

# REDUNDANCY IN SUBSTATION LANs WITH A RAPID SPANNING TREE PROTOCOL

By Michael Galea, Marzio Pozzuoli, RuggedCom Inc. - Industrial Strength Networks

## INTRODUCTION

Ethernet local area networks (LANs) are steadily gaining more acceptance in substation automation applications where the LAN has become an integral part of the protection and control system. Most protective relaying manufacturers now offer Ethernet ports on their relays in both fiber optical and copper media.

Information models and exchange methods currently being defined as part of the IEC 61850 standard allow for real-time control (e.g. Trip/Block) messages to be sent across the LAN between relays or other intelligent electronic devices (IEDs).

Furthermore, there is a trend towards multicasting sampled data of current and

voltage parameters over a 100Mbps Ethernet LAN as defined in IEC 61850-9-2 (Process Bus). With Ethernet LANs playing such a critical role in protection and control systems, new standards such as IEEE 802.1W Rapid Spanning Tree Protocol which are used to implement network redundancy and ring architectures are of critical importance.

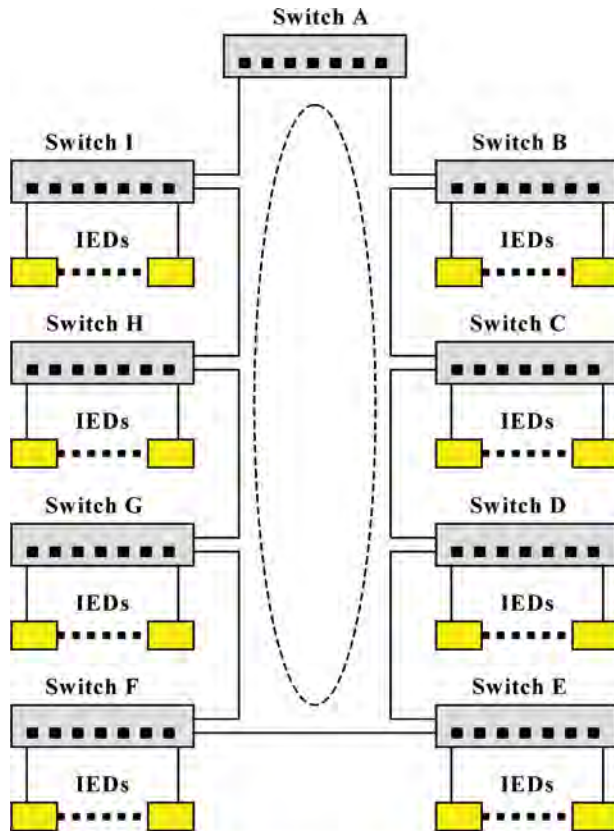


Figure 1: Simple Ring Architecture

nation MAC address is unknown, the switch will "flood" the frame out all of its ports.

If switches in the network are connected in a loop a 'broadcast storm' will result where a single broadcast frame will circulate endlessly. This condition consumes all available bandwidth on the loop, making the network unusable. The Spanning Tree Protocol is used to prevent this situation.

## BRIEF HISTORY OF SPANNING TREE PROTOCOL (STP) AND RAPID STP (RSTP)

The Spanning Tree Protocol (IEEE 802.1D) was designed to solve the fundamental problem of traffic loops. The key idea in STP is to prune (looping) links in order to reduce the network topology to

that of a tree. The resulting tree "spans" (i.e. connects) all switches, but eliminates loops. The steps in order to best accomplish this process are:

1. Allow all switches to send messages to each other that convey their identity and link "cost".

2. Elect a single switch among all the switches in the network to be a "root", or central switch.

3. Let all other switches calculate the direction and cost of the shortest path back to the root using messages received from switches closer to the root.

Each switch must have only one "best" way to forward frames to the root.

4. If two switches servicing the same LAN exchange messages with each other, the one with the lowest cost to the root will service the LAN. The other switch will discard all frames received from that LAN, thus opening the link and blocking a traffic loop.

The STP protocol has proved to be the tried and tested method for providing path redundancy while eliminating loops. The STP protocol does suffer from a number of drawbacks that limit its applicability, namely:

- STP has lengthy failover and recovery times. When a link fails in STP, a backup link to the root requires at least 30 seconds to recognize that it is the best (or only) path to the root and become usable.

- When a failed link returns to service, information about the "better" route will instantly cause a backup link to start blocking. But the portion of the network below the link that is returning to service will be isolated (for about 30 seconds) until that link becomes forwarding.


- Another problem with STP is that it requires that all links must pass through a lengthy period of address learning, even if the link is a point-to-point link to a device such as an IED (e.g. Relay, RTU, PLC).

## BACKGROUND


Ethernet switches operate by forwarding traffic between their ports. The switch examines each Ethernet frame and records (learns) its MAC address and the port upon which it resides. When a frame arrives for a given MAC address, the switch "knows" on which outgoing port to send it. If a frame arrives and its desti-

Continued on Page 26

# Finpoint's 14th Annual Circuit Breaker Test & Maintenance Training Conference



Join Us In Pittsburgh This October!

