Switching Paradigms

Switches are often the unsung heroes of our business networks. They are acquired, setup, and hum away in the closet connecting vital components in our network. If they are running, and running well, we do not tend to pay much attention to them. It is one of those “if it isn’t broke, don’t fix it” types of mentalities. Especially in flat networks, that is networks without network segmentation, even less attention is paid to what equipment is being used, what it takes to replace it, and what sort of trouble an outage can cause.

However, in this attitude of not paying attention to what is going on with switching infrastructure, we can miss out on many key mechanisms that can lead to better business continuity, better and faster network connectivity, and lowered total cost of ownership when maintaining, repairing, replacing switches, and efficiencies in day-to-day management. This holds especially true for networks that utilize segmentation to achieve better functionality and security goals.

Quality Plus Consulting has experience with a broad range of switching manufacturers being applied in just as broad of use cases. Some of the switches we manage service small home or office networks with just a few users. Other networks contain a multitude of switches servicing hundreds of users spread out over acres of physical space. And even with the vast difference in demands between these two scenarios, we have found many key paradigms that need to be adopted, or switched to, when engineering network designs and selecting switches based on the lowest total cost of ownership.

Switching from Acquisition Cost to Total Cost of Ownership

The first paradigm switch that needs to happen is moving from **Acquisition Cost** to looking at **Total Cost of Ownership** The acquisition costs, meaning the initial hardware, warranty, and support contract cost, are a very small part of the cost of equipment. **A much larger cost to implementing a hardware solution, such as a switch, is the labor cost to program, maintain, and make changes to that equipment. Ongoing day-to-day configuration changes must be able to be done quickly and with the confidence that they will not create network outages.**

The only way to know the total cost of ownership is to know key data such as:

- Does this hardware meet our companies long term needs?
- What is the cost profile for maintaining this hardware from year to year?
- How easy is it to do initial setup and programming, and once this is completed, how easy is it to make changes?
- Can the expected configuration changes be made without causing network outages?
  See the Netgear switch example in section “NOS makes all the difference”.
- What is the cost profile of ongoing licensing and support for the hardware?

- What is the total useful life of the solution? In the case of many switches, how is this switch going to help our company over the next eight years that is in service?

All of these questions can only be answered by working with a supplier that can engineer a solution that answers all of these questions BEFORE selecting a product. Working with a reputable company that can supply the right product and evaluate it for fit ahead of time is critical.

Once these questions are answered, it becomes very easy to see how a cheap solution can be very costly. By only looking at low acquisition costs, a company can acquire hardware that needs to be frequently replaced. Initial implementation and replacement are the most expensive phases of a project. **Cheap hardware that only lasts three years vs. robust hardware that can be replaced every eight years lends the cheap hardware to costing more in the long run.** In addition, cheap hardware that does not have a mature, vetted configuration interface can cost more in management time. **If it is more complicated to make changes to the device, then simple tasks that would take five minutes on properly selected hardware can take more than an hour on cheaper hardware and will likely be concurrent with total network outages at the site.** (Refer to the Netgear and 200 series switch examples.)

These are just a few of the examples of how total cost of ownership is the most important cost to look at over and above acquisition costs. Other factors that play in are the cost of warranties and support contracts. A low-cost device that costs thousands of dollars annually for support contracts becomes much more expensive than just the acquisition price. **Also, if the warranty and support contracts are cost prohibitive and not maintained, internal IT Support shoulders the entire burden of getting the system back up and running in the event of a failure.** This means that the company could pick up the cost of replacing the item at full cost prematurely in the event of a major failure. Warranties do not cost the same as buying the device a second time and are a cheap insurance policy over an eight-year period. Likewise, shouldering the time burden of having to do support without the help of the manufacturer is a further costly endeavor when support contracts can be affordable as a second cheap insurance policy when spread out over eight years. It is also important to note that most consulting firms will not work on network equipment that is not covered by a current GTAC contract due to there being zero ability to gain access to current firmware and manufacturer support. It is important that the right hardware warranty and GTAC support contract be combined in order to have an effective support and recovery system for the equipment.

