Single-Site to Cross-Site RoCE Network

Sheffield Co-design Day 2026

Salih MSA
Senior Research Infrastructure Engineer (Cloud, Networking)

Architecture, Infrastructure and Systems
Advanced Research Computing
IT Services
University of Birmingham







2026-09-16

Overview

- Background
- Design
- Implementation
- Lessons learned

Background

Data centres

- On-prem services
- 99% of hardware located in main research data centre (dubbed *'PGDC'*)



Networking

- Ethernet:
  - 100Gbe
    - 1Gbe for OOBM
    - 25Gbe/10Gbe for hosts
    - 100Gbe for ISLs
  - Two-tier spine leaf topology composed of NVIDIA Spectrum switches
  - Each leaf has 1x uplink to each of the 2x spines (2x uplinks total)
  - Internal traffic: BGP underlay, EVPN-VxLAN overlay, decentralised & symmetric routing on leaf switches
  - External traffic: bridged link to FortiNet perimeter firewalls from exit switches

Networking

- InfiniBand:
  - HDR
    - HDR100 for CPU compute
    - HDR for GPU compute + storage servers + infrastructure servers + ISLs
  - Two-tier fat tree topology composed of NVIDIA Quantum switches
  - Each edge has 2x uplinks to each of the 2x cores (4x uplinks total) (blocking)
  - Internal traffic only: non-routed (effectively) flat network

Storage

- IBM Tiger Shark File System / General Parallel File System / Spectrum Scale / Storage Scale
- Lenovo Distributed Storage Solution (DSS-G) (GPFS GNR)
- Data transfer between clients and storage over InfiniBand RDMA
- Used and accessed in multiple ways



Disaster recovery concerns

- Senior ARC management went on an ‘away day’
- Most notable take-away was to have another **active** copy of storage outside *PGDC*

Design

Requirements

- Connect storage in either site, to clients in either site
- Could withstand either of the sites being down
- Performant

Action plan

- Improve upon existing Ethernet network
- Add new spines to *ERDC*
- Cross-connect existing leaf switches to both sets of spines
- 200Gbe capable switches
- Lower blocking ratio for high-throughput storage NICs
- RoCE support

Implementation

Spines

- Added 2x 200Gbe 64 port switches in *PGDC* and migrated connections from old 2x 100Gbe 32 port switches
- Added 2x 200Gbe 64 port switches in *ERDC*
- Infrastructure leaf switches connected to all spines in both *PGDC* and *ERDC*
- Compute & OOBM leaf switches connected only to spines in local data centre
- Infrastructure leaf switches dedicated for high throughput with double the number of uplinks

RoCE

- Most configuration focussed on QoS changes to support lossless RoCE
- Generated multiple switch configurations and coordinated disruptive changes
  - Platform-specific QoS settings
  - Standard QoS settings
  - LLDP DCBX TLVs
- Configured NICs and RDMA stacks on servers
  - Enabling PFC
  - Trusting DSCP
  - Configuring the ‘traffic class’

Lessons learned

Automating RoCE configurations on switches

- RoCE configurations managed via NVUE
- Missed some configuration files out
- Not sure what some files do
- Comparatively poor documentation

High power support on spines

- Limited class 8 ports $\rightarrow$ Limited long range 200Gbe optics
- Restriction not mentioned in the main data sheet
- Always look at fine print of documentation!
- Long term will need to find new spine switches

High power support on spines

From NVIDIA SN4000 Switch Hardware User Manual:

Ports	Maximum High Power Support
All ports	5W
49-50, 53-54, 57-58, 61-62	6.5W

Q&A