- 
- A day at the Pennsylvania Breaker and Pennsylvania Transformer Technology facilities
  - Over a dozen useful substation and switchgear presentations
  - An SF6 handling, management, and emissions reduction seminar
  - A circuit breaker timing and analysis seminar
  - Supplier exhibits every evening in the hospitality rooms

**October 1-5, 2007**

**Omni William Penn Hotel  
Pittsburgh, Pennsylvania**

For further information and on-line registration, please visit  
[www.circuitbreakerconference.com](http://www.circuitbreakerconference.com)

### ENTER RSTP (IEEE 802.1W)

RSTP solves STP's problem with failover time by a number of means. Whereas STP switches store only the best path to the root switch, RSTP switches store all potential paths. When links fail, RSTP has pre-calculated routes to fall back upon. Additionally, unlike STP switches, an RSTP switch will respond to another switch that advertises an inferior or incorrect route to the root switch. This information allows the switch with incorrect information to be rapidly trained.

RSTP solves STP's problem with lengthy recovery time by introducing a new procedure called proposing-agreeing. Proposing and agreeing works after a better path to the root is restored by "shuffling" the restored part of the network one hop at a time towards the network edge. This method also enables the network to come up quickly at inception.

RSTP also introduces a method for quickly bringing up ports at the edge of the network, while still protecting them against loops. If the port is designated as an "edge" type of port, RSTP will continue to send configuration messages out the port (in order to detect loops) but will allow traffic to flow as soon as the port rises. In the event of a loop, some looped traffic may flow before RSTP quickly seals the network. PCs, IEDs and RTUs connected via edge ports can send traffic without the extensive delays imposed by RSTP.

### RINGS

A ring topology offers built-in redundancy and is often the most economical in terms of interconnection costs. Two popular methods of implementing rings are collapsed backbone and distributed switch. The distributed switch method, or simple ring (See Figure 1), is employed when network connected IEDs are geographically distributed.

The IEDs at each location are aggregated onto switches, which are organized into a ring. The connections between switches in the ring may be made using dual redundant links to obviate the possibility of failure at a fiber, connector or port level. The collapsed backbone method (See Figure 2) is usually employed when a large number of network connected IEDs are located in close proximity to one another. The IEDs are aggregated onto switches, the switches organized into a number of rings and all rings terminated in a common root node.




Figure 3: Tree Architecture

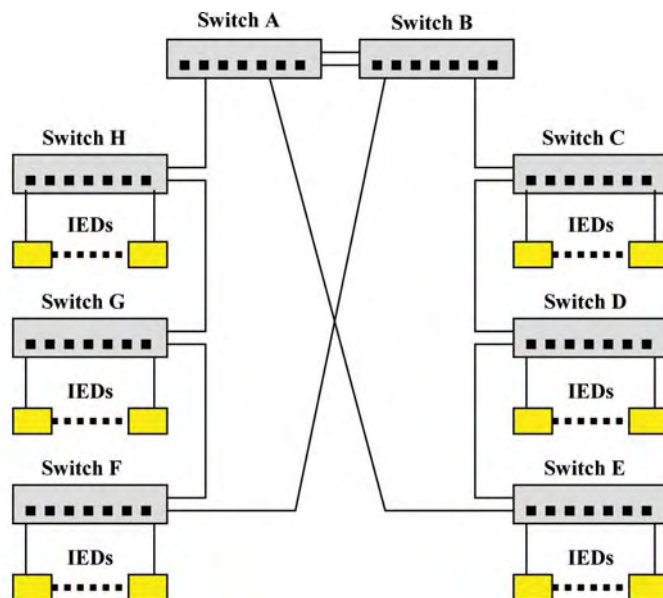


Figure 2: Collapsed Backbone Ring Architecture

Quite often, the network topology is a mixture of both methods, such as a ring of rings. Traffic in a ring tends to be balanced. The ring will open itself with an equal number of switches on either side of the root switch given an odd number of switches in the ring.

Latency in ring networks tends to be greater than in tree networks (See Figure 3) as there are usually more hops to pass through in order to go anywhere useful.

The worst case occurs when switches on either end of the blocked link at the "bottom" of the ring need to forward to each other. In this case, traffic must flow through every switch in the ring.

Ring networks offer only slightly slower failover and recovery times than tree networks. The worst case link failure in ring networks occurs on a port at the root. In this failure case, half of the switches in the ring must retrain themselves to face their root port in a completely opposite direction after a link failure or recovery. The other half of the network must reverse the direction of transmission to switches in the failing half.

The size of the ring is, in theory, limited by the RSTP switch diameter which assumes a pessimistic transit delay of one second per switch. In practice, the maximum number of switches in an optimized ring occurs when the number of priority switch levels has been exhausted. This limits the size of the ring to 31 switches. Rings of more than 31 switches are still possible but will failover and recover in a slower fashion.

### FAILOVER AND RECOVERY PERFORMANCE IN RINGS

Figure 4 presents a network of nine RuggedCom RuggedSwitch Ethernet switches organized in a ring topology. The figure details the sequence of steps to heal the ring after the link between switches A and B fails.

Initially, switch B has information