Switching to a Better Understanding of Sizing

Switch size, or port count, is extremely important when considering when selecting a switch. The reason for this is that even though there are switches with expandable port counts, these are not in the realm of hardware affordable to or needed by SMB clients. **There are many engineering and design considerations that must be considered for sizing a switch.** Here is a chart that shows the sort of calculations that are used when determining the needed gigabit copper port count:

Network Resource:	Ports Needed:
Staff Offices (cross connected to wall jacks)	4 ports
IT Office	8 ports
Servers (per server)	4-12 ports

Firebox (or other core router / firewall)	4 ports
Wireless Access Points	2-4 ports
Printers	2-4 per building
Fiber Cross Connectivity	2- 4 (SFP+ Ports per switch)
Stacking ports	2 per switch

After reviewing this chart, you can see that anything smaller than a 24-port switch falls into a special use case. **Most of the time, it just makes sense to get a 48-port switch because the company will use more than 24 copper ports over the expected life spans of the switch.** It may seem a bit costly to assume that you will need four ports for every office, but keep in mind, the number of offices is a good indicator for the number of devices needed. Each person may need, at minimum, a computer, and an IP phone. Now also, factor in IoT devices that may be in use in the office. And consider that maybe a larger office that used to have a single user may be used as an office for two users. An IP printer may be installed in an office, or a visitor may need to plug into a network jack instead of using wireless. Without properly padding the number of ports needed, a 24-port switch can very quickly be filled.

An important note on port count is that a 48-port switch does not cost twice as much as a 24-port switch. And a 24-port switch does not cost double the cost of a 12-port switch. Often, the entry level 12-port switch will be the base price of the model line, then the 24-port will be an additional 25% more and the 48-port will be an additional 10% above the 24-port. With those costs in minds**, it is rarely ever economical to buy anything other than a 48-port switch.**

Though they fall into exceptions more than rules, a few considerations may result in a smaller switch size. First, most modern 48 port switches have gotten to the point where they need a four-post rack or four points of contact to support when mounting. Sometimes, a four-post rack is either not affordable, or size constraints limit using a four-post rack. Also, there are circumstances when a 24-port switch is all that will ever be needed to support the network. Certain small offices or residential installments will never grow past needing 24 ports of connectivity to the network. **Again, working with a supplier that can engineer a total solution that takes current AND future needs into consideration is of utmost importance to lowering total cost of ownership.**

Switching to Better Management

It is imperative that any switch used in business has a properly configured Out Of Band Management port (OOBM Port). An OOBM port can allow connection to the switch (both physical on-site and remotely) when no other management method works. If the switch is suffering from configuration issues that render the VLAN interfaces inoperable for management, the OOBM port can be used to connect to the switch, diagnose problems, make config changes, and if needed, restore the config to the last good working stated. **Without the OOBM, any loss of higher-level function of the switch becomes costly and time consuming.** There are much more constructive things to be doing with one's time than sitting next to a switch on a console session for hours trying to get it to work. Money spent paying a tech to do this on-site vs restore the switch to operation remotely is also better spent on other things.

But the OOBM port is just one small piece of this paradigm switch to better management. A properly configured Layer-2 network that connects all OOBM ports to a port on the perimeter security appliance is needed. Furthermore, a machine that can always access this Layer-2 network is needed for remote or

on-site management. Usually, the OOBM network can be setup using the *existing* layer 2 dumb switches that a company already has. **These configurations during initial implementation pay for themselves if they are used only a few times over the course of the many years that the switches will be in operation through reduced downtime and labor costs.**

Beyond the OOBM port, Command Line Interface (CLI) ease of use contributes greatly to lowering the total cost of ownership. **A primary reason for this is that by using a well vetted, mature switch operating system with an easy CLI, configuration changes, backups, and firmware upgrades all become quicker and simpler.** The quicker and simpler the process is, the less time it takes and, and therefore, the less money in labor it takes to manage the switch. Some manufactures have not updated their CLI in more than a decade to make it easier to use, such as Cisco. In the case of HP / Aruba, they have taken out management features and shifted everything so far into the ease of use territory that they have removed important functionality or buried needed functions in such a way that it is difficult and time consuming to use them. Even with QPCs go to manufacturer, some of their switch lines are money pits simply because management takes so much extra time that the labor cost quickly eats the price difference between an entry level, lower cost switch, and a moderate priced switch that has a good working NOS. See the appendix section on a comparison between Cisco and EXOS for a code comparison example.

In all decisions discussed in this article, you will see this theme repeatedly. Acquisition costs are a small piece of the entire picture. **By switching the paradigm to total cost of ownership, the end user can then make an informed decision and select a product that may have a higher acquisition cost, but pays for itself over and over on saved labor costs.**

Switching to the Right Hardware

Selecting the right hardware features – outside of our earlier discussion on port count – is the most important consideration when acquiring a switch. There are many technologies, such as Layer 3 capabilities, SFP ports, 10 Gbps throughput, and PoE that cannot be added to a switch later. **Many people sell themselves short by purchasing a switch with hardware limitations to save on perceived acquisition costs.** This very small amount saved becomes very costly down the road when an entire switch replacement occurs because a needed feature was left out of the original switch.

