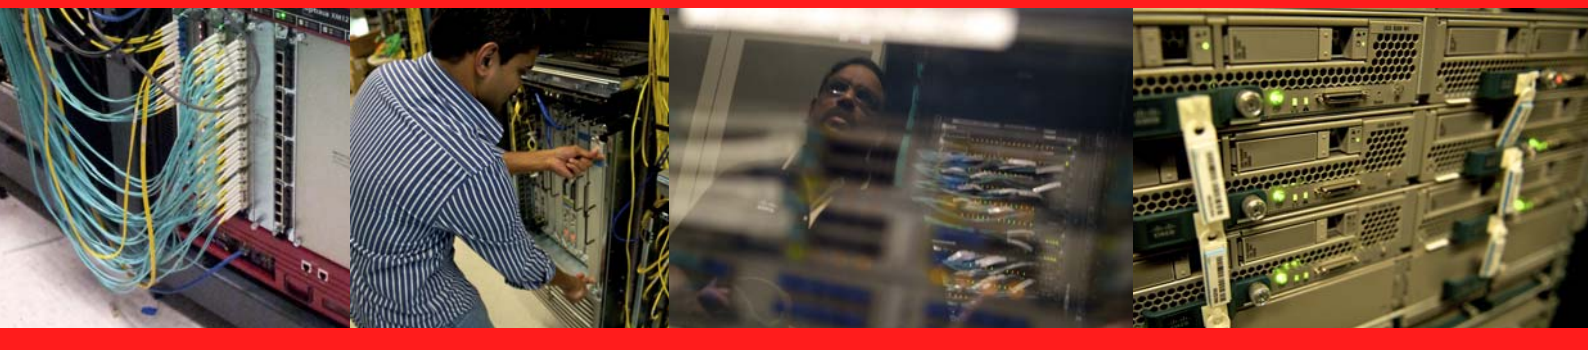

EANTC's Independent Test of Cisco's CloudVerse Architecture

Part 2: Cloud Intelligent Networks

Photos taken by Kim Ringeisen (Flying Monkey Photography)

Introduction from Light Reading

There's nothing new about IPv6. Network equipment providers and service providers have been preparing to migrate from IPv4 for years. But now that IPv4 address depletion is a reality, enterprises and service providers have to address how they'll migrate themselves and what they can do for their customers. (See LR Live: IPv6 Transition Decisions Loom).

That migration will include supporting both IPv4 and IPv6 traffic simultaneously and, in some cases, using carrier-grade network address translation (NAT) as a migration strategy. As more enterprises move to cloud services, they need to be confident that the services providers they are trusting with their most valuable business have a real plan for their networks.

We've covered several aspects of the move to IPv6 and the concerns involved over the past few months:

IPv6 Security: 5 Things You Need to Know

- The Ugly Side of IPv6: Carrier-Grade NAT
- LR Live: The Real Rationale for IPv6
- Cisco to Service Providers: Get Moving on IPv6
- Is IPv6 Finally on the Verge?
- IPv6 Global Event Gets Real
- Transitioning to IPv6

Cisco, like other vendors, is looking for validation that it has several of the pieces needed by service providers to proclaim confidently that they have a plan in place for the future of their IP infrastructure. Here are the sections for this part of the Cisco Cloud Mega-test, where the emphasis is squarely on showing operators how the network core and data center could scale while moving to IPv6:

- Stateful NAT64 Performance
- IPv6 Rapid Deployment (RD) Performance
- IPv6 Dual Stack Performance
- Network Positioning System
- DHCPv6 in the Cloud
- Conclusion: Cloud Intelligent Networks

Table of Contents

About EANTC

EANTC is an independent test lab founded in 1991 and based in Berlin, Germany, conducting vendor-neutral proof of concept and acceptance tests for service providers, governments and large enterprises. EANTC has been testing data center solutions since the early 2000s for both online publications and interoperability and service providers.

EANTC's role in this program was to define the test topics in detail, communicate with Cisco, coordinate with the test equipment vendor (Ixia), and conduct the tests at the vendors' locations. EANTC engineers then extensively documented the results. Cisco submitted their products to a rigorous test in a controlled environment contractually defined. For this independent test, EANTC exclusively reported to Light Reading. Cisco test did not review the individual reports before their release. Cisco had a right to veto publication of the test results as a whole, but not to veto individual test cases.

— Carsten Rossenhövel is Managing Director of the European Advanced Networking Test Center AG (EANTC) , an independent test lab in Berlin. EANTC offers vendor-neutral network test services for manufacturers, service providers, governments and large enterprises. Carsten heads EANTC's manufacturer testing and certification group and interoperability test events. He has over 20 years of experience in data networks and testing.

Jonathan Morin, EANTC, managed the project, worked with vendors and co-authored the article.

Stateful NAT64 Performance

EXECUTIVE SUMMARY: Cisco's CRS-1 loaded with four CGSE cards successfully translated IPv6 traffic to IPv4 at 4 million translations per second. The same system scaled up to 78.4Gbit/s at a total of 67,107,840 translations with almost no loss.

While the industry embraces IPv6 now more than ever, it also recognizes that IPv4 services are not going away soon. The Internet is an obvious example where IPv4 addresses are going to be used for years to come. Cloud applications will use those addresses as well.

While data centers will have different IP migration strategies, they will likely look to serve both IPv4- and IPv6-based customers. Long-term strategies will include native IPv6 throughout the data center, but in the short term a complete IPv6 strategy might not be practical.

For this reason service providers and cloud operators are likely to find themselves needing to deploy Network Address Translation (NAT) from IPv6 users to IPv4 services (NAT64). Let's say an enterprise is building a brand-new large-scale office and wants to use unique IP addressing. The carrier could provide this adventurous customer with IPv6 addresses to use for internal hosts and servers. In order to communicate with the Internet, which at this point is still IPv4 heavy, the carrier could install a

NAT64 device somewhere between the customer and their services to translate the IPv6 addressing to IPv4 before sending the datagrams to the Internet. Another example is the rollout of mobile services en masse using IPv6, to customers who still plan to access IPv4 services, including cloud services.

