

■ CLEAR CHOICE TEST 10 GIG ACCESS SWITCHES

10 Gig access switches: Not just packet-pushers anymore

Extensive testing of seven leading switches turns up major differences in multicast, security, manageability

BY DAVID NEWMAN, NETWORK LAB ALLIANCE MEMBER

Pity the humble access switch. These packet pushers usually work so well they're stuffed into wiring closets and promptly forgotten. Packet in, packet out. End of story. Or is it? If the results of *Network World's* latest switch tests are any guide, network managers may need a new lexicon just to make buying decisions. Our tests found seven next-generation switches bristle with features that don't exist in many previous models — not just physical features such as 10 Gig Ethernet uplinks but also 802.1X-based network access control (NAC) authentication, enhanced multicast support, storm control, denial-of-service (DoS) protection and IPv6 support.

We assessed these switches — all of which sported 48 10/100/1000Mbps ports and two 10 Gig ports — in 10 areas, encompassing Layer 2 and Layer 3 IPv4 unicast and multicast performance, Layer 2 multicast group capacity, 802.1X support, storm control, management and usability, power consumption, and features. (See "How we did it" at www.nwdocfinder.com/4121.)

Overall, we found big differences in support and stability in products tested from Alcatel-Lucent, Cisco, Dell, D-Link, Extreme Networks, Foundry Networks and HP. For example:

- Multicast throughput and latency varied widely, but more basic issues such as group capacity and even system stability were bigger differentiators in our tests. It took multiple software builds from some vendors just to get through industry-standard multicast tests, and then only using very different group counts.
- All switches supported 802.1X authentication, but there were major variations in the level of granularity of access control. Not every switch supported some common-use cases, and two switches forwarded unauthenticated traffic when operating in multi-auth mode, a security issue.
- All devices had "storm control" features to help mitigate DoS attacks, but these varied widely in terms of rate control and signature detection.
- IPv6 support remains a work in progress. Some switches fully support IPv6; others can route IPv6 packets but can't be managed over IPv6; yet others lack support for IPv6 routing protocols.

No switch excelled in all of the many areas we examined, making it difficult to pick winners across the board. Most switches do fine on simple forwarding of Ethernet and IPv4 unicast traffic. If that's all that matters to you, pick a switch on price and usability.

We wouldn't recommend that, though. Increasingly other areas matter more, including security, multicast and IPv6 — and that's where real variations among products exist. The Cisco Catalyst 3750E is the most fea-

ture-complete device we tested, though HP's ProCurve 3500yl, Extreme's Summit X450 and Foundry's FastIron X448 also fared well in most areas.

Because access switches do more than previous-generation products, the first step in picking a product is determining which features matter most — Layer 2 vs. Layer 3, IPv4 vs. IPv6, unicast vs. multicast, managed vs. unmanaged, on-board security vs. no security — and choosing the device that did the best job in these areas.

Unicast performance

Once upon a time, Layer 2 unicast performance tests would have produced by far the most important results, but that's changed. Measuring unicast throughput on all ports, once considered the acid test for access switches, is no longer a major differentiator. Even in the most stressful test case — with a Spirent TestCenter traffic generator blasting minimum-length 64-byte frames at all switch ports — throughput was at or very close to line rate for all switches except D-Link's DGS-3650.

We observed similar results when measuring throughput for 256- and 1,518-byte frames, both in Layer 2 (switched) and Layer 3 (IP forwarded) configurations. Throughput isn't the differentiator it once was.

After we completed testing, D-Link objected to our methodology, saying it isn't indicative of real-world conditions. We take D-Link's point, and hope no network manager would consider running a production network at 99% utilization or above. But we've heard this before many times and believe it misses the point. No one ever represented that industry-standard throughput testing practices use real-world traffic patterns (never mind that reality differs vastly from network to network). Rather, the goal is to determine the limits of switch performance.