Some of the logic that QPC uses when selecting and engineering network designs around switch features is as follows:

- If the customer is using VLANs for segmentation, which is standard for all QPC clients, then all switches in the network must have easy to manage Layer 3 capabilities.

- If the customer is ever going to need more than one switch, then stackable SFP+ ports are a requirement on the switch. Stacking flexibility between models and over fiber is necessary for ongoing flexibility and low TCO.

- If the customer is ever going to need more than one switch, then SFP+ ports AND stacking ports need to be capable of at least 10 Gbps throughput.

- If the customer needs to support 10 Gbps connectivity to a server, then SFP+ ports supporting that transfer speed are required.

- If the customer needs PoE for devices, then the switch needs to have PoE provided and have sufficient power to provide PoE without external PoE injectors for loads 30W and under. PoE injectors are not recommended for the reasons specified in the informational section below on PoE injectors.

- If the customers switch has PoE, then PoE is provided on all ports eliminating the hassle of switch port inflexibility.

To briefly explain each of these points, Layer 3 switching which provides VLANs for network segmentation helps QPC provide additional security to a network through programming and policies. **This can often be much simpler and more effective than other security measures and allows technologies such as intra-VLAN packet inspection that are impossible without at least Layer 3 lite**. There are a few exceptions to the Layer 3 switching rules, but this would likely be only for a select few edge switches that connect downstream of a core Layer 3 switch where the desire is to have all ports presented there identical and no real intelligent operations such as QoS, LLDP, loop protection, or diagnostics.

The need for SFP+ ports is due to the flexibility of media being used to transmit signals**. Being able to use multi-mode fiber, or even single-mode fiber, to run a 10 Gbps signal over great distances increases the quality and efficiency of a network.** Also, there are use cases explained later where SFP+ connectivity is used to provide a better throughput between servers and switching infrastructure that can greatly reduce the number of NICs needed in a server which further reduces ports consumed on the switch.

An additional caveat to this SFP+ port requirement is the ability to stack switches over 10 Gbps ports. The switching technology that QPC has standardized on can stack, not only over the 10 Gbps ports for switches located in proximity, but across fiber distances that maintain a 10 Gbps throughput. Stacking in this way allows the switches to connect to each at the speeds reaching virtually that of the switch backplanes**. Put simply, you can have a switch next to your server, and stack it with a switch 900 feet away, and a device plugged into the switch 900 feet away connects to the server at the same speed as a device plugged into the switch next to the server.**

Lastly, if PoE, or PoE+, is a requirement for devices it should be provided by the switch. Using the switches that QPC has standardized on, there is not only the ability to support PoE+ without an injector, there is also enough power budget to support devices on every port without the additional of a Redundant Power Supply (RPS). An additional consideration of PoE is the quality of PoE that is being provided. **The use of non-LLDP PoE injectors greatly reduces the life expectancy of devices which receive full PoE power all the time.** By using switches with advanced LLDP, the life of the device is extended because the device talks to the switch and only receives the bare minimum power that it needs. In this way, power requirements are reduced saving money on electrical utility costs. **Not to mention, having PoE built into the switch avoids having to buy another appliance or PoE injectors for each device. This not only simplifies setup, but it also reduces total cost of ownership by increasing supportability.**

Switching from Link Aggregation to Stacking

When cross connecting between two switches, there are two different technologies that can be used to achieve the best possible results: Link Aggregation Groups and Stacking. Link Aggregation Groups (LAGs) take two or more ports on one switch and connect them to two or more ports on another switch to increase the bandwidth available for transfer between the two switches and provide redundancy in the case of port or cable failure. **However, connecting two or four 1Gbs ports in a LAG does not give you 2-4 Gbps total bandwidth. This is because LAGs are session based and streams cannot be split across more than one link. Because of that, the total throughput would be 2-4 Gbs, but the maximum transfer speed would be limited to 1 Gbs.** Link Aggregation Groups also have a drawback of consuming multiple ports on each switch. A four port LAG consumes four ports on each switch which ends up costing a total of eight ports within the network.

