© iStockphoto.com/Fernando Alonso Herrero

REDUNDANCY IN INDUSTRIAL NETWORKS

PART 1

Glenn Johnson, Editor

The costs of failure in today's industrial networks can be very high, making network redundancy essential.

The failure of individual components in factory automation, processing and substation applications are inevitable and can never be totally avoided - so they have to be handled in a way that minimises impact on the system. While high availability can be achieved by using completely redundant systems, such as duplicated sensors, actuators, controllers and networks, it is usually far too expensive a proposition to implement.

Some parts of the system can provide resilience, however, and one such element is the networking component. While many technologies used in plants are designed to be robust, networking components are wholly electronic and rely on cables and wireless links, all of which can be more easily damaged or interrupted in some way, so the capability to design a 'self-healing' network is important.

With the increasing use of ethernet as a communication technology in plants and factories, it is possible to take advantage of ethernet redundancy technologies to provide a fault-tolerant network. Most ethernet switches and routers today support various types of redundancy mechanisms that only require some additional cabling and software configuration to implement, and which provide a standby and failover mechanism to secondary network paths.

Network redundancy can be achieved at both the data link layer (Layer 2) and the network layer (Layer 3), with Layer 2 redundancy being provided by switches within a TCP/IP subnet, and Layer 3 redundancy generally being provided by routers, routing traffic via different TCP/IP subnets. Naturally, routing means higher overhead and lower performance, so in this article we will focus only on standardised Layer 2 redundancy techniques. This is not to say that Layer 3 redundancy is not useful in industrial networks in appropriate situations, but this article will focus mainly on redundancy within a single network in which high performance recovery is a must.

But there are choices to be made - differing redundancy protocols and designs will provide different levels of protection and performance. So it is necessary to understand the differences to determine what is sufficient for the particular application. For example, can the process tolerate a delay of a few seconds while the network redundancy 'heals' a fault, or is millisecond response required? Some ethernet hardware may support different redundancy technologies, so choosing the right technology to support your needs is important - as is the architecture of the network as a whole if you want to successfully implement a fast failover capability.

Ethernet does not tolerate loops

It is a basic requirement of a functioning ethernet network that there are not any loops. Loops result in data frames circulating endlessly, flooding the network. So all

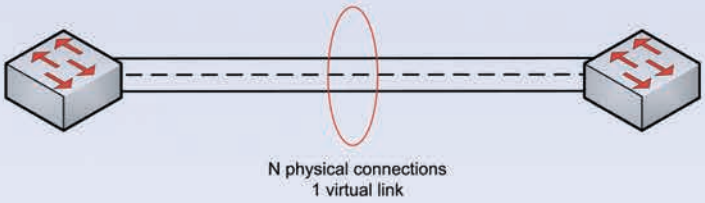

Figure 1: Link aggregation allows links between switches to be bundled to increase bandwidth. Redundancy is improved if the links have different physical paths.

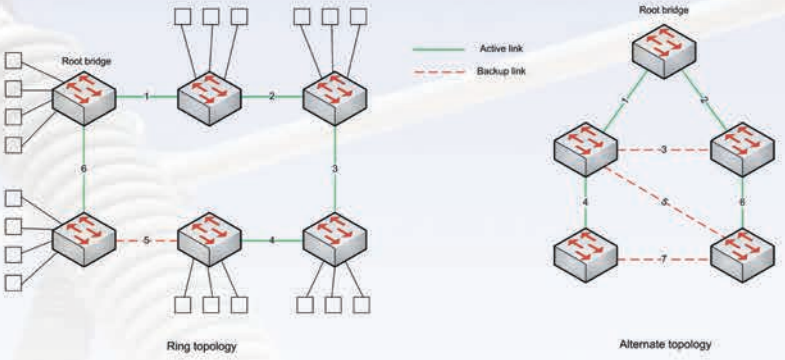

(left) Figure 2: Spanning tree protocols (STP and RSTP) create a tree of connections between switches, disabling connections that would form loops.

