

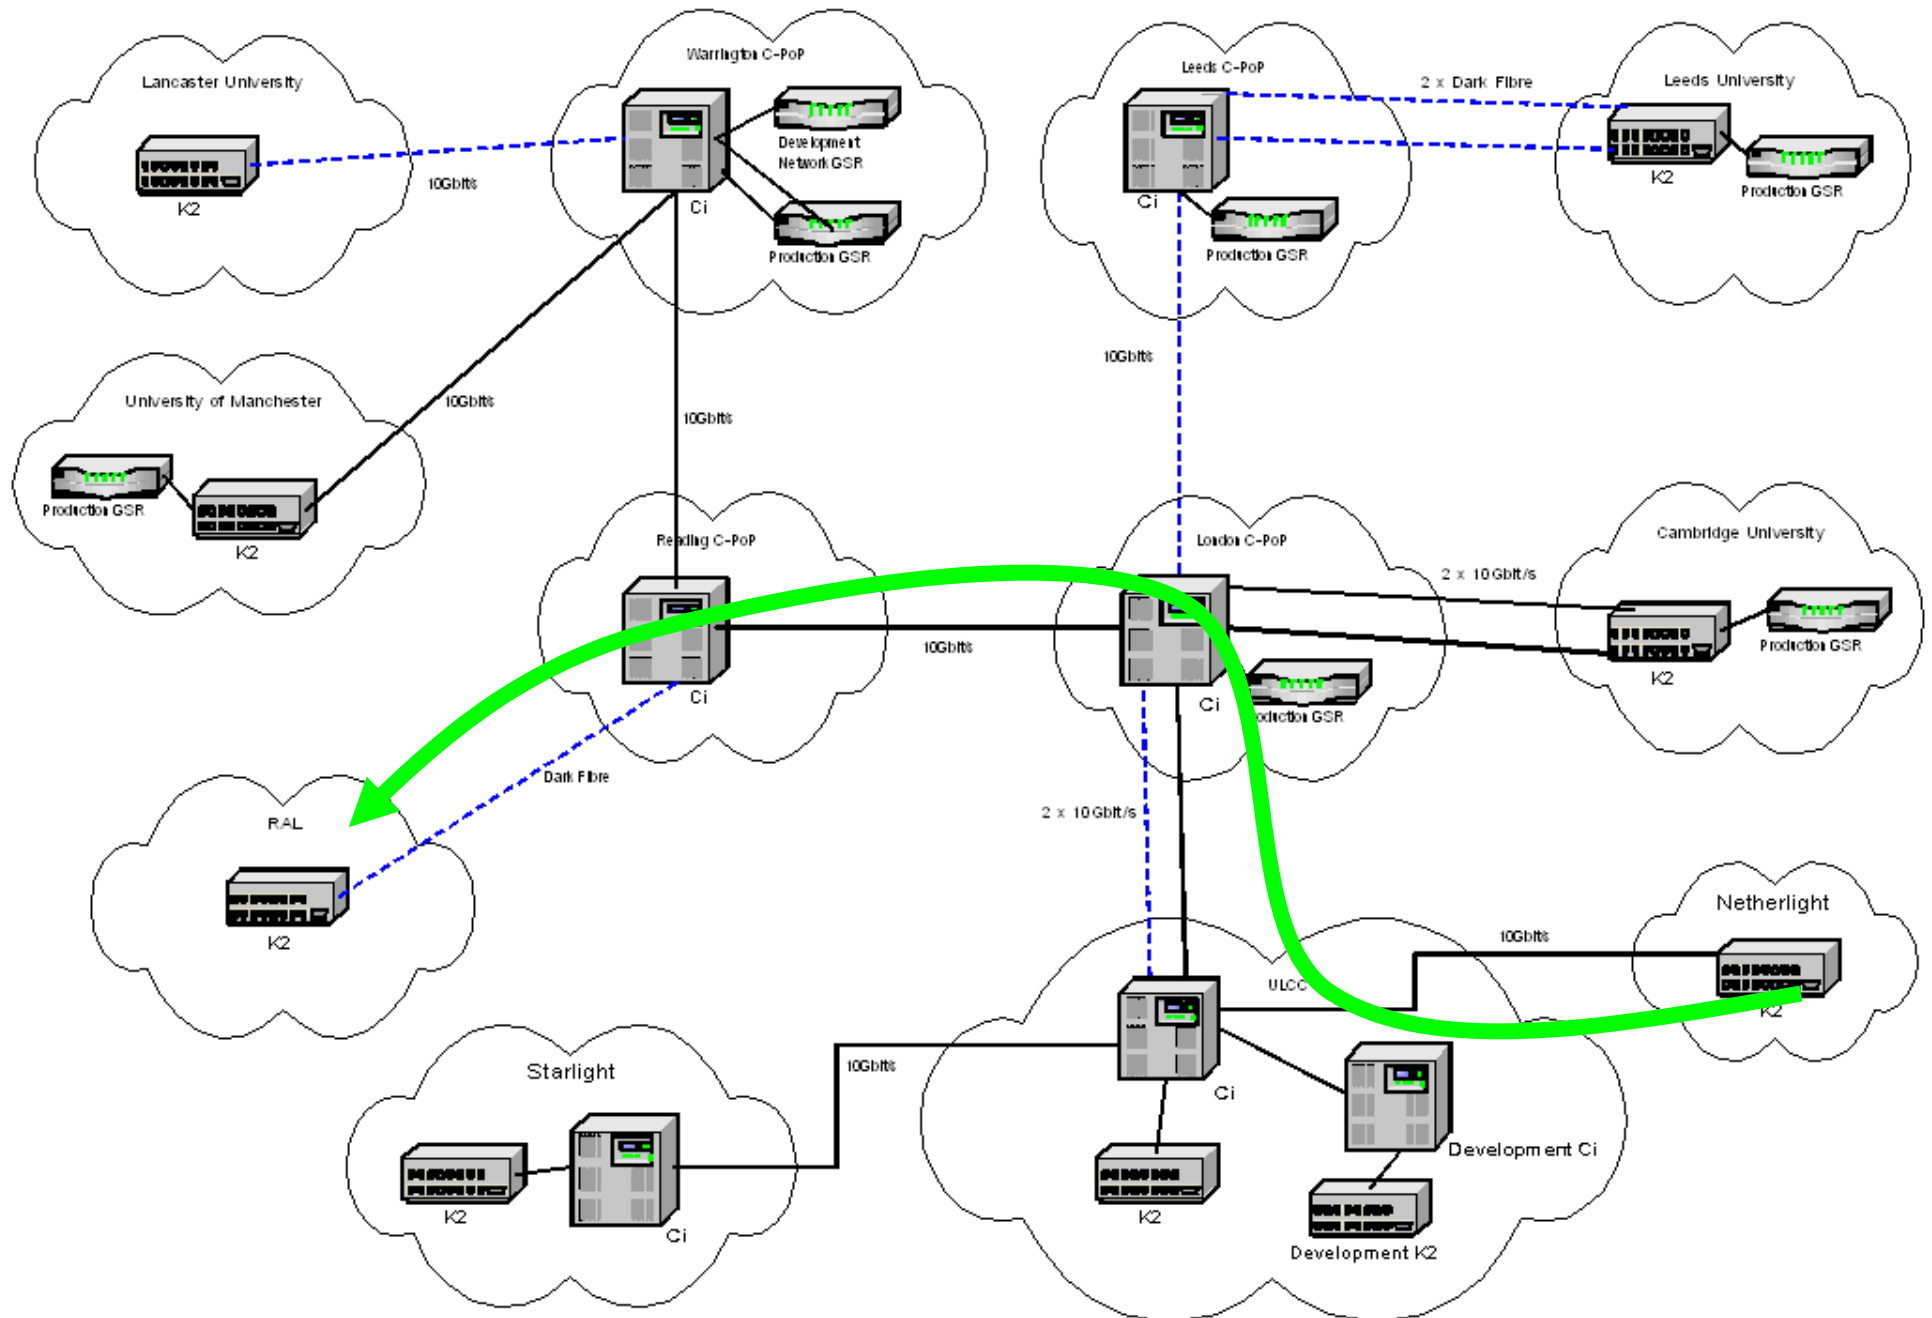


GridPP
UK Computing for