INTEROP LABS

Living Large: A Guide to Wireless LAN Scalability

Wireless LANs are no longer an optional part of the enterprise technology toolkit. As organizations move from small, isolated WLAN trials to enterprise-wide rollouts, scalability becomes the key concern. How many clients will a wireless network support? How fast will clients' traffic move, and with what kinds of response times? How does capacity planning differ between wired and wireless networks?

To help answer scalability questions, this white paper offers an overview of enterprise WLAN scalability issues. For newcomers to radio-frequency (RF) issues, the single most important requirement is understanding the differences between wireless and wired networks. Other concerns include test methodology design; impact on the wired network; and differences between various IEEE 802.11 types.

It's not "just Ethernet"

Wireless LANs differ profoundly from their wired counterparts in ways that have a huge impact on scalability. It isn't just that wireless networks use different physical and link-layer technologies than wired Ethernet LANs. Wireless and wired networks also differ in terms of media access, interference sources, error handling, and power requirements. All these factors affect scalability to some degree.

While virtually all Ethernet networks today use switched full-duplex connections, wireless LANs instead used a shared medium: the air. To deal with contention among clients, the IEEE 802.11 standards define backoff and timing methods that differ from 802.3 Ethernet methods. Further, a single wireless LAN may carry multiple flows operating at different rates, something that isn't possible with wired Ethernet. A lower-speed flow needs access to the medium for a relatively long interval, and this will affect the performance of a faster flow.

Wireless LANs also must contend with interference sources not present in wired networks. Devices based on the 802.11b, .11g, and the forthcoming .11n specs must share the 2.4-GHz frequency range with cordless phones, Bluetooth devices, microwave ovens, and amateur radios. The 5.8-GHz range used by 802.11a is much less prone to interference, but it's still shared by some cordless phones and microwave radio transmitters.

It's a good idea to survey the RF environment and select wireless LAN channels with minimal interference. Note that a single initial site survey is not sufficient for this purpose; interference sources come and go, and the wireless network infrastructure should be able to adapt accordingly.

Modern Ethernet is virtually error-free, whereas wireless LANs run over a highly error-prone, lossy environment. To compensate, wireless LANs add a connection-oriented acknowledgement scheme for every frame. Wireless LAN access points (APs) also broadcast "beacon" frames 10 times per second, adding overhead. Network designers must take this overhead into account in determining capacity limits for users and application bandwidth.

Power consumption may be a nonissue on wired networks, but it matters for both clients and network infrastructure in wireless LANs. Mobile clients such as laptop PCs and phones often use power-saving mechanisms. These mechanisms extend battery life, but they also affect application performance (think of delay-sensitive VoIP traffic) and the speed at which clients can roam between APs.

On the infrastructure side, many Ethernet switches support the 802.3af power over Ethernet (PoE) standard and supply current to APs, eliminating the need for dedicated power outlets. However, some low-end and midrange PoE switches supply only enough power to support half their ports, limiting the number of APs a switch will support. Network managers can quickly determine power budgets from vendors' data sheets; just divide the switch's overall wattage by the port count. Assume a need for 15 watts per port, as specified by 802.3af.

Test, test, test

While it's common practice to run any new technology through trials before implementing it in production, a few special considerations apply when it comes to wireless LAN scalability. These include test bed design, AP count, client association capacity and rate, and the test metrics in use.

Scalability and performance testing must meet three requirements: It must be repeatable, stressful, and meaningful. Of these, the "meaningful" requirement is the hardest to define, since each enterprise has a unique mix of applications. However, it's the most important requirement as well. A test methodology should attempt to model actual user counts and application traffic patterns to the closest degree possible.

Given the spectrum contention issues inherent in wireless LAN testing, control over the RF environment is a must for ensuring repeatable results. One option is to conduct tests at a remote location that's relatively "clean" in RF terms, or to run tests inside a specially designed RF chamber called a Faraday cage.

A more practical option – especially when testing multiple APs that are likely to contend with one another – is to use purpose-built test equipment that models client behavior, and connects to each AP using special shielded cables. While this might seem unrealistic at first glance, this approach ensures that the tester can introduce RF impairments and offer test traffic in a controlled, repeatable way.

Another consideration is the number of clients expected to use each AP. InteropLabs tests have shown that 10 clients associated with a single AP will experience very different forwarding rates, packet loss, and latency, than will a single client. Further, some APs and wireless LAN controllers only support a limited number of client associations per AP. Concurrent client capacity is one among several key metrics that any scalability test plan should include.

Other key metrics include forwarding rate[1], latency, packet loss, rate vs. range, and single- and multi-client roaming times. For networks carrying delay-sensitive VoIP or video traffic, testing QoS mechanisms such as IEEE 802.11e or WMM is also essential.

Walk on the wired side

Comprehensive WLAN scalability testing also should examine the impact on wired networks and devices. Raw bandwidth consumption is one issue; resource usage of wired-side infrastructure such as Radius and DHCP servers is another.

Gigabit Ethernet might seem to offer ample capacity in comparison with an 802.11b/a/g network. That's true if the comparison point is a single AP, but enterprise networks often put traffic from dozens or hundreds of APs onto gigabit Ethernet, potentially leading to congestion. Test results suggest that fewer than 50 fully loaded 802.11g APs can oversubscribe a gigabit Ethernet segment.

Services offered from the wired side also must scale as the wireless population grows. Virtually all wireless clients expect to learn their network configuration from a DHCP server, usually located on a wired network. Many enterprises also use Wi-Fi Protected Access (WPA) and/or 802.1x mechanisms to authenticate users, requiring the use of a Radius server. For DHCP and other services, client connection rate can be an issue. Indeed, in InteropLabs testing engineers reduced client DHCP request rates from 10 per second to 5 per second to avoid overloading a DHCP server.

802.11a: A better choice

For applications that support it, 802.11a is generally a better choice than 802.11b/g in terms of performance and scalability. Data rates and VoIP audio quality tend to be higher with .11a than with .11b/g, all else being equal, due to .11a's use of the less crowded 5-GHz frequency band and different timing mechanisms. And 802.11a may even be preferable to forthcoming 802.11n devices in some situations, since it doesn't have to share spectrum with older, slower devices.

There are some tradeoffs in going with 802.11a. Higher-frequency signals are more easily absorbed by walls and other objects, reducing 802.11a's overall range and leading to a requirement for more APs to cover a given area. Also, VoIP handsets supporting 802.11a are practically nonexistent.

Still, InteropLabs testing suggests that 802.11a technology can deliver higher VoIP audio quality than 802.11g. One caveat is to ensure that all stations operate at the same rate. Testing showed rate mismatches caused quality scores to fall by 10 percent or more.

Putting it all together

As noted, wireless and wired LANs scale differently because the underlying technology is different. With an understanding of these differences and meaningful testing, network managers can ensure that wireless LANs will be just as robust and scalable as any wired segment in the enterprise, now matter how large.

[1] "Throughput" may be a commonly used metric, but it is not meaningful for wireless LAN testing. As defined in RFC 1242, throughput is a zero-loss rate; in contrast, wireless LANs are inherently lossy. Forwarding rate, defined in RFC 2285, is a more appropriate metric for assessing WLAN performance. The latency measurements defined in RFC 1242 are equally valid on wireless and wired networks.