports** provide a redundant link to the same segment as a designated port. Backup ports are also initially in discarding state and do not forward any traffic unless they become designated ports.

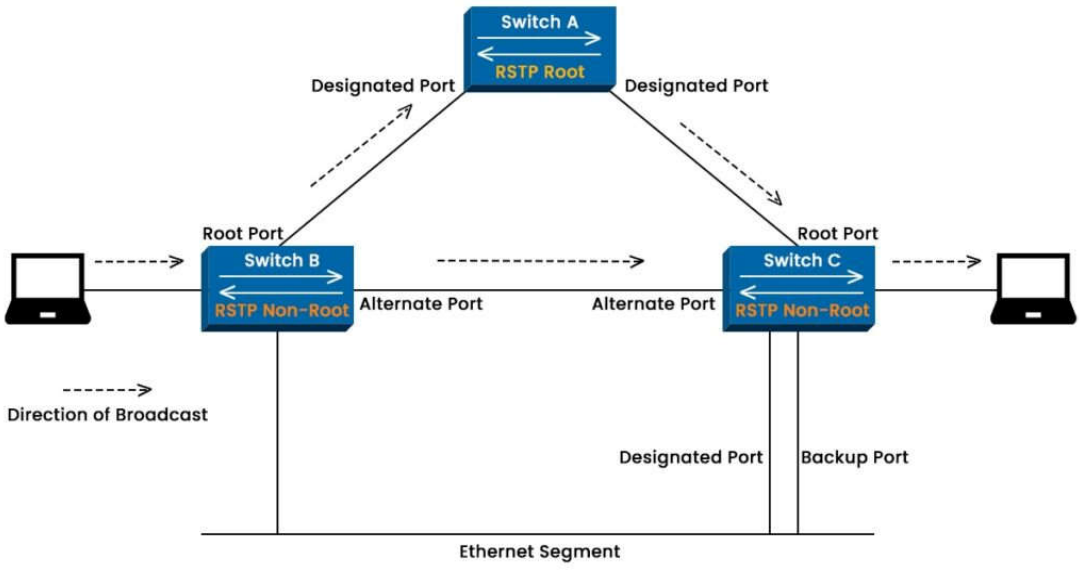

Port States in RSTP

RSTP port states are the different roles and statuses that a port on a bridge can be in, depending on its role and status in the network topology. RSTP has three port states:

- **Discarding**: The port is not forwarding traffic or learning MAC addresses.
- **Learning**: The port is not forwarding traffic but is learning MAC addresses.
- **Forwarding**: The port is forwarding traffic and learning MAC addresses.

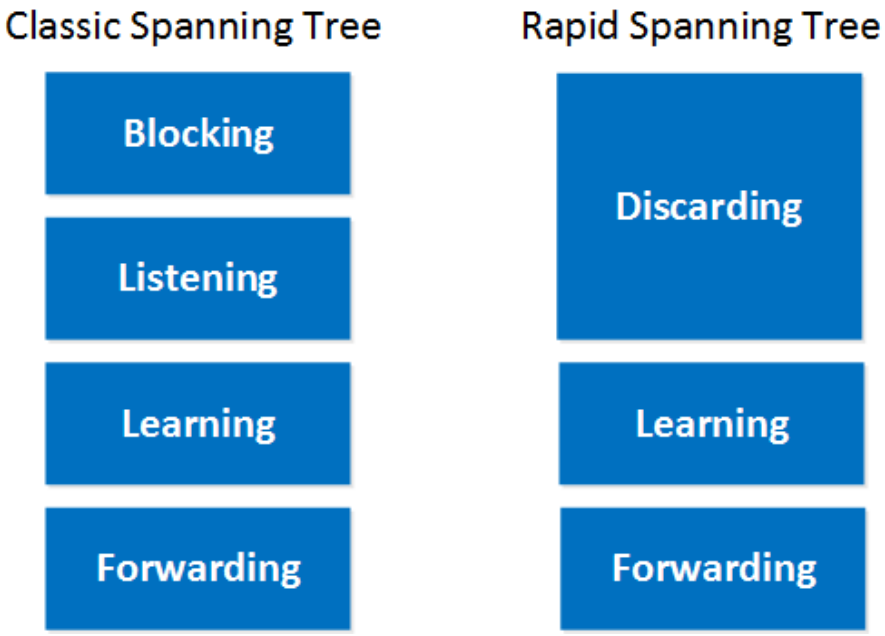

RSTP uses these port states to manage the network topology and prevent loops.

Advantages of RSTP

- **Faster convergence**: RSTP converges significantly faster than STP, typically in 5-10 seconds compared to 40-50 seconds for STP. This is important for modern networks where even brief outages can be disruptive.
- **Improved network performance**: RSTP maximizes bandwidth utilization and minimizes packet loss by quickly transitioning ports to the forwarding state and selecting optimal paths for traffic forwarding.
- **Simple implementation and configuration**: RSTP is easy to implement and configure, and it maintains a similar logical topology to STP. This makes it a good choice for networks that are already using STP.
- **Scalability**: RSTP supports networks of all sizes, multiple VLANs, and multiple spanning trees (MST) for load balancing and performance improvement.

Disadvantages of RSTP

- **Increased resource consumption**: RSTP requires more processing power and memory than STP. This may be a concern for older or less powerful network switches.
- **Potential for suboptimal path selection**: RSTP's fast convergence may lead to less optimal path choices in certain scenarios. This can be mitigated by using MST or by tuning the RSTP parameters.
- **Dependency on network hardware**: RSTP's effectiveness relies on the capabilities and implementation of the underlying network equipment. It is important to ensure that all switches in the network support RSTP and are configured correctly.

Overall, RSTP is a significant improvement over STP in terms of convergence speed, performance, and scalability. However, it is important to be aware of the increased resource consumption and potential for suboptimal path selection. RSTP is a good choice for most networks, but it is important to evaluate the specific needs of the network before deploying it.

Rapid Spanning Tree Protocol (RSTP) can be implemented on Cisco switches on a per-VLAN basis in the form of Rapid PVST+. This means that each VLAN can have its own spanning tree instance, which can improve performance and scalability.