April 2026
Geoff Huston

NZNOG 2026

NZNOG 2026 was held in Christchurch in March 2026. The networking community continues to grow, and the attendee count of 260 was the largest so far. New Zealand is a moderately small island nation in the South Pacific, and it appears to make up for its geographic isolation in its openness for innovation and experimentation. The national community has a long track record of innovation, both in technology and in the underlying investment models for its network infrastructure.

Here's a summary of some of the sessions that I found to be of interest.

Akamai's Approach to Content Distribution

Akamai is a veteran in the Internet content distribution world. Its original model of placing managed content servers in the racks of consumer retail ISPs launched in the late 1990's, and in this way Akamai was one of the earliest Content Distribution Networks (CDNs) in the Internet, taking the earlier opportunistic web caching services and transforming the model by placing the content source at the edge of the network adjacent to the users who consume the content. Akamai's CDN service was subsequently expanded into general cloud computing and security service provision.

Akamai's original placement model inside consumer ISP networks was subsequently augmented by deploying content servers at various exchanges and even some transit networks. Akamai now operate over 4,000 Edge POPs and connect to 1,200 of these access networks. The Edge Pops operate in a cache mode, where requests that cannot be served from the edge tier caches are referred back towards the larger mid-tier servers, who can then pull the content from the origin server it is it not already in their local cache.

Akamai have generally used the Domain Name System to "map" a user to an Akamai server. When any one of the Akami nameservers are queried for an Akamai-managed name, the nameserver attempts to triangulate the assumed location of the recursive DNS resolver making the query to the set of potential Akamai content servers, and it returns the address of what Akamai's resolver calculates as the optimal content server as the result of the DNS query. The DNS response has a low TTL to push new Akamai clients to invoke the same triangulation exercise.

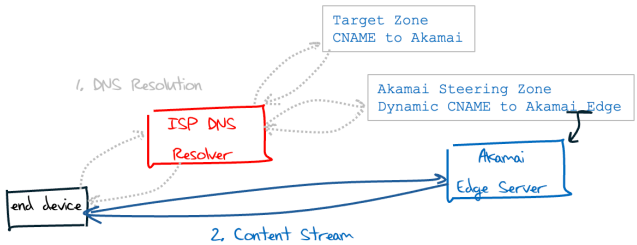

Figure 1 – Akamai's DNS-based Content Steering

The assumption behind this approach is that a user's location and their DNS recursive resolver are relatively close, as the query passed from a user's recursive resolver to the authoritative name server does not contain any details of the end user. In many cases this assumption of the correlation between a user

and their DNS resolver is a workable assumption that provides adequate results. Where it breaks down is where the user's query is directed towards an open DNS resolver, such as Google's 8.8.8.8 or Cloudflare's 1.1.1.1 Open DNS service platforms. There is a DNS-based solution to this. In 2016 the IETF published RFC 7871, "Client Subnet in DNS Queries". This Explicit Client Subnet (ECS) query option is used by the recursive resolver to attach the encompassing subnet of the client's address to the query that the recursive resolver passes to the authoritative server. The authoritative server can use this ECS data in its response to indicate the network scope for which this response is intended. The recursive resolver attaches the subnet to the response and can then use its locally cached copy when there is a match for both the query name and the query source IP subnet in the cache. This is a significant shift away from the inherent privacy properties of the DNS, as passing an indication of the original querier's identity to authoritative servers is not a conventional behaviour in the DNS framework. The fact that this Client Subnet option can be attached to a query made by the recursive resolver without the permission of, or even the knowledge of, the original end user is also a problem.

As RFC 7871 notes: "If we were just beginning to design this mechanism, and not documenting existing protocol, it is unlikely that we would have done things exactly this way. … We recommend that the feature be turned off by default in all nameserver software, and that operators only enable it explicitly in those circumstances where it provides a clear benefit for their clients. We also encourage the deployment of means to allow users to make use of the opt-out provided. Finally, we recommend that others avoid techniques that may introduce additional metadata in future work, as it may damage user trust." It's unclear if anyone has listened to this advice.