Multicast group capacity

If unicast performance didn't differentiate products, multicast performance certainly did. We assessed multicast by measuring group capacity, and Layer 2 and Layer 3 multicast throughput and latency. Multicast group counts turned out to be major differentiators, not just in the capacity tests but also in the throughput and latency tests.

The goal of the group capacity tests was to determine the maximum number of Internet Group Management Protocol Version 3 multicast groups each switch could handle. This is a key measure of multicast scalability: The more groups a switch can track, the more users can do with multicast.

Because this is an access switch test, we configured each device in Layer 2-only mode and enabled IGMP snooping. Then we configured the

CLEAR CHOICE TEST 10 GIG ACCESS SWITCHES

Spirent TestCenter traffic generator/analyzer to join some number of groups, and measured whether the switch would forward traffic to all groups without flooding (see “Breaking the standards,” www.nwdocfinder.com/4032).

The results reveal lots of variation among products, with group capacity ranging from nearly 1,500 for HP’s ProCurve to less than 70 for Dell’s PowerConnect (see graphic, this page). For enterprises that need 70 or fewer multicast groups for the life of the switch, this isn’t an important distinction; for everyone else — this includes most midsize and large enterprises, and many small ones as well — group counts do matter.

The capacity test focused only on maximum group count. When it came to measuring throughput and latency, the group counts supported by each switch were lower in some cases than others. (See “Some switches support lower multicast counts at Layer 3” at www.nwdocfinder.com/4122.)

In part the difference is explained by switch configurations. We measured Layer 2 throughput and latency using more or less the same topology as in the group capacity tests. In the Layer 3 tests we enabled Protocol-Independent Multicast (PIM), a multicast routing protocol, essentially putting a router on every port. Judging from the supported group counts where less than half the switches hit the 500 group-count mark, this is far more stressful on the device under test.

It is important to note, though, that it took multiple software builds for some vendors to obtain these group-count results. Our initial multicast tests of the Alcatel-Lucent, Dell, D-Link and Foundry switches with 500 groups led to lockups or reboots. All these vendors supplied software updates that led to more stable switches. However, as the results show, not all could be tested with 500 groups. If a switch could not hit the 500-group mark we had outlined for throughput and latency testing, we tested Layer 2 and Layer 3 multicast throughput and latency at the switch’s maximum group capacity.

HP’s ProCurve did support 500 groups, but with a twist: In Layer 3 testing, it could use only two virtual LANs (VLAN), IP subnets and PIM router instances, compared with 49 on all other devices. This limitation would rule out the use of this ProCurve switch in situations where more than two subnets and multicast routing instances are needed.

Several vendors observed that few customers support 500 multicast groups at the edges of their networks. But we can argue that conditions may be changing. In some industries, notably financial services, it’s common to support dozens to hundreds of multicast group subscriptions for stock-quote applications. Multicast scalability may not be a top priority in choosing network devices yet, but it could become more important.

Switch jitters

Latency, the length of time a switch buffers a frame, also is a key metric, more important than throughput for such real-time applications as voice and video. In fact, multicast throughput turned out to be a nonissue in our tests, with all products moving packets within 0.5% of line rate.

For unicast traffic, differences between products handling midsize frames were relatively minor, but average and maximum unicast latencies differed widely when switches handled minimum- and maximum-length frames. (See “Unicast latency for switches varies more with larger packets” at www.nwdocfinder.com/4123.) In particular, Foundry’s X448 exhibited unusually high average and maximum delays when handling large frames. The vendor says it hasn’t seen this result in other tests, but it occurred more than once in our lab.

Multicast latencies varied much more, with a 500-fold difference between the lowest and highest result — both from HP’s ProCurve switch (see graphic above). A big delta between average and maximum latency may indicate an issue with jitter, or latency variation, which can have an adverse effect on delay-sensitive applications such as voice and video. The HP and Alcatel-Lucent switches exhibit much greater variation than other switches between average and maximum multicast latency, with spreads of hundreds or thousands of microseconds. In contrast, all other switches held up traffic at most 1 to 4 microsec.

