

No networking protocol has experienced the success of TCP/IP. Thirty years ago, it seemed possible that IBM might have fended it off with its SAA family of protocols or we might have ended up using the Open Systems Interconnection (OSI) family of protocols. But, thanks to the opening up of the US Arpanet once TCP/IP had been developed, the other options were quickly swept aside.

IPv4 itself is so successful that long after the internet was meant to have run out of addresses, IPv6 penetration remains stubbornly below 30 per cent globally, according to Google's analysis and growth has been tailing off after reaching an apparent asymptote in early 2018.

Address-translation techniques appear to have staved off the threat to IPv4. And yet researchers and a handful of networking infrastructure suppliers are planning the unthinkable: to replace IP.

All of the options on the table start from the same place - a plan that removes the idea of compute nodes with unique addresses serving up a smorgasbord of data and services. In its place, under a protocol that employs named-data networking (NDN) every piece of data gets its own address. And that data can live anywhere on the internet that has room for it. There is no need to look up an associated IP address using the Domain Name System (DNS), at least in the purest form of NDN. This ability may be what makes NDN, ultimately, viable as an alternative to conventional IP.

#### Experimental NDN systems

The experimental NDN systems constructed so far work through "interest packets". Each of these packets is a request for data attached to a name that's constructed according to some predefined namespace hierarchy. And there can be many hierarchies: there is no need to shoehorn everything into one long

# Oh data, the places you'll go

Both autonomous driving and the IoT could lead to major changes in how the internet works, but IP is a stubborn incumbent. By **Chris Edwards**

list of dotted addresses.

One application where NDN may gain a foothold is in vehicle-to-vehicle (V2V) communication with its promised ability to sever the link between location and machine ID. The naming hierarchy might be based on the road network rather than the vehicles travelling around it. To use it, a car trying to look at road conditions ahead could harness data from whichever vehicle the network considers to be closest to the point of interest.

In contrast, the IP-based alternative demands a constantly updated centralised map of vehicle addresses, which could easily add unwanted latency to a fast-moving environment. With NDN, the car's ID - which would also be its IP address in today's network - does not matter. A different application would not encode the location but some other aspect of the data, such as a group of frames in a long video.

To get data from A to B, NDN routers, like their IP counterparts, maintain forwarding tables. NDN has two types. One is a pending-interest table (PIT), which keeps a record of any requests for which the router has not seen a reply. Returned data uses the PIT entries to find the breadcrumb trail back to whoever made the request. Van Jacobsen, one of the original IP developers who then came up with the initial ideas for NDN in the mid-2000s, argues the design architecture is much better

**"IPv4 itself is so successful that long after the internet was meant to have run out of addresses, IPv6 penetration remains below 30 per cent globally."**

for multicasting for IP as a router can forward to many ports at once using the PIT entries.

Because chunks of data are no longer located by their IP address, they can be held anywhere on the network and potentially inside content stores managed by the routers themselves, allowing the data to migrate from the original server closer to regions of high request activity. The source data elements themselves are found by looking up portions of the name in a forwarding table and then relaying the request to that upstream port. Each piece of data is encrypted individually to protect it although there may be issues with privacy. Anyone monitoring the network can see the requests for named data passing back and forth, assuming the systems that move into production do not find some way of obfuscating that information.

#### Problems likely to surface

Other problems are likely to surface as trials expand. During a panel session at an ICN workshop in Kansas City last year, UCLA professor Lixia Zhang pointed to the period after TCP/IP first appeared on the Arpanet as one where the protocol stack was tuned gradually over time to deal with problems that emerged as more users joined the network.

Even with small-scale trials performed so far issues have emerged with locating the sources of data.

A pure NDN implementation would

flood the network with broadcast packets announcing the availability of various pieces of data. Zhang and colleagues proposed a distributed directory analogous IP's DNS to deal with the problem. However, experiments have shown that subtle conflicts in how objects are named in a highly distributed system can cause problems. Zhang says a lot more work is needed to research good naming strategies for use with NDN.

Then there is the question of deployment on real-world devices. The development toolkits provided for many operating systems naturally assume the presence of TCP/IP networks and applications may need to be reworked to work efficiently with NDN. Because of this, much of the current R&D work, including several EU-funded Horizon 2020 projects, are looking at ways of making NDN or other forms of ICN coexist with IP.

One option is to simply deploy NDN on top of IP, similar to the tunneling used for other protocols today, but it's far from the only approach.

Two years ago, shortly after buying the content-centric networking

technology from Xerox's PARC research organisation, Cisco said it had started work on what it calls Hybrid ICN (hICN), claiming that its approach solves the compatibility problem by absorbing data-centric methods into IPv6 itself.

The hICN method encodes the names of data into the dotted-number form of IP addresses, taking advantage of the massive expansion of the address space made possible by IPv6. Because of IPv6's slow growth, Cisco has still had to make its hICN work with the older version of the protocol, which means interconnecting IPv6 islands with IPv4.

The rise of software-defined networking (SDN) provides other options do not involve porting a naming system onto IP addresses. SDN switches in the network watch for the packets from these different networks and invoke customised routines to handle each one.

The Horizon 2020 POINT project has also used SDN to implement what researchers call NDN as an underlayer. The project ran a couple of small trials based around a treasure-hunt game in

Bristol in the summer of 2017. In the open trial, players downloaded an app to their Android devices, which used the standard IP stack. But, connected to a private Wi-fi network, requests for video were handled by SDN switches that mapped the HTTP requests onto an ICN protocol. This made it easier to serve up video from the nearest router that had the necessary content in its cache, reducing overall latency.

### Video delivery

One of the early promises of NDN was to deal with the problem of delivering video and other high-bandwidth data to users spread around the world.

However, this is an application that led to the creation of IP-based network caches – content delivery networks (CDNs) – which do not need a novel protocol to operate. Instead, large markets for IP-free ICN may lie in applications that produce far less data but already have problems with IP. These are devices that connect using low-energy protocols such as Lora and 6Lowpan that have severe restrictions on packet size.

One example is US startup Operant Solar, which aims to provide wireless modules for rooftop solar farms using a form of NDN developed with UCLA. The company claims NDN makes it easier to build mesh networks and route around failures. However, as full names could easily consume multiple packets on low-end IoT networks production protocols will most likely compress them for delivery to a gateway that then regenerates the full name that can proceed into the ICN backbone.

Although ICN may make the IoT work better, it still faces an uphill struggle. The rise of CDNs and network-address translation has demonstrated the ability of the IP world to evolve and head off challenges to a protocol with some major flaws. ICN could easily be a victim of another evolution in applications that sit on top of IP that deal with its mobility issues.

Below: The POINT architecture, similar to that used in trials in Bristol, feeds packets from IP networks into a core operated by SDN switches and management systems that convert the headers into versions suitable for name-based routing.