Particle Physics

RAL Plans for SC2

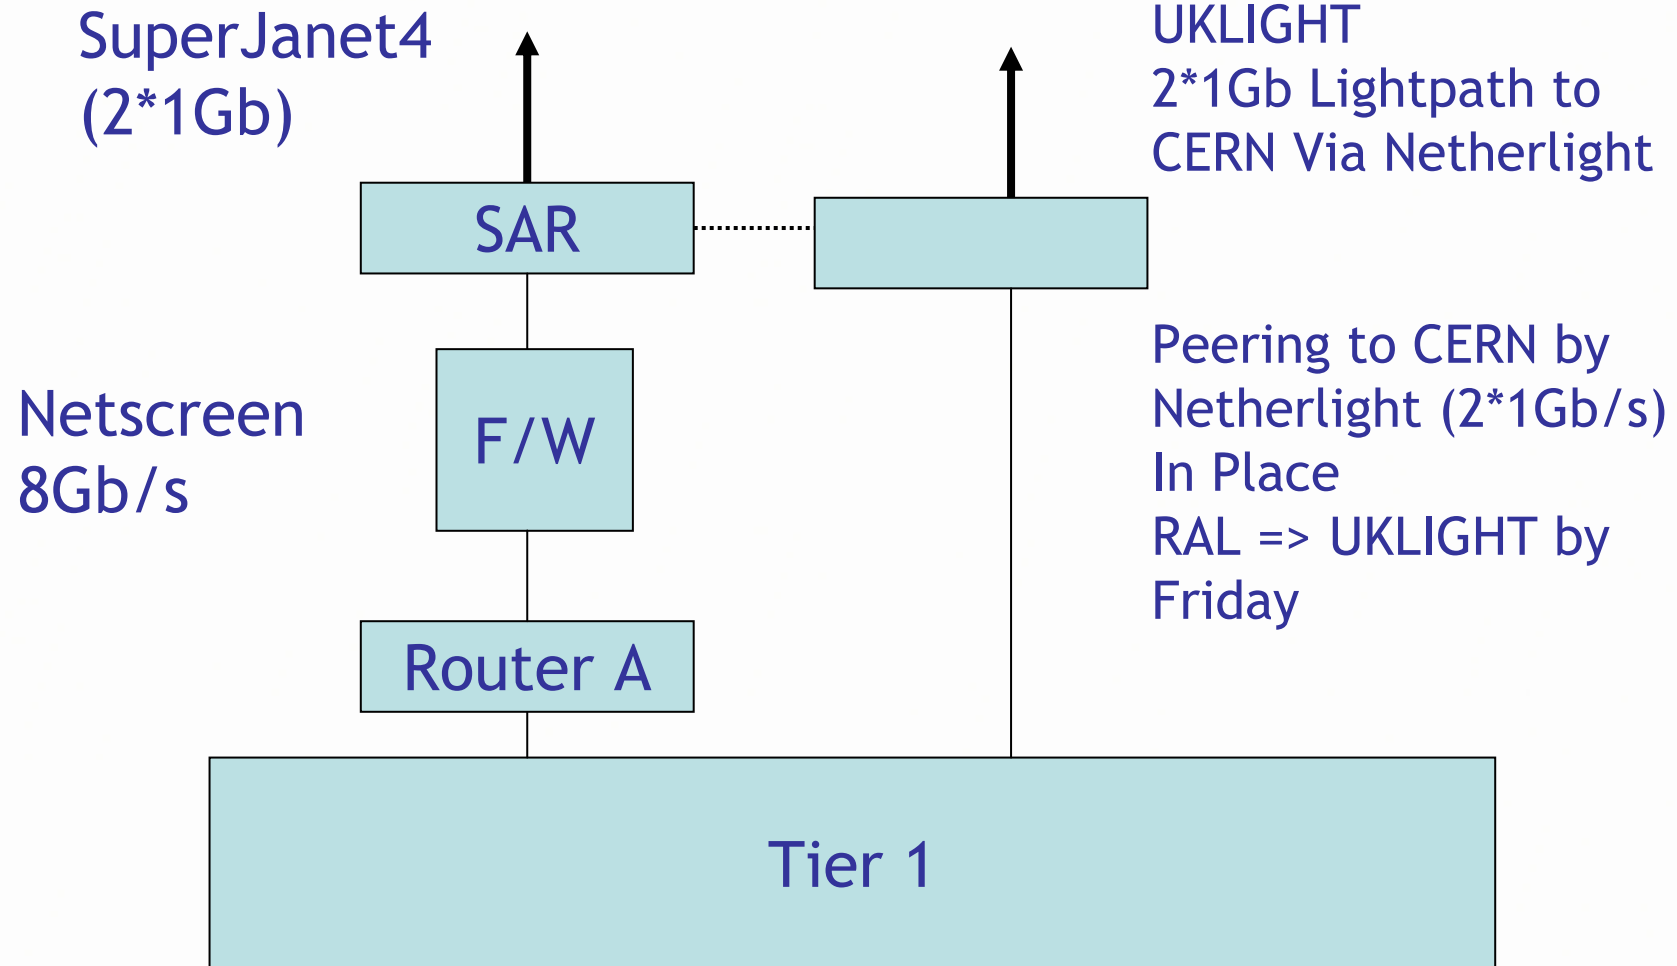
Andrew Sansum
Service Challenge Meeting
24 February 2005

