

When Should I Use a Managed Ethernet Switch?

We have often been asked this question. The answer is somewhat involved but we will offer some insights in this paper.

First let us review what an unmanaged switch is. We also call these **plug-and-play** switches because you essentially plug them in and they work. No configuration is necessary for the unmanaged switch. The primary feature of the unmanaged switch is to segment the traffic to only those devices in the conversation. This gives you much more efficient use of your network bandwidth when compared to hubs which do not segment traffic. When you use a hub, all devices hear all messages.

The unmanaged switch can also auto-negotiate its data rate and flow control settings on a port-by-port basis with attached devices. This allows each port to communicate at its optimum setting. If a port needs to communicate at 10 Mbps, the other ports can still communicate at 100 Mbps.

Many unmanaged switches support Auto-MDIX, also known as auto-crossover. This allows the use of either straight-through or cross-over cables between switches or between switches and end devices (for example, PCs). Thus, the user can select a constant cable pinout and not be concerned about whether the attached device is an intermediate switch or an end device.



Figure 1 — An Unmanaged Switch

Most unmanaged switches provide some simple diagnostics by having individual port LEDs which can indicate link, activity, data rate and possibly even the duplex status of the port.

This leads to the question “what is a managed switch”. The managed switch contains all of the functionality of the unmanaged switch. The added functionality varies among manufacturers. The additional functionality typically includes **SNMP, IGMP snooping, port mirroring, redundancy, VLAN, fault relay, port trunking, rate control, port locking, QoS**, and configurable network parameters.

The list of functionality can be overwhelming. After reading this list you may think these are nice features, but why should I learn to use the functionality within a managed switch? Also, this leads to the question, “**When** should I use a managed switch?”



Figure 2 — A Managed Switch

If you have a larger network, you may want to use managed switches.

The larger the network, the more difficult it will be to diagnose and repair problems. Managed switches have features that can help to diagnose network problems. For example, Ethernet switches provide optimum bandwidth utilization — but because they isolate traffic to only these devices

Contemporary Control Systems, Inc. • 2431 Curtiss Street • Downers Grove, Illinois 60515 • USA
Telephone 1-630-963-7070 **Fax** 1-630-963-0109 **E-mail** info@ccontrols.com **Web** www.ccontrols.com, www.CTRLink.com

Contemporary Controls Ltd • Sovereign Court Two • University of Warwick Science Park •
 Sir William Lyons Road • Coventry CV4 7EZ UK

Telephone +44 (0)24 7641 3786 **Fax** +44 (0)24 7641 3923 **E-mail** info@ccontrols.co.uk **Web** www.ccontrols.co.uk

in the conversation, this can make diagnosing problems difficult. **Port mirroring** allows the traffic of one or more ports to be copied to a designated port for viewing with protocol analyzers like Ethereal® or Wireshark®.

Many managed switches capture diagnostic information and make this available to you via **SNMP (Simple Network Management Protocol)**. SNMP is compatible with many network diagnostic tools. It allows you to view all of the managed switches via one network management application and diagnose problems via this application. Managed switches generally capture diagnostic/statistic information on a port-by-port basis, such as the number and type of messages transmitted/received on each port, the number and type of errors seen on each port, the configuration of the switch, the status of various features within the switch and much more. This information is made available to you via SNMP. On some managed switches you can also view this diagnostic information with your web browser (for example, Windows® Internet Explorer). See Figure 3 below for an example of frame statistics captured by a managed switch and made available via a web browser.

Port Packet Statistics:	
Unicast Packets Received	53146
Unicast Packets Sent	29810
Multicast Packets Received	0
Multicast Packets Sent	130
Broadcast Packets Received	0
Broadcast Packets Sent	9829
Dropped Packets	0
Oversize Packets	0
Undersize Packets	0
Fragments	0
Jabbers	0
Collisions	0
Deferred Transmissions	0

Figure 3 — Managed Switch Port Statistics

Some protocols such as EtherNet/IP™ can transmit large numbers of multicast messages per second.

This traffic load can overwhelm certain end devices which are not designed to handle this type of load. Managed switch features such as **IGMP snooping** and **IGMP query** can automatically isolate multicast traffic to the appropriate ports. Even if an end device is designed for EtherNet/IP networks, it may be designed with the assumption that IGMP snooping is being utilized on the network. It may not be designed to receive an overwhelming amount of unwanted multicast traffic.

For some people downtime can be very expensive.