The Alcatel-Lucent and HP switches also exhibited much higher latency for multicast than unicast. Conversely, Foundry’s X448 did far better with large-frame latency when handling multicast traffic. The traffic topologies differed in the unicast and multicast tests, making the comparison a bit unfair, but given that switches move unicast and multicast alike in silicon we were surprised to see any differences.

Authentication: Six scenarios, seven stories

Many switches today support 802.1X authentication, a building block in NAC. The key question is what kind of access authenticated users can expect. In the six scenarios we developed for this project, we uncovered major differences among products in terms of the conditions under which they’ll grant access, as well as what sort of access they’ll permit.

In the first 802.1X scenario, a client (or supplicant, in 802.1X-speak) gets authenticated, and the switch places the client into a statically defined VLAN. All switches passed this basic test, in which the switch connected Juniper Odyssey supplicants to a Juniper Steel-Belted Radius server (see “Switches vary on 802.1X authentication support, page 4).

The second scenario, involving multi-auth, turned out to be the most problematic, with failures from the Cisco and Dell switches. In this sce-

NETRESULTS

Product	Cisco Catalyst 3750E-48PD-EF Series Switch	ProCurve Switch 3500yl	Summit X450a-48t
Vendor	Cisco www.cisco.com	ProCurve Networking by HP www.procurve.com	Extreme Networks www.extremenetworks.com
Price*	\$33,980	\$16,096	\$14,480
Pros	Very extensive feature set; strong multicast scalability and performance	Strong unicast and multicast throughput and latency; highest Layer 2 multicast scalability	Strong unicast and multicast throughput and latency; passed all 802.1X test cases; extensive feature list
Cons	Forwarded unauthenticated data in one 802.1X case	Limited Layer 3 multicast and IPv6 support in version tested	Limited multicast scalability; factory reset left some personally identifiable information
Score	4.49	4.46	4.35

*Price as tested for switch with at least 48 10/100/1000 Ethernet ports; two 10G Ethernet ports; two 10GBase-SR transceivers; and all necessary software for IPv4 and IPv6 unicast and multicast traffic handling.

CLEAR CHOICE TEST 10 GIG ACCESS SWITCHES

nario, there are multiple users attached to a single switch port, and each must be authenticated before being granted network access. We attached multiple users using an unmanaged hub (a common-use case in many corporate conference rooms where there's only one Ethernet drop). Other uses for multi-auth include IP phones (which sometimes have a two-port switch to attach a PC through the phone) and wireless LAN (WLAN) access points (especially thin access points, which attach to a switch/controller and field associations from multiple wireless clients).

Most switches — other than the one from Extreme — require that multi-auth be explicitly configured.

After doing so, the Cisco and Dell switches authenticated the first user — but then allowed traffic from the second and subsequent users onto the network without authentication. The physical-world analogy of this behavior is “badge tailgating,” in which someone with a door badge enters an office building and others follow while the door is open. The security implications are obvious.

Cisco says it strongly discourages customers from using multi-auth except for certain uses, such as an IP phone with a PC attached, and then encourages customers to segregate traffic onto different VLANs.

Strictly speaking, multi-auth is a violation of the IEEE's 802.1X standard. The spec's media access control (MAC) relay function (the port access entity) includes a logical switch that's on or off. There's no provision for a sort of “selective on/off” state that permits some frames but denies others (see “Breaking the standards,” www.nwdocfinder.com/4032).

Still, because there are common-use cases for multi-auth, it's fairly widely supported. The danger, as our test results show, is that network managers may be lulled into a false sense of security, erroneously believing that enabling 802.1X will result in authentication for all traffic.

The third scenario, involving dynamic VLANs, was far more straightforward. This one modeled networks in which roving laptop users may plug into any switch port at random. The goal was for the switch to dynamically assign a switch port into a given VLAN after authentication.

All switches but one passed this test; the lone exception was Dell's PowerConnect 6248, which doesn't support dynamic VLAN assignment. Extreme's X450 goes the other way: Not only did it pass this scenario, but it allowed the supplicant to be placed into multiple untagged VLANs.

