

Systems Integration Track: Network Topology

Author Jim W. Hammond – Technical Consultant, GarrettCom


Keywords: Ring, Topology, RSTP, Ethernet, Network, VLAN

Industrial networks, like most large networks, are a combination of topologies including mesh, star, and ring. At the edges of a network, star and ring topologies dominate, but the choices of how to manage these sub networks are critical when redundancy, self-healing, and security are factored in. In addition, the choice of network components and their bundled software can enhance or degrade performance.

There have been many advances in industrial Ethernet at the edge of a network that can greatly improve redundancy and availability, and also provide support for not only sensors and intelligent devices, but also video monitors, card and badge readers, and other devices that promote secure installations even in remote locations. The choices of the best topology and the best components to provide this functionality are the purpose of this paper. The first part of the paper discusses topologies and the development of redundancy to support the higher availability demanded today. The final part provides the payoff in application examples.

TOPOLOGY: AN OVERVIEW

The three basic topologies that dominate LAN network designs are *mesh*, *ring*, and *star*. *Tree* topology or branching tree is often formed by placing routers at network segments, or using two-port bridges to decide which branch to take. Each has certain advantages based on how the network is used,



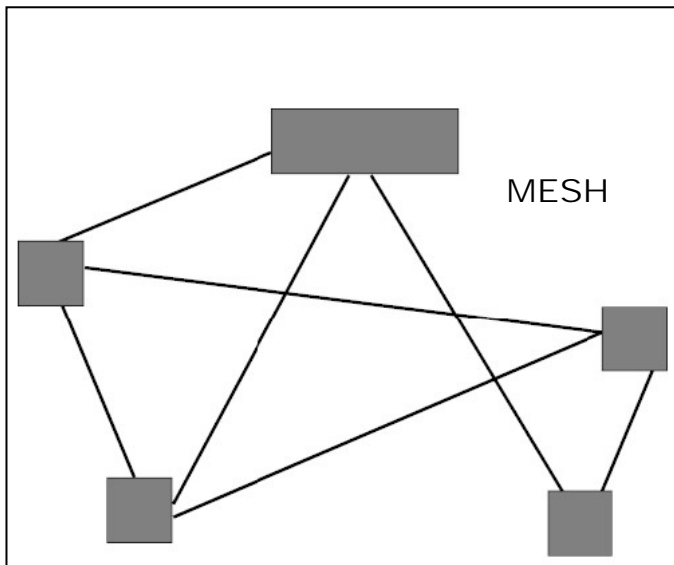
redundancy and reliability issues, and security requirements. Larger networks, interconnected networks, and wide area networks (WAN) are often a combination of a mesh topology with star and ring subnets. Since few networks have evolved from a homogeneous design, integration is one of the key features of today's networks.

The original ring topology used a hub to pass data from the first connected device to the next connected port with the hub's internal wiring providing the topology. Now Ethernet switches are used to join a series of daisy-chained devices together. This can provide an alternate path through the ring as will be

discussed later. Daisy-chaining is an example of a single-ended or “broken” ring. This is analogous to a line topology which is a series of point-to-point links.

WAN stars are obsolete, but LAN stars are numerous. Typically, a hub provides the interconnections, replaced with switches when higher availability and bandwidth are needed. Ethernet evolved as the protocol of choice replacing Token Bus and proprietary protocols as the benefits of higher bandwidth, lower media and component costs, and full duplex Ethernet became widely recognized.

In contrast to a ring hub, a star hub views each device as connected together as on a bus with data sent to all active devices simultaneously. The physical view of both types of hub appears to be similar; it is the internal wiring that is different. The early rings used Token Ring protocols while stars used Token Bus or Ethernet, but again Ethernet’s lower cost and higher bandwidths began to dominate. Ethernet switches can support either topology.



The mesh topology includes the fully-connected mesh, and the much more common partially-connected mesh. Since there are multiple routes between nodes, routing tables are maintained to ensure the fastest route is chosen, and to calculate an alternate route when the primary route goes down. TCP/IP is now the standard for routing and control of WAN/LAN networks with Ethernet the layer 2 (L2) protocol of choice. However, legacy equipment still abounds and protocol conversion, most often provided by routers is required to fully control access to all nodes and devices.

That brings us to the edge of the network where all these topologies may exist, but where routing is generally simpler, and control, uptime, and security are primary factors. This is an area of networking where careful design, the selection of the topology that will best serve the remote devices, and the right protocol set will insure the highest availability and security of connections. The flexibility of Ethernet, the advantages of a common protocol set, the elimination of protocol conversions, and Ethernet’s continual enhancement have made it the choice in the vast majority of LANs. As legacy equipment is upgraded or replaced, Ethernet becomes the logical choice.