On August 11, 2007 a “sputtering network interface adapter” shut down Los Angeles Airport for many hours. A managed switch feature known as **rate control** (or rate limiting) could have prevented this situation. Rate control allows you to set the maximum bandwidth of each port of the switch — independent of the auto-negotiated rate of 10 Mbps or 100 Mbps. For example, let’s say we gave the port which connected to this “sputtering” device a maximum rate of 2 Mbps. Then when the sputtering occurred, the worst problem it could cause would be to send 2 Mbps of unwanted traffic on the network. It is likely most devices could handle this level of traffic with little effect. Also, in normal circumstances 2 Mbps of bandwidth will be acceptable for most communications between non-control devices. Most people find that 2 Mbps of bandwidth for Internet communications is quite acceptable. Figure 4 shows a PC sending excessive traffic, but due to rate limiting, the control system only sees 2 Mbps of traffic. This maximum rate limit can be set to any level under the maximum bandwidth of the link (for example, 100 Mbps).



Also managed switch features such as SNMP, port mirroring and switch diagnostics can be used to diagnose problems and lessen downtime.




Figure 4 — Sputtering device

Redundancy is another managed switch feature which can help repair a network quickly after a network cable break. This is found in managed switch features like **STP**, **RSTP** or **proprietary ring** protocols such as **RapidRing®**. Generally, these systems have an additional cable segment that acts as a backup link. Figure 5 shows 4 links and 4 switches. One link is not needed except to act as the backup. Such systems use the backup link after a primary link has broken. Switching to backup is done automatically by the Ethernet switches. This can be accomplished in under 1 second in most cases.




Figure 5 — A Redundant Ring

If you are interconnecting your control network to your office network you may want to use the features of a managed switch.

Ethernet and TCP/IP are now widely used in the control network and have been widely used in office networks for many years. This makes it easy to interconnect these systems, but you are now exposed to the issues in the office network. The office network will generally have more Ethernet devices than the control network — however, its efficiency may not be as critical. If it takes you 5 seconds extra to bring up a webpage or 1 minute longer to print a document from your office computer, will you complain to your IT department? Probably not. But these types of network issues could be detrimental to your control network because it cannot wait an extra 5 seconds to perform its communications. Managed switch features already discussed, such as rate limiting, can also be effective here. Figure 6 shows an office network with a problem which is sending a large amount of traffic into the interconnected control network. By using rate control, we again can limit the effect of this traffic to a maximum of 2 Mbps (or any desired level) and allow the control network to function normally. Otherwise, the office network issue could force the control network to act abnormally or shut down completely.




Figure 6 — Keep Office Network Problems from Control Network

Another potential issue is sending control traffic to the office network. Control protocols such as EtherNet/IP can generate many multicast messages per second. We had one customer comment that this traffic was getting into the office network and



disturbing the executives at his company (probably not a good career move). However, with **IGMP snooping/IGMP query**, this traffic will be blocked from going to the office network automatically and normal traffic between the two networks can occur without detrimental effect.

Another managed switch feature which can be useful when interconnecting the office and control networks is **overlapped VLANs**. Below, in Figure 7, a SCADA system communicates with a control network as well as the office network. Here we have setup two VLANs. One VLAN is the control network and the other is the office network. The SCADA resides in both. This is known as an overlapped VLAN. The SCADA can still communicate with the control network to collect data and with the office network to allow office employees to view control system status. However, any issues on the office network will not affect the control network (only the SCADA) as the office and control networks are in separate VLANs and they cannot communicate directly.




Figure 7 — VLANs Keep Traffic Where It Belongs

Another managed switch feature which can be useful when interconnecting the office and control networks is **port security** (or port locking). This allows you to specify which Ethernet devices can communicate through the switch on a port-by-port basis. Communications from a device that has not been authorized will be refused.

This can help limit the exposure of the control devices to the office network. For example in Figure 8 below, we have a large office network connected to a small control network. We can restrict all office devices except computers A & B from communicating with the control network. This will lessen the chance of control system problems from occurring due to office network problems. This will also help limit the amount of normal outside communications being received by the control devices, allowing them to perform their normal functions.




Figure 8 — Port Locking Allows Access to Specific Devices

At the beginning of this paper we asked the question, “When should I use a managed switch?” Basically it would be needed in the following cases:

- When you have a large network you will probably need managed switches simply to be able to diagnose problems.
- When you have multicast traffic on your network (for example when using EtherNet/IP).
- When downtime is expensive.
- When you are interconnecting office and control networks.

Managed Ethernet switches provide many features and you may find other reasons to use them beyond those described here.

For more information on managed Ethernet switches please go to the Industrial Ethernet University www.industrialEthernetU.com or to www.CTRLlink.com.