There is also an issue with the location granularity of the Akamai response. For networks that operate over a wide geographic area Akamai have a poten tial problem. It's not a case of publishing the location of the client IP addresses, as when Akamai are attempting to triangulate the user location they are doing so by using the IP address of the recursive resolver used to pass the query to Akamai, not the client's IP address.

Akamai do not operate their own backbone network, as the flow from the original content servers to the mid-tier caches to the edge caches occurs over the public Internet. Like the DNS itself, Akamai gathers its performance leverage through the use of content caching close to the edge. This means that is it not reliant on the routing system, as is the case with anycast content distribution networks. Akamai's major issue is that it has limited control over the DNS's retention and potential redistribution of DNS responses. Its use of short TTLs and ECS attempts to limit the reuse of DNS responses and force new content clients to go through the same Akamai DNS-based location triangulation process.

In delivering content, once Akamai has made a steering cache decision for delivery of an item it appears to be a fixed decision. Other CDNs have gone a step further in attempting to optimise the end user experience. For example, the steering approach used by Google's YouTube is a hybrid approach, using both DNS steering and breaking the content into chunks that are variously served from nearby content caches. The service uses DNS steering to route the client to a collection of front-end service units that are located close to the client, and then they periodically "dither" the feed of individual content chunks across the other candidate service units to ensure that best performing service delivery unit is being used for the majority of the content traffic.

This generic approach of DNS steering in the service delivery world has some profound implications for the network, in aspects of architecture, design, robustness and utility. If we no longer rely on carriage services to sustain service quality, then we are no longer reliant on the routing system to produce optimal outcomes in terms of path quality. The true value of adding robust security to the routing environment dissipates in accordance with our reduced dependence on routing as a service platform.

Arista and an alphabet soup of LAN emulators

If there is a single consistent story of the past few decades in networking, it's a story of increasing capability of connected devices and a comparable drop in the demand for complex services from the

network that they connect to. Value, in its various guises, is moving out of the network and into the edge devices and the applications that they host. It just does not make sense these days to outsource to the network the role of setting up and maintain closed communities of connectivity (Virtual Private Networks), when the capabilities of connected devices and applications readily exceed those available to the network provider.

The response from the vendors of network equipment is to add to the alphabet soup of EVPN VXLAN, EVPN-VPWS, EVPN-VPWS-FXC, EVPN MPLS, VXLAN over IPSEC, SR, SR-TE, SRv6, and on and on. It strikes me as an act of desperation on the part of the network equipment vendors to try and persuade network operators to purchase this gear in the forlorn hope that customers will value such complex service offerings, to the extent that they are willing to pay a price premium for the service. When sufficient numbers of customers fail to materialise, they conclude that it was a lack of features and functions that was the cause of the market failure of a service, and they furiously work to add more complexity in the embedded features for their equipment and more letters to the names of their offerings!

The leakage of enterprise customers from running their own private servers and private networks, virtual or otherwise, to simply pushing all this back into a cloud provider has become a stampede for the enterprise network exit doors, and no amount of alphabet soup from vendors such as Arista can stop this!

RADIUS Still Lives!

In the words of the presenter Alan DeKok "Radius is the protocol that will never die." RADIUS, or to give it its full name "Remote Authentication Dial-In User Service" had its moment in the days of dial-up access services, where a user would make contact with a network access service and provide a username and password. Radius was supposedly overtaken by Diameter as a general-purpose Authentication, Access and Accounting service. But Radius is ubiquitous, open sourced, and, interestingly, baked into IEEE 802.1X standards, and for these reasons it lives on!

It's not that Radius is perfec. Far from it. RFC 5050 from 2007 listed a number of issues and fixes and there is an ongoing effort to list issues and fixes, and there are some Radius implementations contain these fixes. But the generic issue with freeware is that it's simultaneously everyone's problem and nobody's problem. Unlike vendorware, there is no controlling entity who is maintaining coherency, tracking bugs, working on incremental changes to the code base that fix the issues without disturbing all other aspects of the operation of the code, and distributing updates to folk who are using the code, in both servers (easier) and embedded clients (mind-bendingly hard at times!). With many code bases of the specification, maintained in many difference ways, including not at all, the effort to maintain the various instances of the radius code often exceeds the total capacity available.