INDUSTRIAL ETHERNET

The integration of existing sub-networks into a homogeneous network that supports effective routing, redundant links, and beefed up security can now call upon multiple vendor components and protocol enhancements to more effectively manage and monitor industrial applications. Industrial Ethernet has emerged as the leading protocol for supporting an integrated, reliable, and secure

network. Thus, Ethernet managed and un-managed switches, including edge switches, coupled with the right topology provide the best solution to the control and support of remote sites such as power substations and unattended sites.

Industrial Ethernet networking has inherent advantages. By utilizing a standards-based solution that supports multi-vendor implementations, Industrial Ethernet users enjoy highly reliable systems with rapid ring recoveries, reduced costs of deployment, and a guaranteed upgrade strategy as needs evolve. Ethernet makes possible redundant and self-healing networks for 24/7 uptime.

PROVIDING REDUNDANCY

One of the reasons Ethernet is the preferred protocol for redundant industrial applications is the plentiful supply of industrial-grade switches and hubs running at 10/100 Mb/sec and higher speeds that provide more than adequate bandwidth.

Redundancy in a mesh topology requires multiple communications links to each node so an alternate path can be utilized if the primary link goes down. In WANs, this represents a costly investment in media, routers, and transmission rates. In a LAN environment, the media costs and the physical location of end nodes or devices for providing redundant links can outweigh the benefits of redundant mesh networks.

Where many devices are co-located, a star topology, created using a hub or, preferably, a switch can reduce costs while providing high bandwidth, however, it also represents a single point of failure. Some Ethernet edge switches feature dual-homing whereby two connections, a standard and backup port, provide redundant paths. However, as distances between devices increase, a ring topology is generally a better choice.

Ring subnets use a "daisy-chain" or sequential point-to-point topology which is optimal for minimizing the cabling expenses that dominate overall installation cost. In most cases, routing the end of the cable string back to the switch that manages the daisy-chained units is fairly easy. Ring structures easily provide redundant capabilities when spanning tree protocols are used.

THE SPANNING TREE PROTOCOLS

The IEEE 802.1d standard Spanning Tree Protocol (STP) was the original standard for Ethernet fault recovery. STP resolved redundant physical connections to maintain the operation of Ethernet LANs that only allow one path for a packet transfer at a given time. The Spanning Tree Protocol is now being replaced with Rapid Spanning Tree Protocol (RSTP).

STP was deemed too slow for most industrial applications. Some vendors offer proprietary alternatives, however RSTP is considerably faster than STP and has become widely accepted. Redundant LAN configurations can be constructed in a variety of topologies, and RSTP is designed to work with all of them. While mesh configurations are a more general topological case, ring

configurations for redundancy are especially useful and cost-effective in industrial LAN systems, and will be treated in more detail. Alternatives or enhancements to RSTP will also be discussed in a general way.

RAPID SPANNING TREE PROTOCOL (RSTP)

Rapid Spanning Tree Protocol (RSTP) is defined in IEEE 802.1w and like STP, was designed to avoid loops in an Ethernet network. RSTP provides for faster spanning tree convergence after a topology change, as quickly as one second, instead of the 30 second timeout of STP.

Rapid Spanning Tree replaces the STP settling period with an active handshake between switches (bridges) that guarantee topology information will be rapidly propagated through the network. RSTP converges in less than one second to six seconds. RSTP offers a number of other significant innovations. These include:

- Topology changes in STP must be passed to the root bridge before they can be propagated to the network. Topology changes in RSTP can be originated from and acted upon by any designated switch (bridge), leading to more rapid propagation of address information.
- STP recognizes one state - blocking for ports that should not forward any data or information. RSTP explicitly recognizes two states or blocking roles - alternate and backup port, including them in computations of when to learn and forward and when to block.
- STP relays configuration messages received on the root port going out of its designated ports. If an STP switch (bridge) fails to receive a message from its neighbor it cannot be sure where along the path to the root a failure occurred. RSTP switches (bridges) generate their own configuration messages, even if they fail to receive one from the root bridge. This leads to quicker failure detection.
- RSTP offers edge port recognition, allowing ports at the edge of the network to forward frames immediately after activation while at the same time protecting them against loops.
- RSTP allows configuration messages to age more quickly speeding up recursive route discovery.

RSTP has three operational states: discarding, learning and forwarding.

When a port is first enabled it is in the *discarding* state. The port does not learn addresses in this state and does not participate in frame transfer. The port looks for STP traffic in order to determine its role in the network. When it is determined that the port will play an active role in the network, the state will change to *learning*. The port learns addresses in this state but does not participate in frame transfer.