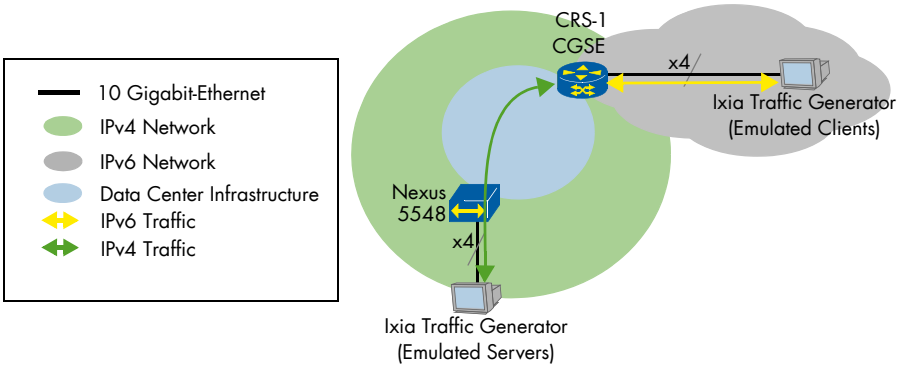

Figure 1: NAT64 Traffic Flow

Cisco claimed to be ready for these scenarios – delivering IPv4 services to IPv6 customers – at scale. Since we have already reported results on Cisco's stateless NAT 64 capabilities we wanted to use this opportunity to verify Cisco's stateful NAT64 performance claims – that by placing four Carrier-Grade Services Engine (CSGE) modules into a single CRS-1, we could scale up to 60 million NAT64 translations, at 4 million translations per second, all while transmitting up to 80Gbit/s of data.

Would any carrier need this performance? Probably not anytime soon, but we have learned that those who purchase large-scale core routers want to know that they can use their significant financial investment for a while.

Given the scale, we looked to verify each metric separately. Even with this divide-and-conquer approach, NAT can become complex to test. Cisco explained, and showed, that when their NAT64 implementation chooses an IPv4 address to map to an incoming IPv6 request, it is done at random. Now imagine manually configuring the tester for 60 million mappings, when all 60 million incoming requests are given random IPv4 addresses – clearly this was not the way to go.

One alternative that we considered was to use stateful traffic using Ixia's IxLoad application, but emulating up to 60 million sessions would have required a significant amount of very high-performance test equipment – again, not really a workable option. The solution we used involved Ixia's IxNetwork generating stateless traffic, with the appropriate TCP fields set to emulate a stateful session (TCP SYN/TCP ACKs). Since Cisco's implementation randomly assigned TCP port numbers and IPv4 addresses to incoming IPv6 requests, we schemed to simply exhaust the entire pool of resources on the CRS-1. This way we were able to predict which addresses and ports would be used – it would be all of them. If your head is spinning, we hope the following diagram will help.

To summarize, we sent client traffic from 1,024 IPv6 addresses – each of whom opened 65,535 TCP sessions. In fact, this brought us to a total of 67,107,840 translations on the CRS-1. We sent traffic in return toward all 960 IPv4 addresses, each with all 65,535 TCP port numbers, as was configured in the CRS-1 pool. All traffic used IMIX frame sizes – 122:7, 512:4, 1500:1 (106 in place of 122 on the IPv4 side) at a rate of 38.4Gbit/s toward the clients and 40Gbit/s toward the servers, all across four 10-Gigabit Ethernet links. Once the configuration was pre-staged and verified to be working, we could breathe a sigh of relief.

As we started the official test run we recorded only a small amount of loss – 0.002 percent on eight of the 16 flows configured from the IPv6 emulated clients toward the IPv4 emulated servers. The other four of such flows ran with no loss, and no flows in the return direction observed any loss either. Considering that we had planned to only test 60 million translations rather than 67,107,840, the loss was considered very minimal. We also verified, using the CRS-1 Command Line Interface (CLI) that all expected translations appeared in the enormous translation table. We also measured latency. The maximum latency values were not very surprising given the translation work to be done by the CRS-1, but in general, given that the latency also included the seven other devices in the test bed, the average latency was quite low.

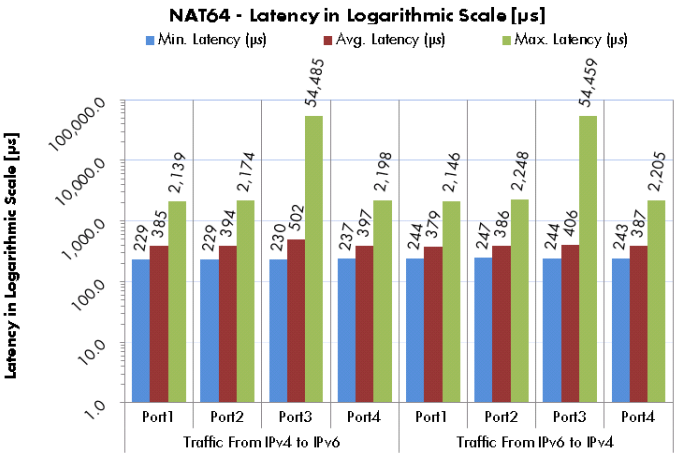

Figure 2: NAT64 Latency at 67,107,840 Translations

Next was performance. How quickly could these translations be built in hardware? Now that our test methodology was proven, we felt safe clearing the NAT table on the CRS-1. After doing so, we lowered all frame sizes to 150 bytes so we could increase the frame rate to 4 million frames per second – 1 million frames per second on each of the four 10-Gigabit Ethernet ports. In order to add realism to the test we configured IxNetwork to randomly assign TCP ports to the IPv6 flows, so that they were not sequential. This however required that we also lower the total number of ports to 13,824, bringing the number of translations to 56,622,848 in total. We ran the test for two minutes without loss.

After some pretty long nights of some complex configuration, we had finally established a test that was able to verify the rate, translation capacity, and throughput of Cisco's NAT64 solution. Impressive.

IPv6 Rapid Deployment Performance

EXECUTIVE SUMMARY: Cisco's CRS-1 forwarded 79.6Gbit/s across 1 million IPv6 RD tunnels to aid in quickly migrating customer networks.

There are several technologies designed to to ease the migration from IPv4 to IPv6. The "right" one depends on the use case, but it often boils down to a simpler question: Which part of the network will be based on IPv4 and which part will be based on IPv6? We've reported Cisco's Stateful NAT64 Performance to document its use of the technology required when there are IPv4 endpoints communicating with IPv6 endpoints.