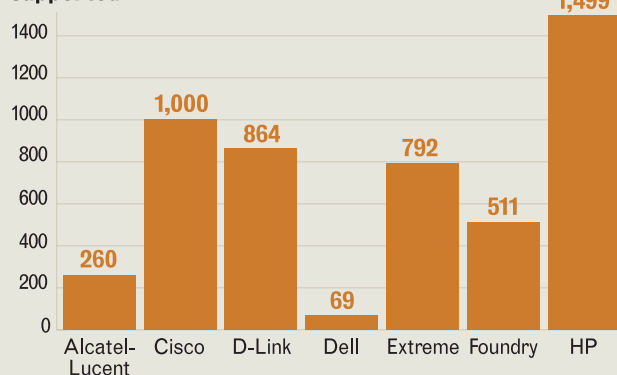
In the fourth scenario, we determined whether the switch could dynamically enable an access control list (ACL) upon authentication, governing where the client can go. As with dynamic VLAN allocation, dynamic ACLs can be useful with mobile work forces, where employees should gain access to specific resources regardless of location.

The Cisco, Extreme, Foundry and HP switches all support this feature.

Tracking multicast group capacity

As a key measure of multicast scalability, this group capacity test determined the maximum number of IGMPv3 multicast groups each switch could handle.

IGMPv3 groups supported



We needed to use an undocumented syntax to get dynamic ACLs to work with the HP switch, but the vendor says this has been corrected in currently shipping software (we did not verify this). Switches from Alcatel-Lucent, D-Link and Dell do not support this feature.

So far, all the 802.1X scenarios have covered situations in which authentication succeeded. In our fifth scenario, we deliberately failed authentication to determine whether switches would place a client into a guest or restricted VLAN. This is a common requirement, not just for enterprise employees who mistype a password but also for visitors and contractors who may not have authentication credentials. All switches tested offer a guest VLAN capability without issue.

In our final test scenario, we looked for the switch to concurrently support both 802.1X clients and non-802.1X clients. For better or worse, 802.1X isn't yet pervasive. There are large numbers of networked devices, such as printers, that do not have 802.1X supplicant software. For this, it's desirable to have a feature Cisco calls “MAC authentication bypass.”

All switches we tested, except those from D-Link and Dell, support fallback to MAC authentication with a non-802.1X client. D-Link's DGS-3650

FastIron Edge X Series 448+2XG-PREM

Foundry Networks
www.foundrynet.com

\$15,985

Strong multicast performance; passed all 802.1X test cases; extensive feature list

Higher power consumption; larger form factor compared with other switches

4.23

OmniSwitch 6850 Model OS6850-48X

Alcatel-Lucent
www1.alcatel-lucent.com

\$13,685

Strong unicast performance; lowest idle power consumption

Limited multicast scalability; verbose CLI; factory reset left some personally identifiable information

4.05

PowerConnect 6248P

Dell
www.dell.com

\$5,779

Strong unicast and multicast throughput and latency

Limited multicast scalability; limited 802.1X support; forwarded unauthenticated data in one 802.1X case

3.58

DGS-3650

D-Link Systems
www.dlink.com

\$8,841

Strong multicast throughput and latency

Much lower unicast throughput and latency than other switches; limited 802.1X support; limited storm-control granularity

3.55

CLEAR CHOICE TEST 10 GIG ACCESS SWITCHES

Switches vary on 802.1X authentication support

As a basic building block to many network access control (NAC) schemes, 802.1X authentication support is required in any modern day access switch. We tested these switches in six 802.1X authentication scenarios and the level of success was all over the map.