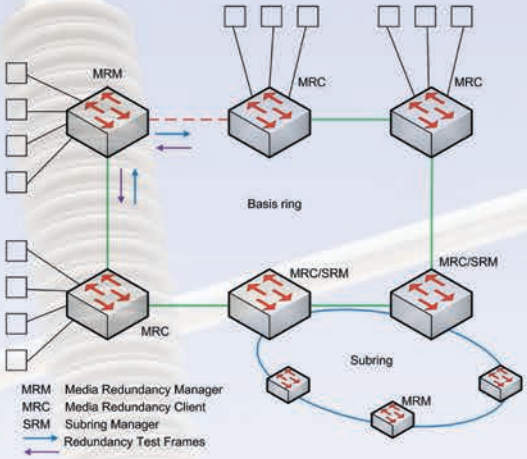

Figure 3: In MRP, switches react to received reconfiguration frames from the Media Redundancy Manager (MRM). Subrings are also supported through Subring Managers (SRMs).

ethernet networks need to be implemented to make sure there is only a single path between any two devices.

For redundancy, however, there must be an alternative path available, in case the primary path becomes unavailable. For this to work, it must be possible to have multiple physical paths between devices, but to make sure that only one path is active at any one time.

The main way this has been achieved is through monitoring the communication paths, detecting failures and switching to the backup path if the main communication path fails. There are several protocols that can achieve this functionality, but they vary in their performance. All changeover mechanisms of this type depend on detecting the fault, then reconfiguring the network to a new topology (alternate paths) to re-establish communication - and these steps all take time. The protocols available on the market can differ greatly in their failover speed, which is in turn also affected by the size and design of the network.

Link aggregation

A simple form of redundancy is link aggregation, or link redundancy (Figure 1). Link Aggregation Control Protocol (IEEE 802.1ad) provides the ability to bundle groups of switch ports between switches to form one link with the aggregated bandwidth of the individual links. In the event that a single connection fails, the remaining links keep working with reduced bandwidth. To best take advantage of link redundancy, it is most effective if the physical links (cables) are routed via different paths, to minimise the risk of multiple link failures.

Spanning trees

One of the first protocols developed to implement redundancy was the Spanning Tree Protocol (STP) that was developed in the early 1990s. Designed for failover in IT networks, the failover time for this protocol can be as long as 10 seconds, but can handle different network topologies, including mesh networks. Apart from the slow failover time, it also has a limitation in the number of switches between endpoints in the network, due to the time required to converge on a new configuration. Although larger networks can be built, depending on the topology, the original RFC for STP recommended that the number of hops (the number of bridges or switches between any two devices) should be no more than seven.

Spanning tree protocols work by creating a tree of connections between switches and by disabling all the connections that are not part of the tree (and that would form loops), as shown in Figure 2. Special frames called Bridge Protocol Data Units (BPDUs) are used to communicate between switches and to set up optimal paths in the network, with one switch defined as the 'root bridge' for the tree (by default the switch with the lowest MAC address, but can be manually defined). When the topology changes, Topology Change Notification BPDUs are used to announce the change, resulting in a recalculation of the spanning tree, and the activation of backup paths to re-establish the network.

STP has generally been replaced by the Rapid Spanning Tree Protocol (RSTP), an improved version of STP that was defined by the IEEE 802.1 working group in 1998. RSTP networks support a larger number of switches (20 in a path) and the typical failover time is around one second. Regardless of the failover time, however, neither STP nor RSTP can provide *deterministic* failover. The failover time will vary depending on the particular implemented topology and the location of the individual failure. Restricting RSTP to simple ring networks and with careful configuration, it has been shown to be possible to keep failover times down to around 100 ms, however.

The main benefit of a spanning tree protocol is that depending on the design, it is possible to design a network that is resilient to more than one simultaneous link failure. For example, the loop configuration in Figure 2 can recover completely from only one failure (a weakness of loop topologies). If a second link were to fail (both links 3 and 5), then a switch or even a whole section of the loop would be isolated. In the alternative configuration of Figure 2 (a partial mesh), there are multiple backup links, and this allows, for example, two links to fail (such as links 4 and 6), and the network should reconfigure to allow the network to keep working (assuming in the example that the two failures disconnected