What about when a new customer plans to access IPv6 based services, with an IPv6 address, but is connected to an IPv4 access network? Expecting these scenarios to occur frequently, the IETF defined technology for IPv6 Rapid Deployment (IPv6 RD). The technology is pretty straightforward – encapsulate IPv6 packets with IPv4 headers, no control plane is required. The idea is for residential gateway routers to implement IPv6 RD on the customer side, and for these IPv6 RD tunnels to converge on a more powerful platform in the provider's network, which can then de-capsulate the packets from IPv4 and route them based on their IPv6 destinations.

The implications are that there would be additional overhead in the IPv4 network, and the gateways used would still have IPv4 addresses, but this allows operators to ramp up IPv6 deployments even where the access or aggregation network is IPv4 based – hence "rapid deployment."

Cisco says it would like to be ready to support its customers in any migration scenario, therefore IPv6 RD had to be included. We wanted to use the same setup we had already created using the CRS-1, loaded with four CGSE modules, and connected with 4 x 10 Gigabit Ethernet interfaces to again reach line rate, but the more interesting question was how many residential gateways we could scale to. Cisco claimed that 1 million tunnels would not be a problem. In order to test such a high number of residential users, we emulated them using our Ixia equipment running IxNetwork. Behind the 1 million residential gateways, evenly spread across the four physical ports, we emulated 10 million users – 10 users behind each gateway.

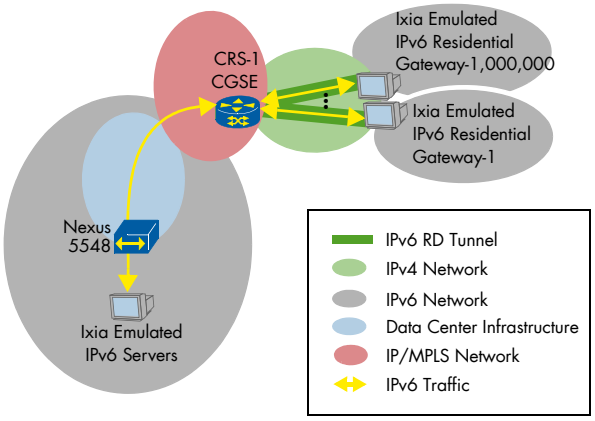

Figure 3: IPv6 RD Test Setup

We sent bidirectional traffic in pairs between these 10 million users and 20,000 emulated IPv6-based servers, using IMIX frame sizes: 122:7, 512:4, 1500:1. We had to account for the overhead incurred by the IPv6 RD encapsulation, so we were focusing on the data rates on the IPv6 interfaces. These native IPv6 interfaces connected via Cisco's Nexus 5548 transmitted a full 10Gbit/s, and received 9.9Gbit/s each. In total we generated 79.6Gbit/s of traffic. The good news for Cisco was that no frames were lost. Given the encapsulation/decapsulation, we were also interested in recording latency, expecting it to be lower than NAT64 latency since it was not stateful. The results are shown below.

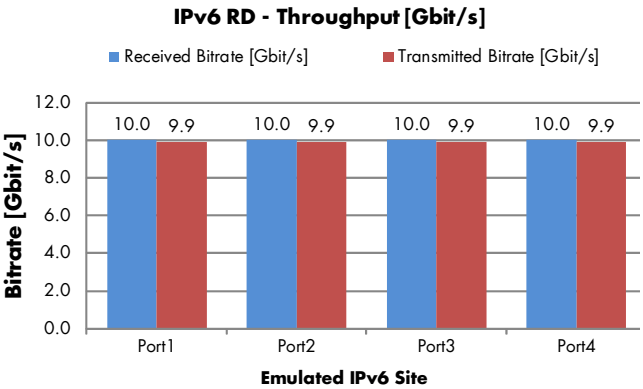

Figure 4: Throughput Results

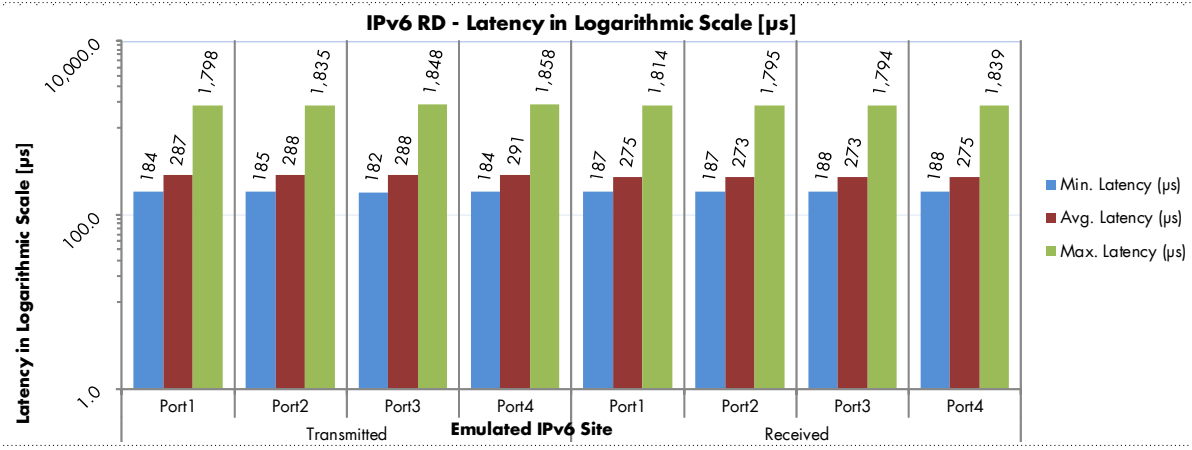

Figure 5: Latency Results

Finding no issues, Cisco proved that they are ready to help operators deploy IPv6 rapidly, scaling to a million customers, each with the potential of having multiple users, all communicating simultaneously.

IPv6 Dual Stack Performance

EXECUTIVE SUMMARY: Cisco's core and data center infrastructure successfully forwarded 96Gbit/s of bidirectional dual stack traffic of 50 percent IPv4 and 50 percent IPv6, without losing a single packet.

Operators can obviously not flip on IPv6 throughout an end-to-end infrastructure in a day, but it is safe to say that all major carriers have the end goal of someday enabling IPv6 on their entire infrastructure. Technologies such as NAT64 and IPv6 RD ease the migration process by stitching and tunneling IPv4 and IPv6 networks together, but many operators are planning not only to support IPv6 everywhere, but to support IPv4 everywhere.