To overcome the built-in limits to LAGs, QPC has adopted the paradigm of stacking as the primary means for cross connecting switches. **Stacking switches has one advantage over LAGs because the stacking ports are ports outside of the standard, configurable ethernet ports on the switches.** This conserves the network ports on the switch on each side of the stack. A second benefit of stacking is that by utilizing 10 Gbps ports and transceivers, connections made by devices utilizing any switch in the stack run at virtually the same speed as the backplane of the switch**. It essentially makes one big switch out of several stack member switches.**

Most important in the paradigm switch into stacking vs LAGs is manageability. **It cannot be understated that taking the time to configure a switch for stacking, even if the first switch is purchased and the second switch is not needed for two or three years, is likely the single best decision to be made when selecting and setting up switches to reduce total cost of ownership**. Up to eight stacked switches are all managed as if they are one switch. With eight LAG connected switches, each switch must be managed as its own device leading to eight times the labor when it comes to firmware updates. **Replacing a switch in a stack when using the switches QPC recommends is typically an hour worth of time from the time that the new switch is in the rack to it being fully engaged as the replacement stack member.** It is not this simple and quick in a standalone switch because of having to restore from a configuration backup assuming identical models.

Server connectivity, and specifically, server Link Aggregation Groups, can also be simplified to reduce ports when a proper switch is selected. To get enough throughput for a VM host server, 4, 8 or even 12 ports may be needed for the NICs on the server. Not only does this tie up switch ports, but there are also higher implementation costs because it takes time to trace and label all ports in the physical as well as the virtual environment**. Should a server need high throughput, a switch with adequate ports to do two 10 Gbps connections between the server and the switch is much easier to manage than 12 separate NICs and cables.**

The NOS makes all the difference

NOS is the network operating system of the switch. One cannot assess a NOS from a user manual or a datasheet. A person can review the datasheet for an HP Comware NOS switch, and it says that it supports link aggregation. When you dig into its actual real functionality, it is found to lack viable LAG

diagnostics. If there are no diagnostics, it cannot be discerned why a LAG config is not working. LAG only works if all the ports and connected devices on both ends fully agree. Missing diagnostics is a deal killer.

One of the fallouts of the damage from the HP split into HPE (Comware) and Aruba (Provision/Procurve) was that HPE messed up the NOS in the Comware switches and Aruba completely busted the webUI management as well as the stability of the Provision switches. In both situations, the support costs for both switch NOS flavors exploded to an unmanageable level. Unreliable switching equipment is completely unacceptable and cannot be tolerated.

It can be difficult to do a like-for-like comparison between Extreme, Aruba, HPE, and Cisco because there are things that Extreme EXOS switches can do that no other can. In a single switch comparison of a layer 3 capable 48 port PoE switch with stacking capabilities and an equivalent hardware warranty with GTAC contract, the Extreme switch is the least expensive to acquire.

It is absolutely imperative that organizations purchase switching equipment ONLY from high quality network engineering consulting firms who have vetted and tested the switching equipment thoroughly. In-house IT staff do not have access to the materials and methodologies to be able to vet equipment on their own.

Even IT consulting firms cannot rely upon the internal teams at distributors who are supposed to know the equipment and have access to it in order to simulate configurations. QPC works with many distributors. We have tried repeatedly to work with these distributors to use their internal highly certified network equipment experts regarding assessing configurations, features, functionalities, or even to gain remote access to the equipment to test things. The result has always been lack of access to equipment or lack of knowledge on the behalf of people that work for the distributor in their network engineering department. In one case, we were literally told by the top certified network engineer at TechData that he did not use the CLI for switch management in general so could not comment on what its capabilities were on a new model switch we were enquiring about.

Extreme came out with their 200 series switches which were supposed to be a direct competitor to the midmarket switches. They were supposed to be less expensive than an EXOS switch. In practice, we have found them to be insanely expensive. The reason is because management of them is extremely unreliable, and every time a change is made, there is some network outage. Configurations must be removed before they can be put back in using different commands or a different config simply because the NOS is a Frankenstein of Comware and Cisco. Any amount of money supposedly “saved” during the initial procurement of the switch was completely lit on fire during the first two maintenance cycles for the switches. Any switch that cannot have its configuration changed on the fly without creating a network outage is an abomination.

QPC worked with a client on a network migration project with an internal IT manager who decided to “save his company money” by buying Netgear layer 3 lite switches. “Lite” is a designation that describes when VLAN tagging can be done, but there is no ability in the switch to actually be a full layer 3 router. A standard area were most switches blow up in their TCO analysis is when they are used as anything beyond a dumb layer 2 device. This was most definitely the case with these “cost saving” Netgear switches.

