Ethernet in industry – challenges and opportunities

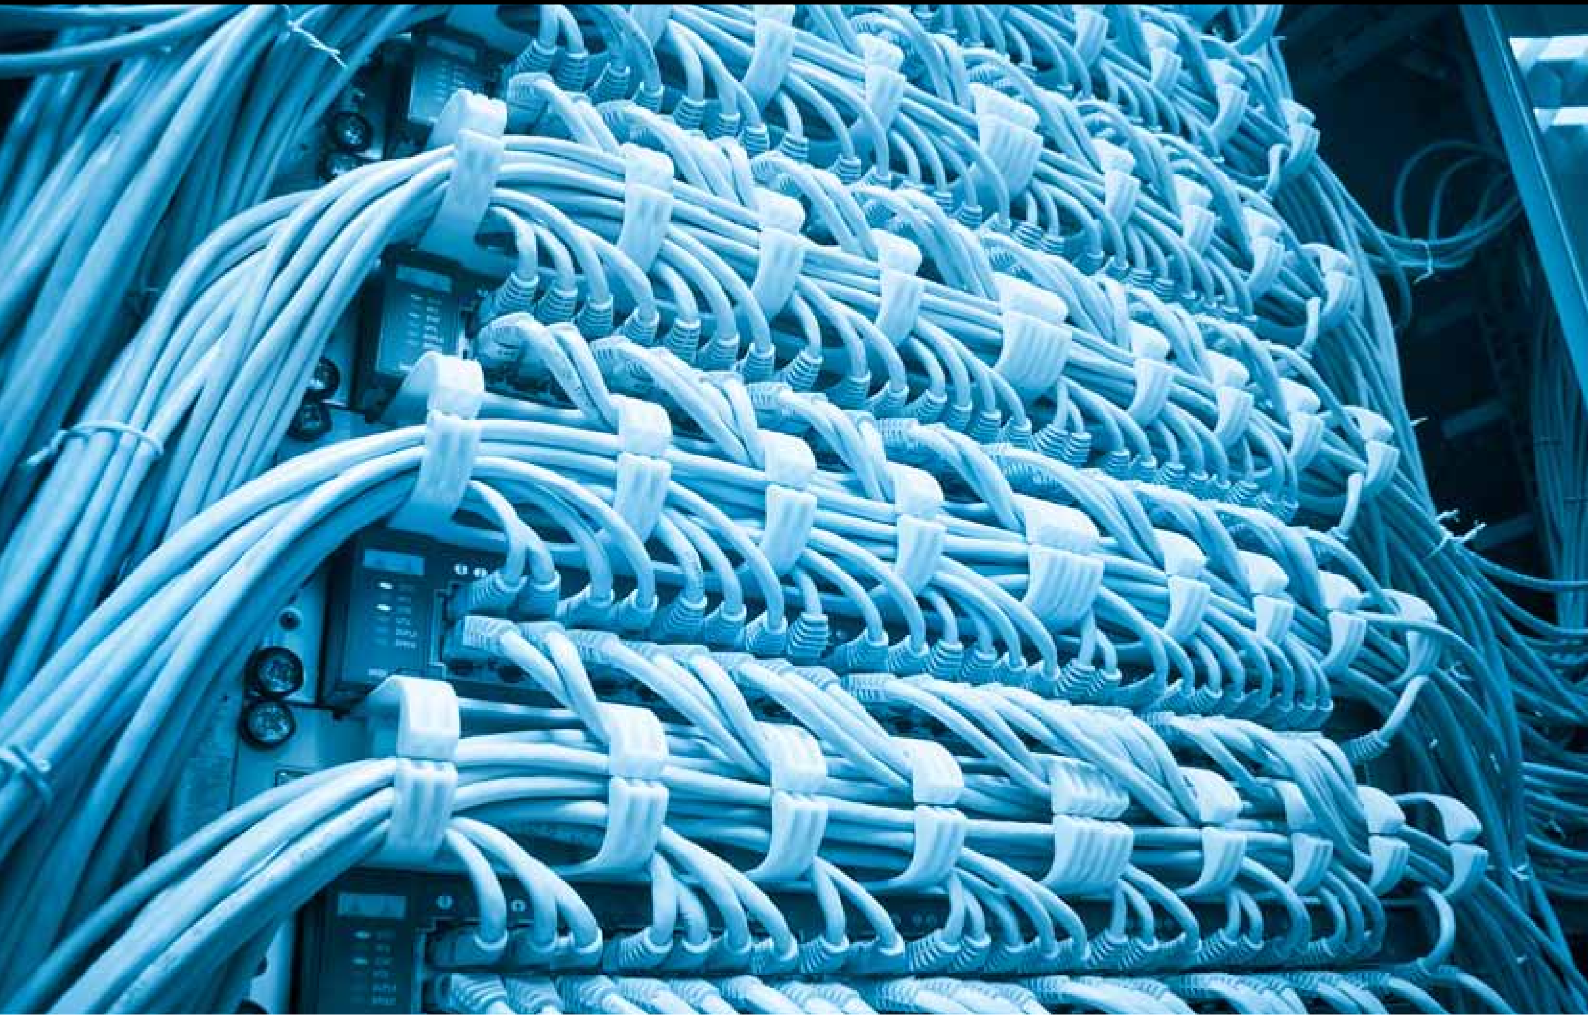

Contents

ebook

processonline .com.au

Automation, control & instrumentation

Redundancy in industrial networks

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 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

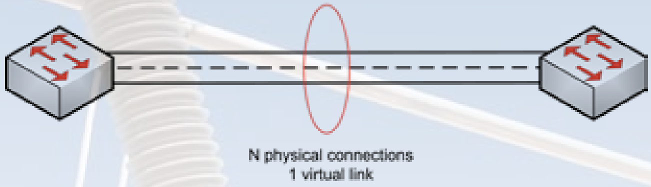

Figure 1: Link aggregation allows links between switches to be bundled to increase bandwidth. Redundancy is improved if the links have different physical paths.

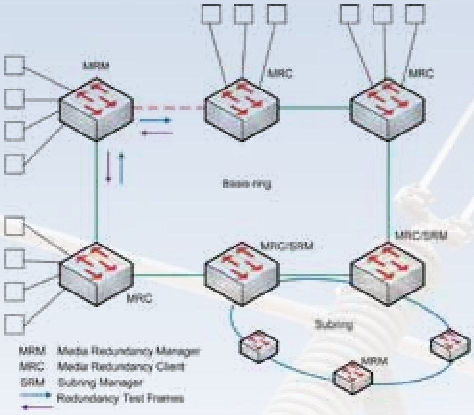

