

Ethernet Switches Performance & Types

Ethernet LAN Switch Market Classes

LAN switches have been generally categorised in the following three market classes:

1. The Node/Workgroup Switch
2. The Segment Switch
3. The Backbone Switch

The Node/Workgroup Switch

The node switch (also generally referred to as a "workgroup" switch) is intended to provide increased bandwidth and speed for a single workgroup (typically with a small number of 10Base-T nodes). These workgroups usually have only one server and each

client may have a direct connection to the switch. Some of these switches allow for upto 4000 addresses in their address tables. Hence they can also be successfully used as low-end segment switches. Some switches in this category have one or two dedicated fat pipes.

The Segment Switch :Segment switches interconnect workgroups. On each port of a segment switch, you may connect a node/workgroup switch or a shared media hub. Segment switches have larger address tables, are more flexible in their port configurations, and support higher throughput speeds. (All switches have address tables. The address table has two columns. In the first column all the port numbers are listed and in the second column, the ethernet addresses of nodes available on the respective port are listed.) On each port of a segment switch, you have a workgroup switch or a hub which again have nodes connected to each of its ports. So the number of nodes indirectly connected to each port of a Segment switch is high which means larger address tables and higher memory requirement to hold them.

The Backbone Switch: Backbone switches are intended to interconnect a large site and provide connectivity to remote services. The backbone switch is highly modular and provides connections to ATM and FDDI. Often, they have options for built-in routing services.

Store & Forward Vs Cut Through Switches

The store-and-forward technology buffers data packets in order to perform a complete error check and filtering, and then forwards the packet out of the switch. The more traditional cut-through technology immediately forwards a packet out of the switch to decrease latency (the amount of time the switch takes to process a packet), but typically at the cost of reduced error checking within the switch. The store-and-forward method is required to move a frame from a low-speed 10 Mbps LAN to a high-speed LAN. This method is increasingly being used more for error checking and for dual speed (10 Mbps/100 Mbps) data forwarding. Since most switches today provide 10 Mbit/s to 100 Mbit/s switching, they will be in store-and-forward mode.

ASIC Design

Switches can be either Application Specific Integrated Circuit (ASIC) or processor-based. Processor-based switches are built with existing industry-standard processors with the switching done in software. ASIC switches are more of a combination of hardware and firmware with the Ethernet switching procedures entirely encapsulated within the ASIC. ASIC switches are sometimes preferred over processor-based switches because they are much faster.

Shared Backplane

Ideally, switches would implement a pure cross-point matrix architecture. A matrix is basically a single ASIC that meshes multiple communication paths with each port having a dedicated path to every other port. Unfortunately, the pure matrix architecture does not lend itself well to expandability, flexibility, or cross technology switching (for example, 10Base-T to 100Base T switching). A pure cross point matrix also requires a lot of circuitry.

The other alternative to cross point architecture is a high-speed shared backplane architecture. With a high-speed shared backplane, a switch shares its backplane with all domains connected to it. Unlike a pure matrix switch, there is a shared data path from point A to point B. The shared backplane is fast enough to process all data requests, making certain there is no contention on the attached segments. The backplane speed could be typically 1Gbps. In shared backplane technologies, high-speed backplanes can easily handle multiple simultaneous conversations.


Virtual LAN Implementation

A switch's architecture should also be able to support the creation of Virtual LANs or VLANs. VLANs are groups of users that are defined based on their logical function rather than their physical location. For instance, criteria for a VLAN membership may be that a user use a common protocol or be part of the same department function. Regardless of where the device/user is physically located, workgroups can be defined based on their logical function via simple port

configuration. VLANs offer increased flexibility, performance, and filtering. Network users assigned to a VLAN are considered a separate broadcast domain so that packets are only forwarded between ports that are designated for the same VLAN. With VLANs, cross domain broadcast traffic is minimised and bandwidth is saved by not allowing packets to flood throughout the network. A VLAN restricts broadcast traffic to a given domain. They create barriers between the domains by grouping ports together into separate VLANs. These barriers, or actually, "lack of communication paths" help to provide security between the domains. Security mechanisms can be enforced between the separate VLANs.

