

February 2025
Geoff Huston

Network Automation and AI at NANOG 93

I recently attended NANOG 93, in Atlanta in the first week of February (<https://nanog.org/events/nanog-93/agenda/>). The dominant theme of the presentations this time around was the combination of automation of network command and control and the application of Artificial Intelligence tools to this control function. The interest in AI appears to have heightened of late, and while the hype levels are impressive even for an industry that can get totally fixated on hype, the deliverables so far still appear to fall somewhat short. I suspect that we have yet to peak out on the levels of hyperbole on the boundless applications of AI to this area, and the greater the level of hype the greater the bar of the capability of deliverables that are being placed on this effort right now. The minimum level of functionality expected of the application of AI to today's networks is rising day by day.

Understandably, this is a source of tension between suppliers (vendors) and customers as to the extent of customer-driven customisation of the network equipment. A limited repertoire of configuration settings is preferred by the vendor, on the basis of limiting the amount of rope that the customer has access to, to prevent the equipment from being configured into unstable or destructive operational states. At the same time, there is strong push in the industry all the way to customer-programmed routers. We have seen elements of this in switches with P4 and Intel's Tofino switch ASIC offerings, and programmable interface drivers in general purpose compute platforms, so it's no surprise to see the same pressures emerge in the routing space between standalone packaged routing functionality and elastic platforms that require significant levels of added software to function.

At the same time, there is still the tension in the networking space as to "whose network is it anyway?" in terms of command-and-control functions that are intended to govern the behaviour of the network and the means of sharing the common resource across a number of concurrent network applications. About the best illustration I can provide is the original CSMA/CD Ethernet as one extreme end of this model. The common cable had no command-and-control capability as it was just a passive cable. The sharing of the network across the competing demands was a matter for each connected device and the application. The other extreme was the old carrier virtual circuit system where two end points of the network were interconnected by a synchronous virtual circuit that was implemented as a set of stateful rules in a sequence of carrier switches that were installed essentially by hand.

We've seen various networking technologies swing to and fro between fully stateless systems that exhibit a degree of self-regulation (and hence have reduced reliance on centralised command and control) and systems that recreate end-to-end virtual circuits by the imposition of network state through command and control. Network operators and their vendors have a keen interest in maintaining value in the network platform, and the perceived path to achieve this is through amassing of function into the network platform (QoS, and VPNs are classic examples of this), while the increasing functionality loaded into end systems, and a strong push towards hiding end-application behaviour through encompassing encryption show a strong counter-force that implicitly assigns a simple undistinguished commodity role to network carriage.

In this space AI is being recruited by both parties. The network-centric view sees AI as a means of increasing the capability, accuracy and scale function of command-and-control driven network platforms,

working closely with automation of these functions. At the same time AI functionality is being loaded into the application space in a myriad of ways. Some see this as a means of increasing the capability of end systems in the networked world in the same manner as a conventional traffic network: passive roads and human drivers evolving to place the AI function into driving functions, rather than into a more nebulous concept of AI roads!

Network Automation

The opening presentation for the meeting, by Kentik's Justin Ryburn, was on the topic of network automation and AI. I'm reminded of the observations about self-driving cars, that it's harder than it might've looked at the start! But the theme of this presentation was the assertion that the technology platform for automation of a network's command and control system is not the hardest part. The more challenging part is culture change both with a workforce that has built their career around mastering human interfaces to the network control systems and the vendors who build to what customers claim that they want to buy. So, in 2025 we persist in using Command Line Interfaces to network equipment that has a continuous heritage to the command line interface of a PDP-11 minicomputer running RSX-11 in the 1970's! Change is hard.

