


Upgrading to STM-256

Rob Evans

JANET(UK)

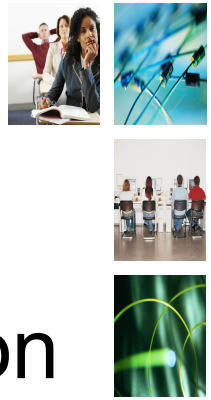

What is JANET?

- UK's National Research and Education Network (NREN)
- Connects educational establishments to each other, the Internet, and other R&E networks around the world
- Serves schools, colleges, universities and research facilities
- Potential userbase of 18 million

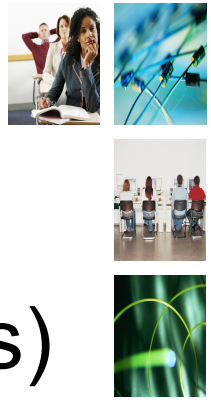

What is JANET?

- Backbone and regional networks (RNs)
  - Sites connect to RNs
  - RNs connect to backbone through two diverse fibre routes
- Two PoPs in London Docklands with most of the external connections (transit, peering)
- This talk concentrates on the backbone

JANET (Fibre)

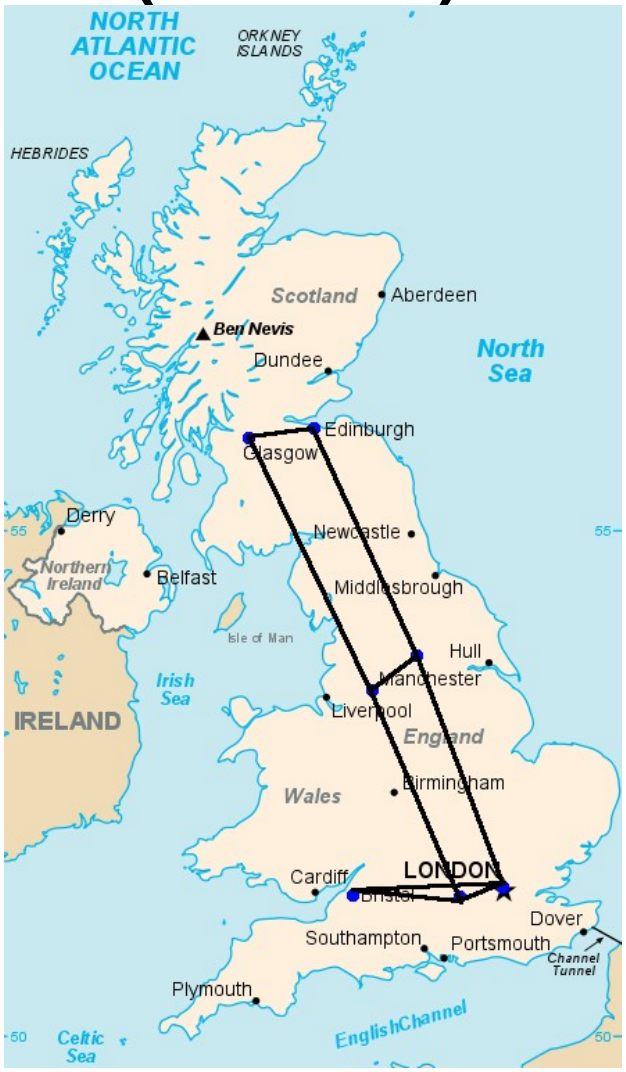
NORTH ATLANTIC OCEAN
ORKNEY ISLANDS
HEBRIDES
Scotland
Aberdeen
Ben Nevis
Dundee
North Sea
Edinburgh
Glasgow
Derry
Newcastle
Northern Ireland
Belfast
Middlesbrough
Isle of Man
Hull
Irish Sea
Manchester
IRELAND
Liverpool
England
Wales
Birmingham
Cardiff
LONDON
Bristol
Dover
Southampton
Portsmouth
Channel Tunnel
Plymouth
Celtic Sea
EnglishChannel
55
50

JANET (Fibre)

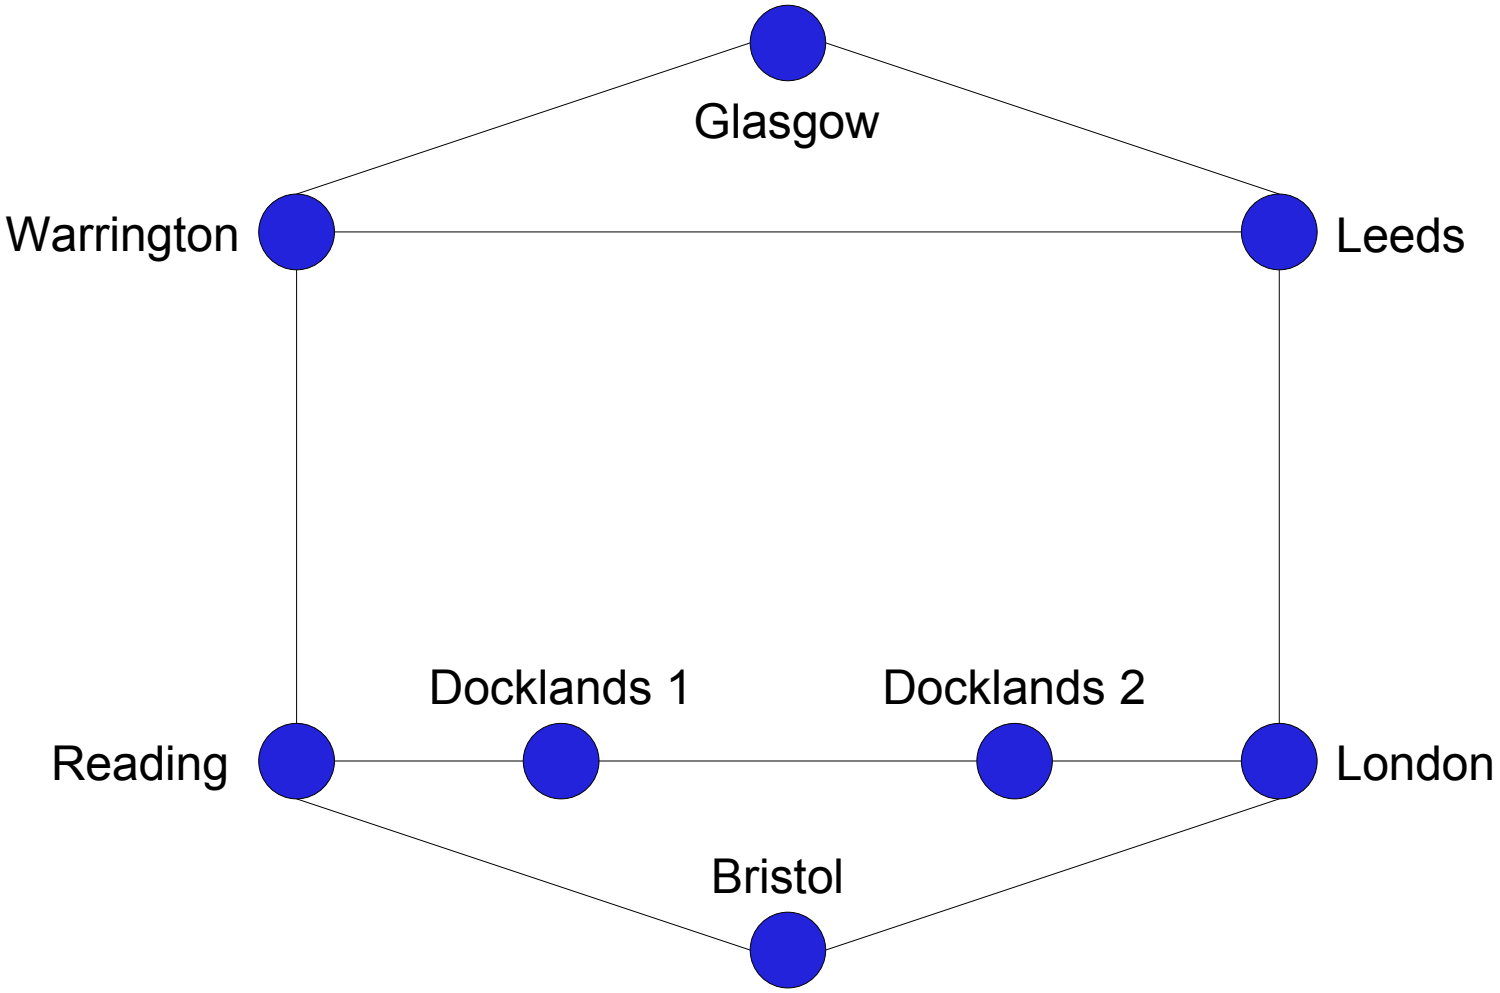
Glasgow
Warrington
Leeds
Docklands 1
Docklands 2
Reading
London
Bristol