In port-based VLANs, workstations assigned to a VLAN are identified by the port they are connected to. A port-based VLAN is a grouping of ports that belong to a single broadcast domain. For example, ports 1 through 5 may be assigned to VLAN 1, while ports 6 through 10 may be assigned to VLAN 2. Therefore, traffic (broadcast, multicasts, unicasts) from workstations on

Multiple VLANs Within a Single Switch



- Isolates broadcasts and multicasts
 - (can route between VLANs internally)
 - Full management of each virtual LAN

5009-1
12/95 Mac C

ports 1 through 5 is restricted to only ports 1 through 5. Likewise, traffic from ports 6 through 10 is contained within ports 6 through 10. Port-based VLANs not only help to reduce broadcast traffic, but they can be used to enhance security by isolating traffic.

Port Trunking: It is the ability to use switched ports in unison. This approach effectively connects switches across multiple ports, resulting in total bandwidth between switches equal to the sum of the bandwidth provided by individual interswitch links.

Port trunking is used to allow several ports to be connected together and treated as a single, higher speed port. This enables two switches to be interconnected by multiple links with all those links acting as one higher-speed link. For example, with port trunking, you can link two switches together with two 100 Base T ports. This would actually double your switch-to-switch performance compared to connecting the two switches with a single 100 Base T port.

Evaluating Switch Performance : The performance criteria when evaluating switches relates to the speed and efficiency with which a switch delivers data. Performance measurements include: latency, throughput, and its ability to deal with congestion.

Latency

As mentioned earlier, latency is the time a switch takes to process a packet. It is the amount of time between when a switch receives a unit of data and when that unit of data is forwarded out from the switch. Latency is measured differently depending on whether its for a device using the store-and forward method or for a device using the cut-through method. Latency for a store-and-forward device is measured by LIFO, while for a cut-through device, it is measured by FIFO.

LIFO means last (bit) in, first (bit) out. This is the length of time it takes from the moment the last bit of a frame is received in a port, until the moment of time, when the first bit of a frame is sent out a destination port. Typical LIFO latency is 3.4 micro seconds.

FIFO means first (bit) in, first (bit) out. This is the length of the time it takes from the moment the first bit of a frame is received on a port until the point of which the first bit is sent out to the destination port.

Because the methods for measuring latency differ, it is important that you compare latency for store-and-forward devices to store-and-forward devices and latency for cut-through devices to cut-through devices.

Throughput & Packet Loss Rate

Throughput is the rate of data transfer the switch can sustain without packet loss. While latency measures the delay of a single frame, throughput measures the number of packets, or frames per second without packet loss. In a switch, throughput is typically measured in packets per second (PPS), but may be also referred to as frames per second (FPS).

Packet Loss Rate (PLR) is the percentage of packets that the switch does not forward within a time window of when the data was sent. Under PLR, a packet is

termed "lost" if it doesn't get forwarded within a certain time period, even if the switch ends up forwarding it from one of its buffers outside the time window.

PPS values vary due to the difference of capacity of the switches. To effectively compare capacity between switches, look at precisely the number of switched ports, and the media speed of the ports that are switched. It would be unfair to compare possibly excellent capacity and throughput numbers of a 6 port Ethernet switch to that of a 12 port.

Congestion Control

Congestion occurs any time more packets are destined for a particular segment than that segment can handle. It is when the switch is unable to release the data onto the destination media. The data is then held in buffer memory until it can be sent to the destination. If enough data is buffered, for congestion or other reasons, and the buffers are completely consumed, frames will be dropped since there is no place to store them. Some switches attempt to handle congestion by making a heavily used segment appear to have numerous

collisions, causing all stations on that segment to back off in their sending of packets. This method is called backpressure.

There are other features which could be considered while evaluating an Ethernet Switch:

- ? Flexibility in providing various types of ports like 10 Base T, 100 Base T, 100 VG Any, FDDI etc
- ? Redundant Power Supply Unit
- ? Hot Swappable port modules

Conclusion: This article has dwelled into Ethernet Switch Market Classes. Different features and performance criteria which could be considered while evaluating a switch was also discussed. This will enable you to evaluate your customer needs better and make a better sale.