	Alcatel-Lucent	Cisco	D-Link	Dell	Extreme	Foundry	HP
One user	Pass	Pass	Pass	Pass	Pass	Pass	Pass
Two users	Pass	Fail	Pass	Fail	Pass	Pass	Pass
Dynamic VLAN	Pass	Pass	Pass	Not Supported	Pass	Pass	Pass
Dynamic ACL	Not Supported	Pass	Not Supported	Not Supported	Pass	Pass	Pass ⁽¹⁾
Guest VLAN	Pass	Pass	Pass	Pass	Pass	Pass	Pass
MAC fallback	Pass	Pass	Fail ⁽²⁾	Fail ⁽³⁾	Pass	Pass	Pass

1. Syntax issue with 12.25 software, corrected in current 13.x release.

2. Supports MAC authentication, but not concurrently for 802.1X and non-801.X clients. Switch instead puts failed 802.1X clients and non-802.1X clients into a guest VLAN.

3. Does not support MAC authentication; switch does allow a user-defined number of MAC addresses to be learned dynamically.

supports MAC authentication but not concurrently with 802.1X. Dell's PowerConnect 6248 does not support MAC authentication, although it can restrict access to a user-defined number of MAC addresses.

The Cisco Catalyst 3750E also supports three 802.1X scenarios we didn't test for. It can place non-802.1X clients into a special restricted VLAN, distinct from a guest VLAN for unauthorized or unremediated 802.1X clients. It can automatically fall back to Web-based authentication if 802.1X authentication doesn't occur within a given timeframe. And it can authenticate multiple devices on a port and place each in a different VLAN (this is different than the multi-auth case in which all devices enter the same VLAN). We didn't test any of these additional capabilities.

Management and security

In assessing switch management and security, we sought to answer three questions: Did devices follow current best practices by default? Could users configure switches to follow these best practices? And could switches be wiped clean of any sensitive information before being taken out of deployment?

The "wipe clean" question stems from regulatory requirements in a growing number of industries. For example, the National Institute of Standards and Technology, the U.S. government's standards body, and the credit card industry's Payment Card Industry Data Security Standard (PCI DSS) both require the deletion of any personally identifiable information before disposal.

We assessed reset capabilities by deleting each switch's start-up configuration file after putting it through performance and security tests. For all but the Alcatel-Lucent, Extreme and HP switches, that was enough to wipe the systems clean. HP's ProCurve switch stores passwords separately in flash memory, but these can be deleted by using the front-panel buttons. The procedure is documented, and HP says it's moving toward inclusion of encrypted passwords in the switch configuration file.

The Alcatel-Lucent and Extreme switches retain passwords even after a factory reset. In addition, Extreme's Summit X450 also retains the private SSH key, which could allow an attacker to pose as an authorized device even after the switch has been retired.

We also determined which management methods were enabled by default, and which would need to be enabled or disabled by network managers (see "Tracking support for management and security best practices" at www.nwdocfinder.com/4124).

These best practices include disabling nonsecure management methods such as telnet (supported out of the box over IPv4 by all switches by default), Web and SSHv1. Best practices means accessing the switch only through secure methods such as SSHv2 and/or Secure-HTTP and also logging switch events to a syslog server (a requirement under many

enterprise security policies).

The Cisco Catalyst 3750E adhered the closest to security best practices. However, it supports telnet by default, as do all switches. Also, when enabling SSH the Catalyst supports the nonsecure Version 1 of that protocol (although SSHv1 can be disabled via an additional command).

In general, management over IPv6 isn't as solid as over IPv4. Two switches, from Dell and HP, didn't support IPv6 management on their default VLANs in our tests, although HP says it's currently shipping 13.x software that does support IPv6 on the default VLAN. Also, there were a couple of cases in which options offered with IPv4 weren't available over IPv6. We were unable to configure syslog over IPv6 on the Cisco Catalyst 3750 or Extreme's Summit X450. And the Extreme switch didn't support Web- or SSL-based management over IPv6.

As with multicast and 802.1X, IPv6 support is relatively new in many switches, and support for all features is far from complete. For network managers considering IPv6 deployment, it's not enough to consider whether a switch will forward IPv6 packets; supporting management over IPv6 is critical as well.