JANET (Circuits)

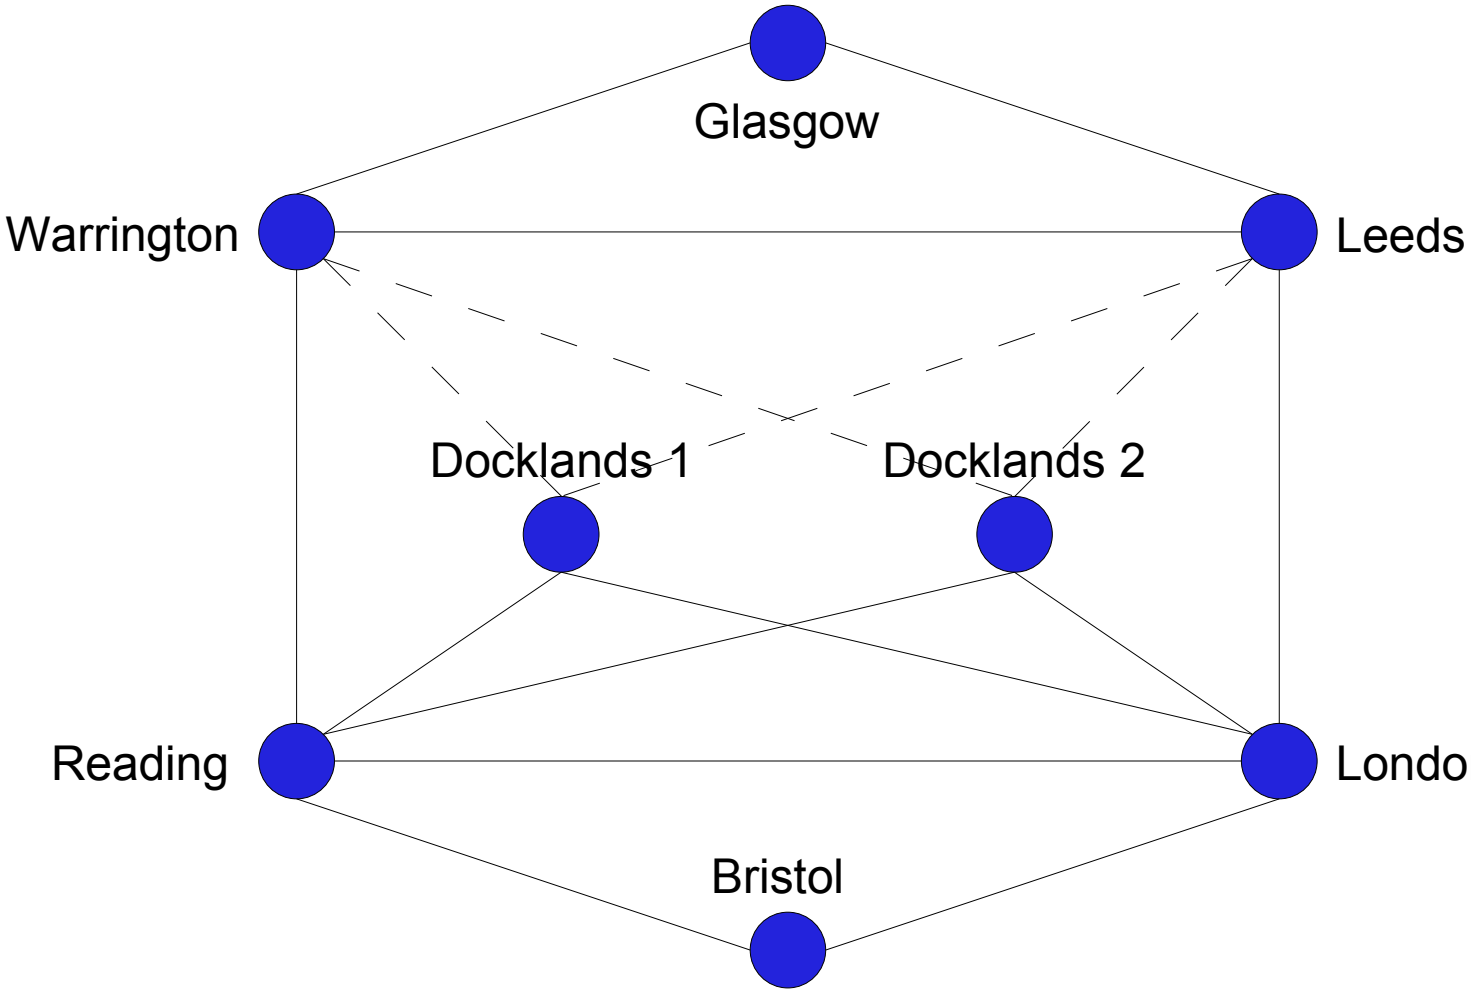

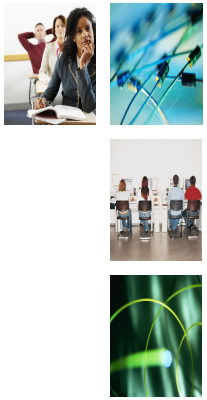

What is JANET?

