

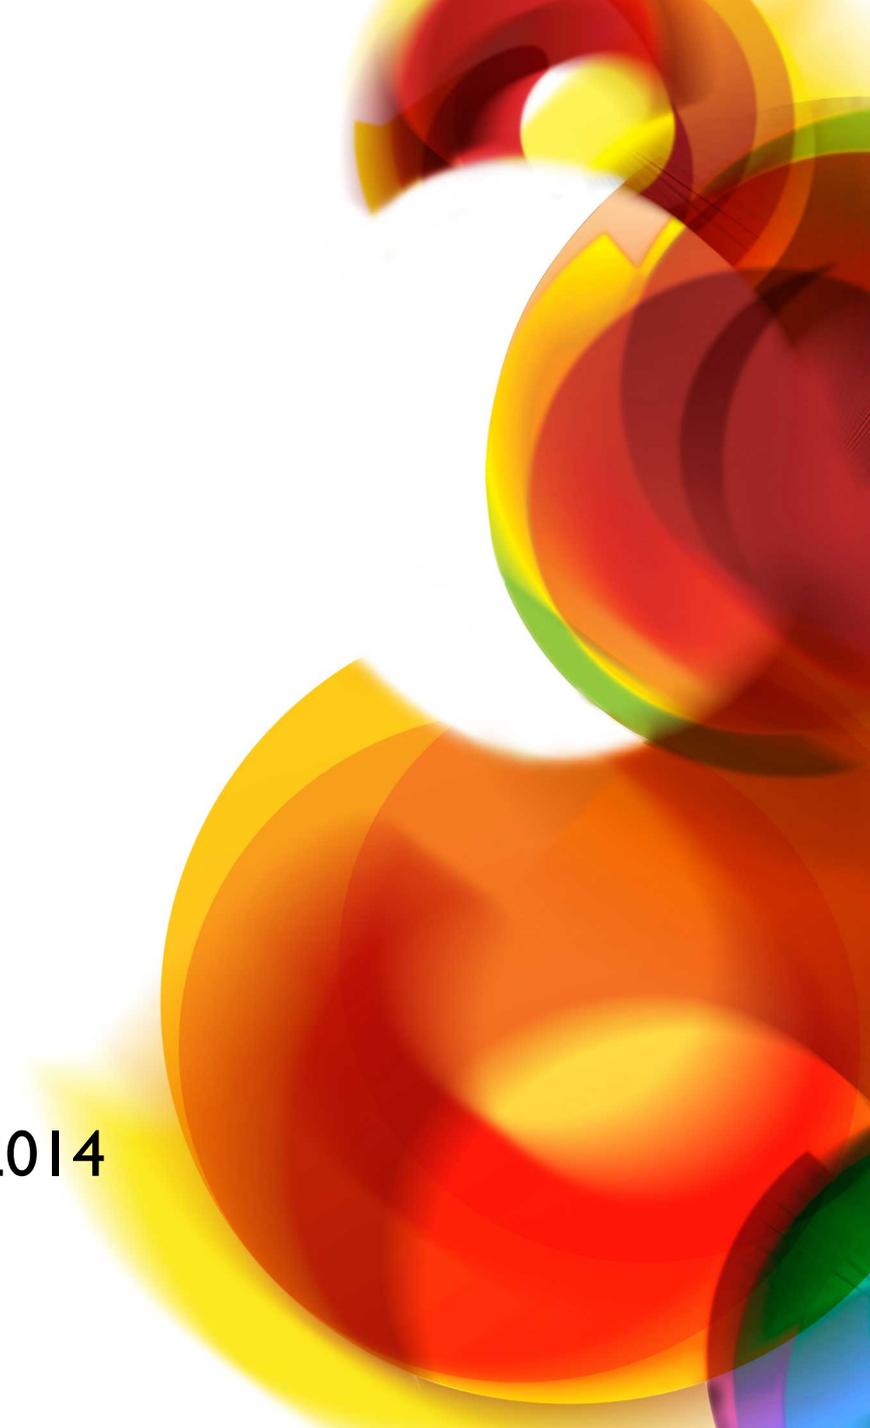
janet

Building the perfect R&E network.

Rob Evans

Janet

Internet2 Technical Exchange, Oct 2014



- Years of
 - Requirements gathering
 - Workshops
 - Capacity planning &c
 - Government engagement (funding)
 - Procurement
 - Competitive dialogue
 - Building
 - Migration
- Dark fibre
- Coherent optical
- Lots of 100GE
- Ethernet over MPLS
- No OTN switching
- Pretty good, but not the perfect NREN.



- A network based on all user requirements, now and for the foreseeable future?
- A network built without making any compromises?
- A network that achieves the right balance between cost, management and features?