Figure 3: In MRP, switches react to received reconfiguration frames from the Media Redundancy Manager (MRM). Subrings are also supported through Subring Managers (SRMs).

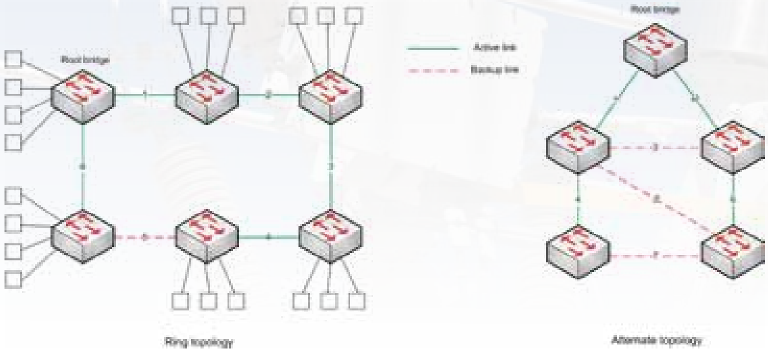

(left) Figure 2: Spanning tree protocols (STP and RSTP) create a tree of connections between switches, disabling connections that would form loops.

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 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

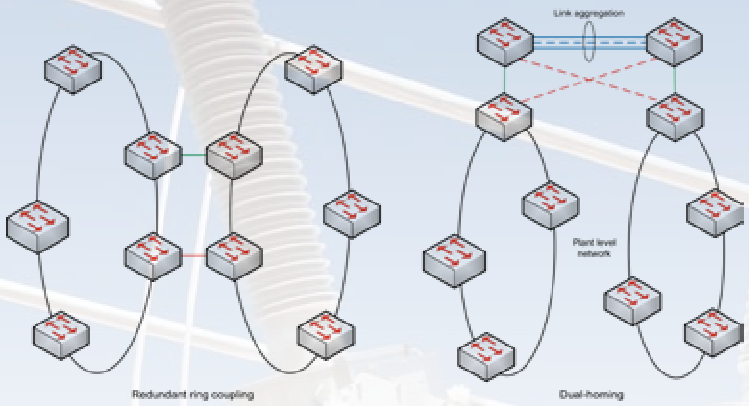

Figure 4: Dual-homing and ring coupling may be used if these protocols are supported by the switches.