Company	Technology	Claimed recovery time	Network size	Topology	URL
Advantech	X-Ring	<10 ms	30	Ring	www.advantech.com
Hirschmann	HiPER Ring	<500 ms	200	Ring	www.belden.com
	Fast HiPER Ring	<60 ms	200	Ring	
Moxa	Turbo Ring	<20 ms	250	Ring	www.moxa.com
N-Tron	N-Ring	~30 ms	250	Ring	www.n-tron.com
ORing	O-RSTP	~20 ms	40	Any	www.oring-networking.com
	O-Ring	<10 ms	250	Ring	
	Open-Ring	Variable	250	Ring	
Rockwell Automation	Cisco REP	20-250 ms	Unknown*	Line/Ring	www.rockwellautomation.com
Weidmüller	Turbo Ring	<20 ms	Unknown*	Ring	www.weidmuller.com.au
Westermo	Cisco REP	<20 ms	200	Ring	www.westermo.com

*Table 1: Examples of proprietary ethernet redundancy offerings. (*Some data unavailable to the author at the time of writing.)*

both ports of a single switch, which would effectively isolate the switch - such as links 4 and 7).

The disadvantage of spanning tree protocols is that while, with careful design, the recovery time can potentially be low, it is also not predictable. The recovery time will depend on the topology, the location of the failure and the number of failures that occur - and the larger the number of switches, the more the recovery time increases.

Media Redundancy Protocol

STP and RSTP are enterprise network protocols supported in all managed ethernet switches. A protocol commonly found in industrial ethernet switches that is designed more for industrial applications is Media Redundancy Protocol (MRP). It is defined in IEC 62439 as an industry standard for high-availability networks and is a standardised version of the HIPER-Ring protocol first released by Hirschmann and Siemens in 1999. It is exclusively for ring networks, but can guarantee deterministic ring failover.

The reason that MRP can have a predetermined recovery time is that it is not a protocol in which all the switches need to reconfigure their forwarding ports hop-by-hop and 'converge' to a new topology, as in Spanning Tree protocols. Instead, one of the switches is configured in the role of Media Redundancy Manager (MRM), which sends frames out of one of its ring ports and receives them on its other ring port, in both directions, while maintaining one port closed to normal data. All other switches act as Media Redundancy Clients (MRCs), and can act on configuration frames received from the MRM, as well as detect and signal link changes on their ring ports (Figure 3).

With MRP, the failover time is nearly independent of the number of switches in the ring, because MRP control frames are forwarded as multicast frames through the ring, and so can be processed nearly simultaneously in all switches, resulting in a maximum reconfiguration time of around 200 ms and a typical time of less than 80 ms.

As stated above, however, ring topologies have the weakness that they cannot tolerate more than one failure.

MRP (along with many proprietary ring technologies) also has the ability to support *subrings*. Depending on the support that is included by your hardware vendor, some switches can be configured as Subring Managers (SRMs), allowing them to take part in two rings. For example, two of the MRC switches in Figure 3 could be configured as SRM switches and connect a subring of additional switches off another of their ports. The two switches then take part in two rings - the original ring being known as the *basis ring*. The subring will need to have at least one other switch, since there needs to be a switch taking the role of MRM for the subring.

It should be pointed out, however, that the subrings need to be configured on different VLANs, so further configuration is required to share traffic between the rings.

Proprietary solutions

Many industrial ethernet switch manufacturers offer their own proprietary redundancy protocols (see Table 1). If you don't mind being 'locked in' to a particular vendor for your network, or at least a part of it, then you may be able to take advantage of redundancy protocols that perform better than RSTP or MRP and offer additional features to enhance the redundancy further. However, if you need interoperability between vendors you will have to either settle for a standardised protocol or design a hybrid architecture in which sections of the network use proprietary redundancy, while others are linked using standard protocols.

In Part 2

The approaches to network redundancy discussed so far have focused on standard network topologies in which there is a single path between any two points. Redundancy depends on the paths being reconfigured in the event of failure and, depending on the protocol used, there may be a trade-off between speed of recovery and the number of concurrent failures that can be recovered.

In Part 2 of this article we will discuss fully redundant network architecture approaches in which total redundancy is achieved using independent paths between any two devices - PRP and HSR.