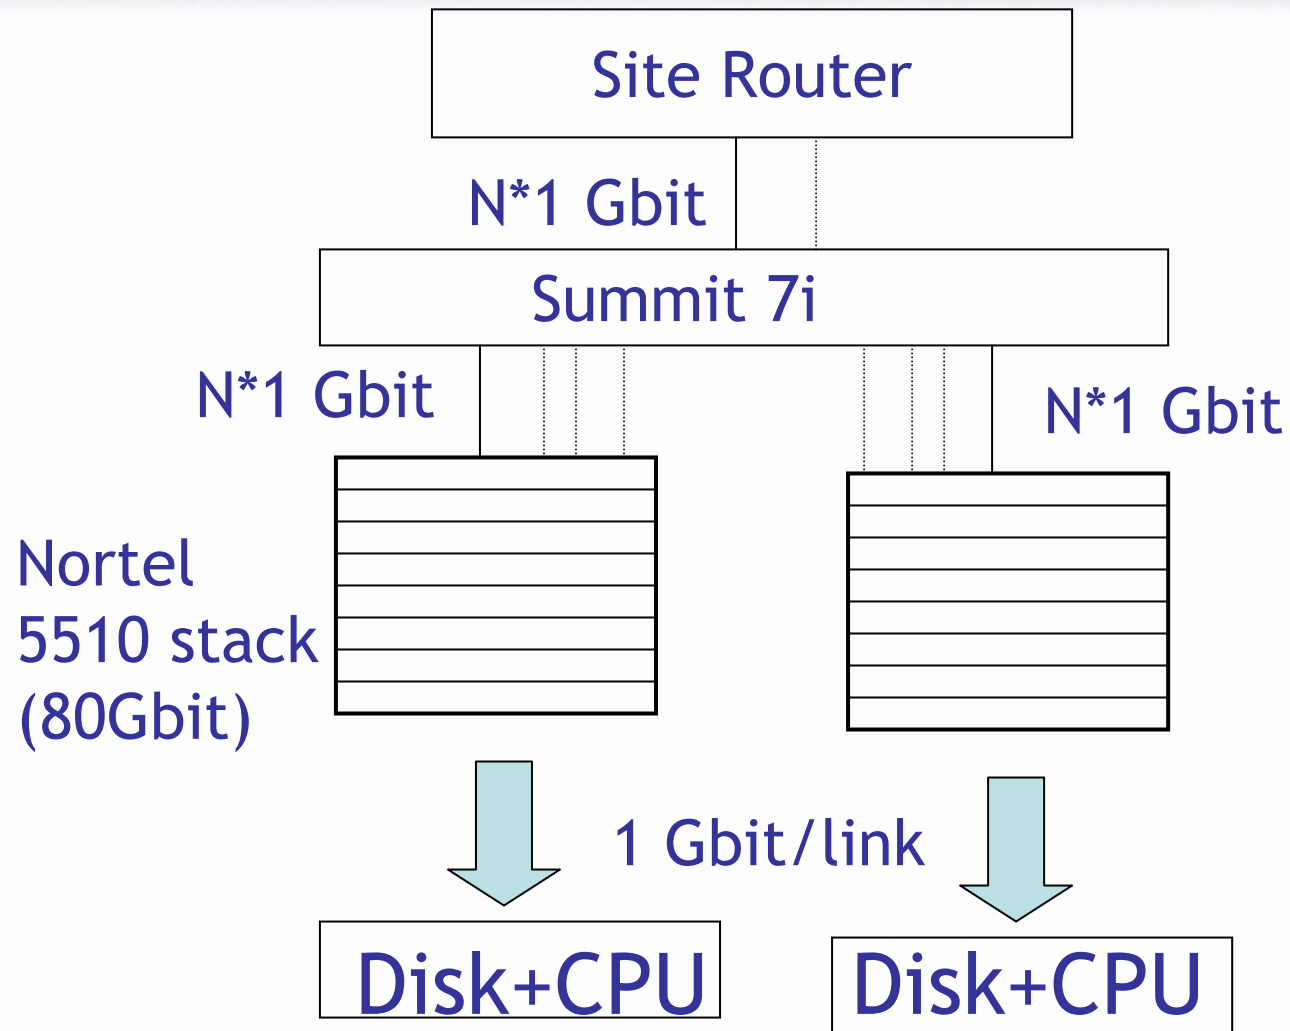
UKlight & the extended development network - Phase 1 & Phase 2