Despite this, Radius lives on. Some implementations are still vulnerable to decades-old bugs, some are well maintained, but operate in subtly different ways. Some code bases try to keep up with the ongoing efforts identifying issues and their resolution, while others do not. Is this different from other widely implemented open tools that have a strong open-source freeware component? Not really.

PON Broadband Access Networks

New Zealand has been deploying fibre optical networks for network access reticulation for more than a decade now, similar to other developed economies. For the Enable NZ enterprise the network was based on 10G trunks to the GPON optical line termination units and 2.5G./1.25G services to the PON splitters that service individual residences.

While a decade ago this may have appears to be a large capacity network, the installation of network capacity in access networks is a self-fulfilling prophesy. This provider embarked on an upgrade moving away from Huawei to a Nokia Access network with a Cisco backbone core, giving the network an order of magnitude upgrade in capacity as well as a more politically acceptable network equipment vendor. This

provides the network with XGSPON capability with symmetric 10G capacity, and 100G internal links from the spine equipment to the OLTs.

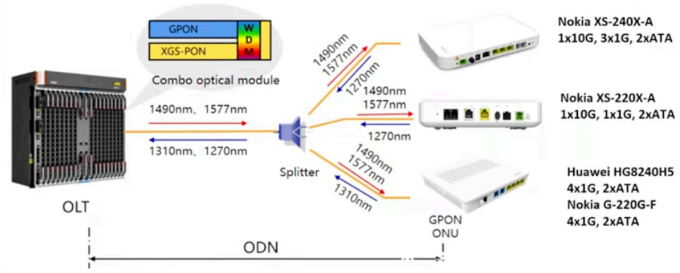

Figure 2 – Upgrade to XGSPON

The upgrade has not been without its issues, with problems of increased latency and packet drop, misalignment of the ONT line termination cards in a multi-vendor environment, fragmentation drop issues and an issue that many IPv6 providers struggle with, DHCPv6 and DHCP relay forwarding.

One of the more telling questions of the presenter on this topic was if they knew when they were selecting vendors what they subsequently learned when operating the platform, would they have still selected a multi-vendor network? The short and sweet answer was "No!"

The Inevitable AI Presentation

These days it seems to be mandatory to have a presentation on the application of AI techniques into network operations. This time it's the treatment of network log processing.

There are some caveats about feeding a Large language Model with the network's configuration and diagnostic report, as the results so far appear to be somewhat variable in quality. The issue as always is that AI is not a deductive tool and does not apply semantic knowledge to the domain it is operating in. AI is a far simpler process of assembling patterns of behaviour and matching new inputs against those observed patterns to infer a most likely next element. Admittedly, it's still early days here but the promise of replacing the role of human operational support with basic pattern matching seems to be another triumph of the current wave of AI hype.

Network Management

As much as some aspects of network technology appear to be rapidly changing, the area of network management appears to have its feet embedded deeply in the mire! Here change is incredibly slow.

The early days of network management in the early 90's saw the emergence of the Simple Network Management Protocol (SNMP) as the ubiquitous management tool. The SNMP model of a managed device as a collection of registers (or counters), with the basic operations of GET, and SET to interrogate and modify the device's operation. There was also the concept of a device-initiated notification when a device's register value changed beyond some threshold value. A structure was imposed on this collection of registers through the definition of a Management Information Base (MIB), allowing the network manager to peer into the operational state of a managed device and infer its operational state by examining the values in the device's MIB. It was a clunky approach, but it permitted the addition of network management hooks into otherwise simple devices, and SNMP quickly became the foundation of most network management systems.

SNMP was accompanied by a text command line interface (CLI). These CLs drew their inspiration from the data initialisation statements used in the C language (among others), where the device was modelled as a set of named records, which could contain variables or other records. A configuration was generated by a text sequence naming these records enumerating the variables and their values.

This approach had its issues, particularly in its clunky nature of viewing a management device as a set of registers, and a very simplistic (naive) security model, but we persisted with it for some decades.