- Dedicated optical equipment
  - Ciena CoreStream Regional & CN4200
- Dedicated fibre
- Optical layer managed by Verizon Business
  - We have read-only access to OnCentre
- Most circuits are SDH
  - We like alarms, diagnostics and error counters provided by SDH

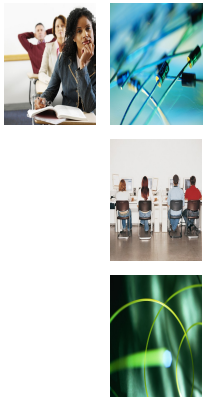

What is JANET?

- Juniper T series routers
  - Quite a few Ciscos too
- IP layer managed in-house
  - IPv4/IPv6 dual stack, of course!
    - Started with a tunnel service in 1997
    - Native since 2003
    - Unicast and multicast.

Why increase speed?

External Traffic (Gbps)

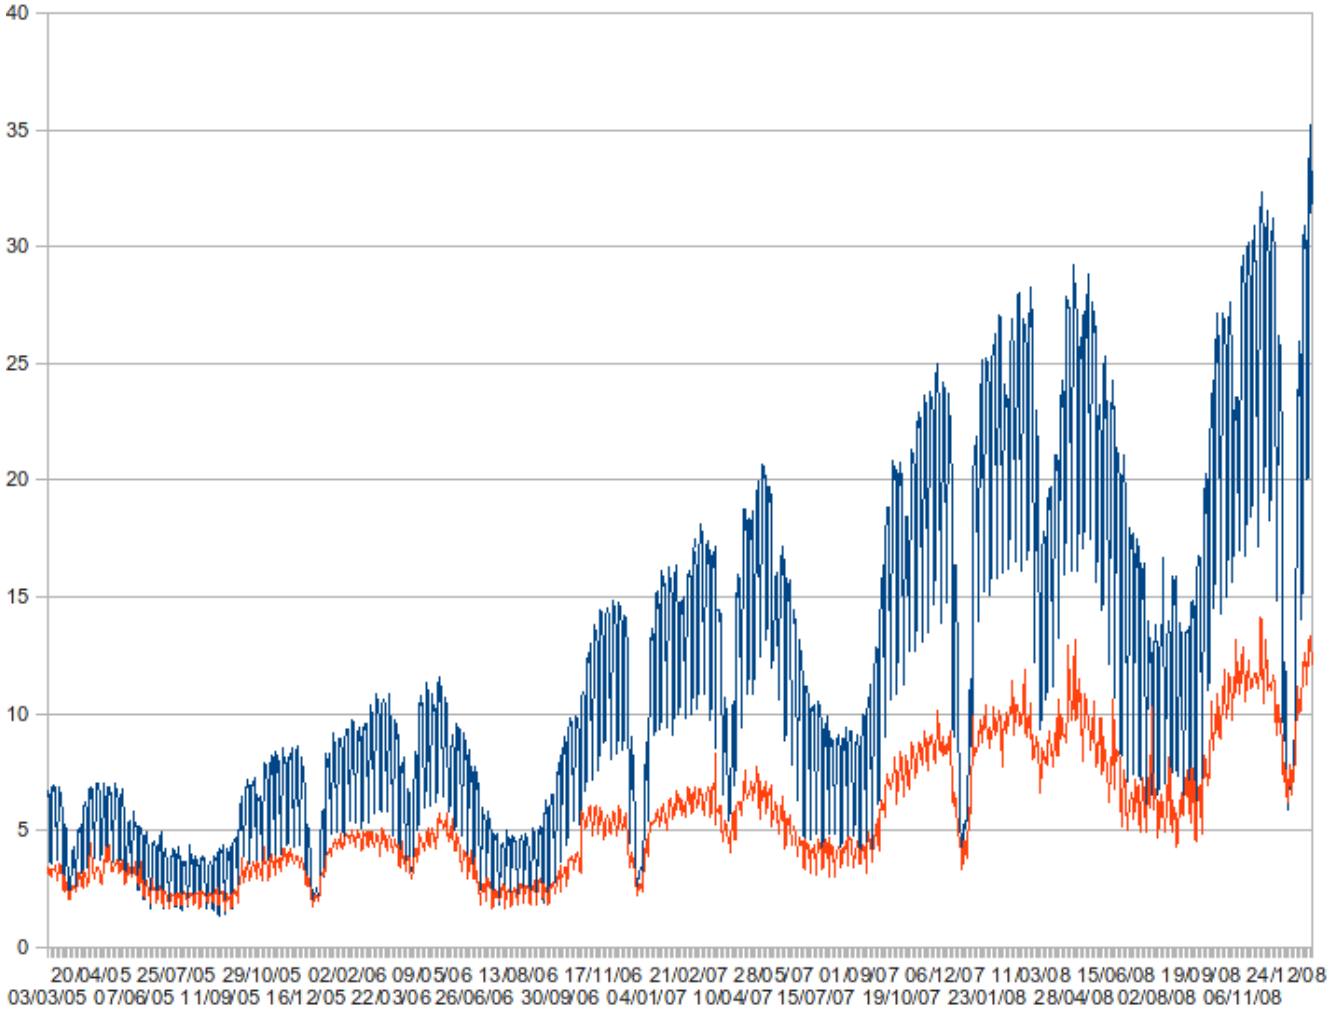

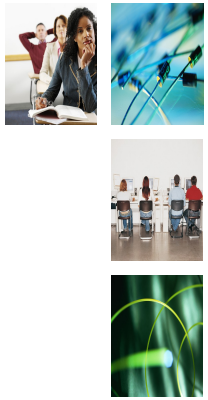

What were the options?

- Trunk multiple links
- Faster circuits

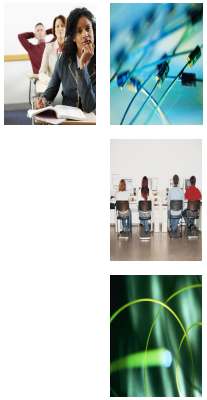

Why choose STM-256?

- Client (router) side:
  - We like simplicity
    - Prefer one link to multiple
  - Some of the scientists generate large flows
    - May overload a particular link in a bundle

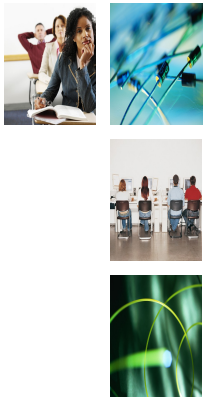

Why choose STM-256?