IPv4 addresses will not be available to all uses in all cases, but the goal of Dual Stack networks is to be ready for it when it is needed, all while supporting IPv6 end-to-end as well. Technically, Dual Stack is simply the concept of supporting both IP versions, and having an IP address assigned for each, on the same physical interface. If you would ask a major core network operator, "Dual Stack everywhere?" you would likely hear responses ranging from "Eventually ..." to "Next year" and everything in between.

This test was the most straightforward of our three IPv6 migration tests. We simply had to send both IPv4 and IPv6 traffic, bidirectionally, end to end. From Cisco's perspective, it was also straightforward, but they still had to configure both IP versions on the CRS-1, CRS-3, ASR 9010 and Nexus 7010 at a minimum, to match our previous IPv6 test setups. Cisco was keen to include more functions in the test.

Traffic labeled as Gold (by VLAN) was routed by the Nexus 7010 to the Catalyst 6500 and its ACE30 module and the ASA5585-X60 firewall appliance according to Cisco. The goal was to transmit 8Gbit/s of Gold traffic and 40Gbit/s of Bronze traffic, per direction, for a total of 96Gbit/s across six 10-Gigabit Ethernet interfaces on each side of the setup. The traffic was sent between 1 million emulated clients and 50,000 emulated servers – both split down the middle, half IPv4, half IPv6. The traffic rate was split in half for each IP version as well. Due to the ASA firewall, all Gold-labeled traffic (16Gbit/s) included UDP headers, else the ASA firewall would have correctly considered the traffic invalid. In fact, the ASA firewall was also the cause of some loss in our first test runs. After making extra

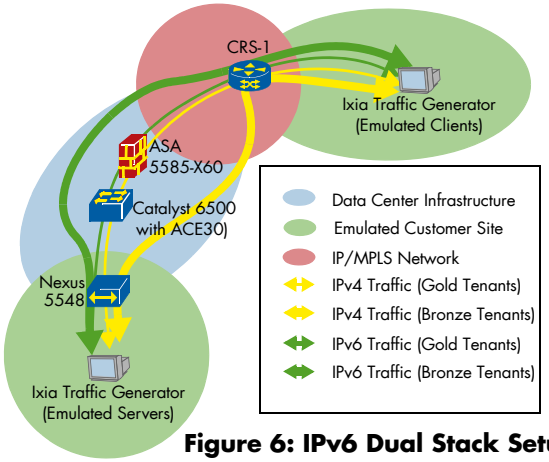

Figure 6: IPv6 Dual Stack Setup

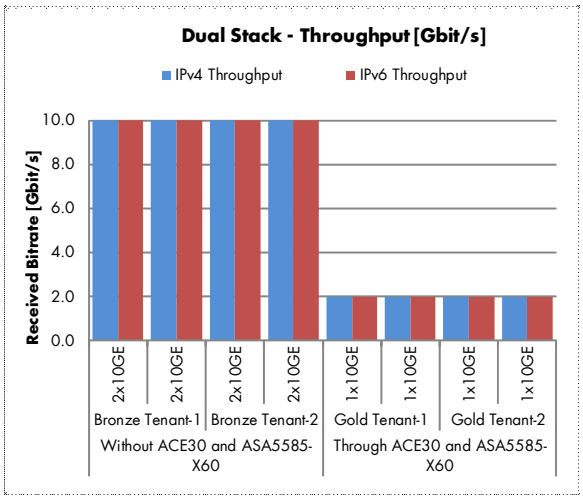

Figure 7: IPv6 Dual Stack Throughput

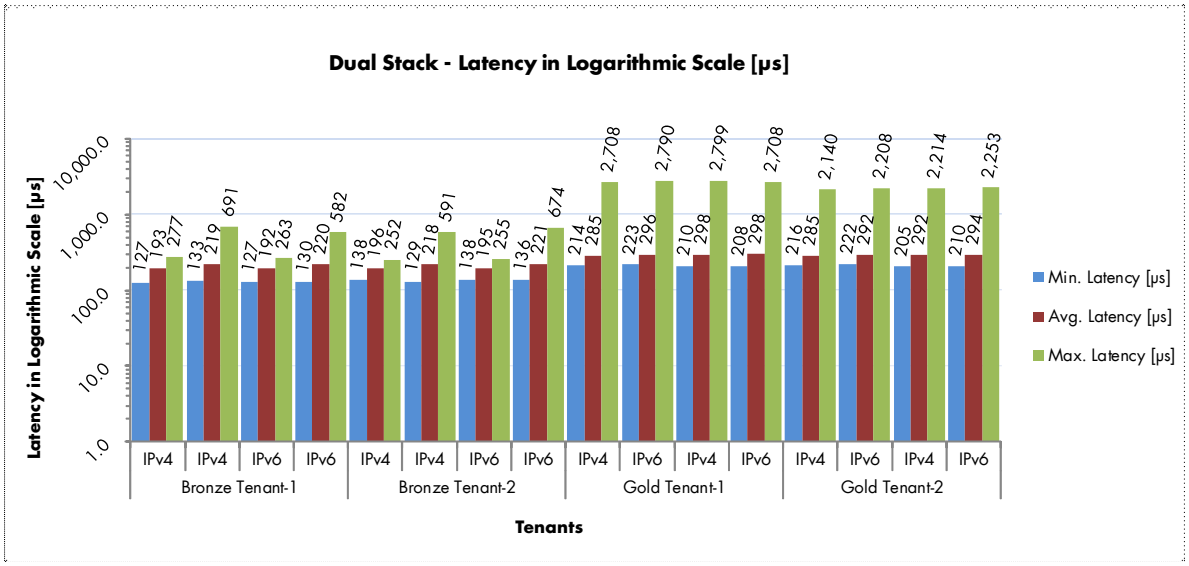

Figure 8: IPv6 Dual Stack Latency

sure that nothing else was going on in the network (at least not that touched the ASA firewall), we were able to get a clean run for five minutes with no loss at line rate. Effectively we reached 96Gbit/s – 48Gbit/s of IPv4 and 48Gbit/s of IPv6, in parallel – 16Gbit/s going through the ASA 5580.

Cisco, as one of the champions of IPv6, has shown that the three major migration scenarios are supported, at scale, in various pieces of their hardware: from the core backbone combination of CRS-1 with CGSE to the full palate routers with dual IPv4/IPv6 stack. For service providers looking at any of these migration scenarios we could attest that in our initial review the pieces are there – now it is time to verify their functionality and performance in context of your network.