In the early 2000's the industry headed into the world of Netconf, which used the concept of Remote Procedure Calls, where the management controller could invoke more complex functions to be executed by the managed device.

The data storage definition was refined in Yang in 2010, a data description language built on a hierarchy of modules, container(s), list(s)s and leaf(s). The value of Yang was that there were bindings defined for many programming languages and platforms, allowing programmatic (API) access to the operating state of a managed device.

In 2015 Google open-sourced "gRPC", using protocol buffers as the description language for request and reply payloads, with the RPC function layered over HTTP/2, which itself could be layered over TLS. There were gRPC interfaces for common network operations, including GNMI for configuration and state manipulation, set/get functions and streaming telemetry. There was gNOI that allowed for common operations on a device, including file operations, ping and traceroute and validation of a link's operational status.

There is certainly an evolutional path going on here, albeit very slowly. The initial approach was "imperative" where particular directives were passed to the managed device in the context of the device's capabilities. The direction is heading to a "declarative" approach of describing an end goal to the network management system and allowing the system to determine the specific actions to achieve that objective.

The cry of frustration in the network management space goes along the lines that if we can already build systems for autonomous self-driving cars, then why can't we build systems for autonomous self-operating networks?

If the answer is that if such a task requires a massive capital investment, then network operations are still just not an attractive investment proposition for comprehensive network automation. As the same time as we plead for better network management tools, we continue to move functionality and value out of the network and load it up into the edge and the application layer. I suspect that the desire for more sophisticated network management tooling is not accompanied with a capability to pay for it! So the inevitable result is that nothing much happens. Slowly.

Optics

In contrast, the area of optical transmission is a very active one.

The initial optical systems used light emitting diodes as light sources and light sensitive diodes as receivers, using a basic on/off optical keying. Successive refinements of both the transmitter and the receiver have allowed this simple signal modulation to reliable achieve reliable signal rate of 10Gbps, and lab results ten times that speed at 100Gbps. Further speed refinements of this modulation form are generally considered to be somewhat unlikely. If we want to achieve higher optical speeds, we can use the same techniques that we used in electrical signal processing, turning to modulation techniques that change the amplitude, phase and polarisation of the carrier light signal.

This functionality can be implemented in pluggable optical transceivers. These are simple devices with three components: a digital signal processor, a receiver and a transmitter (Figure 3).

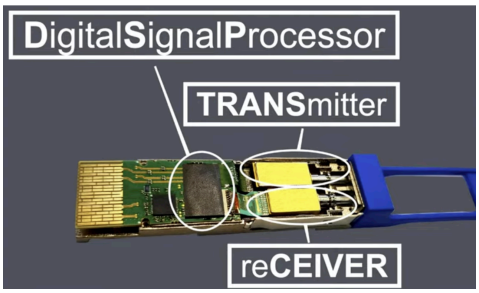

Figure 3 – Pluggable Optical Transciever

The optical transmitter is a solid-state semiconductor laser, operating at room temperature. LED light sources operate as a 1300nm light source with a large line width of some 100nm. A refinement to the LED laser is the Fabry-Pérot (FP) laser, which has a centre wavelength of 1310 nm and a line width of less than 10nm. A Distributed Feedback (DFB) laser is a type of semiconductor laser diode that uses a periodic diffraction grating within the active medium to create a stable, single-wavelength, and narrow-linewidth light source of less than 1nm. Known for high precision, they are widely used in fiber-optic communications, long-reach data transmission (10-25Gb/s). External-cavity diode lasers (ECDLs/ECLs) are highly stable, narrow-linewidth lasers that use external optical elements (gratings, mirrors, or filters) to provide precise wavelength control. They offer superior performance with high power, low noise, and wide tunability (>100 nm). A traditional approach to increase the capacity of a fibre is to use frequency division multiplexing, dividing the total fibre capacity by frequency. This is inefficient as the filters required to pull apart the individual channels need guard bands, or channel separation to prevent channel crosstalk. Another approach is to use a single carrier with a far higher bandwidth.