- Line (optical) side:
  - Don't want to squander wavelengths that researchers may need
- As an R&E network, we're expected to be on the 'leading edge.'
  - Within reasonable constraints
  - Whilst still keeping a stable network, of course.

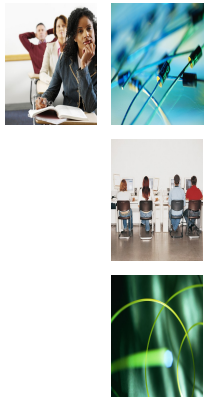

Equipment specification

- Ciena cards made by Stratalight
- First generation card
  - PMD tolerance: 2.1ps DGD
  - Duobinary encoding
- Second generation card
  - PMD tolerance: 2.5ps DGD
    - 8ps with compensation
  - DPSK encoding




Fibre characteristics

		A to B	B to A
	Length (km)	DGD (ps)	DGD (ps)
London – Telehouse	25.10	1.44	0.47
London – Telecity	28.30	1.45	0.49
Reading – Telehouse	116.50	0.75	1.64
Reading – Telecity	113.30	0.74	1.64
London – Leeds	379.68	15.46	13.66
Reading – Warrington	322.85	10.01	11.52
Warrington – Leeds	107.00	0.15	1.31
London – Bristol	252.94	2.03	1.89
Bristol – Reading	150.00	1.10	0.84

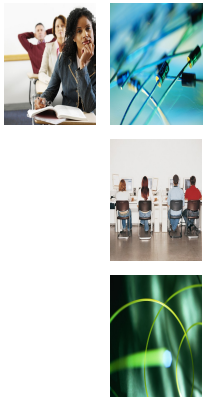

Step 1: Trial

- Split into trial and two phases.
  - Trial: Telehouse to London

Trial

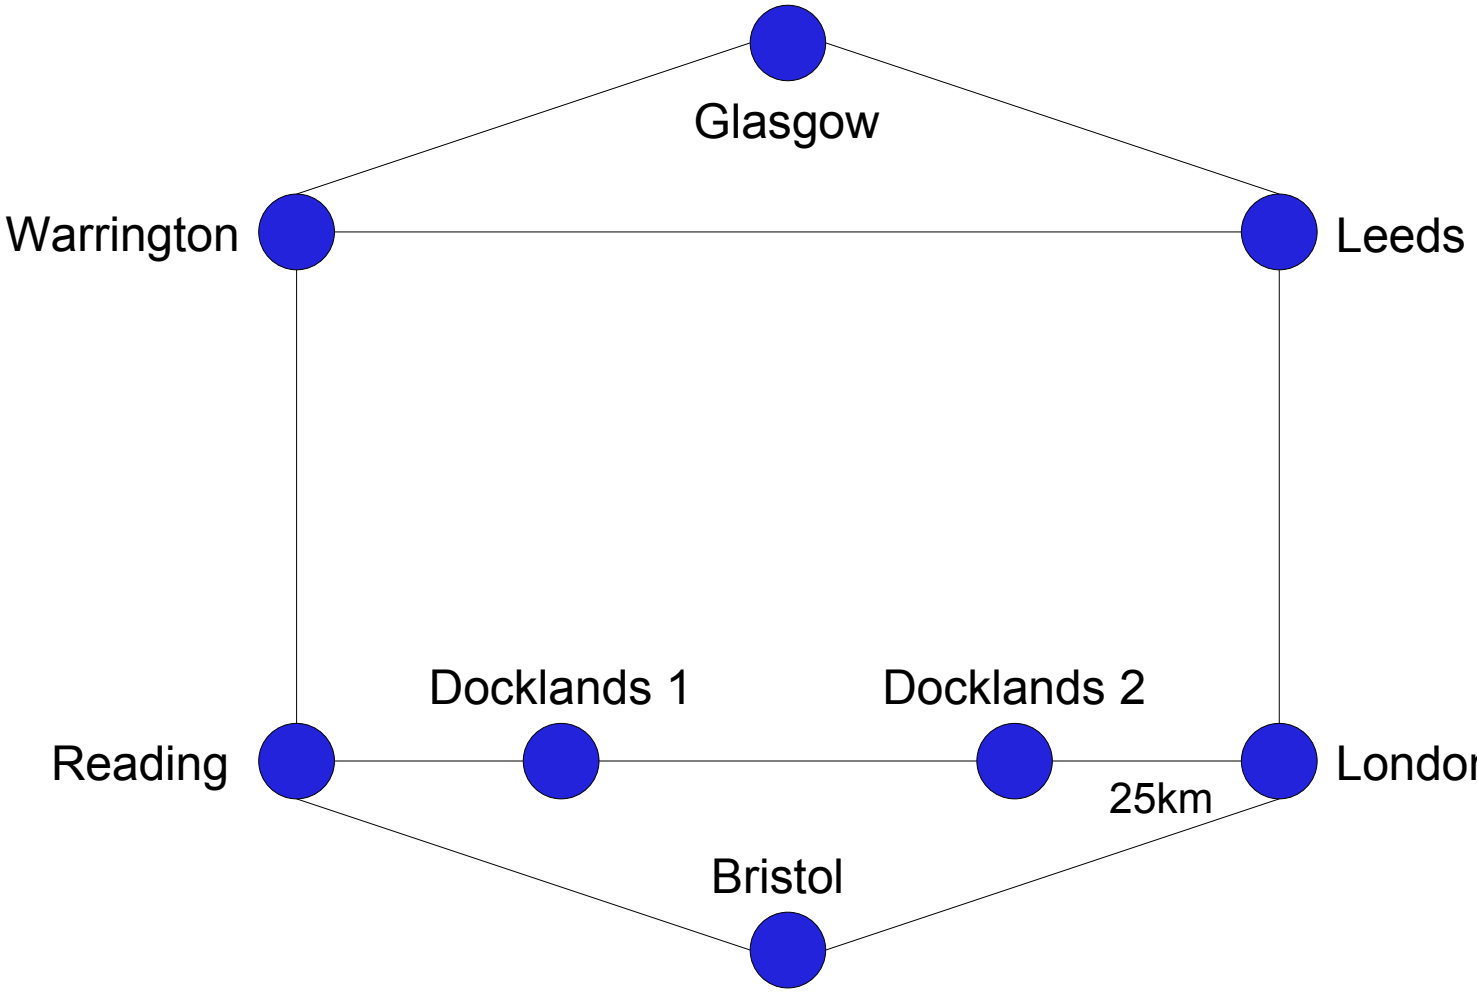
Glasgow
Warrington
Leeds
Docklands 1
Docklands 2
Reading
London
25km
Bristol

Trial

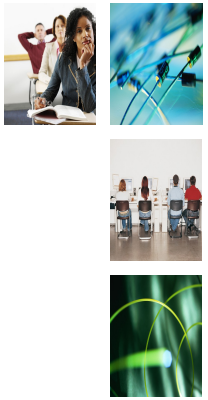

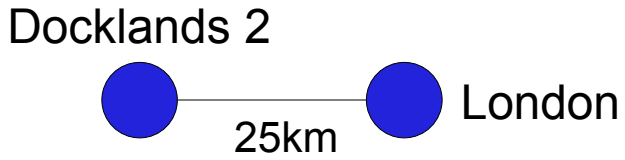

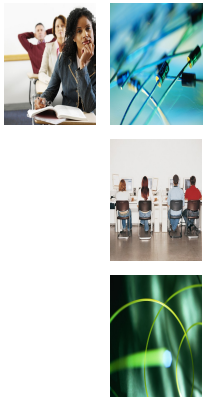

Trial: June 2007

- Started with a separate fibre in same bundle as primary route (~25km)
- Ciena 1st generation cards
- Loaned Juniper interfaces
- Logistics more complicated than technology.