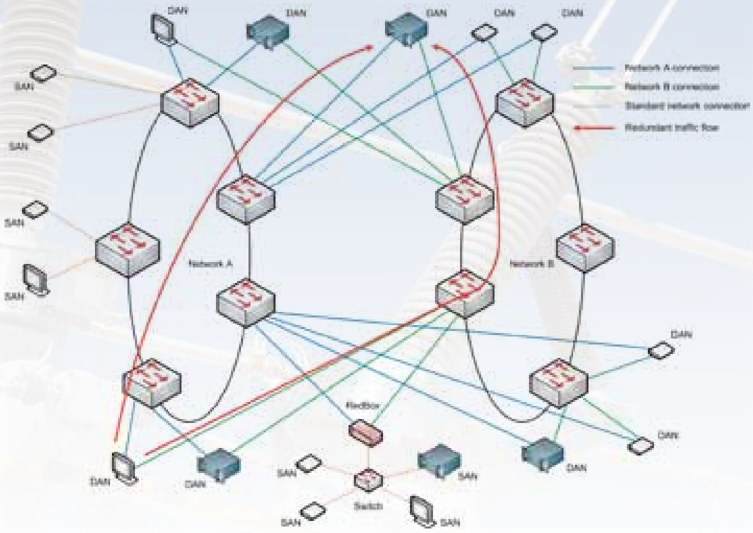

Figure 5: An example of a redundant configuration using PRP. The redundancy box acts like a redundancy proxy for the SANs attached to it.

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. 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.

Fully-redundant solutions

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. Redundancy protocols can also be combined to further enhance network availability.

Other methods of achieving redundancy (usually using proprietary methods) are dual-homing and ring coupling.

Dual-homing and ring coupling

One example of redundancy techniques that are based on proprietary technology is dual-homing protocols, also known as redundant coupling protocols. They usually have a recovery time in the 200 ms range. Although they may be installed as the sole redundancy method, they are more typically used in tandem with other methods. They are used to give redundancy, or to connect ring topologies, to enable redundant links between rings or between other lower level networks and a higher-level network. All data runs through a primary link and, on failure, a backup link is opened. Usually, both the primary and secondary links connect with two separate switches in the lower level network so that there would be no single point of failure. For example, with redundant ring networks, one process area might be put into one ring while another process area would be configured into a separate ring, with all the information directed to a central control station or historian server (Figure 4). Each ring or process would be redundantly coupled back to the main or backbone network so that the flow of information would not be interrupted.

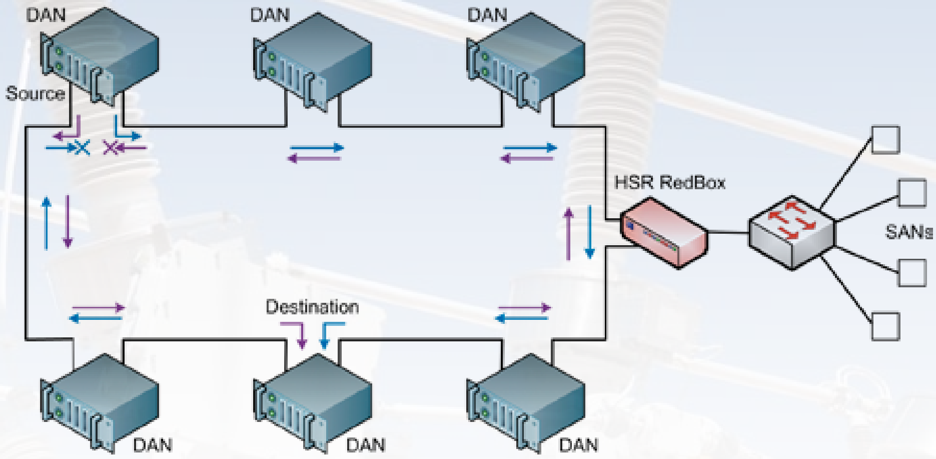

Figure 6: A typical HSR configuration. Single attached nodes cannot be connected directly to the ring, as in PRP.

Zero recovery time technologies

The Layer 2 redundancy methods we have examined so far have three basic weaknesses:

1. They provide redundancy at the switch/network level, but not right down to the end device. If the link between the end device and the switch should fail, then the device will be unable to communicate, the effect of which will depend on the device and the process.
2. They provide redundancy for links, but not for entire switches. If a switch fails, all devices connected to the switch become cut off.
3. They have a finite, and often not deterministic, recovery time, which may be a problem for some high performance applications.

An example of an application where a slow or non-deterministic recovery time can be an issue is in substation automation. Intelligent Electronic Devices (IEDs) is the name used for the technology that has come to replace protection relays and other technology for high voltage circuit control. Today, many of these devices have ethernet interfaces and, in a typical substation environment, communicate with each other and the higher level SCADA system via ethernet using the IEC 61850 protocol. Under this protocol, sample data may be collected up to 256 times per 50 Hz cycle (or 12,800 times per second) and network latency is a significant consideration for network design, under normal operating conditions. These networks are also implemented in an environment where large surges and EMI bursts are commonplace. If it is intended that network failures be accommodated in the design, then the recovery time of standard redundancy protocols may not be fast enough to ensure no loss of important data.