Optical receivers can be light-sensitive diodes where light reaching the diode opens the conductor path. LED receiver switching speeds can be tuned into the gigahertz frequency ranges. If the incoming optical signal is coupled with the output from a local oscillator, then the carrier signal can be removed and phase shifts used by modulation keying can be detected by a pair of received diodes.

So, we have both simple on-off keying (OOK) and phase-amplitude keying to modulate a n optical carrier. For a 400G data stream, using a carrier with 16QAM modulation, then it's necessary to use some 75Ghz optical spectrum bandwidth. 800Gbps requires double this bandwidth and 1.6Tbps double again. Dense wave division multiplexing of multiple OOK signals use 96 channels, each of 10Gbps, mixed into a common bearer with 50Ghz band filters for each channel. Those inter-channel guard bands consume a lot of spectrum on the cable. When we go towards 400G and above, we need to reduce the number of channels down to 12 for example, then we have 400 gigahertz filter blocks. And then we can feed in 400G, 800G, or even 1.6T signals into such coherent systems.

The optical receiver is configured as an Intradyne receiver where the local oscillator is tuned as close to the optical carrier frequency as possible, and within the optical carrier bandwidth. This gives an intermediate frequency of around 5Ghz, which is easily handled within the on-chip capabilities of the digital signal processor. The side-effect of this is that any other carrier signal in the fibre is not detected, as the intradyne receiver only picks up the carrier signal for which the receiver has been tuned.

This allows for a high degree of flexibility where narrow band DWDM signals can be used in one part of a fibre's optical spectrum, while another can be used for these wide-band 400-Ghz spectrum blocks (Figure 4). Coherent Detection techniques in fibre are robust and very impressive, and 400Ghz muxes provide a lot of flexibility in the design of optical networks. 800G Coherent ZR+ DD and OSFP pluggable optics are now available.

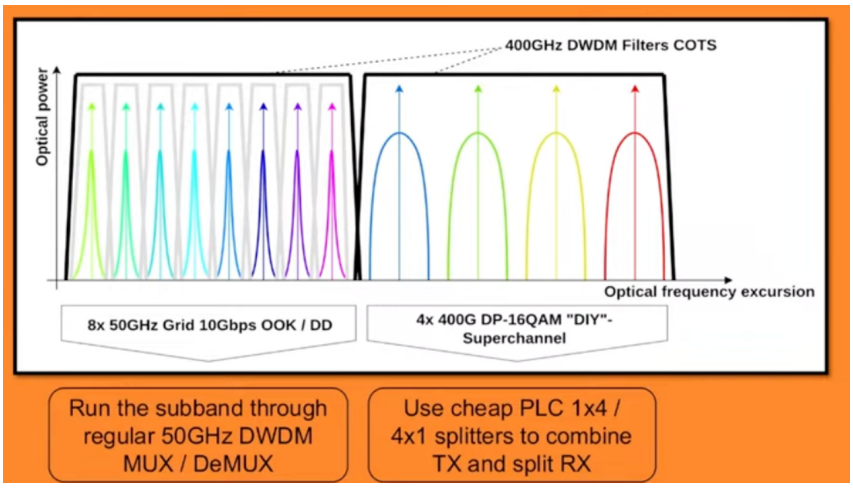

Figure 4 – Combining modulation methods on a single cable.

A LEO service for Mobiles

Starlink is a technically impressive service. Using phased array software-steerable antennae and a 1m dish, looking some 340km – 550km upward to a constellation of some 10,000 spacecraft, it can provide Internet access at speeds of some 200Mbps to almost anywhere on the surface of the earth. Amazon LEO is launching some 3,000 spacecraft into a slightly higher orbital plane of some 600km, apparently offering a service portfolio similar to that of Starlink, which is equally impressive.

Starlink provides a service for mobile phones, but there are some quite severe restrictions on the signal capabilities of their service. This is due to the unfocussed transmission of the mobile device, the limited transmission power and the 350km distance between handset and spacecraft. Reports of available capacity for mobile devices using Starlink are apparently in the region of 5Kbps, so the service is limited to SMS and location services, and to, some extent, voice capability.