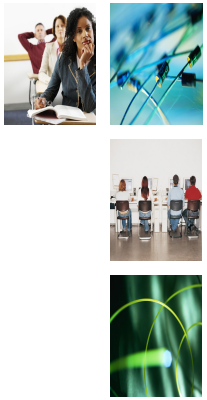

Trial

- One month later, wavelength multiplexed onto backbone fibre
- No problems.




Step 2: Production

- Phase 1: Docklands (Telehouse and Telecity) into London and Reading
  - All could be done with the Ciena equipment
- Phase 2: London and Reading north to Leeds and Warrington
  - Needed to find an alternative solution

Production

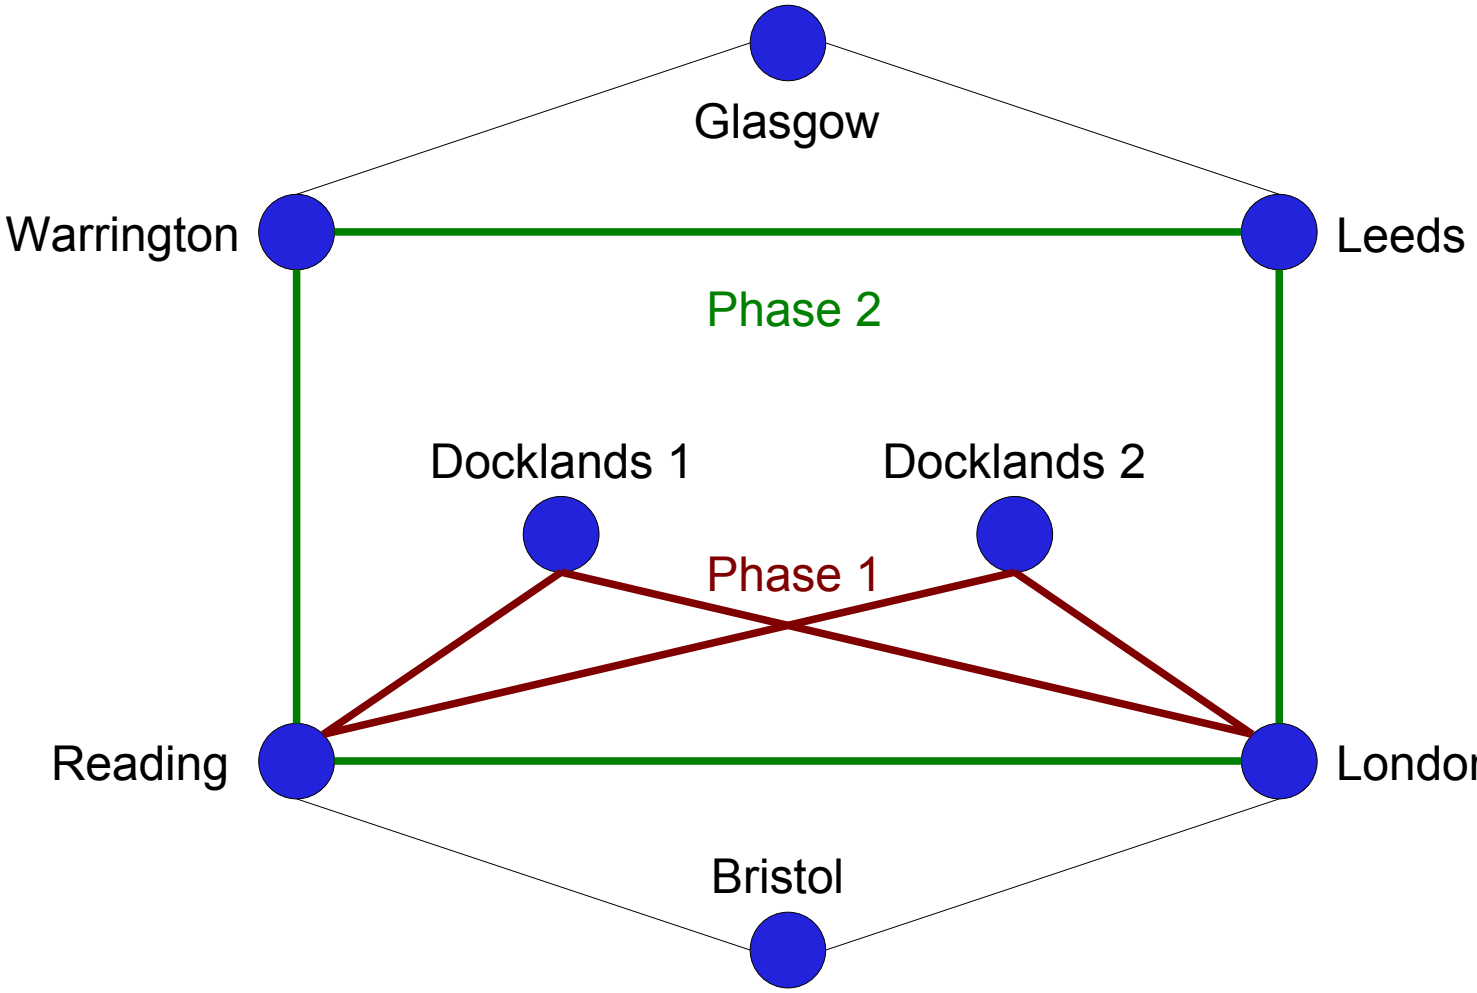
Glasgow
Warrington
Leeds
Phase 2
Docklands 1
Docklands 2
Phase 1
Reading
London
Bristol




Phase 1: June 2008

- 2nd Generation Ciena cards
  - Became available 1Q08
- Required software upgrades to existing Ciena shelves to work with new management software
  - Not always so smooth
- September 2008