In order to overcome these three limitations, two new standardised technologies are available which allow for two independent paths between any two devices, providing complete communication redundancy. They are both specified in IEC 62439-3. The big advantage of both these protocols is zero reconfiguration time, guaranteeing the highest communication availability.

Parallel Redundancy Protocol

Parallel Redundancy Protocol (PRP) is implemented in the end devices and two independent paths are configured to exist between these end devices. The two networks are completely separated and are assumed to be fail-independent. They can have the same topology or be completely different and can also internally implement previously discussed redundancy protocols. The end device does not need to be 'aware' of any of the features of the networks themselves (Figure 5).

A source node with PRP functionality simultaneously sends two copies of a frame, one over each of two ports. The two frames travel through their respective separate networks until they reach a destination node, in the fault-free case, with a certain time skew. The destination node accepts the first frame of a pair and discards the second, taking advantage of a sequence number in each frame that is incremented for each frame sent.

The result is that, as long as one network is operational, the destination always receives one frame. This protocol provides a zero-time recovery and allows the redundancy to be continuously checked to detect failures. The only inefficiency in this design, however, appears to be that the redundancy control information is late in the frame and the message has to be processed in order to determine if it is a duplicate.

For PRP to work it must be implemented in software in the end nodes - the switches are standard devices and do not need to have any PRP functionality. An end device with PRP functionality is a Double Attached Node (DAN), having a connection to both networks.

A standard device with a single network interface (a Single Attached Node, or SAN) can only be attached to one network. Such a device has no redundant path in the event of network failure between it and another SAN. A device called a Redundancy Box (RedBox) can be used, however, to connect standard devices (or networks of standard devices) to both networks.

In many implementations, only the critical devices need be DANs, while non-critical devices can remain as SANs or be connected through a RedBox. The RedBox implements the PRP for all the SANs attached to it as a type of redundancy proxy.

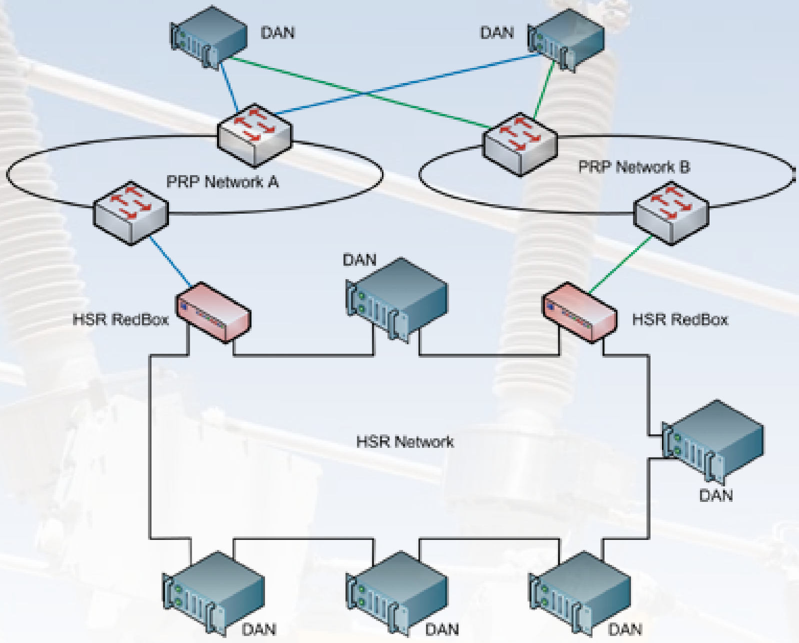

Figure 7: An example of mixing HSR and PRP in a redundant network.

This system works seamlessly provided both networks do not experience a failure at the same time. Availability can be enhanced further by implementing standard redundancy protocols within the two networks, independently.

If we overlook the cost disadvantage of duplicated network hardware, the main cost advantages of PRP are:

- Static redundancy reduces network engineering costs.
- The lower likelihood of network outages reduces operational costs.
- The use of standard ethernet hardware.
- Critical and non-critical systems can exist on a single network, rather than having to implement separate networks.

High Availability Seamless Ring

High Availability Seamless Ring (HSR) is implemented in a ring topology with DANs connected to each other in a ring without dedicated ethernet switches. Nodes within the ring must be HSR-capable switching nodes.