If you want to serve mobile devices in a manner that is similar to the service from terrestrial networks from a space platform, the basic requirement is to increase the signal power and signal sensitivity. That's where AST Space Mobile and their spacecraft come into play. This is a spacecraft with very large antenna, some 223 sq^m in size, where the earth-facing side is an array of phased array units, and the other side is covered by solar panels. AST's deployment plans for 95 of these spacecraft in orbit. Communication with the satellites use the Q/V bands.

With their large antennae they are intending to operate as a space-based base station with direct-to-handset 4G and 5G digital service using the 3GPP standards, using the 900Mhz Band 8 LTE, with conventional voice and data services. AST will use a 5Mhz bandwidth allocation, compared to 5G's allocation of 145Mhz.

Their commercial model is not a Starlink-style direct retail mode, but a wholesale one based on service agreements with Local mobile providers. The local mobile provider operates the ground stations and connects the ground station directly to the terrestrial mobile network.

The New Zealand provider, 2Degrees, is looking at using this AST service for disaster recovery and emergency services, remote communications in search and rescue, and mission critical applications.

It's an interesting set of market tensions that are at play here. Is there a viable market volume for services to mobile handsets alone that can sustain a dedicated service operator such as AST's offering? Or are terrestrial services platforms already so well established that the only exposed markets are niche markets in extremely remote locations with low numbers of subscribers. And is the population of such niche markets sufficiently large to meet the capital return requirements of a of initial investment required to

launch a LEO service? Starlink and Amazon Leo have headed in the other direction, driven by the market for high-speed data services that are price cost competitively with many terrestrial providers. Its mobile services are a secondary thought for these LEO constellation operators, and the service quality available to mobile handsets is far lower than users expect from terrestrial 4G and 5G network. As Spacex heads closer to an IPO, there is a level of support in the investor market that this data-oriented model is the one that is sustainable. It's not clear yet whether the AST offering is another run of the Iridium story, but such an outcome is a distinct risk in its focus on the mobile handset market.

Open Fibre Data Standards

As we have seen over the years, telecommunications infrastructure as a national strategic asset. The open disclosure of the location of cables, buildings housing switching equipment and power stations and generators is all too easily transformed into a target list when hostilities break out. This is by no means a recent concern.

> When Britain entered World War I one of its first acts was to cut 5 submarine cables in the English Channel linking Germany to France, Spain and the Azores, forcing Germany to use radio systems for its international communications. One of the most serious consequences of the cable cutting for Germany was that Britain was able to intercept and decode the Zimmermann telegram. This was an attempt by Germany to make a secret alliance with Mexico who stood to gain United States territory as a result. Without a secure telegraph connection of their own to the Americas, the Germans were allowed to use the US diplomatic telegraph link, which the US believed would assist peace efforts. Unfortunately for the Germans, this supposedly secure route went through Britain and was listened to by British intelligence. The revelation of this German duplicity was partly responsible for the US later entering the war.

One response is to try and keep the location of telecommunications infrastructure a highly restricted data set, and criminalise its deliberate disclosure, particularly if the intent behind the disclosure was hostile or malicious. However, by shrouding the location of this infrastructure in a veil of secrecy it increases the likelihood of accidental disruption, and the precise location of this infrastructure is already well known to capable adversaries. So why not go the other way and make the entire data set open? This is the motivation for the Open Fibre Data Standard program, supported by ISOC and others.