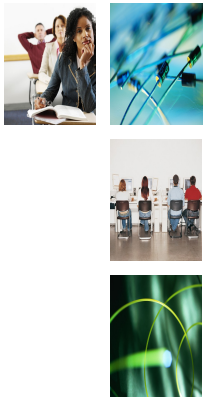

Phase 2: Nov 2008

- Outside specification for Ciena cards
  - Even with external dispersion compensation
- Alternative solution: Nortel
  - Uses Polarisation Multiplexing Differential Quadrature Phase Shift Keying (POLMUX-DQPSK)
- Carried as 'alien wavelength' on CoreStreams

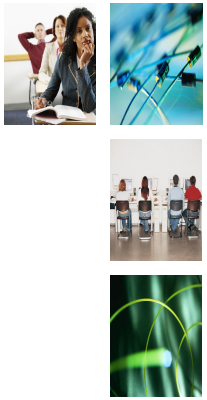

POLMUX-QPSK

- Quadrature Phase Shift Keying
  - Allows transmission of 2 bits/symbol
  - ~1 bit/s/Hz of spectrum
- Polarisation Multiplexing
  - Two QPSK signals sent with orthogonal polarisations
  - ~2 bit/s/Hz of spectrum
- Plus fancy Forward Error Correction
  - 112Gbit/s within 50GHz ITU grid channels

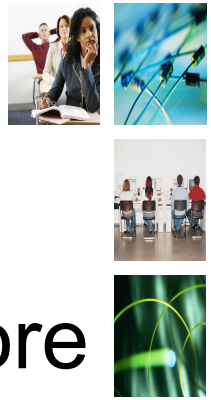

Phase 2

- Additional optical equipment at four core points of presence
  - Expensive
    - “Not our problem,” we'd had 40Gbit/s in our contract.
  - New management systems
  - New management procedures
- December 2008

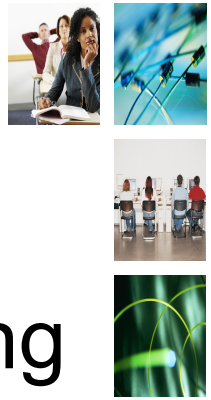

Problems

- Card failure shortly after commissioning
  - Backed off to 4-hour spares replacement
  - ...only no card in stock
    - ...or in the country
      - ...or in the continent
  - Had to be sourced from manufacturing
  - Best part of two weeks

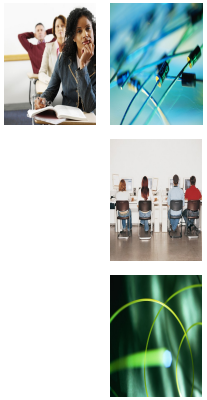

Problems

- Still had STM-64 circuit to back it up
  - Just in case there were teething problems
  - Still had sufficient capacity
- Took several more weeks until sparing was completed
  - Followed up contractually

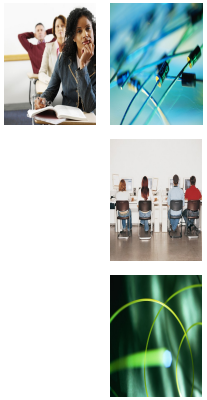

Problems

- Mainly logistical rather than technical
  - Weekly phone conferences for six months

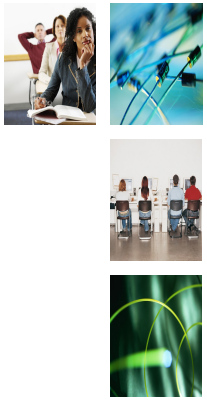

Routing kit

- So far, only talked about circuits
- What about the routers at either end?

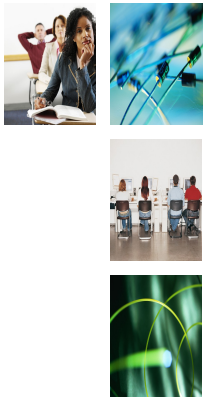

Routing kit

- Started off with Juniper T-640s
  - Eight chassis slots
  - 40Gbit/s per slot
- Some routers would have four STM-256 circuits
  - Half the chassis just for core links
- Upgrade to T-1600s!

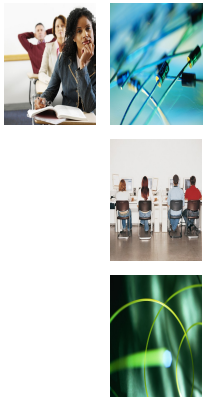

Routing kit

- T-640 to T-1600 is an “in-service” upgrade
  - Needs an extra DC feed per PEM
  - Telehousing providers worried about power and cooling
  - Power rating increased from 6.5kW to 9.1kW
- Lots of talking to convince providers we wouldn't be running at maximum power draw
  - Actual draw ~3kW

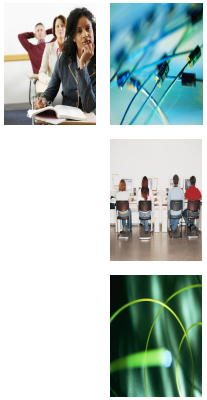

Routing kit

- First time links were brought up, interface card needed to be taken offline and brought online again to work
  - Not happened since
    - Even after circuit outage
    - Until two weeks ago after some reprovisioning work...
- Otherwise, not much to report
  - Generally it “just worked.”

What do we have now?