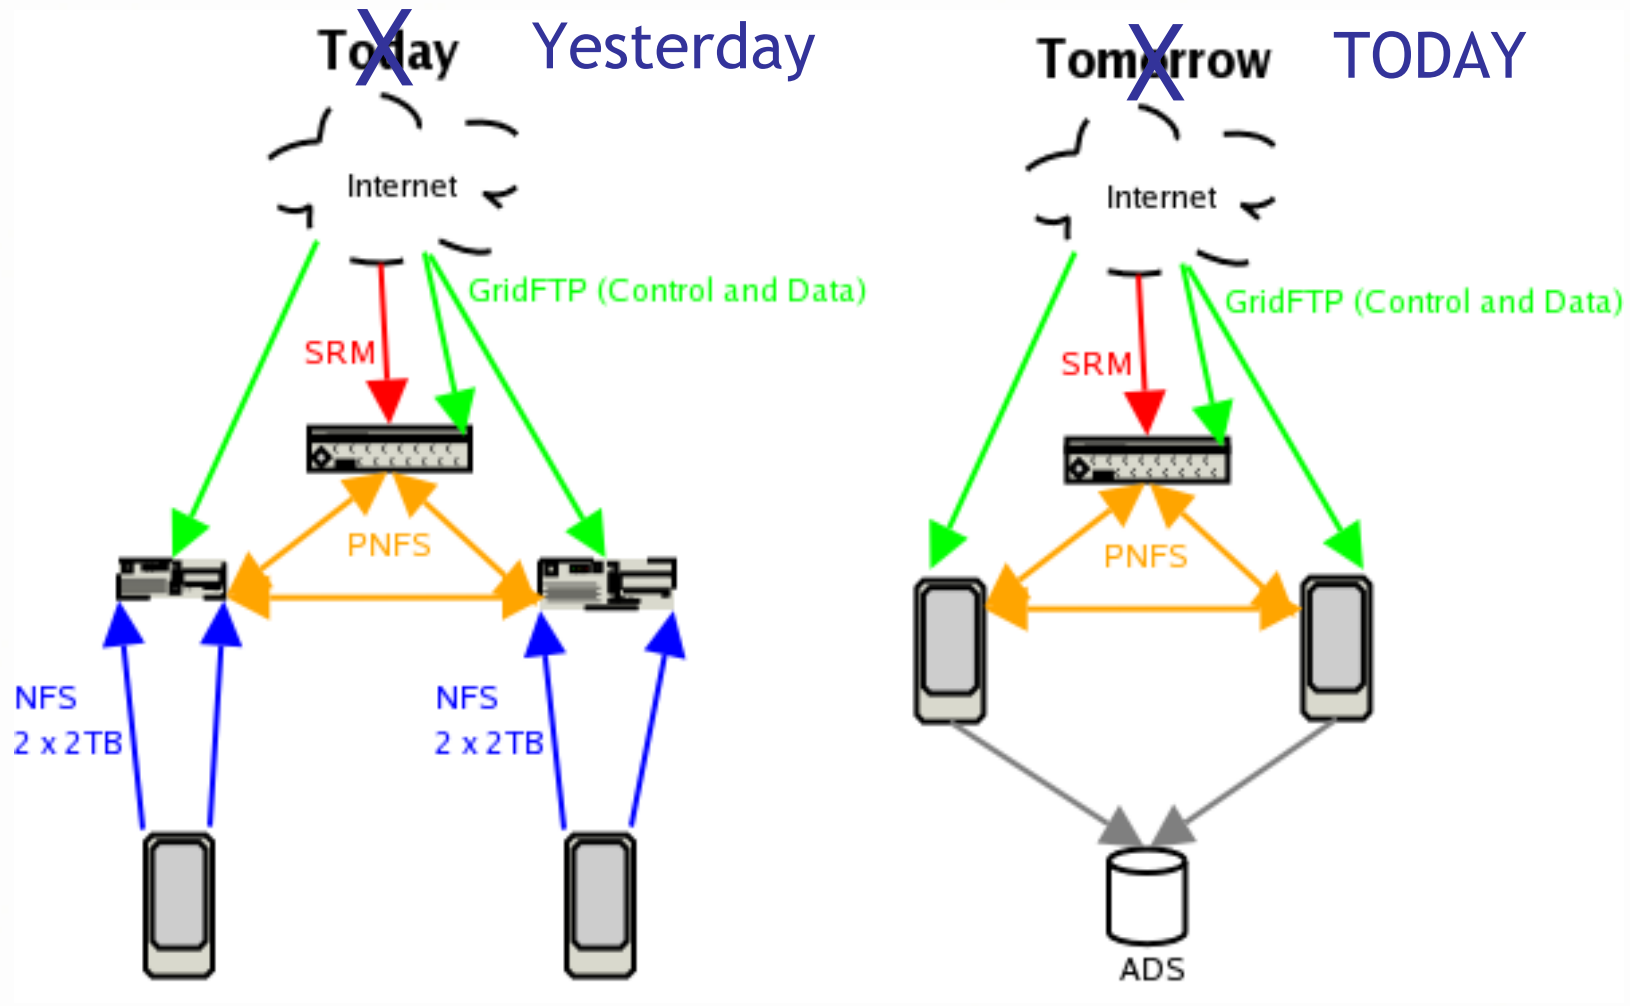
Site Network







dCache Evolution





SC2 Disk Server

2 dCache
Disk pools

2*1Gbit uplink (1 in use)

