5 Questions to ask before buying a Layer 1 switch

Despite the apparent similarity of devices on offer to firms deploying or looking to deploy a Layer 1 switch, there are a number of sometimes subtle but important differences between them. All devices generally offer single nanosecond latency, dynamic patching and port replication, however the availability of features such as media conversion, signal regeneration, statistics and telemetry vary widely across the market.

Asking the following five questions before buying a layer 1 switch will give you a guide to available features and ensure that you are selecting the right switch for your needs.

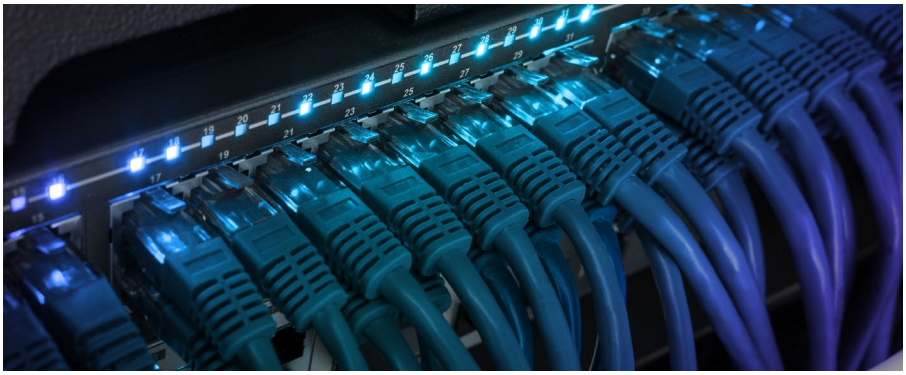

1. Does the switch perform Ethernet signal regeneration?

Typically inserting any device between Ethernet endpoints can have a detrimental effect on the bit error rate of the traffic carried as well as the ability to meet the Ethernet specifications at an electrical level. Multiple media conversions such as connecting a device via an optical transceiver, and patching it to another optical transceiver or direct-attach copper (DAC) cable can degrade the signal significantly. The results may be frequently detected as packet errors on the endpoints or even difficulty with obtaining a stable link.

The solution to this is for the Layer 1 switch ports to perform signal regeneration and clock data recovery. Doing so ensures that the outgoing Ethernet stream contains precisely the same data as the incoming data but with the signal and embedded clock margins equivalent to that of the originating data. In scenarios where a Layer 1 switch is used to fan out and replicate Ethernet streams without signal regeneration the probability of bit errors goes up with the number of replicated streams. With signal regeneration there is no such problem as each replicated signal margin is equivalent to the original data. Ethernet signal regeneration converts a Layer 1 switch from a device that may degrade data passing through it to one that actually reduces bit errors to receiving endpoints; in particular when endpoints have longer connections between them. Another key benefit of Ethernet signal regeneration is that it allows the use of different types of Ethernet media between two endpoints without compromising signal integrity. For example one end of a connection can be fiber and the other copper. The Layer 1 switch can also be used to convert individual links to DWDM wavelengths allowing many to share a single fiber pair.

2. What kind of monitoring/telemetry is provided?

Layer 1 switches are devices that operate on Ethernet links at the physical layer and therefore switch data statically based upon pre-configured rules rather than dynamically. Layer 1 switches need have no knowledge of the content of the Ethernet frame. As such most Layer 1 switches provide very few per port statistics and are often not even able to detect bit or frame errors in the Ethernet frames. Parsing the Ethernet frame and generating statistics requires additional logic. Even the lowest latency Ethernet switches typically require double-digit nanoseconds to do so which is why most Layer 1 devices do not offer them. Some vendors do however offer per port Ethernet statistics such as overall packet and byte counts, packet type counts, Ethernet error counts per type of error and frame size bin counts providing detailed visibility into each endpoint's Ethernet traffic.

As with most Layer 2/3 switches the Layer 1 switch counters are cumulative over time so calculating a time series generally requires periodically polling for a snapshot of the port counters then building it manually. If a Layer 1 switch does offer useful Ethernet statistics it is therefore worth also verifying that the Layer 1 switch has integrated telemetry features to perform more advanced analysis, displays, or even alerting based upon interface statistics.