In a network of RSTP switches (bridges) the time spent in the learning state is usually quite short. After ‘learning’ the switch will place the port in the *forwarding* state. While in this state the port learns addresses and participates in frame transfer.

The older Spanning Tree Protocol takes a fairly long time to resolve all the possible paths and to select the most efficient path through the network. The Rapid Spanning Tree Protocol significantly reduces the amount of time it takes to establish the network path. The result is reduced network downtime and improved network robustness. In addition to faster network reconfiguration, RSTP also implements greater ranges for port path costs to accommodate the higher connection speeds that are being implemented.

VIRTUAL LAN (VLAN)

VLANs are widely used today for both reducing broadcast traffic and providing an added level of security by limiting the size of a shared-traffic domain and managing cross-domain communications. Since crossing a domain involves a routing decision, the security of a given domain can be assured. A VLAN creates separate network segments that can span multiple Ethernet switches. A VLAN is a group of ports designated by the switch as belonging to the same broadcast domain. The IEEE 802.1Q specification establishes a standard method for inserting VLAN membership information into Ethernet frames.

VLANs provide the capability of defining two or more Ethernet segments that co-exist on common hardware. The reason for creating multiple segments in Ethernet is to isolate shared-traffic domains. VLANs can isolate groups of users, or divide up traffic for security, bandwidth management, etc. VLANs need not be in one physical location. They can be spread across geography or topology.

A group of network users (ports) assigned to a VLAN form a broadcast domain. Packets are forwarded only between ports that are designated for the same VLAN. Cross-domain broadcast traffic in the switch is eliminated and bandwidth is saved by not allowing packets to flood to all ports. For many reasons a port may be configured to belong to multiple domains.

Since VLANs are entirely separate segments or traffic domains – they require a routing protocol (e.g., a router or an L3-switch). The routing function can be done internally within an Ethernet switch if it includes the layer 3 protocol. One advantage of an L3-featured Ethernet switch is that it can also support multiple VLANs. The L3 switch can thus route traffic across multiple VLANs easily and provides a cost effective solution if there are multiple VLANs configured.

VLAN ports are defined as *port* VLAN or *tag* VLAN. This functionality is assigned to each port on an Ethernet switch that provides VLAN support. The term “port VLAN” is specific to a switch which assigns a port or group of ports as belonging to a VLAN. When ports belong to multiple VLANs and are specific to a switch, they are defined as port VLANs as well. Port VLANs do not recognize or manipulate the VLAN identifier (VID) information included in an Ethernet frame. The port works “transparently” and propagates any VLAN information.

In tag VLAN, the VLAN identifier (VID) is either inserted or manipulated. This manipulated VLAN tag allows VLAN information to be propagated across devices or switches, allowing VLAN information to span multiple switches.

VLANs, as the name suggests, create virtual LANs administratively. Instead of going to the wiring closet to move a cable to a different LAN segment, the same task can be accomplished remotely by configuring a port on an 802.1Q-compliant switch to belong to a different VLAN. The ability to move end stations to different broadcast domains by setting membership profiles for each port on centrally managed switches is one of the main advantages of 802.1Q VLANs.

802.1Q VLANs aren't limited to one switch. VLANs can span many switches. Sharing VLANs between switches is achieved by inserting a tag with a VLAN identifier (VID) into each frame. A VID must be assigned for each VLAN. By assigning the same VID to VLANs on many switches, one or more VLAN (broadcast domain) can be extended across a large network.

802.1Q-compliant switch ports can be configured to transmit tagged or untagged frames. A tag field containing VLAN information is inserted into an Ethernet frame. If a port has an 802.1Q-compliant device attached (such as another switch), these tagged frames can carry VLAN membership information between switches, thus letting a VLAN span multiple switches. Connections between switches can carry multiple VLAN information, which is called port trunking or 802.1Q trunks.

Private VLANs are private to a given switch in a network, and are usually restricted to a single switch. Private VLANs are constructed for security reasons. For example, if some confidential data were residing on VLAN 5, then only the users connected to the switch on VLAN 5 can access that information. No one else can access that VLAN. Similarly, if another switch had video surveillance equipment on VLAN 20 then only ports with access to VLAN 20 can access the video surveillance information.

SUMMARY

VLANs provide great flexibility in controlling what traffic goes where, either to limit traffic where it isn't wanted (think about all the junk mail you receive), or to avoid sensitive data from reaching the wrong people (think about that crabby Email you sent to the wrong person). VLANs utilize a function of Ethernet switches, reading addresses before sending a frame out, to insure a faster more responsive network with better security and robustness.