Server: dual
2.8Ghz Xeon

7501 chipset

U320 SCSI

1 Lun per device, ext2

EonStore

Eonstore

15+1 RAID 5
250GB SATA

120MB/s

120MB/s



dCache Versions

d-cache-core-1.5.2-33

d-cache-lcg-4.0.0-1

d-cache-opt-1.5.3-15

d-cache-client-1.0-34

vdt_globus_info_server-VDT1.1.14rh7gcc3-4

vdt_globus_essentials-VDT1.1.14rh7gcc3-4.lcg1

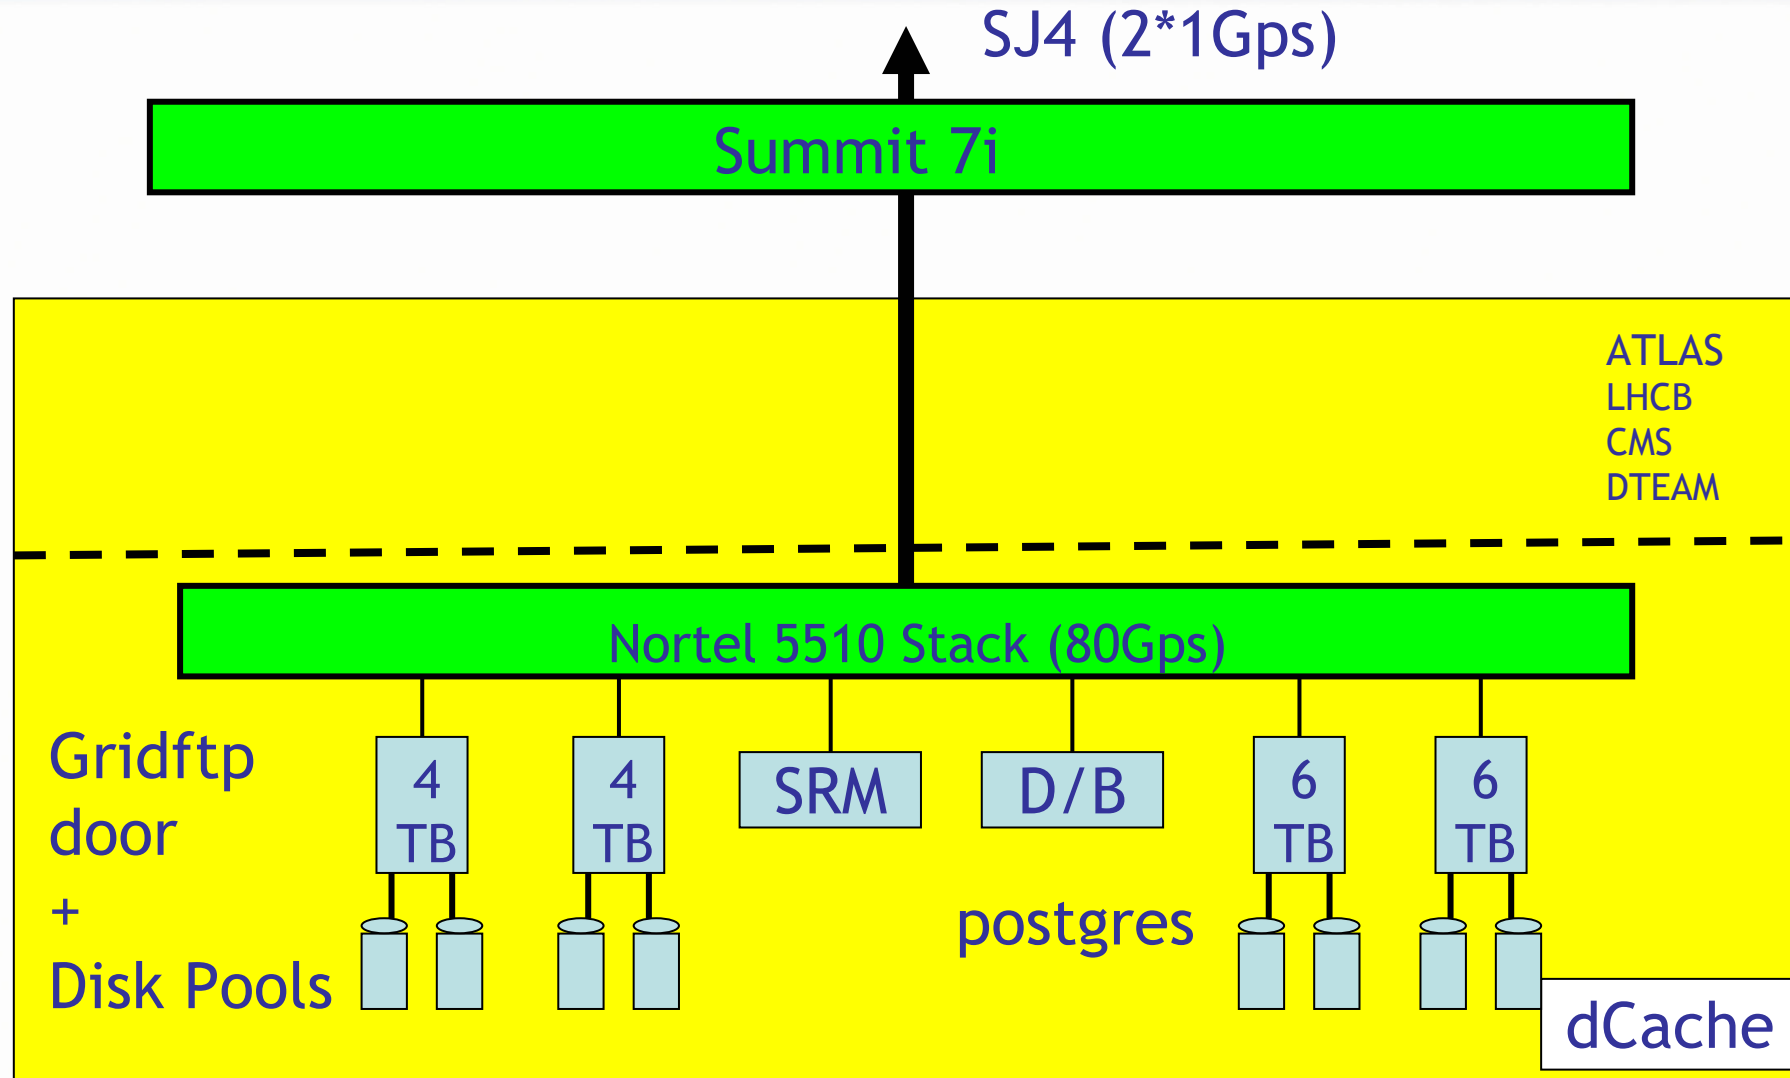