HSR works by passing the frames around the ring in both directions at once, resulting in a halving of the available bandwidth. Unicast frames, when received by the destination node, are removed from the ring and the data passed up to the application on that node. Multicast and broadcast frames, when received, will be forwarded on the other ring port. The sending node is responsible for removing the frame when it has traversed all the way around the ring, to avoid frames circulating forever.

General purpose SANs cannot be attached directly to a HSR, except via a HSR RedBox (Figure 6).

The advantage of HSR rings is that, like PRP, there is seamless failover. Unfortunately, being a ring topology, it cannot recover from multiple failures in a single ring. Being implemented in hardware, its application is in high-speed networks that require instant redundancy for a single failure, such as in substation networks and motion control. There have also been concerns raised in some quarters that the fact that all traffic must go through all devices (twice) means that, in large implementations, the network speed may need to be more than 1 Gbps.

The use of specialised hardware interfaces allows the upper-layer application and protocol stack to be ‘unaware’ of the underlying redundancy topology, but the disadvantage of this is the necessity for this specialised hardware. PRP, on the other hand, does not require specialised hardware as it still uses standard ethernet switching technology. PRP’s overall performance however is dependent on the standard networks it is implemented over.

But, like PRP, HSR provides other cost benefits:

- Static redundancy reduces network engineering costs.
- The lower likelihood of network outages reduces operational costs.
- Critical and non-critical systems can exist on a single network, rather than having to implement separate networks.

HSR also defines a double RedBox known as a ‘QuadBox’ that can be used to link HSR rings together. Complex topologies, including ‘rings of rings’ can be implemented. HSR rings can also be maintained only for the high-speed critical parts of the network (such as for networks of IEDs in substations) and be connected via a RedBox to a standard RSTP or MRP redundant network as a backbone, or even to a PRP network using two RedBoxes - one for each of the two PRP networks. Figure 7 shows an example of mixing HSR and PRP in a redundant network.

Conclusion

In today's automation applications, RSTP and MRP are the redundancy control protocols typically used, or alternatively a range of proprietary protocols (see breakout box). Most, if not all, industrial ethernet switches have RSTP and MRP redundancy control protocols implemented and have proved their worth. These protocols cover most requirements, but there have always been applications that cannot tolerate a failure of even a few milliseconds. Until now there has been no effective way to overcome this problem.

But now, with the availability of PRP and HSR, it is possible to implement zero-changeover, fault-tolerant network architectures. However, both PRP and HSR are very new. While PRP is already in use in some applications, HSR is still very new and is dependent on the development of equipment with HSR interface hardware. ■

The Ethernet evolution in industry

Shaun Loesch, PAC Solution Manager, Industry Business, Schneider Electric

In our day-to-day office and personal environments, we've become accustomed to a wealth of benefits delivered by internet-based technologies. As we move further and further towards a common IP highway, we consider the significant benefits that could be delivered by Ethernet backbones in industrial processes.

Will Ethernet bridge the divide between business and production environments or are we already there?

The simple answer to the above question is that Ethernet technology has evolved to meet the needs of the industrial automation market and its capabilities provide significant advantages compared to the older proprietary networks. Some of these advantages include:

- High-speed communications, 10 Mbit, 100 Mbit and 1 Gbit options.
- Large data packet size coupled with higher speeds improves communications to large I/O drops and intelligent field devices.
- More predictable communications with the introduction of Ethernet I/O scanning in automation systems.
- No need to specially train personnel on proprietary networks as Ethernet is taught in universities and understood by personnel in other industries.
- Simplified configuration and troubleshooting allows management of the entire network from one central location and access to a wide range of existing Ethernet diagnostic tools.

Yet despite such advantages, some organisations continue to express concerns regarding the use of Ethernet at the plant level. These concerns have chiefly focused on nervousness regarding its real-time capabilities and the robustness needed to operate in the harsh environment of a plant floor. While some of these concerns may have been valid previously, developments in Ethernet technology mean they should no longer be an issue. Today's Ethernet solutions directly address the previous limitations expressed by end users.

Ethernet development to suit industrial processes

In the 1990s, vendors recognised that the inherent advantages of Ethernet would make it an attractive fieldbus network and began to build open application protocols based on Ethernet. The protocols used Layers 1 and 2 of the Ethernet stack and a new application layer optimised for automation applications.

The resulting protocols needed to have the flexibility to meet a range of industrial requirements while being easy to use for non-IT personnel within the plant. An additional requirement was for these open protocols to use standard Ethernet hardware technology so users could utilise common off-the-shelf network components. Now large device catalogues allow customers to pick best-in-class devices for their system, and be assured that these devices will work together.