RING TOPOLOGY: THE SOLUTION FOR THE EDGE OF A NETWORK

Ring topologies are very effective at the edges of a network. The daisy-chain connections of ring topology significantly reduce wiring costs. Since signals are reshaped and retransmitted from one device to another, transmission errors are greatly minimized. Operating with Ethernet switches that support Rapid Spanning Tree Protocol (RSTP), they provide redundancy and self-healing functionality to remote sites.

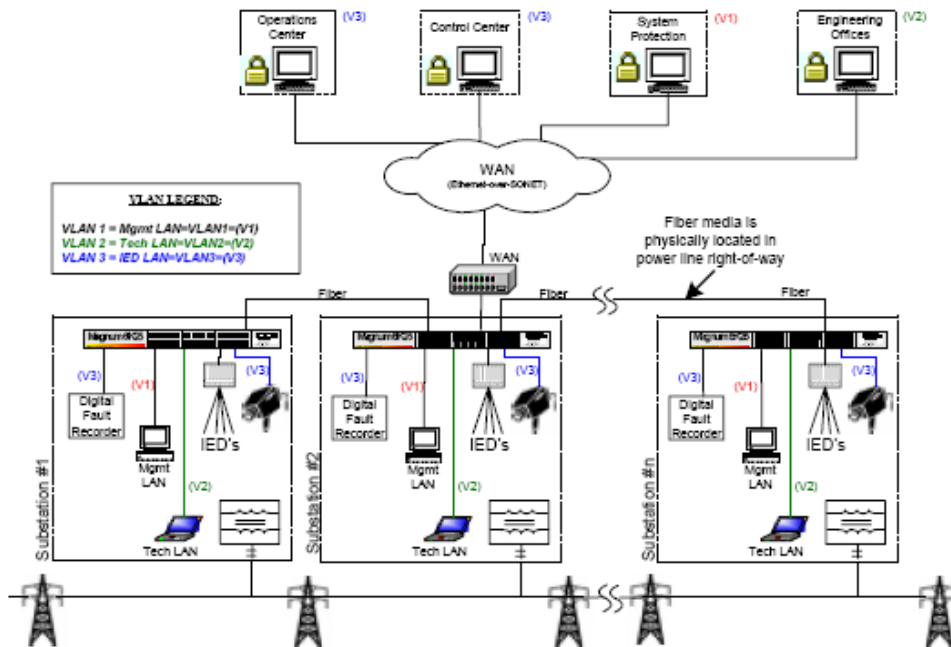
A few vendors also include VLAN support, which provides additional security and limits unwanted traffic, improving the responsiveness of remote devices.

Ethernet switches that support Power over Ethernet (PoE) simplify adding devices such as badge readers, security cameras and other monitoring equipment, which can be critical to sensitive sites such as remote power stations and water and power distribution systems. In addition, PoE can greatly simplify the wiring of the many sensors, monitors, and input devices found in industrial Ethernet environments, while reducing cabling costs.

CASE STUDIES

VLANS MAKE SECURE REMOTE MANAGEMENT POSSIBLE

VELCO is a transmission-only power company that interconnects all Vermont Electric utilities. Its large territory and the need for rapid, secure exchange of information among the substations and operations areas caused the company to consider a broadband standards-based solution to its communications needs. A secure management system over Ethernet for strictly internal device management was critical.



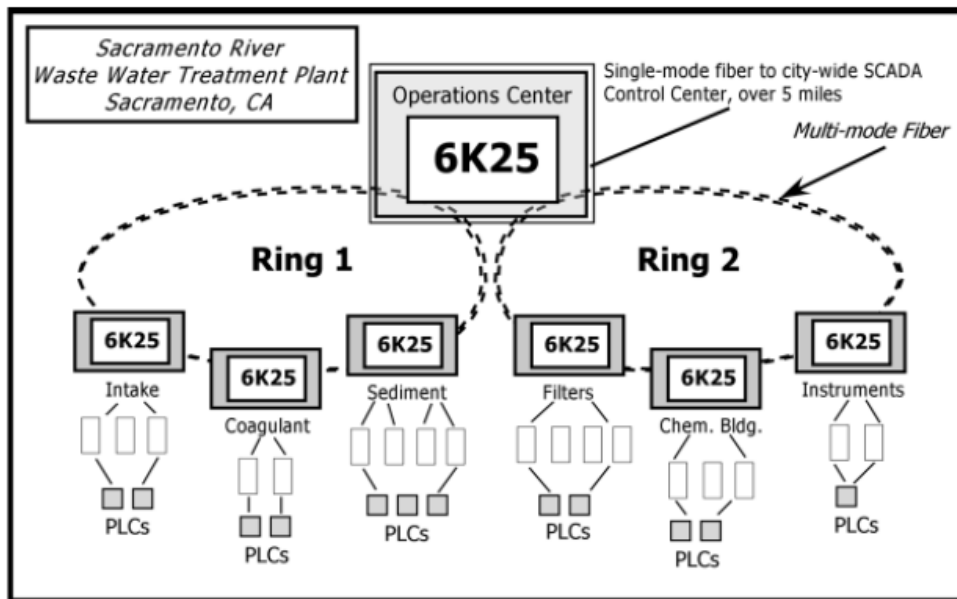
High-speed access, however, is not without its concerns. The accessibility of information over long distances means that communications are vulnerable to eavesdropping and attack. The company needed not only a cost-effective high speed connection, but also a secure one.