Network Positioning System (NPS)

EXECUTIVE SUMMARY: When provided access to multiple data centers with the identical service, Cisco's NPS correctly chose the best performing option for user traffic.

If cloud services are so important, than so too must be the availability of cloud services, which will require cloud providers to use multiple data centers to tackle issues involving scale, reduction in application latency based on geographical proximity, and resource distribution.

If the existence of multiple data centers is a given, resource distribution and customer experience optimization becomes a critical business concern: Will the data center operators distribute load across the data centers? Will they provide customers with the data center that will give them the best experience based on network proximity or performance?

Cisco says it can arm the network with the intelligence to make these decisions. This is the idea behind Cisco's Network Positioning System (NPS).

To see NPS in action we needed two data centers and, as luck would have it, our test setup came equipped with two data centers.

Our intention was to verify that when the same customer requests a service, NPS makes the decision as to where that request would go – Data Center 1, or Data Center 2. When we discussed this idea with Cisco, we prepared to have NPS work based on proximity, but they also explained that NPS was built as a customizable tool. We felt it would be more relevant to see NPS choose data centers based on performance – delay, for example. Cisco agreed with us and got to work.

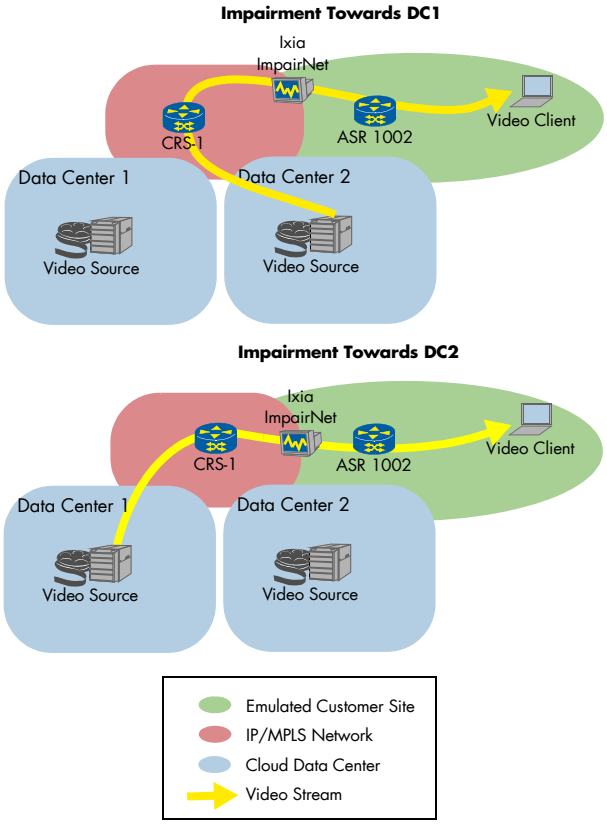

Figure 9: Impairments Illustrated

The NPS system database was incorporated into our customer-facing CRS-1. Cisco then set up some special NPS client software on our client laptop, and configured NPS on the ASR 1002 - the router used as a Customer Edge (CE). The CRS-1's central purpose in the NPS setup was always know which data center is the optimal match for the defined metrics. The client laptop and ASR 1002 polls the CRS-1 with the preferred metrics and uses the CRS-1's response for the customer traffic. In our test case, Cisco set up their IP-SLA measurement probes between an ASR 9010 in each data center, and the CE (ASR 1002) to constantly measure and report the delay to the CRS-1.

Cisco installed two simple video servers in each data center. We connected a laptop client to our ASR 1002, and began requesting video through a Web portal Cisco had setup. In the beginning, the Web portal would almost randomly choose different data centers each time we refreshed. We found this was because the latency measurements were extremely close, and mildly fluctuating. No problem, this meant both video servers were working. Also, we came prepared. We connected Ixia's new shiny ImpairNet impairment generator between the customer edge ASR 1002 and its upstream CRS-1. This was the customer's link to both data centers, but, by using a filter on the impairment generator, we could add delay to all packets for a given destination. We toggled back and forth between adding 50 milliseconds on all the IP-SLA measurement packets to Data Center 1, and then disabling it and adding the delay to Data Center 2. Each time, we observed that when the video client was refreshed it was showing video from a different data center according to the path with the lowest latency.

In addition we verified that NPS would not include data center options that didn't run a service all together. We used "CPU hog" on Data Center 1 to disrupt the video server. The NPS system detected the failure of this virtual server to respond, and the CRS-1 updated its NPS database to no longer include DC-1 as a viable option for this video service. We refreshed our browser and were consistently directed to Data Center 2.

For service providers offering cloud services the ability to optimize the customer experience when accessing geo-redundant or distributed data centers could well be a competitive edge, especially when the cloud services begin to be commoditized. It is impressive to see that functions that required complicated traffic engineering knowledge in the past have been simplified and repackaged for general consumption.

DHCPv6 in the Cloud

EXECUTIVE SUMMARY: Cisco Network Registrar successfully provided IPv6 addresses to up to 18,036 users per second – all from the cloud.

Despite IPv6's initial promise of removing the need for DHCP, the need for central management made DHCPv6 a necessity. We decided to test Cisco's DHCPv6 performance for two reasons. First, Cisco Network Registrar (CNR) has been ported to run in the cloud – that is, the DHCPv6 server runs on virtual servers within Cisco's Unified Computing System (UCS). Second, Cisco had some pretty high-performance claims.

This triggered us to ask why performance numbers for a protocol like DHCP are even important. DHCP address requests are typically quite slow and scattered, and even in a failure event, when users come back online and the servers cannot service all requests, the users will simply retry immediately and will typically get it. Cisco explained that residential service providers have expressed concern to them that DHCP would be the weak link when thousands of customers come back online after a failure event – a failure event that they perhaps don't want customers to know existed.

Since Cisco's CNR was running not on its own appliance but rather on the UCS, the performance can vary depending on what CPUs and memory are on each UCS blade. We used four identical UCS B200-M1 blades for the test. Seven VMs were installed on each of the three UCS blades to run a script emulating users, and a fourth blade had a single VM equipped with Cisco CNR.

At the time of the test, testing DHCPv6 scale was in the Ixia IxLoad roadmap, so we used scripts running on these 21 VMs to emulate the user exchange. Each blade had two 2.93GHz quad-core processors and 98GB of memory. Once Cisco configured 20/64 IPv6 subnets, we could start the client requests.