Unforeseen by internal IT prior to procurement was all the massive outages and increased IT consulting expenses created by the horrible NOS on these Netgear switches. Because of the horrible NOS, anytime any tagging for a trunk had to be changed, it caused a complete network outage. None of the changes could be performed remotely either. So the concept of saving money by buying unvetted low grade network equipment caused internal IT to have to drive to every branch office to physically connect a laptop to a specially programmed port directly on the switch after hours anytime any trunk tagging change needed to be made. In contrast, doing this change in an EXOS switch is less than 5 minutes of remote work with no risk of outage. It can be reliably done at any time during the day as it will not interrupt existing traffic.

The selection of this horrible network equipment also exploded the IT consulting costs on the project. This was because of all the project delays, increased workload for what should be a simple change, and the fact that QPC could not simply just make the change remotely and quickly. If an organization has any desire to control costs or to have a predictable cost profile for a technology solution, it is absolutely imperative that the right technology be selected inclusive of a fully engineered solution thoughtful to the ongoing support model.

Final Thoughts on Switching Paradigms

It can be difficult to switch from a paradigm of looking at only acquisition costs to looking at total cost of ownership. **Even though this document points to most of the reasoning behind this switch in paradigms, there are still many nuances of an individual company's network that go into selecting a switch that leads to the lowest cost of ownership**. The switches that QPC recommends typically have a manufacturer specified usable life for eight years when purchased new. It can be difficult to think ahead eight years to the future needs of a company. But one thing is for certain: if labor costs for implementation, management, repair, and replacement are all known factors, then there is no way that the company does not save money on a well-engineered, proven solution. Saving a few hundred, or even $1,000, on a new piece of hardware that is selected simply because it is cheap will likely run out any savings after the first few support issues. Multiply that over the eight-year lifespan of that equipment, and it quickly becomes evident that selecting the right hardware is a worthwhile investment. **QPC is has the experience and has proven time and again that we save our customers money by investing in the right solutions with lowest total cost of ownership in mind.**

Cisco versus EXOS – code simplicity example

Code simplicity leads to easy and lowered cost management. This is a simple example of adding ports to an existing VLAN in Cisco versus EXOS. This is a very typical example of what someone would do as a regular switch management routine.

Cisco

Enable

Configure

Interface 1/0/23-1/0/24,0/3/1

Switchport mode trunk

vlan tagging 1,10,7

EXOS

Configure vlan default add ports 23-25 tagged

Configure vlan WapMgmt add ports 23-25 tagged

Configure vlan Printers add ports 23-25 tagged

Looking at this example, there is no necessity to go into enable and configure mode switching context to select an interface range context. In Cisco, you cannot configure things unless you already have selected their context. This means you cannot simply issue commands and have them work. You have to always be selected in the right context. With EXOS, you can have your commands in TXT or Excel and just copy and paste them into the SSH session. No need to worry about context selection.

There is no need to tell the switch if it is an access, general, hybrid, or trunk type port. The switch just knows based upon the tagging configured. In EXOS, we call VLANs by their real human meaningful names instead of a number that you have to look up in a spreadsheet all the time. Selection of port ranges is obvious instead of having to look up how it is that Cisco does their port numbers. You waste time issuing commands when the command sets are not logical and obvious.

Considerations for used equipment

Equipment should not be used in a production environment if it is not under hardware warranty contract and a viable support contract that provides recovery for failure in the amount of time that matches the organization's tolerance for outages. For a core switch, the hardware replacement time should not be any longer than next business day.

Used equipment often has questionable provenance. It may be impossible to get a manufacturer support contract on used equipment. The equipment may be scam equipment from China. The equipment may already have an active infection installed on it that is intended to compromise your environment. And used equipment may already be 4 years into its s 8 year expected lifecycle. Investing labor into two migrations over the same time period is more expensive than one migration.

Negatives of PoE injectors

- Often PoE injectors are powered by a connection other than a UPS which can introduce power surges into the switches and connected equipment blowing out equipment that is very expensive.
- PoE injectors do not understand LLDP communications for power sipping and power negotiation like switches do.
- PoE injectors usually blast full power all the time which leads to higher electricity costs and prematurely burned out downstream connected equipment.
- PoE injectors lack remote diagnostics on power monitoring.

PoE injectors may make sense in an environment with very minor (less than 5) PoE devices as long as the devices are all powered by the battery protected and power conditioned portion of a quality UPS. Surge suppressors alone are typically not sensitive enough.