3. Which tranceivers and media types are supported?

The SFP (1 GbE) and SFP+ (10 GbE) standards allow 1 and 10 GbE to be carried in a multitude of media types. Devices may be connected via types of optical fiber or electrical cables. There are multiple standards for Ethernet over optical fiber via the SFP/SFP+ interface with the most common being:

1000BASE-SX over multi-mode fibre up to 550 m
1000BASE-LX/LX10 over single-mode fibre up to 10 km
1000BASE-EX over single-mode fibre up to 55 km
1000BASE-ZX over single-mode fibre up to 90 km
10GBASE-SR over multi-mode fibre up to 400 m
10GBASE-LR over single-mode fibre up to 10 km
10GBASE-ER over single-mode fibre up to 40 km
10GBASE-ZR over single-mode fibre up to 80 km
10GBASE-ZR DWDM over single-mode fibre up to 80 km

The most common electrical SFP/SFP+ interfaces are:

1000BASE-T over twisted-pair cables terminated in RJ-45 connectors in lengths up to 100 m

10GbE direct-attach copper or twinax cables in lengths up to 15 m

A number of SFP+ transceivers support both 1 and 10 GbE modes and can be configured for either. From a Layer 1 switch perspective support for each of the above standards will require that it can recognise the type of transceiver inserted and configure it for the desired class of operation. It also must supply it with adequate power and remove the heat it generates within the SFP/SFP+ cage.

In particular LR, ZR and DWDM transceivers require higher power to be supplied and often specialized cooling within the device. Some optical transceivers can provide information on their light levels however the Layer 1 switch specifically needs to know how to query them; this is an extremely useful feature in troubleshooting errors on the link.

When twinax cabling is used in lengths greater than 3m, signal integrity can be an issue requiring that switches possess the appropriate logic to compensate for the varying electrical properties of the profusion of cables available on the market . Each has individually varying electrical characteristics that may or not meet the requisite Ethernet Alliance's standards (there is no IEEE standard for 10GbE over twinax).

All Layer 1 switches support at least some of the above SFP/SFP+ transceivers, it is important however to verify that they support the transceivers and media types that you currently require or may require at a future date.

When looking to leverage twinax cables, it is definitely worth finding a Layer 1 switch that has signal integrity features, such as integrated bit error rate testing (BERT) that maximize flexibility in the choice of twinax cables while maintaining the lowest possible bit error rate.

4. What management/operations/industry standard protocols does it interact with?

There is a profusion of standards available to interact, directly or programmatically with a network device.

Examples include:

Web-based Graphical User Interface (GUI)

Command-line interface (CLI) via secure shell (SSH), Telnet, serial connection

Local and remote logging via Syslog

JSON-RPC API

Simple network management protocol (SNMP) v1, v2, v3

NETCONF

It is important that those protocols running over the network offer security in the form of an encrypted transport layer such as the Transport Layer Security (TLS) or HTTPS. Integrating with existing authentication databases via protocols such as the Terminal Access Controller Access-Control System (TACACS) or Remote Authentication Dial-In User Service (RADIUS) are also useful features. Devices may offer support for some or all of these protocols.

5. What is its latency and how is it measured?

In general Layer 1 switches deliver single digit nanosecond levels of pass-through latency. Ethernet traffic can pass through them and optionally be fanned out in as low as 4 nanoseconds. Actually measuring and confirming this latency is not a simple exercise. Where are these latency values measured from; the transceiver cages? Do they include a transceiver? What is the quoted latency number; is it a minimum, mean, median or other value? Across which ports on the device? The two with the lowest latency between them? All of them? Was the result obtained indirectly by statistical averaging or directly via high-precision packet capture?

Decision makers are often asked to take these numbers on trust by vendors as testing it for themselves is usually rather difficult mainly due to the aforementioned complexities. There are independent, unbiased specialist entities such as STAC Research who perform these sort of tests. It is certainly worth asking the vendor precisely what the measurement is that they quote and if they have any independent results to back up their numbers.

In summary all Layer 1 switches provide basic Layer 1 functionality however there are key features that vary significantly from product to product. These include questions as to the latency and how it is measured, whether it provides Ethernet signal and clock regeneration, what traffic statistics are available, what media types are supported - and can be converted between - as well as what protocols can be used to manage the switch.

Ask yourself those questions and you are sure to select the right device for your requirements.

Santa Clara—Corporate Headquarters
5453 Great America Parkway,
Santa Clara, CA 95054

Phone: +1-408-547-5500
Fax: +1-408-538-8920
Email: info@arista.com

Ireland—International Headquarters
3130 Atlantic Avenue
Westpark Business Campus
Shannon, Co. Clare
Ireland

Vancouver—R&D Office
9200 Glenlyon Pkwy, Unit 300
Burnaby, British Columbia
Canada V5J 5J8

San Francisco—R&D and Sales Office 1390
Market Street, Suite 800
San Francisco, CA 94102

India—R&D Office
Global Tech Park, Tower A & B, 11th Floor
Marathahalli Outer Ring Road
Devarabeesanahalli Village, Varthur Hobli
Bangalore, India 560103

Singapore—APAC Administrative Office
9 Temasek Boulevard
#29-01, Suntec Tower Two
Singapore 038989

Nashua—R&D Office
10 Tara Boulevard
Nashua, NH 03062