Sharp-eyed readers will notice we haven't covered SNMP management over either IPv4 or IPv6. Problems with our test bed setup prevented us from completing SNMP verifications; however, SNMP support is covered in our features section.

Storm control

Tools to block DoS attacks, once the exclusive purview of intrusion-detection/prevention systems, are now included in most switches' security arsenals. While all switches we tested can classify and block malicious traffic, there are differences in the depth of coverage.

At a high level, "storm control" takes two forms: Rate-controlling traffic and blocking specific attacks. Rate control in turn may be divided into separate commands for throttling unicast, broadcast and multicast traffic, though not all switches support this. For example, it may be desirable to set one drop threshold for unicast traffic (TCP SYNs, say, to block a SYN flood attack) and another threshold for broadcasts (perhaps to avoid overwhelming the switch's CPU).

All switches offer the ability to throttle traffic. The D-Link 3650's rate controls are limited to broadcast and multicast traffic, while Extreme's Summit X450's rate controls specifically target CPU-bound packets. Dell's PowerConnect 6248 Web-based GUI appears to allow rate control over only one class of traffic at a time (unicast, broadcast or multicast), but in practice different classes with different thresholds can be defined by issuing multiple commands. The other switches (and the Dell PowerConnect's CLI) all support individual commands for throttling different traffic classes.

CLEAR CHOICE TEST 10 GIG ACCESS SWITCHES

SCORECARD

Action and weighting	Cisco	HP	Extreme	Foundry	Alcatel-Lucent	Dell	D-Link
Layer-2 unicast performance (15%)	4.50	4.75	4.50	4.17	4.67	4.42	3.00
Layer-3 unicast performance (15%)	4.50	4.67	4.50	4.08	4.58	4.42	3.17
Layer-2 multicast group capacity (10%)	5.00	5.00	4.25	4.00	3.00	2.00	4.25
Layer-2 multicast performance (10%)	4.50	5.00	4.75	4.58	3.50	4.50	4.50
Layer-3 multicast performance (5%)	4.50	3.67	4.75	4.67	4.00	4.58	4.50
802.1X/NAC support (10%)	3.50	4.50	5.00	5.00	3.50	2.00	3.00
Storm control (5%)	5.00	4.50	4.25	4.50	4.50	3.00	2.50
Management and security (10%)	4.00	3.50	2.50	3.75	3.50	2.75	3.50
Power consumption (5%)	4.25	4.00	4.50	2.50	5.00	4.50	4.50
Features (15%)	5.00	4.25	4.50	4.50	4.25	3.50	3.50
Total	4.49	4.46	4.35	4.23	4.05	3.58	3.55

Scoring key: 5: Exceptional; 4: Very good; 3: Average; 2: Below average; 1: Subpar or not available.

Attack signature detection varied widely among switches. Some devices — such as those from Alcatel-Lucent, Dell, Extreme and Foundry — include signatures for between two (Foundry) and 29 (Extreme) well-known forms of attack, and to drop these packets HP's ProCurve 3500 uses an anomaly-based approach it calls "virus throttling" to detect and block malicious traffic. Foundry's FastIron X448 also has hooks that tie into external monitoring tools, such as SFlow monitors or a Snort IDS, that will drop frames when traffic matches a given signature.

No spoofing allowed

Some switches also support antispoofing mechanisms targeting Dynamic Host Configuration Protocol, Address Resolution Protocol and even plain-vanilla IP traffic.

All switches support DHCP snooping (D-Link calls this "DHCP server dynamic binding"), which sets up a binding between an authorized DHCP server's IP and MAC addresses. This helps prevent clients from receiving bogus addresses from a rogue DHCP server.

The Cisco and Foundry switches support IP source guard, which is conceptually similar to DHCP snooping. The switch blocks all traffic until it sees a valid DHCP conversation, then it will allow traffic only from that IP-MAC binding. This helps prevent some man-in-the-middle attacks in which an intruder spoofs a source IP address.