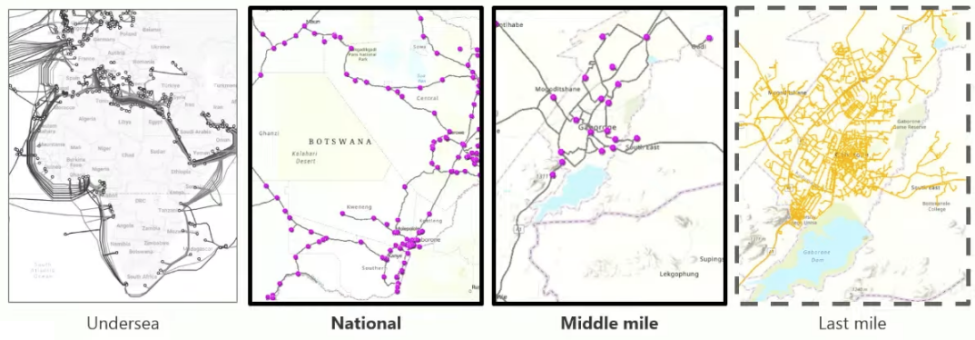

Figure 5 – Progressive granularity of communications infrastreucture

Does this result in a more resilient outcome for users? Do such maps assist in responses to natural disasters? Or does the disclosure of such information act to increase the attack surface for critical infrastructure? Personally, I am of the view that more data is always better. These days we are increasingly reliant on private sector investment to build and maintain comms infrastructure and such loosely

coordinated actions by multiple providers are materially assisted by more accurate and timely information on state of comms infrastructure.

Optical Transport

These days 100G, 200G and even 400G is passé in trunk communications, as we shift optical networks to 800G and 1.6G line speeds in optical communications (Figure 6).

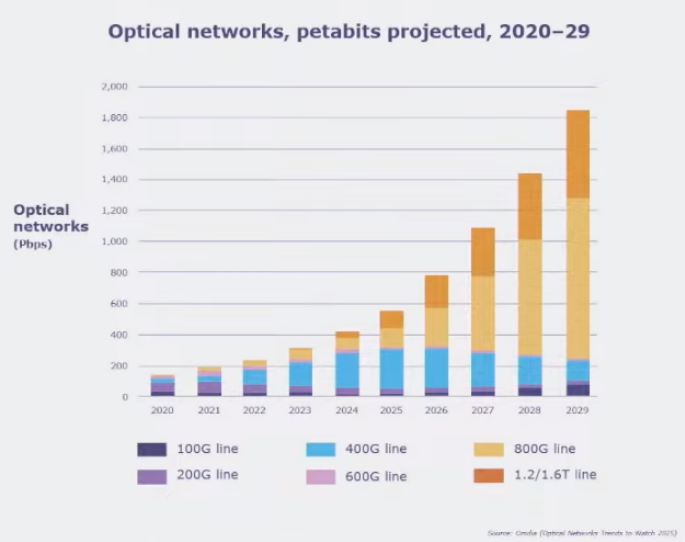

Figure 6 – Optical Carriage Trends

There are a number of technologies that have enabled this trend. The critical shift is a shift from OOK modulation to coherent modulation, using the phase amplitude space to encode more bits per baud, as well as increasing the basic baud rate, and this is supported by more capable digital signal processors. Such capabilities are bound by the gate density from the chip fabrication process and higher capabilities in DSPs are enabled as the industry shifts into 3nm and 2nm chip feature sizes. Today's systems are also capable of performing probabilistic shaping of the points used in the phase amplitude space, so instead of falling back from an 8x8 grid used in a 64 QAM space to a 4x4 grid in 16 QAM space, the DSPs can remove just those points in the QAM space where the line noise prevents accurate decoding. Thin-film lithium niobate (TFLN), Silicon Photonics (SiPh), and Indium Phosphide (InP) represent the three most critical platforms in modern integrated optics, driving high-speed, low-power data communication in the 800 Gbps/1.6T range. Coherent pluggable transponders are also evolving. Their power requirements are increasing, and this creates a heat dissipation problem. Liquid cooled pluggables, which can cool optical modules that operate at 70 – 100 watts, are now on the market for these high-speed longer distance applications.

With a wealth of tradeoffs in the design of a fibre backbone that are available these days, the overall design task becomes more complex. A backbone carrier will have a number of client demands for different capacities and different transmission lengths, and the exercise is to manage the overall fiber capacity and optical transponder settings to maxmise the efficiency of the fibre network and do so without unduly inflating the cost of the network.

NZNOG 2026

This is a small selection of the presentations at NZNOG 2026. The full program can be found at https://www.nznog.org/conferences/nznog-2026/programme, and the streaming video is at https://www.nznog.org/conferences/nznog-2026/streaming.

Disclaimer

The above views do not necessarily represent the views or positions of the Asia Pacific Network Information Centre.

Author

Geoff Huston AM, B.Sc., M.Sc., is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region.

www.potaroo.net