




Science and
Technology
Facilities Council

Networking

Jonathan Churchill
SCD Network Architect

Jonathan.Churchill@stfc.ac.uk



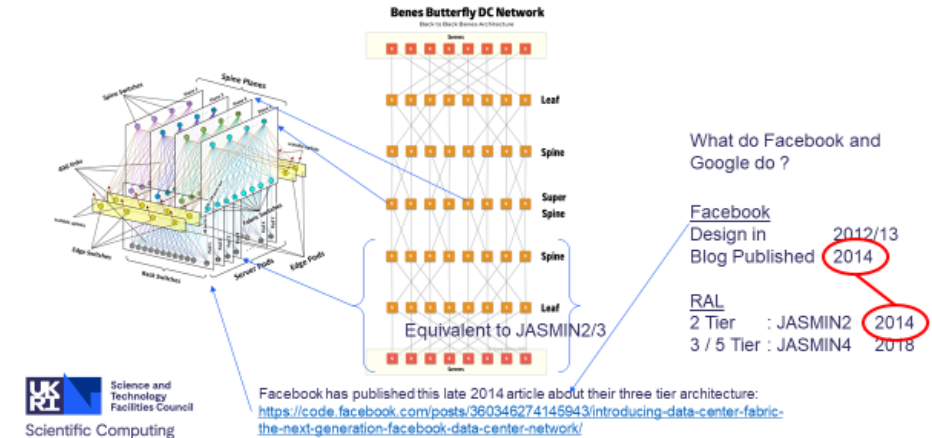
Introduction

- Network Architect ... for Scientific Computing
 - (Philip Garrad is Network Architect for STFC & RAL – Enterprise networks)
- New role for SCD since May 2021
 - Co-ordinates departmental activities and standards.
 - Outreach to: other areas of STFC/UKRI, Janet/JISC, Science (eg SKA SRCnet, LHCone, DUNE, (WLCG) etc
 - Recommends: Hardware & software choices (but projects do procurement)
 - Responsible for the core superspine & exit/firewall networking.

Current Hardware

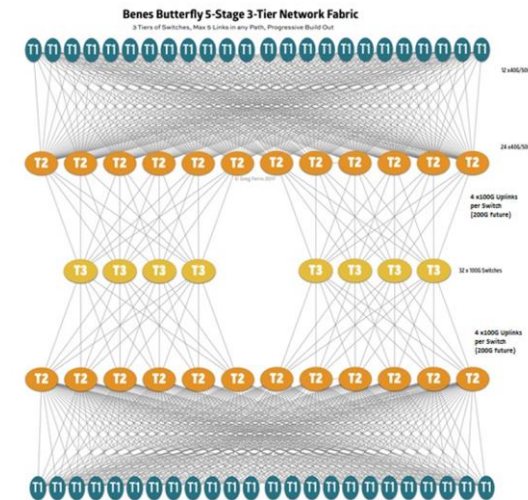
- 10/40/25/50/100G Switch/Routers
 - NVIDIA/Mellanox, Dell Force10
 - Choice mostly driven by procurement
 - 270+ Mellanox (SN2100/2010, SN2410C/B, SN2700, SN4600C)
 - 50+ Dell 'S' & 'Z' (S4048, S4810, S5248, Z9000/9100/9264)
 - Cumulus and Dell OS10 NOS's
 - Converting Onyx to Cumulus for now ... but NVUE ☹ ☹ ☹
 - Hence considering SONiC +/- OS10
 - Ansible (and other linux) config management.
- 10/100/1000Mb "Management" Switches
 - 188 Dell S3048's (replacing Old Nortel/Avaya) <- The 'old' workhorse.
 - Some Netgear dumb unmanaged.
 - HP Aruba 6300M (or Dell N3248TE) (or Mellanox 'AS') <- New purchases.

A Data Centre "Fabric" Network



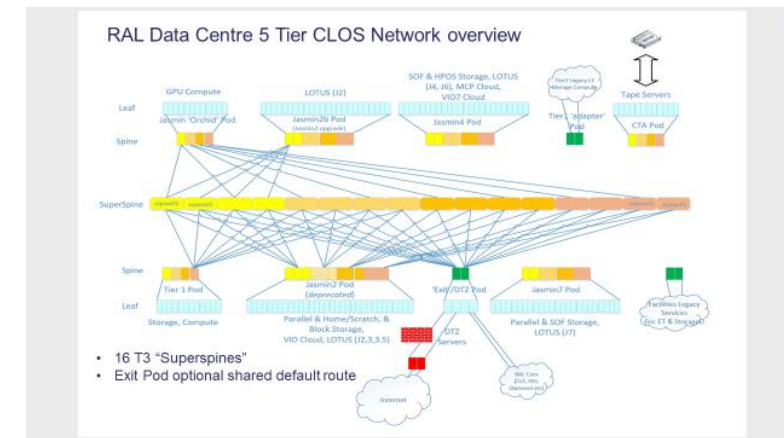
Future Procurements – Server/Pod

- Tier1, STFC Cloud, JASMIN, SCARF etc
 - New purchases driven by price/performance procurement.
 - Assume Mellanox (Cumulus) or Dell Force10 (OS10) this year.
 - But all options are possible that support :
BGP/OSPFv3/MLAG/IPv6/ECN/(RoCEv2)/sFlow/OOB mgmt
 - Mix of 25G and 100G Server connections.
 - Not 200G or 400G this year.
 - Port to port Low latency (low 100's nS)
- Key issues
 - Quick lead times on DAC/AOC cables in small quantities.
 - Availability of long AOC (10/15/20/30M) 100G in “sensible” lead times.



Future Procurements - Superspines

- Superspines upgrade 400G + port capacity increase.
 - Initial Support for:
 - New Pods @ 100G
 - 400G between data centres.
 - Then FY24/25-27/28:
 - New Pods @ 400G - Spine to Superspines 10's 400G AOCs
 - Campus data centre ring @ 400G (100's 2km SM optics)
 - Target is 16x 64 port 100/400G routers.
 - Backwards compatible with
 - Existing AOC 100G cables (QSFP28)
 - Existing 2km SM LC Optics (QSFP28)
 - Plus 400G SM transceivers and 400G AOC's





Science and
Technology
Facilities Council

Questions?



Science and
Technology
Facilities Council

Backup