The Cisco, Extreme, Foundry and HP switches also support "dynamic ARP inspection," which will drop any packet with previously unseen IP-MAC bindings. This is useful in preventing some man-in-the-middle attacks, in which an attacker poses as a previously seen station and redirects traffic through a different switch port using a new MAC address.

Usability

Any assessment of switch usability is subjective. While there are some objective measures that can be applied (for example, it might take 17 steps to enable SSH on one switch and five on another), usability assessments ultimately come down to what's most comfortable for the user.

For most of the industry, "comfort" means a command-line interface (CLI) that is or closely resembles Cisco IOS. It hasn't escaped the attention of Cisco competitors that more network engineers are trained in IOS than any other CLI. In this test, the Dell, Foundry and HP CLIs were very IOS-like. HP's was probably the closest, with Foundry's close behind (although they use different syntaxes for referring to physical and virtual interfaces). The Dell CLI was inconsistent in a few places. For example, some commands refer to an interface with an Ethernet prefix and others don't.

The Alcatel-Lucent, D-Link and Extreme switches use homegrown CLIs.

Perhaps it's our greater familiarity with it, but we found Extreme's XOS CLI by far the easiest of these to learn and navigate. It also offers some useful monitoring features we didn't see in other switches, such as the ability to monitor port statistics, even across multiple ports, in real time.

The CLI in Alcatel-Lucent's OmniSwitch saves configurations twice, in "working" and "certified" directories. This feature can be very useful in testing new configurations, because it lets network managers roll back to a known good configuration in case of error.

One aspect of the OmniSwitch CLI we didn't like: Unlike all other switches tested, it can't execute the shortest unambiguous version of a command. For example, while most switches will understand that "sh run" is an abbreviation for "show running-configuration," the OmniSwitch instead must receive the longhand version of "show configuration active." The fact that the OmniSwitch has tab completion for commands is only partial compensation; switch configuration would run faster if the CLI accepted abbreviated commands, like all others.

The D-Link DGS-3650's CLI configuration syntax is verbose, sometimes too much so. As with many other CLIs, typing the tab key will display options for completing a command. Unlike all others, the switch places a full string on the command line, which the user has to erase. For example, typing "config vlan <tab>" places the string "config tab <vlan_name 32>" on the command line, and the user must delete "<vlan_name 32>" before continuing. While it's useful to know a VLAN name can be 32 characters long, the need to erase strings got tiresome after a while.

We also did a quick review of vendors' documentation. While all documentation adequately described the commands available on each switch, they differed in explaining the basic technology behind each command, and why users would want to use (or not use) that technology. Dell's and D-Link's documentation included relatively little technology background. Of the others, we considered the Cisco, Extreme and Foundry documentation to offer the most complete technology tutorials. HP's documentation also is first-rate but doesn't cover as many features as some of the other switches, especially for IP multicast.

One usability area we did not assess was the Web-based management of each system. We freely admit a bias for the CLI. While we're sure there are plenty of fans of graphical management, we're not among them.

Switch features

While this test's key takeaway may be the big differences in new features, the good news is that, with a very few exceptions, all switches support the same basic Layer 2/Layer 3 functions. They're all 1U systems with 48 10/100/1000Mbps ports and at least two 10 Gig Ethernet uplinks

CLEAR CHOICE TEST 10 GIG ACCESS SWITCHES

(except Foundry's FastIron X448, which is 1.5U high). They all offer basic Layer 2 and Layer 3 IPv4 forwarding features, and full support for VLANs, 802.3ad link aggregation and Layer 2 and Layer 3 QoS controls. All even re-mark diff-serv codepoints (DSCP), a best practice when classifying traffic for QoS treatment. (It's not a good idea to trust incoming DSCPs.)

Differences start to appear beyond these basics. For example, Foundry's FastIron X448 is the only one not stackable (Foundry has other stackable products but supplied the X448 for this project). MAC address capacity ranges from 8,192 for Dell's PowerConnect 6248 to more than 64,000 on HP's ProCurve. And the Alcatel-Lucent and D-Link switches were the only two not yet supporting the IEEE's 802.1AB Link Layer Discovery Protocol (LLDP), a relatively new standard describing how link partners can exchange capabilities information.