The move to Ethernet-enabled PACs

With many businesses currently focusing on maximising output from existing installations, one key means of achieving this improved performance and overcoming some of their challenges is by making better use of the vast amounts of information which exist within their operations.

The move towards the connected enterprise requires a technology architecture that is capable of moving large volumes of data and information from the many connected devices found across the operation. This data is moved to the higher level applications and systems used for visualisation and analysis. At the centre of this technology architecture is the programmable automation controller (PAC). The traditional role of the PAC and its predecessor, the PLC, has been to monitor and control the devices, equipment, applications and processes found within the industrial operation.

The new generation of PAC implements functions and services which support:

- secure and efficient process automation
- Ethernet transparency, distributed intelligence
- links with business applications
- web integration
- interoperability
- device communication using Ethernet and web standards

The new generation PAC with Ethernet at its core not only ensures that its performance exceeds the demands placed on it, both now and in the future, but that it achieves this while maintaining high levels of security. Cybersecurity threats from external or internal sources are issues confronting all manufacturing companies today, and deliberate or accidental breaches to system integrity have the potential to impact not only profits but also people and the environment.

PACs, big data analysis and energy management

With a global focus on sustainability, we now find energy measurement information available from a large number of sources. Power meters, smart devices and process instruments are the most common ones. All of these sources of information need to be brought together and combined with process data in order to achieve effective energy management. The highest quality of this data is obtained from smart devices and instrumentation, so the control system needs to provide open interfaces to each of these different devices and have the ability to time stamp (to milliseconds) the data which they are collecting (both electrical and process data). The combination of high-quality process and energy information shifts the data value from energy dashboard variables to being an information source which can indicate faults within the process that cannot otherwise be detected.

Current levels of energy management are mostly focused on energy consumption and looking at peak demand or power factor. These measures are the ones which need to be used in order to have the most direct impact on the processes to control our energy consumption. Beyond this, there is more analysis which can be based on energy data that provides predictive rather than reactive information about the process. To complete this analysis there needs to be increased access to more information that is available within the smarter drives and power meters that are available today.

The information available in smart devices and meters is growing, but they are typically located on proprietary networks or on open networks dedicated for PAC controls. These networks ensure that appropriate priority is given to the control messages that enable great process automation. But they force engineers to select a subset of the data which is required for energy management and copy it into the PACs in order to transfer a set of data to energy management and other systems. This replicated data is limited in its scope because of the need for the PAC, rather than acting as a mailbox.

To unlock all this automation and energy management data for big data analysis, without compromising the control focus required, there needs to be open access from remote systems to this wealth of data. Any transfer solution places unnecessary load on the controller or reduces the data available. The best solution is to allow direct Ethernet connection between analysis applications and the smart devices.

With many businesses currently focusing on maximising output from existing installations, one key means of achieving this improved performance and overcoming some of their challenges is by making better use of the vast amounts of information which exist within their operations.

To obtain the various multidisciplinary functionalities needed to run a plant, process end users in industries such as water and wastewater, food and beverage, hydropower, metals and mining, as well as in the cement and glass industries, require secure, reliable interoperability among their automation products. As a hub for both real-time control and information, PACs can benefit from being designed with an open Ethernet backbone to optimise connectivity and communications, increase bandwidth and provide a high level of security.

PACs typically provide complete automation, real-time information and motion control functionality using a single programming and engineering tool and a single programming language. PACs also provide transparent access across all parameters and functions, along with easy integration to the enterprise though the use of internet and other IT standards.

As the needs of process end users continue to evolve to meet their ever-increasing challenges for productivity, flexibility, efficiency and profitability, the designs of PACs have also evolved. PACs must leverage the latest and most powerful silicon offerings in hardware to increase robustness and the reliability of the memory. PACs must also provide a high memory capacity to avoid creating bottlenecks.

Evolution from the PAC to the ePAC

With today's process plants requiring more rapid changeover capabilities, it is critical to be able to change automation configurations and architectures on the fly, without stopping the process. PACs must also have an architecture geared for maximising production flexibility, data and information transparency, and openness for diagnostics performed both locally and remotely. This has led to the next evolution of the PAC, known as the ePAC.

Harnessing the best parts of Ethernet and PAC technologies, the ePAC offers users an even more adaptable platform to integrate with their existing hardware. Companies no longer have to take a costly rip-and-replace approach to deploying new solutions and now have an easier way of migrating operations to the platform. The end result allows companies to pick and choose the right services for them from multiple vendors, instead of relying on a one-size-fits-all approach.