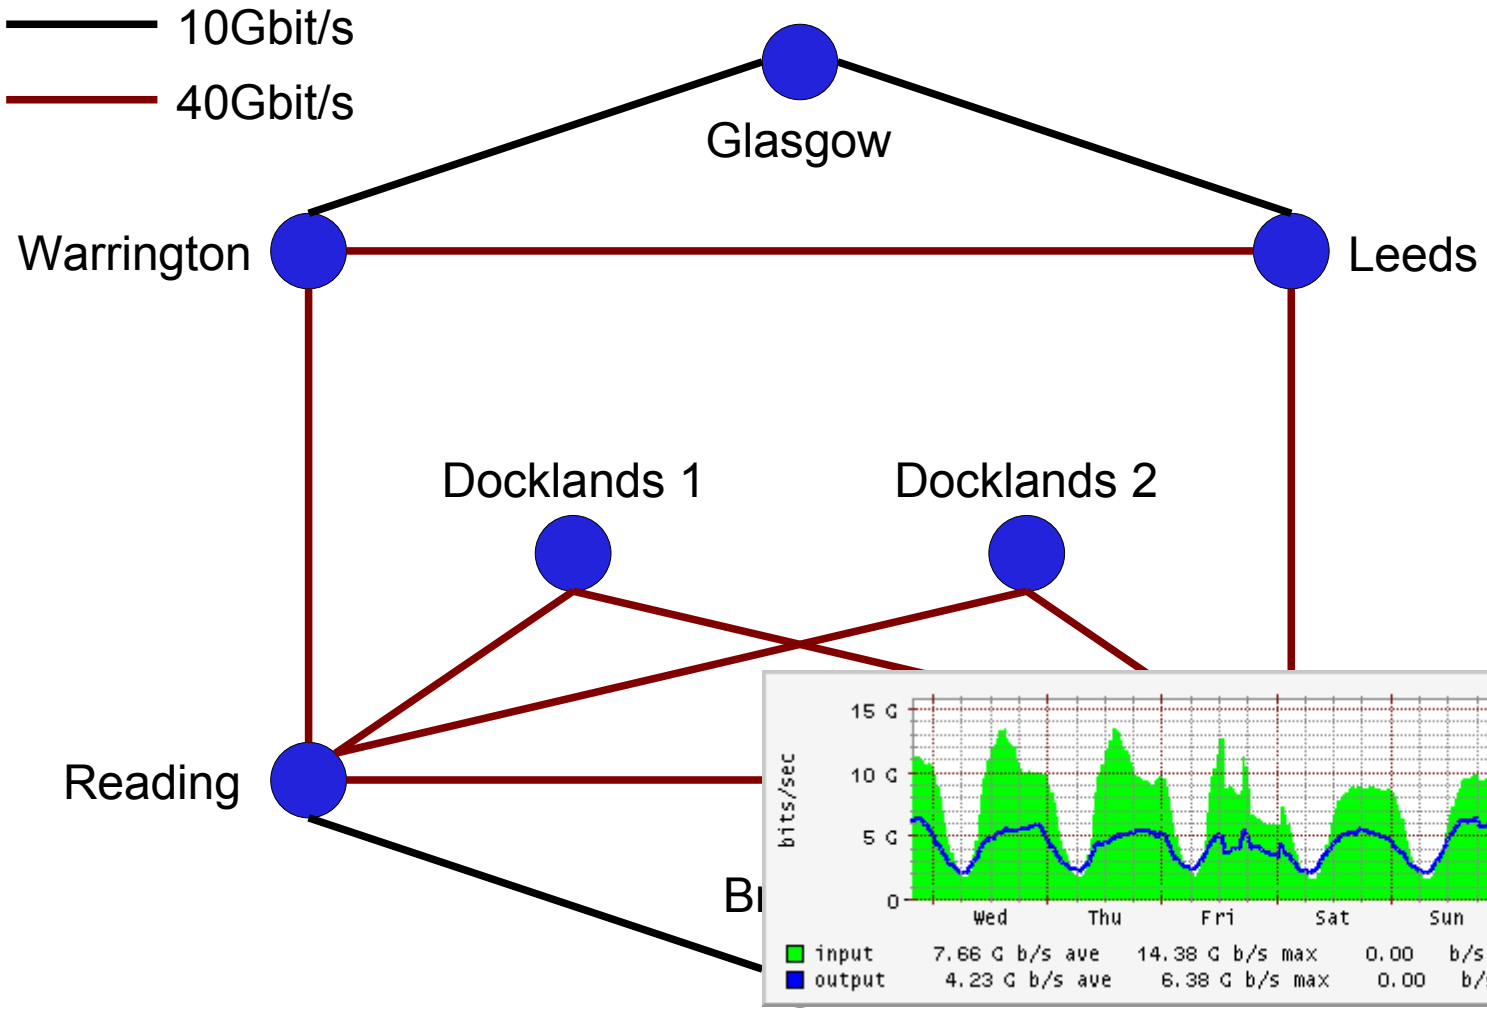

US Inauguration

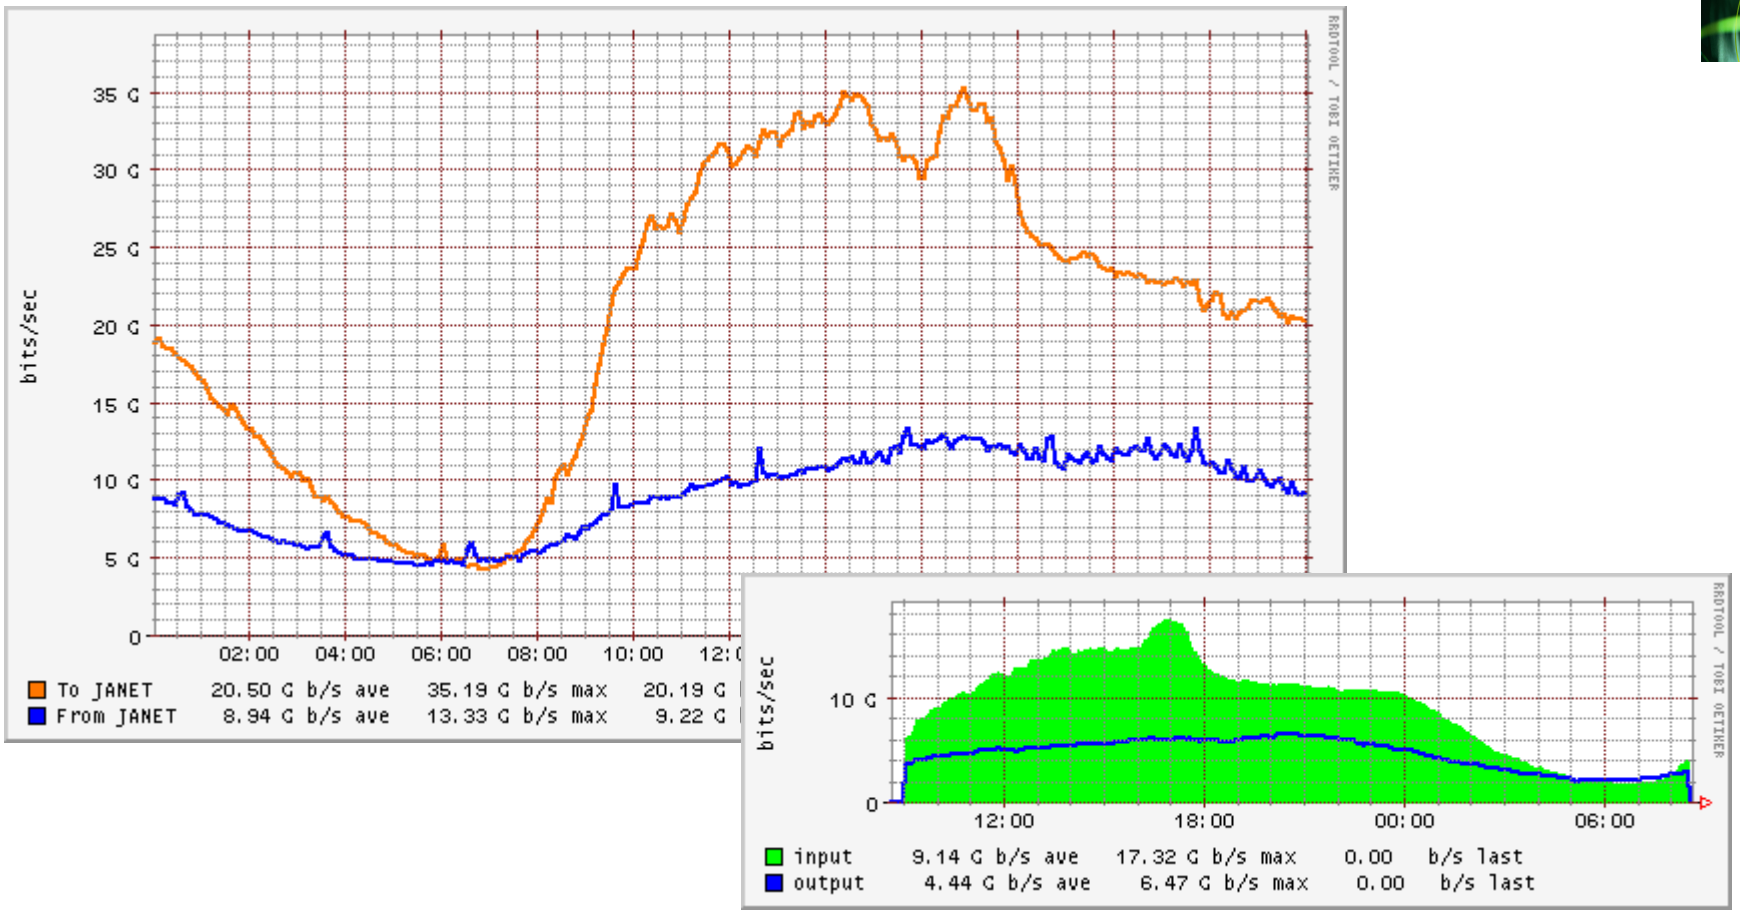

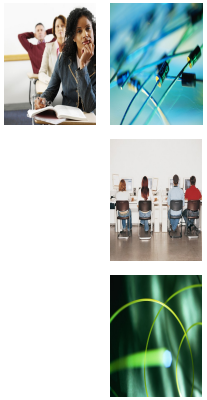

Should you go 40G?

- I can't tell you that, but...
- 40G has very strict fibre requirements
  - Aim is for 100G requirements to be closer to 10G requirements
  - 'Coherent detection' will help things further

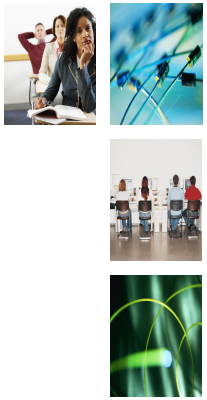

Should you go 40G?

- 40G SDH prices haven't dropped significantly
  - Some large deployments in 2008 may help
    - Maybe not, these have mainly been 'line side' deployments, multiplexing 1GE/2.5G/10G circuits onto 40G wavelength
- SDH reaching the end of its life
- 'Carrier Ethernet' the way forward

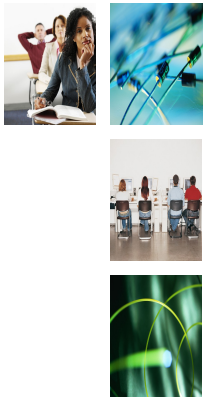

Should you go 40G?

- 100G standardisation marches on
