A new way to connect to the internet

Authors

Dave Wilson, HEAnet. dave.wilson@heanet.ie
Ole Frendved Hansen, DeIC. ole.frendved.hansen@deic.dk
Peter Lavin, Trinity College Dublin. lavinp@cs.tcd.ie
Pau Minoves, i2cat. pau.minoves@i2cat.net
Bo Peng, University of Bristol. bo.peng@bristol.ac.uk

Keywords

IaaS, NaaS, virtual router, ip as a service, network virtualisation, provisioning automation, router deployment, connection deployment, OpenNaaS, Mantychore, NREN

Abstract

We have hit some sort of limit. Despite the community's best efforts in abstracting out the service from the network, the amount of physical equipment being deployed - metal boxes consuming power and space - doesn't seem to be going down. If anything, it's rising.

What is all this equipment doing? Take a typical setup on the boundary of a university network where it meets its NREN. Both the university and the NREN will probably deploy a router each, so they can manage the boundary between their networks. For resilience, you then need to deploy two more. Then the NREN may have non-routed services to deliver, so that will require at least a switch - two for resilience - and then the optical transmission equipment that underlies it all. Moreover, the university also has to enforce its own security policy, requiring them to deploy a firewall - make that two firewalls, one for backup.

This "equipment mountain" is a far cry from the vision of abstracting the IP network as just one of a number of services neatly delivered to the client. And it's a problem, not just because of the absolute cost of the equipment and energy in this one use case. Any attempts to make network services cheaper and easier to deploy - and any benefits we could gain from innovative new network setups - are stymied because the additional equipment costs at each end have come to outweigh the benefit.

How can we overcome this? Well, some parts of the mountain seem intractable - the physical characteristics of the fibre in the ground require a certain level of power to light it, for example. And the other parts of the mountain are the product of hard-learned lessons. We've seen the impact on reliability and manageability of concentrating too many different duties on one router.

The essential work that a router does is to separate networks from each other, so that the customer can operate its network independently of the NREN and other customers, and neither is impacted by changes in the other.

It seems hopeless. We can automate so much of our provisioning, but we still need to order, buy and deploy all these devices in order to get a network that we can actually manage. A router cannot serve two masters; BGP, OSPF, Access Control Lists, route filters and MSDP are all tools that an operator uses to maintain the separation between one network and another.

But those tools also have one other, very important, characteristic in common. They're all done in software. You can virtualise software.

We now have tremendous experience in our community of virtualisation, and the very same benefits that we see on server software can be brought to routing software as well. A device can indeed serve two masters, as long as a virtualisation layer is able to provide that separation.

This seems like the answer to our problem of the equipment mountain. If we can virtualise some of that routing equipment, then we don't need to put so many physical boxes at each site. We save cost, space and energy by aggregating the virtual routers centrally, just as we've been doing for years with servers.

If we virtualise on-site equipment, then we can give customers direct access to “their” routers, which become a true part of their network. The NREN can continue to manage the BGP sessions, while the customer can make their own technology choices on their own internal network and integrate the virtual routers fully with that network. We can go so far as to abstract out parts of the setup, so that the customer has visibility of the operation of their upstream connections without the requirement that they must become familiar with BGP.

It even opens the door to those innovative new services that we've not yet been able to implement, or - even better - that we didn't know our clients had imagined, because the logistics of all the extra equipment made them impractical.

Except for one thing. Virtualisation can save us some cost today, but it doesn't make a dent on operational complexity. Just by freeing ourselves of the need to match physical equipment to every service doesn't mean we can manage every possible configuration that can be imagined.

Fortunately, this is a problem that's been solved as well - by automation. When you automate your provisioning, it doesn't just make things quicker and cheaper over the long term, but it minimises the problems of inventories and monitoring going out of sync. It even eases the burden of change control, because the impact of an automated process that has been run thousands of timed before is much easier to narrow to a tight level of confidence than that of even the most expert and reliable human operator.

So, the software we have developed, called OpenNaaS, takes this very approach. This is a open source framework and the Mantychore project supplies a broader community around the software. HEAnet is one of the participants, and it is rolling out a virtual CPE service based on OpenNaaS. Project partners are applying OpenNaaS in other fields like High Definition Video, Cloud Computing and Connecting Multiple Networks.

This talk will describe both our technical and operational experiences deploying OpenNaaS into services and the experience of developing software that has multiple objectives for a number of different users.

But it will also describe one other thing. Replacing existing equipment is a worthy technical goal and the benefits are clear. However, the project doesn’t stop here. Because when you have taken the steps of virtualisation and automation, one final step is within reach - and it is a very interesting one indeed.

This service - and not just the replacement of physical equipment but the general virtualisation of the layer 3 network - is just one building block of a general vision of being able to build networks on demand. But the other building blocks have already been built, and because they are also automated, we can integrate with them.

OpenNaaS uses Bandwidth on Demand for its interdomain connectivity. This gives it access to circuits that touch every NREN in GEANT, and where those NRENs also implement Bandwidth on Demand, it provides true end-to-end connectivity between endpoints in those NRENs.

But the most interesting applications - the services that we actually deploy - are not built with bandwidth alone. Neither are they built with routers alone. When building most services requires an expensive purchase and a long wait in order to deploy equipment, it is hard to get the full benefit of a Bandwidth on Demand system. But when all parts of the network can be spun up on demand, this changes.

As NRENs, we have been laying the building blocks for this sort of solution for over a decade. Far from being limited to virtualising on-site equipment, we now have a solution for deploying full, virtualised networks which can be managed wholly, partially or not at all by end users with authority delegated from their institutions and NRENs.

Vitae

Ole Frendved Hansen has since 2006 been with the Danish NREN, DeIC (former operated by UNI-C). In Mantychore FP7 research project he is Work Package leader and takes the role of software owner in development of OpenNaaS. Consulting the use cases and bringing forth the requirements. Engaged in GEANT3 project, SA3, providing the eduroam service.

Peter Lavin is a PhD student in the Computer Architecture and Grids research group at Trinity College Dublin. His research interests include integration of virtual networks with grid and cloud infrastructure. He also works in the area of resource allocation in heterogeneous computer resources such as cloud, grid and volunteered computing, and in particular where resources have different social and economic backgrounds. Peter holds a BSc(Eng) in mechanical engineering and a MSc in computer science from Trinity College Dublin.

Pau Minoves graduated on Telecommunications Engineering for the Universitat Politècnica de Catalunya (UPC, 2005) and received an Erasmus Mundus Master in Software Engineering by the Blekinge Technical High School (BTH, 2008) and Universidad Politecnica de Madrid (UPM, 2007). He is currently the DANA's Professional Services Manager, as well as Technical Coordinator and Work Package leader of the MANTYCHORE FP7 research project. Since 2011, he performs as software engineering Associate Professor at Universitat Pompeu Fabra.

Bo Peng received BSc from Xidian University, Xi'an, China, and his MPhil and Ph.D from University of Leeds in Electronic and Electrical Engineering in 2006 and 2010, respectively. He is currently a senior research officer with the High Performance Networks Group, University of Bristol. He has been involved in several national and EU funded projects. His research interests are in high-performance network architecture and technology, software defined networks and virtualisation in heterogeneous wired and wireless networks. Bo is a member of the IEEE and the IET.

Dave Wilson has worked for HEAnet, Ireland's National Research & Education Network since 1996. Dave's work centers around planning the national IP network, including a close focus on areas such as BGP policy and IPv6, and he led the procurement and deployment of the latest iteration of this network. Dave also participates actively in a number of international fora, including RIPE, GÉANT3 and in the Irish IPv6 task force.