We did two types of test runs – some where the DHCP server still had the user (MAC address) to IP address mapping cached, and some tests where we cleared the cache to emulate a situation that users were being added for the first time, or there was an outage that was long enough for the server to timeout its entries. When testing for renewals, the system was pre-cached with 400,000 leases and the emulated clients would send a total of 60,000 requests. When testing for new lease request rates, the cache was cleared, and the clients would send

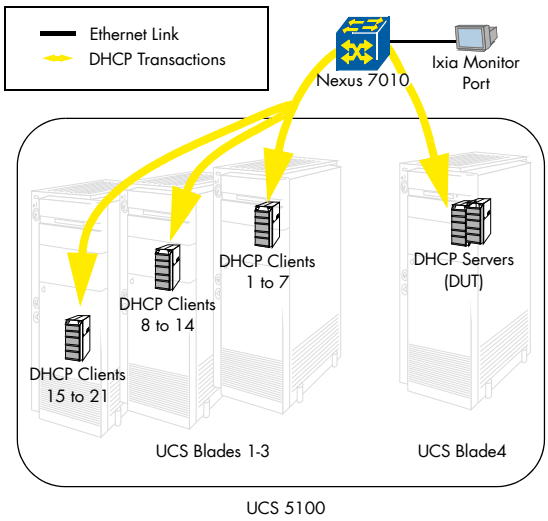

Figure 10: DHCPv6 Test Setup

the same 60,000 requests.

To verify that the scripts were working as expected we captured traffic through a SPAN port on the Nexus 7010 as shown in our setup diagram. We compared what we saw from the capture with what the server was reporting. The graphs show a varying rate of requests and offers, and when it is compared to the rate reported by the DHCP server, it is about the average value from the graph. This is what we were expecting to see.

The capture consistently left out exactly 25 percent of the offers. Cisco explained that these were being hashed onto different links – some internal, to the UCS across the Nexus 1000v. We proved it by looking at the port counters on the Nexus 1000v, seeing exactly 120,000 packets in and 120,000 packets out – one solicit and request coming in for each of the 60,000 transactions, and one advertise and reply going out for each of the transactions. Proving the servers stats were accurate, we repeated each test type three more times for consistency, and looked at the performance as reported by the server.

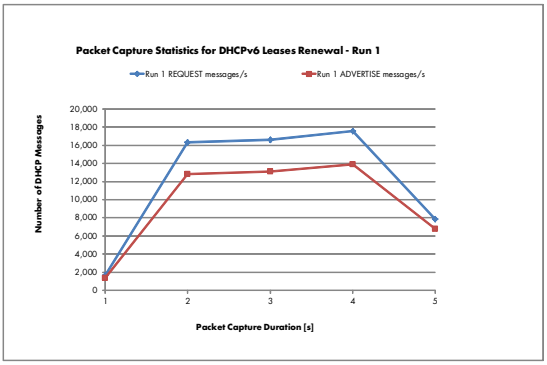

Figure 11: DHCP Messages per Second for Run 1 (Lease Renewals)

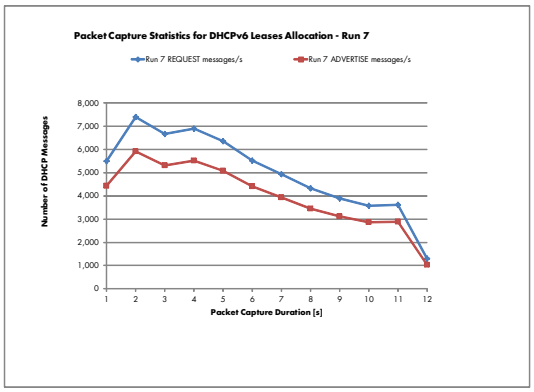

Figure 12: DHCP Messages per Second for Run 7 (New Lease Allocations)-

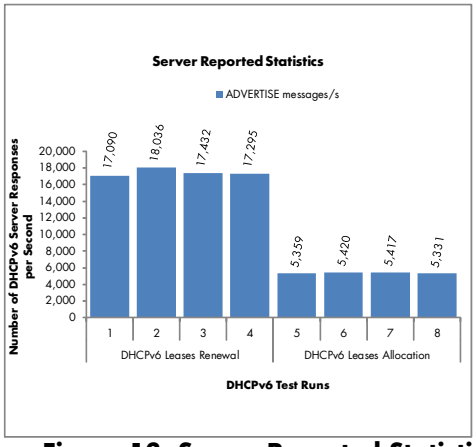

Figure 13: Server Reported Statistics

After our deep dive inspecting the results, and observing the consistency, we felt confident that operators looking to produce these results could surely do so with this setup. For those looking for this kind of performance, we have listed the hardware specifications to get there, but such performance from a DHCP server is not always required in all environments. Finally, the cloud. The setup we used demonstrated that even such basic IT requirements as DHCP can be outsourced to the cloud. While it may be hard to imagine now, even these base services will need to run in the cloud in order for enterprises to reduce IT investments.

PCRF IN THE CLOUD

EXECUTIVE SUMMARY: Cisco's ASR 5000 successfully retrieved policies from the PCRF in the cloud, and throttled customer traffic accordingly.

In 2010 we conducted a comprehensive test of Cisco's mobile solution including mobile core and mobile backhaul (See Testing Cisco's Mobile Core, Data Center & Business Services and Testing Cisco's Next-Gen Mobile Network). At the time, we used a third-party Policy and Charging Rules Function (PCRF) as Cisco had not implemented its own. Now, not only did Cisco have an early version of their PCRF for us to test, but it came with a very timely message – it was ready to be run in the cloud.

Mobile carriers need PCRF to dictate the rules their subscribers must follow when using the network. These rules could include data allowance, mobility and roaming to name but a few. The Policing and Charging function has been defined both for 3G and Long Term Evolution (LTE) scenarios by the 3rd Generation Partnership Project (3GPP) , and is typically done by a dedicated system with access to subscriber information databases, charging systems and mobile gateways.

In this sense, we think Cisco was wise to port its PCRF to its Unified Computing System, either to be run locally or in a cloud. By doing this, mobile operators could benefit from the flexibility and agility of the cloud, and Cisco has a new use for the UCS systems. With the flexibility of running the PCRF in the cloud comes questions of how many subscribers

can it support and what kind of new mobile core topologies could be erected using this idea. Since the system was brand new and since such scalability tests are extremely time consuming, we focused on initial functionality proof points.