I also think that part of this problem lies in an incremental approach to network automation. The resistance to change from the workforce is an issue when automation continues work alongside human-controlled network command and control. I have yet to hear of a case study where the transition to a fully automated network was via a single "flag day" where automation was installed comprehensively to replace every human-controlled network operational function. I'm not even sure any network operator would behave such a seemingly reckless manner today! I guess that behind all the hype about how network automation is ready to roll out today, lies a large-scale lingering doubt about the true capabilities of such systems, and a consequent reluctance to pull all human operators out of the network control room. I can't see this changing anytime soon.

gRPC

This shift toward network automation is associated with a shift in the model of interaction between an automated controller and the network elements it is intended to manage. Much of the network management still runs the venerable SNMP protocol, which is, as the word "simple" suggests is a naive polling protocol. Nokia have developed a network management framework based on generalised Remote Procedure Calls (gRPC).

To me this approach makes a whole lot of sense in the context of network management. It quickly leads to the adoption of a common object model by both the controllers and the objects being managed in this way. This allows the typical tasks of configuration management, troubleshooting and streaming telemetry applications to be implemented within a single RPC framework. They have also implemented a set of operational tools, including ping, traceroute, file get and put. There are also a set of services relating to authentication and credentials, and another set of services relating to forwarding and route management.




Figure 1 – GRPC Framework (From https://storage.googleapis.com/site-media-prod/meetings/NANOG93/5292/20250202_Laichi_Grpc_Services_Under_v1.pdf)

The gRPC framework is developed on top of *protobuf* (<https://protobuf.dev/>), rather than JSON. Protobuf is faster and allows a more compact representation of serialised data. It also uses the *gRPC framework*

(<https://grpc.io/>), implemented on top of an HTTP 2.0 transport. This latter aspect surprises me, as this a serial transport base that can encounter head of line blocking. I had thought that QUIC would've been a more appropriate choice for transport here, as QUIC allows for multiple RPC threads that operate asynchronously - that is without inter-thread dependences and no head-of-line blocking.

IAB Workshop

A general view of the state of the industry with respect to network management tools and techniques was provided in a report from the December 2024 IAB workshop on the Next Era of Network Management Operations (<https://notes.ietf.org/nemops-workshop-next-steps>). The survey report that was presented at this workshop provides an interesting snapshot as to how networks are actually operated today.

There has been a lot of effort put into various automation tools to manage networks, including NAPALM, Terraform, Ansible, and of course the persistent use of Command Line Interfaces (cli). The heartening news appears to be that the use of ansible now exceeds that of using the device cli, but these two tools are the two most popular configuration management tools being used in today's networks (Figure 2).