ARC has observed a major trend for process end users to employ open networking technologies, such as EtherNet/IP, and to be able to take advantage of an open integration environment, higher information bandwidth, standardisation, cost savings, the flexibility to physically move portions of their processing and increased data visibility at all levels. The increasing need for distributed intelligence makes networking critically important. This market driver is leading to increased adoption of ePACs with built-in Ethernet backbones, especially for connectivity to either on-premise or cloud-based enterprise applications, such as enterprise resource planning (ERP), manufacturing execution systems (MES), enterprise asset management (EAM) and supply chain management (SCM).

Today's connected applications demand tighter integration and more information, with a higher expectation that the control system will initiate communication, update the controller at the device level in real time and serve up potentially massive quantities of information. Automation platforms with a built-in Ethernet backbone help meet these requirements in a highly flexible manner because they can support instant access, regardless of hierarchy, and avoid the limitations of proprietary software interfaces and protocols.

Network-centric ePACs with a built-in Ethernet backbone are accelerating the trend towards distributed I/O, providing process end users with significant cabling cost savings and reductions in installation, start-up and commissioning costs. Using Ethernet cables to replace I/O extension cables and fieldbus cables can result in significant cabling cost savings. Ethernet cables are also much less expensive than even standard coaxial cables. In addition, the use of single optical fibres to connect long-distance remote drops and devices can also result in significant cabling cost savings. ePACs create new opportunities for both traditional in-rack applications as well as for distributed I/O.

At the control level, process end users seek an increasingly more flexible, expandable, interchangeable and reliable control platform that ideally covers the widest range of required applications. They're looking for the ability to easily interface their control platforms with both fixed/wired and mobile/wireless HMIs, serial devices, motors, thermocouples, analog and digital I/Os, and other equipment and devices. With control room and rack space at a premium, they want automation platforms with the smallest practical footprint. And since power consumption affects both electricity and air-conditioning costs, they're looking for more energy-efficient solutions.

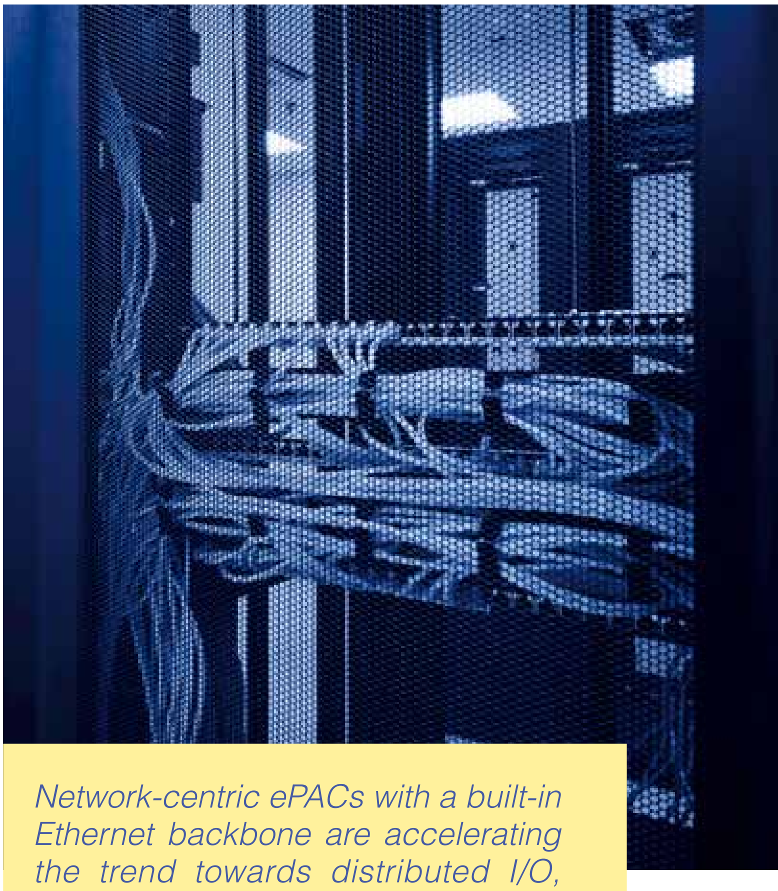

Network-centric ePACs with a built-in Ethernet backbone are accelerating the trend towards distributed I/O, providing process end users with significant cabling cost savings and reductions in installation, start-up and commissioning costs.

Taking the journey

System modernisation issues are becoming more important than ever as industries increasingly move to automate and analyse big data. Several years ago, ARC estimated that worldwide, installed process automation systems worth about $65 billion are reaching the end of their useful life, with most over 20 years of age.

When evaluating automation system modernisation projects, end users should seek solutions that minimise downtime and risk, while providing a tangible business value proposition that will have a real economic impact on their business. In many instances, ePACs will represent a viable, easily cost-justifiable modernisation solution.