Cisco claimed that its ASR 5000-based mobile core could use the PCRF to implement throttling for different customer tiers, so we set out to test just that. Cisco's mobile setup was in a different lab than our cloud test bed and we decided not to move it. If the PCRF can run in the cloud, then it should certainly be able to run in our cloud test bed to a remote mobile core and test setup. Ixia helped us to bring an extra XM2 tester over to the building where Cisco's ASR 5000 was and we set up the test. Cisco configured a single ASR 5000 to do the work of the Packet Gateway (PGW) and Serving Gateway (SGW) in a Long Term Evolution (LTE) scenario.

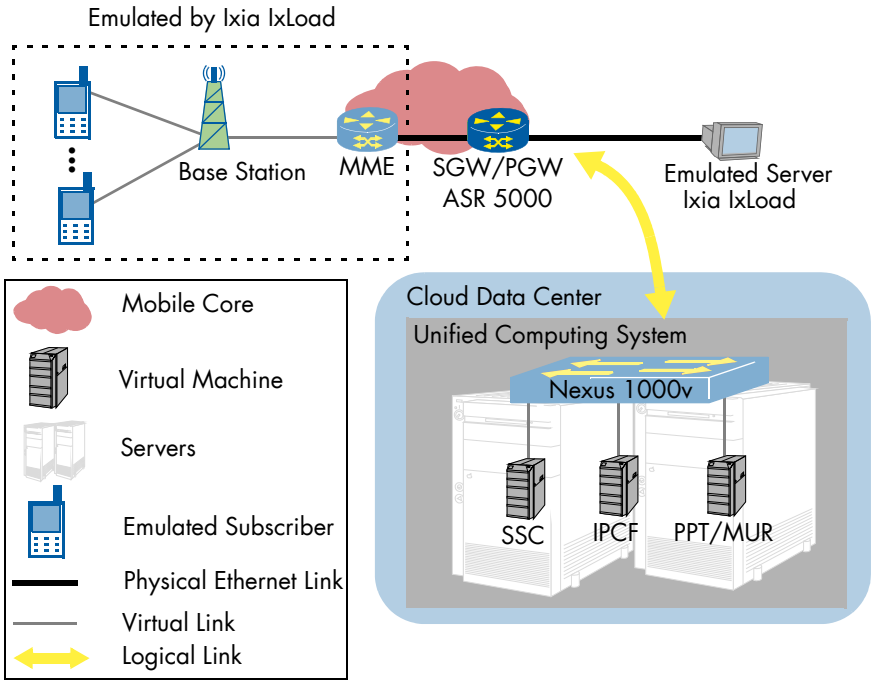

Figure 14: PCRF Test Setup

Ixia's IxLoad was used to emulate the Mobility Management Entity (MME) on one port connected to the ASR 5000, with the base station and clients behind it, and the emulated Web servers with content on a second port, also connected to the ASR 5000. In the cloud, the PCRF was set up with three virtual machines. One had Cisco's Inteligent Policy and Control Function (IPCF – Cisco's implementation of the 3GPP-defined PCRF) installed; the second ran Cisco's Subscriber Service Controller (SSC), which held the database of subscriber data, and the third virtual machine ran Cisco's Policy Provisioning Tool (PPT) and the Mobility Unified Reporting (MUR) tool.

Before we looked at throttling, as a sanity check, we ensured that we could establish both default and dedicated bearers to up to 50 subscribers. Since only data traffic was going to be used in this test we only configured default bearer per subscriber.

To test the throttling feature we configured three subscribers – one bronze, one silver, one gold. Cisco's ASR 5000 and PCRF categorized them based on IMSI ranges. Each subscriber was configured to create an HTTP session with the emulated server, attempting to reach as high a data rate as possible. Each subscriber type had a different bandwidth policy assigned: Gold subscribers received 4Mbit/s per bearer, Silver subscribers received 3Mbit/s and Bronze subscribers received 2Mbit/s. Each subscriber had two additional rules assigned. The first rule was a traffic volume limit of 50MB. Once this limit was reached, each subscriber bandwidth was throttled some more: Gold subscribers were throttled back to to 2Mbit/s, Silver to 1.5Mbit/s, and Bronze to 1Mbit/s. We cleared the volume usage on each subscriber and tested each one at a time. The graph below shows that each subscriber was throttled approximately as expected. The behavior of each line shows that the ASR 5000 would allow a burst before dropping, and Ixia's TCP sessions slowly learning to home in on the rate it could consistently get.

Once the test was complete, Cisco mentioned it is also working on enabling dynamic policies – the reconfiguration of how the ASR 5000 throttles traffic based on some condition. One of such conditions was when a specific Access Point Name (APN) crosses a bandwidth threshold as a percentage of how much bandwidth the ASR 5000 was seeing in total. Another dynamic policy was to limit specific protocol if traffic from this protocol exceeds a given percentage amongst the total traffic, which could be used to throttle P2P and YouTube traffic, for example.