“having all the required or desirable elements, qualities, or characteristics; as good as it is possible to be”



What is a perfect NREN?

- High capacity IP provisioning
 - Flexible, statistically multiplexed packet switching
 - Cope with peaks in demand
 - Demanding users
 - Media events
- Rapidly provisioned circuits
 - No bandwidth guarantees, but not congested
- Rapid provision of dedicated bandwidth
 - Mbps, Gbps and many Gbps
- Reliable
 - Multiple failures!
- Meet the needs of production and network research

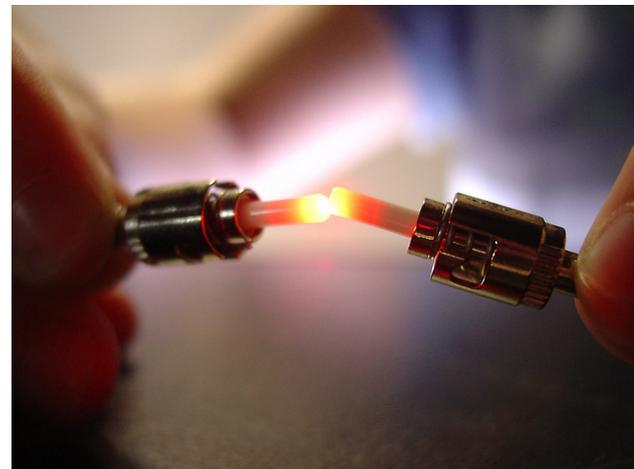


How might that network be built?

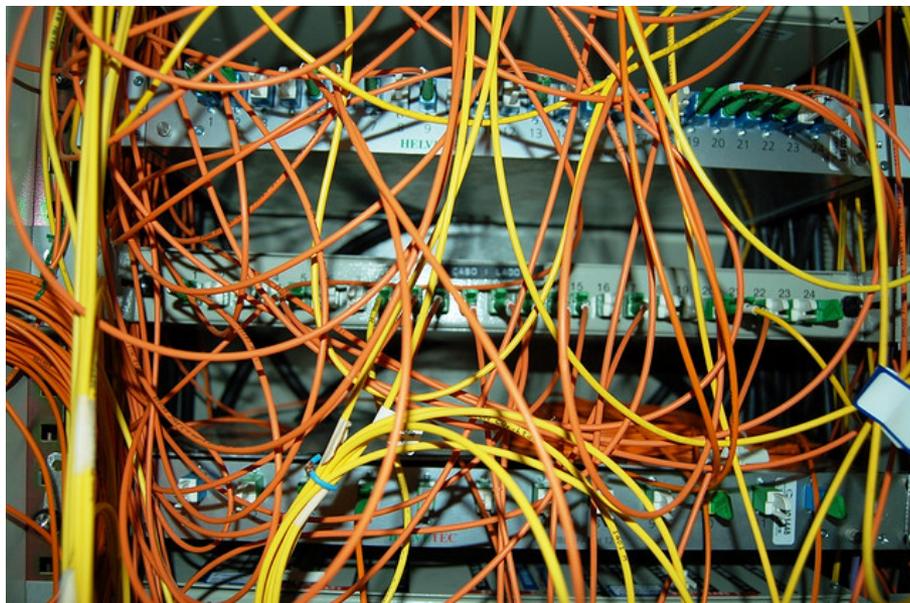
- Dark fibre
- Optical layer
- Switched layer
- Routed layer
- Network management
 - Rapid build of new bandwidth
 - Automated provisioning
 - Network intelligence



- Do you need it?
 - Good for high capacity
 - Visibility of underlying network
 - Requires a long-term contract
 - Cheaper if you don't need capacity
 - Optical spectrum sharing alternative?
- Can you get it?
 - Fibre market in UK is fairly limited, what's it like elsewhere?
 - 'Fibre glut' of early 2000s is history
 - Telcos driving existing fibre harder rather than laying new fibre
- Will you always be able to get it?
 - Unless new fibre is laid, market is unlikely to open up
 - Telcos prefer to sell services, need to be in a strong position to buy fibre.
 - Will you be able to expand a dark fibre network in five years?



- How much of a mesh?
 - Many routes creates a resilient network
 - More degrees required on ROADMs
 - How do you use all the fibre?
 - Switching and protection at the optical layer?
 - Provision IP paths over all the directions?
 - Lots of router ports!
 - Complex IGP topology
 - Simpler may be better



- How 'reconfigurable' do your ROADMs need to be?
 - Less is more
 - Colourless
 - Directionless
 - Contentionless
 - Gridless
 - MEMS or LCOS
- Capacity requirements
 - Are you going to be using a large number of the 88 (96?) channels?
 - Will you want to dynamically reroute wavelengths around the network?
- Where are you building resilience?
 - Protected optical circuits, or at the IP layer?
 - WSSs won't reroute circuits instantly.