  - We've started looking at it for technology trials soon
    - Line side available before client side
    - Considering deployment in 2011-2012
  - But... will your chassis route/switch 100G/slot?
- 40G is here now




Credits

- Verizon Business worked very hard with us to make this happen
- Ciena
- Nortel
- Alcatel-Lucent and Juniper

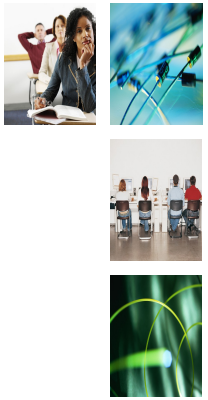

Questions?

POLMUX-DQPSK

Quadrature Phase Shift Keying

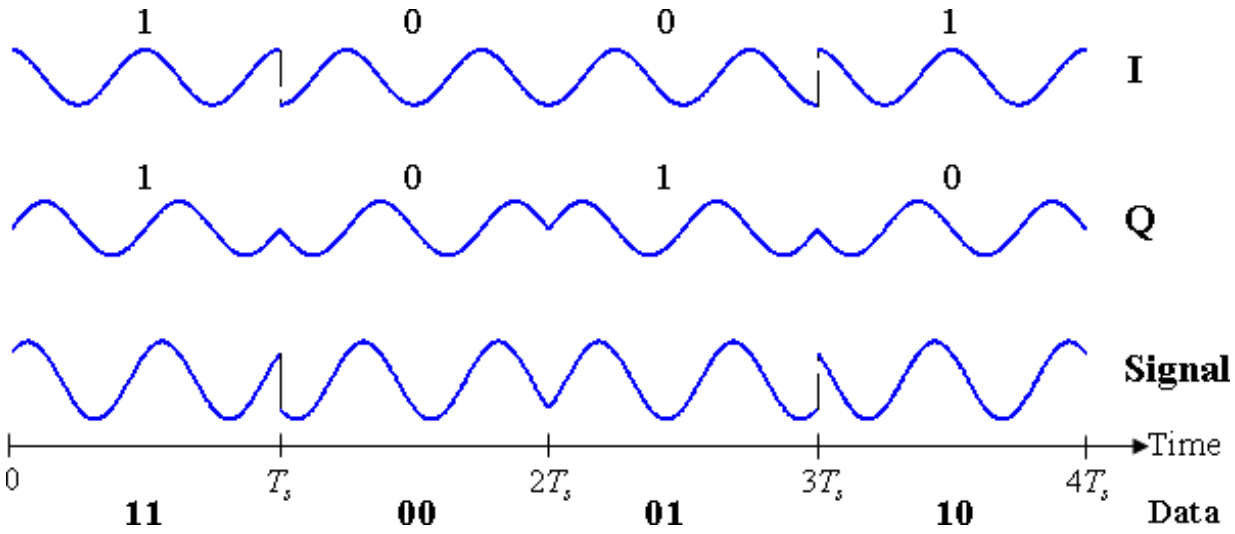

2 bits per symbol, ~ 1bit/s/Hz

Diagram by wikipedia:User:Splash

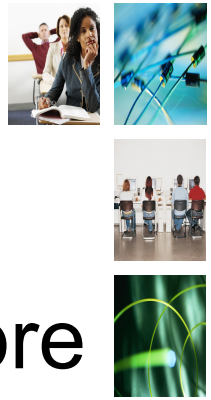

POLMUX-DQPSK

- Two DQPSK signals sent down the fibre at the same frequency with different polarisations
- 2 x 2 bits per symbol = 4 bits per symbol
  - ~2 bit/s/Hz
- Can carry 100Gbit/s signals (112Gbit/s for ODU-4) within neighbouring 50GHz ITU grid channels