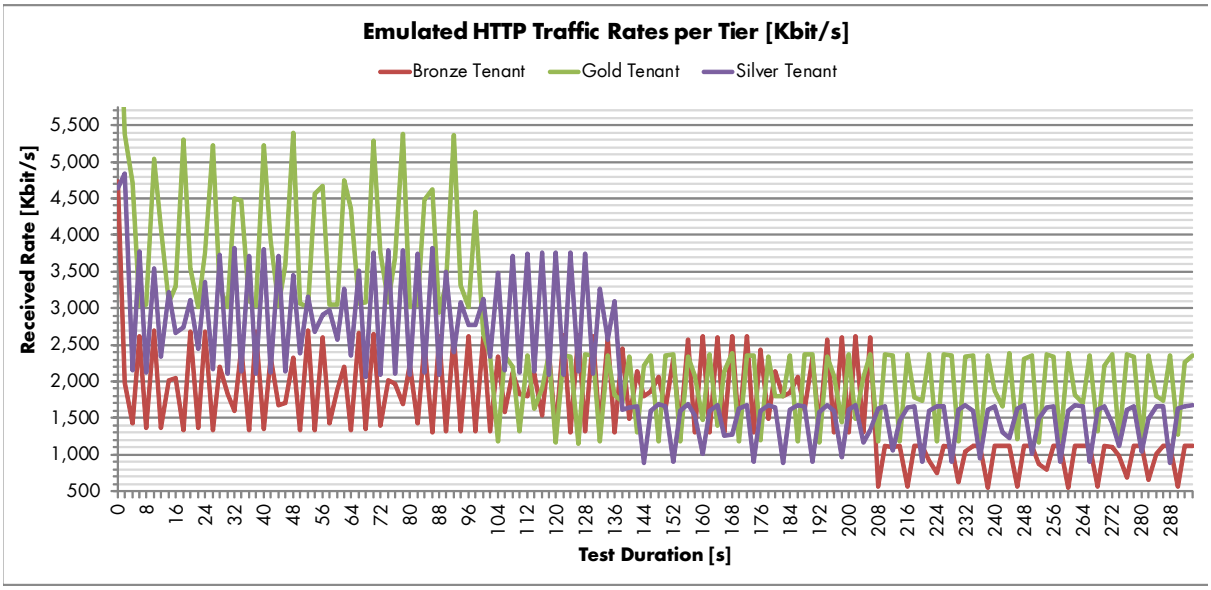

Figure 15: Throttling Based on Service Tier

Cisco explained that operators have asked for such features. One example we heard from Cisco was that operators would like to be able to limit peer-to-peer traffic, dynamically ensuring that it never reaches a high percentage of the total traffic in the mobile network and that it doesn't reach a high data rate. Such functions could also be performed in the mobile core firewall or DPI devices for example, but putting them directly into the mobile gateway enables mobile operators to register the offender (since the gateway has an IMSI and account association). Interesting, powerful, and potentially a can of worms, depending on how it's used.

These functions are also where the MUR and PPT come into play. Cisco explained that the MUR should normally poll live traffic statistics from the ASR 5000 and the PPT will send the new configuration to the ASR 5000 if they see the conditions met. At the time of the test, the ASR 5000 polling was not yet implemented so Cisco was using some in house scripts for their own testing to manually update the MUR with traffic statistics. In this concept demonstration, we observed that when these scripts were used in accordance with the APNs or protocols we configured with the Ixia equipment, the throttling rates indeed changed.

We validated that the PCRF worked from its installation in the data center. It controlled the mobile gateway located across campus and applied policies to subscribers both statically and dynamically. The question on mobile service provider minds is very often: "Will it scale?" This question is left unanswered at the moment since a scaling test, in the policy and control area, is a completely different beast, one that we did invite Cisco to take on. Meanwhile we also welcome Cisco's ideas for using the PCRF in the cloud – ideas that increase the potential scalability, and optimize both agility and access to the data.

Conclusion: Cloud Intelligent Networks

Under the header of intelligent networks we tested Cisco's solutions for IPv6 migration as well as a method to optimize customer access to services based on automatic parameterization of the network. These tests was important because while IPv6 solutions have been developed and tested since the early- to mid-2000s, two IPv6-related aspects are new. For one, high profile service providers (Phil: Maybe a link to recent Verizon or DT or BT announcements about IPv6 is a good link here) are now paying attention. in the past, operators have often put "IPv6 Migration" on the very end of their TODO list, but these days they can no longer afford such luxury – the very real depletion of IPv4 address space was a wake-up call. This brings us to the second reason as to why this is an important test section – answering the question of how to strategically migrate. We hope that this public report will help operators build a blueprint, answering which technologies will scale, and how the various IPv6 migration pieces could be put together for a comprehensive story.

Does IPv6 network infrastructure relate to cloud services? Absolutely. The point is that these migration questions affect all services, and as various service providers and large IT operators are now considering plans for cloud services, now is also the time to get serious about IPv6. Enterprises ready to move to the cloud are putting their trust in their service provider and the latter, by having a comprehensive strategy for the future of their IP infrastructure, could validate this trust.

So what did those pieces look like and how did they perform?

Cisco's NAT64 implementation was put to quite a bit of scrutiny at a high scale. This was not a simple task, we learned, after several long nights. The four CSGE blades inserted into the CRS-1 were able to translate more than 67 million IPv6 user sessions into IPv4 sessions, while forwarding 78.4 Gbit/s with only some minor traffic loss. When testing the rate at which the CRS-1 could create new translations, we were able to successfully validate that one million new customers could be supported per second.

TABLE 1. Cisco Devices Tested

Cisco Device Tested
Cisco ASA 5585-X60
Cisco ASR 1002
Cisco ASR 9010
Cisco Catalyst 6500 (with ACE30)
Cisco CRS-1
Cisco CRS-3
Cisco Nexus 7010
Cisco Nexus 7018

By encapsulating IPv6 packets with IPv4 headers, the CRS-1 was also able to forward IPv6 traffic through IPv4 networks, reducing the requirement for translation services while taking advantage of IPv6 rapid deployment tunneling. We validated that the solution could support one million such tunnels as well as support a combined traffic rate of 79.6 Gbit/s with no user traffic loss.

And since no IPv6 migration is complete without a strategy for the days in which the network will have to support both IPv4 and IPv6, natively, side by side, we verified that Cisco's network infrastructure was able to service a total of 96 Gbit/s of Dual Stack traffic – IPv4 and IPv6 based user traffic in parallel.

Putting concerns to rest regarding disaster recovery, Cisco showed that their Prime Network Registrar, running on Cisco's Unified Computing System, could quickly respond to up to 60,000 IPv6 address requests.

Finally, we verified that Cisco's Network Positioning System (NPS) correctly directed customer traffic to the data center that would give them the best performance at the time of the request, when measured by Cisco's IP-SLA performance monitoring tool.

These tests focused mostly on Cisco's network core and data center equipment, including the CRS-1, CRS-3, ASR 9010, Nexus 7010, and ASA 5580, where we feel Cisco's results are quite good. Of course, this is not the entire picture - customer premise and residential gateway equipment is just as important to test in terms of interoperability when it comes to IPv6 readiness, but here we wanted to show operators how the network core and data center could scale, and hence focused on exactly those network areas. We look forward to further tests of the much broader topic of IPv6 migration. This concludes our network focused tests of Cisco's CloudVerse architecture, look out next week for a report diving into some more data center applications, particularly Videoscape.

EANTC AG
European Advanced Networking Test Center

Salzufer 14
10587 Berlin, Germany
Tel: +49 30 3180595-0
info@eantc.de
http://www.eantc.com

Light Reading
A Division of United Business Media TechWeb

240 West 35th Street, 8th floor
New York, NY 10001, USA

http://www.lightreading.com