

Ethernet switches: why time matters

Ethernet switches: why time matters

Introduction

A managed Ethernet switch can offer significant advantages over an unmanaged switch in terms of flexibility, availability and security. However: organizations will often compromise on these benefits in favor of the faster start-up times offered by an unmanaged switch. Now, work undertaken by the network engineering team at Abaco Systems means that compromise may no longer be necessary: it is possible to have the best of both worlds.

Background

Measuring the start-up time of an Ethernet switch is, in theory, straightforward. Start the virtual stopwatch at the moment the switch is powered on, and stop it when the first packet is switched. The answer will be some number of seconds.

That, however, is only the theory. The reality, especially in the case of a managed Ethernet switch, can be very different – and sometimes surprising. Why is that?

It's important to understand the differences between an unmanaged Ethernet switch and its managed equivalent.

Unmanaged vs. managed switches

An unmanaged switch is built around a very specialist silicon chipset (ASIC) called a switch fabric. This makes very fast decisions about where to send a packet coming into it. The switch can look up the destination Ethernet address and decide which port to “forward” the packet out of. This is a very simple decision, and is done very fast – in fact, at “wire speed” in most cases. Unmanaged switches work well in simple networks, where most operations are handled at the Ethernet layer – “Layer 2” in OSI (Open Systems Interconnection) terminology.

A managed switch is still based around a similar switch fabric, but one which is capable of making more complex forwarding decisions. A processor is added alongside the switch fabric. The fabric can still make those simple, fast decisions - but the processor and associated software can additionally set up new forwarding tables for new routes and take care of protocol interactions. Management also allows controls for security, statistics gathering and so on, but those are not relevant here.



Redundancy

A simple example of the kind of function a managed Ethernet switch will be needed for is in managing multiple links. This is a fundamental requirement when designing redundancy into a network – providing more than one possible route from A to B. Redundancy means that, when something goes wrong - maybe someone pulls a cable - the traffic can still get through.


Redundancy is important based on the potential for a network failure - in military and aerospace applications due to enemy attack and in automotive applications due to a possible accident.

However: having multiple links can cause a range of problems to an Ethernet network. Among these is the classic “broadcast storm” in which traffic goes out via one of the redundant links, then instantly comes back in on the other one, and then goes out the first, then.... All the traffic locks up because of the loop.

But these problems are well understood, and are generally handled by using a managed switch. Managed switches support many features – and handling of redundant networks is a good example. Typically, this management of redundancy makes use of the “Spanning Tree Protocol” (a topic beyond the scope of this white paper). The simplest thing to consider here is that the management software on the CPU needs to block transmission on some of the links to prevent these broadcast storms.



Ethernet switches: why time matters



Dynamic Routing

Another factor influencing what level of management is needed on the switch is related to how simple or complex the network is going to be. Simple networks only need L2 (MAC-address) switching – no need for decision to be made on the IP address – and no need for management. However, there is an increasing need for networks to do what used to be called “routing”, but is now handled as L3 switching. (i.e. the switch decides where to send the packet using the IP address.)

When L3 decisions need to be made, some level of switch management is required. But it can be limited – if you can guarantee that your network will never change, then “static” routing can be used. Thus, the management is used to set up L3 routing tables, and these are used from then on. This is the approach some suppliers provide in their management software – static routing only.

In real life, we find most of our customers need more than this. They need “dynamic” routing – where the tables which decide on L3 forwarding are controlled and updated by “routing protocols” – be that RIP, OSPF, BGP or such like. With the routing tables the switch management entity talks to its peers and they decide on the best routing between them.

The good news here is that OpenWare is a fully switch management suite, and so supports fully dynamic routing, including all the popular routing protocols. And, it does this while still keeping its start-up time very low!

“Could we get a managed switch fully functional from power on in less than 30 seconds?”



Start-up times

So: to return to the issue of start-up times. For a managed switch, this requires the CPU environment to start up, perform Built-In Test (BIT) and load the management software – which then reads the configuration and initializes the switch fabric and interface devices.

The requirement for an operating system and a set of management software programs to crank up, then initialize and configure a complex ASIC, means that a significant amount of time can pass before the switch can become fully operational.

On some managed switches, this time has been measured in minutes. For a product based on Abaco Systems’ OpenWare switch management system, and assuming a relatively simple network, a figure of around 30 seconds is typical.

In many cases, the time taken to start a switch really doesn’t matter – even if that time is measured in minutes. In many commercial environments, for example, the switch is powered on when the equipment is installed, and is probably never powered off again. If it is, it will be in a situation where changes are being made, and so a few minutes disruption is of no consequence.

Mission-critical applications

However: there are many mission-critical applications that need the functionality of a managed Ethernet switch – but that need the switch to be 100% functional in an even faster time than 30 seconds. In such environments, there may well be the need to turn the switch off when it is not in use in order to conserve power and/or minimize heat dissipation. The requirement, though, is that it should be 100% functional as rapidly as possible once turned on again.



Ethernet switches: why time matters

Imagine, for example, an environment in which there are multiple end nodes. The system could be designed such that multiple single board computers (SBCs) communicate with a central storage device over the network, from which they load their file systems. In such cases, the network needs to be operational before the SBCs can boot.

Alternatively, the environment might be one in which the only external vision available to a vehicle is via IP cameras, feeding monitors for the driver, or where that driver is entirely reliant on other sensor-derived data. Until the network is operational, the driver is effectively “blind”. Clearly, the time between power being applied to the Ethernet switch and images and/or information appearing on the screen is of paramount importance.

The GBX411 Ethernet switch

It was with such mission-critical switch applications in mind that the network engineering team at Abaco resolved to address this need to minimize wait time for the new Abaco Systems GBX411, a 3U VPX fully managed Ethernet switch that supports Gigabit and 10 Gigabit communications, in various configurations. It is based on a leading-edge switch fabric, and a small PowerPC processor. It runs OpenWare, Abaco’s in-house-developed switch management software suite, a comprehensive set of software to configure and control the switch.

OpenWare makes substantial use of open source software for its processing capability, including the boot-loader and operating system environment.

The advances in barebox (the chosen boot loader) and Linux recently gave the team the ability to tailor the start-up time, taking a number of combined approaches, including:

“Yes, we could.”



- trimming everything possible from the boot sequence
- migrating the start-up scheme from SysVinit to systemd
- tailoring the exact order of the numerous tasks required during start-up

Best of both worlds

The reduction in the time taken for a managed Ethernet switch to become operational as a result of these approaches was significant, with a “power-on to first frame switched” time of approximately 15 seconds being achieved. Not only did this represent a substantial 50% saving in time – but it also meant that a managed switch was capable of becoming fully operational in a time close to that achievable with an unmanaged switch.

The implication of this is, of course, that customers can achieve the multiple benefits of a managed switch in terms of flexibility, availability and security – but with the start time closer to that of an unmanaged switch.

And: the team believes that similar reductions may be possible in other products in Abaco’s switch portfolio.

WE INNOVATE. WE DELIVER. **YOU SUCCEED.**

Americas: 866-OK-ABACO or +1-866-652-2226

Asia & Oceania: +81-3-5544-3973

Europe, Africa, & Middle East: +44 (0) 1327-359444

Locate an Abaco Systems Sales Representative visit: abaco.com/products/sales

abaco.com | @AbacoSys

©2017 Abaco Systems. All Rights Reserved. All other brands, names or trademarks are property of their respective owners. Specifications are subject to change without notice.