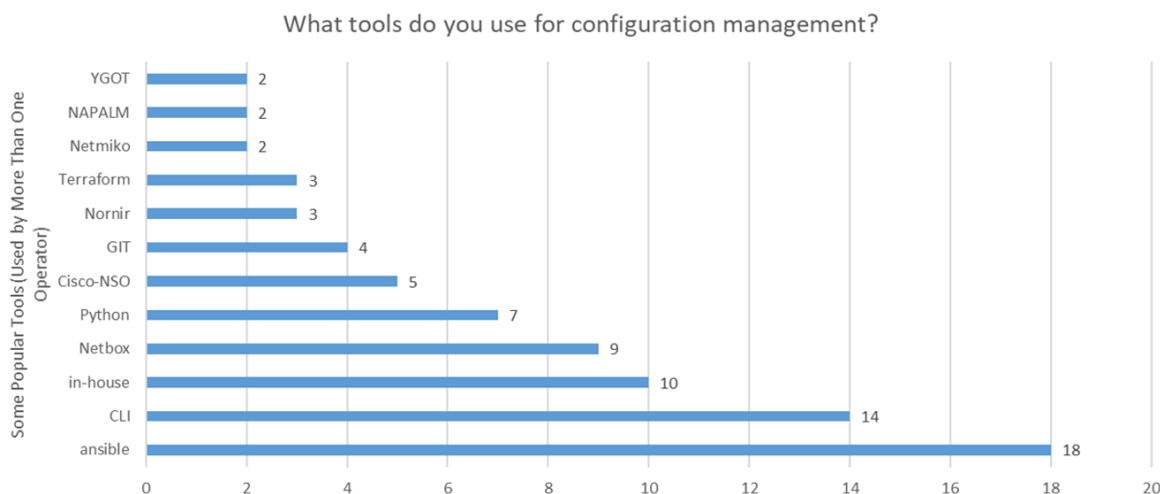


Figure 2 – Network Management Tool Use – From https://storage.googleapis.com/site-media-prod/meetings/NANOG93/5348/20250202_Krishnan_Summary_And_Next_v1.pdf

SNMP is still a very widely used platform for these tools, as well as NETCONF. For network monitoring snmp remains a firm favourite, while YANG-Push and Netconf have yet to gain a large base of adopters in this space (Figure 3).

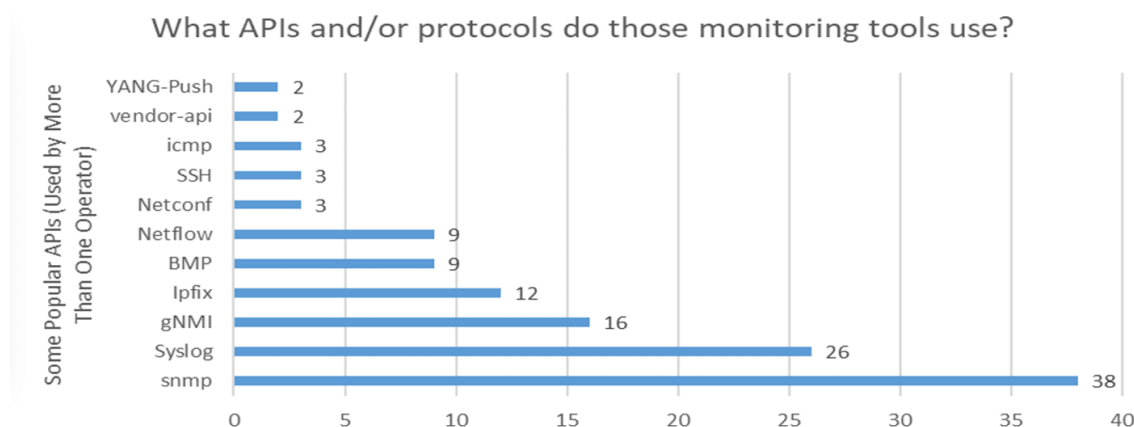


Figure 3 – Network Monitoring Tool Use – From https://storage.googleapis.com/site-media-prod/meetings/NANOG93/5348/20250202_Krishnan_Summary_And_Next_v1.pdf

It has often been observed in the transition to an IPv6 network that the plethora of dual stack transition approaches that were developed by the IETF at the time served more to confuse that market than to

facilitate a rapid and efficient transition. Similarly, the plethora of approaches to automation of network management has managed to fragment and confuse the space, and CLIs and SNMP remains popular because of its simplicity, ease of use, and uniform availability. It's not clear that this situation is going to be changing anytime soon.

Networking for AI

AI is not just a currently fashionable approach to the role of network operations. The computing characteristics to train the Large Language Models (LLMs) are intense and the trend in the compute cost (until the recent announcement of DeepSeek) has been up and to the right. Today's models require hundreds, heading to thousands of GPUs to train these LLMs (Table 1).

Model Name	Model Size (Billion Parameters)	Dataset Size (Billion Tokens)	Training ZETA FLOPS (10 ^{^21})	Number of GPUs
OPT	175	300	430	474
LLaMA	65	1400	600	662
LLaMA 2	34	2000	400	441
LLAMA 3	70	2000	800	882
GPT-3	175	300	420	463
GPT-4 (est)	1500	2600	31200	34392

Table 1 – Compute Requirements for LLMs

In AI data centres there is a need for a “backend” switching fabric interconnecting connecting the servers that contain the GPUs. In many ways this environment is operating in a manner more closely aligned to remote memory access, as distinct from remote procedure calls.