- Coherent
 - Offers 40G and 100G over long distances
 - Single 10G across an uncompensated network can be expensive
- Thick or thin?
 - Transponders in the optical equipment, ‘grey’ light to the routers
 - OTN optics in the routers
 - Several years behind the optical vendors
 - Coherent requires a lot of DSP
 - Penalty of ‘grey’ light optics and OEO
 - Especially with the cost of 100GE optics



Which circuits to offer higher layers?

- 100GE is a lot of traffic to fail in one go
 - 100GBASE-LR4 optics are still expensive, but use normal single-mode fibre.
 - 100GBASE-SR10 are cheaper but require special cabling (multi-core multimode cables with MPO connectors)
 - 10x10MSA may have support issues.
- Parallel 10GE?
 - No single flow larger than 10G
 - Lots of cabling
 - If you require all the capacity, then any one of the optics failing in a bundle might mean you have to take the whole bundle down



- What do I mean by ‘switched layer?’
 - Something between the optical layer and the IP layer
 - Circuits
 - The things we had in X.25 before we binned them.
 - The things we had in ATM before we binned them.
 - The things we will forever reinvent.
- What will you be using it to provide?
 - Dedicated capacity circuits? How much capacity?
 - Is the fundamental statistically multiplexed aggregation layer of the network here or at IP?
 - ‘Pool of available bandwidth’
- Virtual networks
 - Network research



- Where should it be?
 - In the ROADMs?
 - OTN, ethernet?
 - Dedicated boxes?
 - Yet more rack space and power
 - In the routers?
 - EoMPLS
- The more places you build in switching capacity, the more layers you need a bandwidth overhead to provide it



-
- How many backbone routers?
 - Don't want to carry traffic further than is necessary.
 - Don't want to process the packet at lots of hops.
 - How far out do you manage?
 - Options for collapsed CE/PE using virtualised functions?
 - I don't need to say IPv6, do I?



- Scaling
 - Power, space
 - Vendors suggesting more boxes rather than always meeting demand with one box



- Content delivery
 - Lots of traffic comes from usual sources
 - Google, Akamai, Limelight, Netflix
 - House servers?
 - Netflix can serve several 30Gbps from a IRU server
 - Can distribute them around network to reduce carrying traffic
 - Local nodes may be cheap/free, but ‘pay’ for power, rack space, time spent arranging broken servers to be replaced.
 - Do you need Google servers, Akamai servers, Limelight servers, Netflix servers?
What about the next big CDN?
 - Just peer?
 - May not be available in all locations
 - Need to carry traffic around the network
 - Easy to add new peers
- ‘Cloud’ access and peering



- QoS?
 - “You don’t need QoS if you build an adequate network” (Owen DeLong commenting on “A reasonable discussion on Net Neutrality”)
 - If you’re preferring some traffic, that means you’re deprioritising everything else.
 - How do you do admission control fairly so only the VoIP traffic is prioritised, but it’s open to all VoIP providers?
 - Network neutrality



- Where does it fit in the wide-area network?
 - A tool for point-to-point circuits?
- Which layer?
- Testbed or more?
- Multi-domain SDN
 - East-west interfaces



- Capacity needs to be lit on the optical layer to use it at higher layers
 - Or at least sitting around...
 - Expensive pre-provisioning
 - May never be used
 - Who has the money?
 - Technology may be obsolete before it is used
 - Flexigrid and dynamic encoding may help here!



- Know what you're selling
 - Bandwidth on demand? Or a circuit on demand?
- Heterogeneous networks
 - Vendor-specific solutions are bad
 - How much experience of network operations does your development team have?
- Who is the user? How do we get to them?



- How smart should the network be?
- Historically, distributed control planes
 - BGP, LDP, OSPF, IS-IS.
- Towards more centralised control
 - SDN
 - MPLS-TE using centralised software to provision/route LSPs
- Is this a good thing?
 - Trust
 - Simplicity



- Now you have your dark fibre, optical layer, switched layer, routed layer and automated circuit provisioning, how do you see what's going on?
- Each layer may be from a different vendor.
- Expensive commercial tool
 - Does it do what you want?
- Open source
 - Will require customisation
 - Will require interface modules
 - Beware: for the optical equipment they may be expensive!



Questions, discussion...

Janet, Lumen House
Library Avenue, Harwell Oxford
Didcot, Oxfordshire
t: +44 (0) 1235 822200
f: +44 (0) 1235 822399
e: Service@ja.net