Ethernet managed switches employing fiber interfaces are installed in substations throughout the service area, and are tied into the communications network through a WAN port dedicated to

substation automation. The switches come with a powerful, highly secure and robust remote access management capability that was needed to protect its resources and customers. The switches provide encrypted authentication and access over Ethernet from remote locations. VLANs were implemented to enhance security and minimize unnecessary traffic.

UPGRADING TO REDUNDANT ETHERNET

The Sacramento River Water Treatment Plants account for 85 percent of Sacramento's water supply. Because the current intake structure was out of compliance with current environmental standards, the decision was made to update the networking system as well. In developing the upgrade plan, the city wanted to use state-of-the-art technology that ensured a platform from which future upgrades could be expected to be integrated easily.



The new installations are based on a redundant ring topology, and utilize fiber media to support the distances required in managing the various SCADA (Supervisory Control and Data Acquisition) systems located throughout the city. The EMI noise immunity of fiber cabling increases operating reliability of the network. The installation uses the less expensive copper cabling for some local devices and computers within control system racks. Cost was an issue as the city wanted to make its funds stretch as far as possible in upgrading the system.

The contractor recommended Ethernet managed switches utilizing SNMP Network Management Software. The switch features a unique modular port structure that enabled the customer to define on a port-by-port basis the medium to use, providing immediate cost savings because of the efficiency of the units, and eliminating the need for media converters.

To ensure network reliability in the ring portion of the network, a sophisticated-ring redundancy manager, based on the industry standard 802.1d Spanning Tree Protocol, provided support for dual

rings (see illustration above). As an added bonus, a management feature of the switches allowed the city to monitor the health of the SCADA network as well as the health of the network system at the treatment plants without the necessity of full-time networking monitoring professionals at each site. Immediate feedback when the monitoring system itself is in trouble can avert shutdowns and reduce the costs of managing and maintaining the networks.

Where proprietary serial lines are expensive and can max out in the multi-Kilobit (Kb) range, Ethernet protocols support 10 Mb to 100 Mb per second, with migration paths to the 1 Gigabit (Gb) range and beyond. Cost benefits are equally impressive. Proprietary serial network interfaces average 40 times the cost of an Ethernet interface while providing 1/1000th of the bandwidth. Proprietary media converters for serial lines cost more than 10 times the price of an Ethernet media converter – and the need for them is eliminated by the switch’s modular port structure and built-in fiber ports.

With the Ethernet network in place, the city will be able to easily implement technology advances such as video surveillance over the network with a minimum of effort and cost. While not in the current plans, the ability to implement these features in the future gives the city the ability to expand the remote monitoring and management of sites with a reduced staff.

BIBLIOGRAPHY

IEEE 802.1d and IEEE 802.1w Standards

Networking as a 2nd Language; Understanding Spanning Tree Protocol -- the Fundamental Bridging Algorithm, Michael Norton, O'Reilly Network, 03/30/01

Achieving Fault-Tolerance with PC-Based Control, David W. Cawlfeld, ISA Automation & Control Subsystems Committee

IEEE Standards for Local and Metropolitan Area Networks: Draft Standard for Virtual Bridged Local Area Networks, P802.1q-rev (D4) 2005.

“Redundancy with Standards in Industrial Ethernet LANs”, Frank Madren, RTC Magazine, October 2003, <http://www.rtcmagazine.com/home/article.php?id=100156>

Vermont Electric Power Builds Secure Network Using Ethernet: An Industrial Ethernet Application. http://www.garrettcom.com/techsupport/appnotes/velco_appnote.pdf

Sacramento Utilities Department Chooses Magnum 6K25 Switches and S-Ring Redundancy Manager for Water Treatment Plant: An Industrial Ethernet Application. http://www.garrettcom.com/techsupport/appnotes/sacto_water_treatment.pdf