As with the evolution of the processing environment to achieve higher throughput some decades ago, the techniques of parallelism, pipelining and segmentation can be used in the datacentre to improve overall performance. In very general terms, a local AI network is a large set of GPUs and a set of memory banks, and the network is cast into the role of a distributed common bus to interconnect GPUs to memory in a similar vein that the backplane of the old mainframe computer design was used to connect a number of processing engines to a set of storage banks. Here Remote Memory Access performance is critical, and the approach, in very general terms, is to enable the network to pass packet payloads direct to the hardware modules that write to memory and similarly assemble packets to send by reading directly from memory. The current way to achieve this is by RoCE, or Remote Direct Memory Access (RDMA) over Converged Ethernet. Network-intensive applications like AI, networked storage or cluster computing need a network infrastructure with a high bandwidth and low latency. The advantages of RoCE over other network application programming interfaces such as the socket abstraction are lower latency, lower processing overheads and higher bandwidth. However, the concept of using the network for remote memory (Remote Memory Access, or RMA) where it is possible to transfer from one device's memory to another without involving the operating systems of the two devices is not a new one, and RMA has been around for some 25 years, often seen in parallel computing clusters. However, these approaches had their shortcomings, including the lack of explicit support for multi-pathing within the network (which constrained the network service to strict in-order delivery), poor recovery from single packet failure and unwieldy congestion control mechanisms.

Ultra Ethernet is not a new media layer for Ethernet but is intended to be an evolution for the RMA protocol for larger and more demanding network clusters. The key changes are the support for multi-pathing in the network and an associated relaxing of the strict in-order packet delivery requirement. It also uses rapid loss recovery and improved congestion control. In short, it appears to apply the behaviours of today's high performance multi-path transport-level protocols to the RMA world. The idea is to use a protocol that can make effective use of the mesh topology used in high performance data centre networks to allow the scaling of the interconnection of GPUs with storage, computation and distribution networks.

A critical part of this approach lies in multi-path support which is based on dropping the requirement for strict in-order packet delivery. Each packet is tagged with its ultimate memory address, allowing arriving packets to be placed directly into memory without any order-based blocking. However, this does place a higher burden on loss detection within the overall RMA architecture. Ultra Ethernet replaces silent packet discard with a form of signalled discard, similar to the ECN signalling mechanisms used in L4S in TCP. If a packet cannot be queued in a switch because the queue is fully occupied, the packet is trimmed to a 64-octet header snippet and this snippet is placed into a high priority send queue. Reception of such a trimmed packet causes the receiver to explicitly request re-transmission from the sender for the missing packet, which is analogous to selective acknowledgment (SACK) signalling used in TCP.

The framework also proposes the use of fast startup, including adding data payloads to the initial handshake used to establish a flow. This approach eliminates the delay of a round-trip handshake before transmitting, as the connection is established on-demand by the first data packet.

The challenge for many networking protocols is to be sufficiently generic to be useful in a broad diversity of environments. If you are permitted to look at optimising the performance of the protocol in a far more limited scenario then it's possible to introduce further optimisations for the protocol. That is what is going on for UltraEthernet, which takes the more generic approach of RMA over IP and makes some specific assumptions about the environment and the payload profile that are specific to high performance data centres used for data-intensive processing, as seen in AI applications. It's an interesting approach to scaling in data centres, as it does not attempt to alter the underlying Ethernet behaviours, but pulls in much of the experience gained from high performance TCP and applying it directly to an Ethernet packet RMA management library.

NANOG 93

The agenda, presentation material and recordings of sessions can all be found at the NANOG 93 web site: <https://nanog.org/events/nanog-93/>

Disclaimer

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

Author

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

www.potaroo.net