ARC recommends process end users follow a stepwise approach that allows them to evolve the components of their legacy systems that have the greatest impact on their processing operations, while preserving the components that have not yet outlived their useful life. Process end users require an approach that leverages automated tools and a range of services targeted at reducing or even eliminating the downtime required to complete a modernisation project. When modernising from PAC to ePAC, end users should consider the benefits of selecting an ePAC that embeds all legacy technology in its microprocessor to help ensure compliance with older technology.

Three reasons to use Ethernet as your industrial communication protocol

Jérome Petit, Schneider

I would like to discuss the benefits of using Ethernet as a backbone for automation control systems. It's a technology that enables controllers to connect to this powerful medium.

There are plenty of benefits to using Ethernet as an industrial communication protocol, but I will focus on the following three: Ethernet openness, transparency and flexibility.

Openness

For me, Ethernet openness is the ability to mix different application protocols on the same media. These different application protocols provide different services to best fit with a business's operational needs.

For industrial communication, the ODVA (Open DeviceNet Vendors Association) specifies Modbus and Ethernet industrial protocols to enable communication between controllers or between controllers and operational visualisation where determinism is a must.

On the other hand, during operations there are other needs such as diagnostics, time synchronisation or IT connection. Using Ethernet will allow businesses to use existing standards such as web based diagnostics over HTTP, time synchronisation with NTP and IT connections over web services.

From a business standpoint this openness means a cost-effective solution and freedom of manufacturer selection.

Transparency

Ethernet transparency is also a key point and a great technology benefit. The drawback of this transparency is the administrative requirement to ensure security of your network to avoid intrusions. This being said, being able to securely access your data from everywhere on your intranet has lots of advantages.

Process automation has a huge impact on the energy consumption costs for a company. Having Ethernet on process units at a lower plant level will help to access data straight away, without the need of middleware. Data consolidation will also enable businesses to manage and forecast their energy consumption, and at the same time decrease it without infrastructure evolution.

I pointed out the need to connect the control system to the IT world; Ethernet is the IT protocol so the control system is, de facto, integrated with your IT intranet. The transparency of Ethernet will help businesses to have agile operations by directly connecting manufacturing execution systems with no additional development costs.

There are plenty of other advantages to accessing data from control systems at higher levels of the enterprise such as asset management, production data, remote maintenance, evolution and much more. As I pointed out, there is an administration cost to manage the security of your installation and to maintain it. But Ethernet will avoid the need to develop specialised middleware for energy management or MES connections, for example. The maintenance will only occur at the network level and does not need to be dedicated to specific areas by specialised teams.

Flexibility

Nowadays, Ethernet allows different topology over different media. This flexibility helps to complement plant topology at the right cost.

The bus topology will provide low-cost connection with a daisy chain feature where availability is not mandatory. The use of fibre optic cable will enable a long-distance network while retaining a high bandwidth.

When designing the topology, it can help to have distributed device from a single part of the plant. Star topology can be used to secure network devices with critical data, thereby avoiding a 'man in the middle' attack (where a cyber attack is aimed at the communication between the endpoints on a network).

The ring topology is a typical architecture in automation, as most of the time network availability is a requirement and designed to allow at least with one fault tolerance.

The ability to mix those topologies within a plant helps operations managers adapt network layouts within physical and availability requirements.

To conclude, an Ethernet backbone increases the agility of an enterprise's operations. It enables evolution without any change on the infrastructure itself. That being said, we understand that having a controller that is fully Ethernet enabled will reinforce all of those advantages.

Watch this video to find out more about an Ethernet enabled PAC.

Jérome Petit is a highly experienced automation engineer with almost 20 years' experience across customer and business support. He has worked with customers in a variety of market sectors, including water and wastewater, mining minerals & metals and food & beverage. His specialities include automation architecture, asset management and optimisation and energy efficiency.

resources
from our sponsor

Schneider Electric

From steel in the 19th century, to electrical distribution and automation in the 20th and energy management in the 21st, Schneider Electric has always been driven by an international, innovative and responsible mindset to shape the transformation of the industry it was evolving in.

To find out more, visit **www.schneider-electric.com**

Read more about:

The ARC White Paper on the world's first ePAC:

https://www.schneider-electric.com/tools/registration/promo/au/en/getpromo/41646P/

Schneider Electric Industry Business
Pacific Head Office
78 Waterloo Road
Macquarie Park
NSW 2113
Australia
02 9125 8000

another ebook from
www.processonline.com.au

published by
wfm
WESTWICK-FARROW MEDIA