Power over Ethernet (PoE), often used at the edges of enterprise networks to drive IP phones and WLAN access points, is another differentiator. Every vendor in this test sells PoE-capable switches, but only Cisco, Dell, Extreme, Foundry and HP supplied PoE gear. The Cisco Catalyst 3750E is the only device tested capable of delivering power to all 48 downlink ports simultaneously. The others require an external power supply to do so. (We didn't measure PoE power consumption; these figures are from the vendors' responses to our features questionnaire.)

IPv6 support also varies widely. For most enterprises, spotty IPv6 support may not matter — at least not today. But there's a strong and growing probability that IPv6 will matter before switches end their depreciation cycles in three to five years. Even for enterprises with no IPv6 now in place, it's still very much worth considering.

Any switch configured in Layer 2 mode can forward IPv6 packets because it doesn't know or care about Layer 3 headers. When configured as Layer 3 forwarding mode, all switches tested except HP's ProCurve move IPv6 packets between subnets (at least in the version we tested; HP says current 13.x releases do support IPv6 but we didn't verify that).

That's not the end of the IPv6 story, though. It's important to distinguish between forwarding (moving packets between subnets using direct or static routes) and routing (running a routing protocol to learn dynamically where to send packets). The D-Link, Extreme and HP switches do not support the major enterprise IPv6 routing protocols, Open Shortest Path First Version 3 and RIP next generation. And, as noted, there are major differences in switch management methods over IPv6.

As for multicast over IPv6, the Dell and HP switches don't support either version of multicast listener discovery, IPv6's functional equivalent of Internet group management protocol in IPv4. D-Link's DGS-3650 supports Multicast Listener Discovery Versions 1 and 2.

Power consumption

With large data centers' electric bills topping \$1 million a month, power

consumption is a major concern. Using Fluke clamp meters, we measured each switch's power draw when idle and again when its control and data planes were fully loaded. (See "Tracking power consumption: How low can the switch go?" at www.nwdocfinder.com/4125.)

The results show a roughly threefold difference between the most miserly and power-hungry device, but most switches used similar amounts of power, drawing anywhere between 128 and 154 watts when fully loaded. Alcatel-Lucent's OmniSwitch 6850 wins bragging rights for the most efficient device when idle, using just 79 watts. Extreme's Summit X450 was the most efficient when fully loaded, requiring 128 watts.

Foundry's FastIron X448 was an exception. It uses 255 watts when idle and 316 watts fully loaded, more than double that of other switches. At 1.5 rack units, it's also slightly larger than all other switches, which take up one rack unit apiece.

Wrapping up

There are plenty of differences among switches, especially when it comes to newer features. Just because basic functions long ago entered commodity status doesn't mean the switch wars are settled. As our test results show, new additions such as multicast, 802.1X and security are making access switching interesting all over again.

Newman is president of Network Test, an independent test lab in Westlake Village, Calif. He can be reached at dnewman@networktest.com. Fellow Lab Alliance member Rodney Thayer contributed to this testing.

THANKS

Network World gratefully acknowledges the test equipment vendors that supported this project. Spirent Communications supplied its Spirent TestCenter Gigabit and 10 Gigabit generator/analyzer, and senior software engineer Timmons C. Player updated Spirent ScriptMaster for use in multicast testing. Juniper Networks provided Steel-Belted Radius Enterprise Edition 6.1, an IC 6000 network access server and Odyssey 802.1X client software for our 802.1X NAC tests. Juniper engineers Denzil Wessels and Christian Macdonald provided extensive assistance with test bed setup. Thanks too to Fluke Corp., which provided Fluke 322 and 335 clamp meters for measuring power consumption.



Cisco Systems, Inc.
170 West Tasman Drive,
San Jose, CA 95134-1706
USA