vdt_globus_sdk-VDT1.1.14rh7gcc3-4.lcg1

vdt_globus_info_essentials-VDT1.1.14rh7gcc3-4



GridPP
UK Computing for Particle Physics

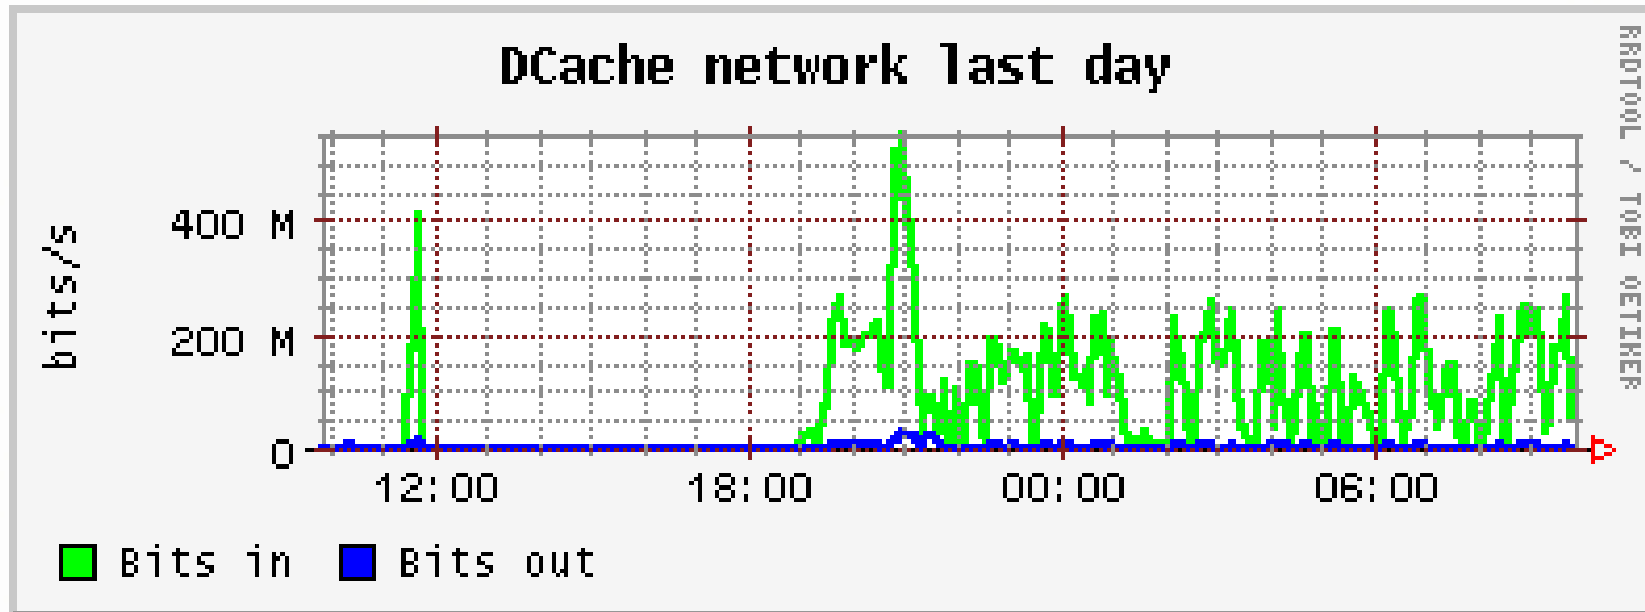
Production dCache SRM





GridPP
UK Computing for Particle Physics

CMS dCache Network Load




At peak: 500Mb/s sustained until RAL caught up with CERN stager. Not yet optimised - RAL lightly loaded



GridPP
UK Computing for Particle Physics


CMS Production Network Load




At peak: 500Mb/s sustained (full end to end including cataloguing etc) until RAL caught up with CERN stager.
Not yet optimised - RAL lightly loaded



SRM Server Load



Postgres CPU load very high - mainly waiting on update.
Needs tuning?

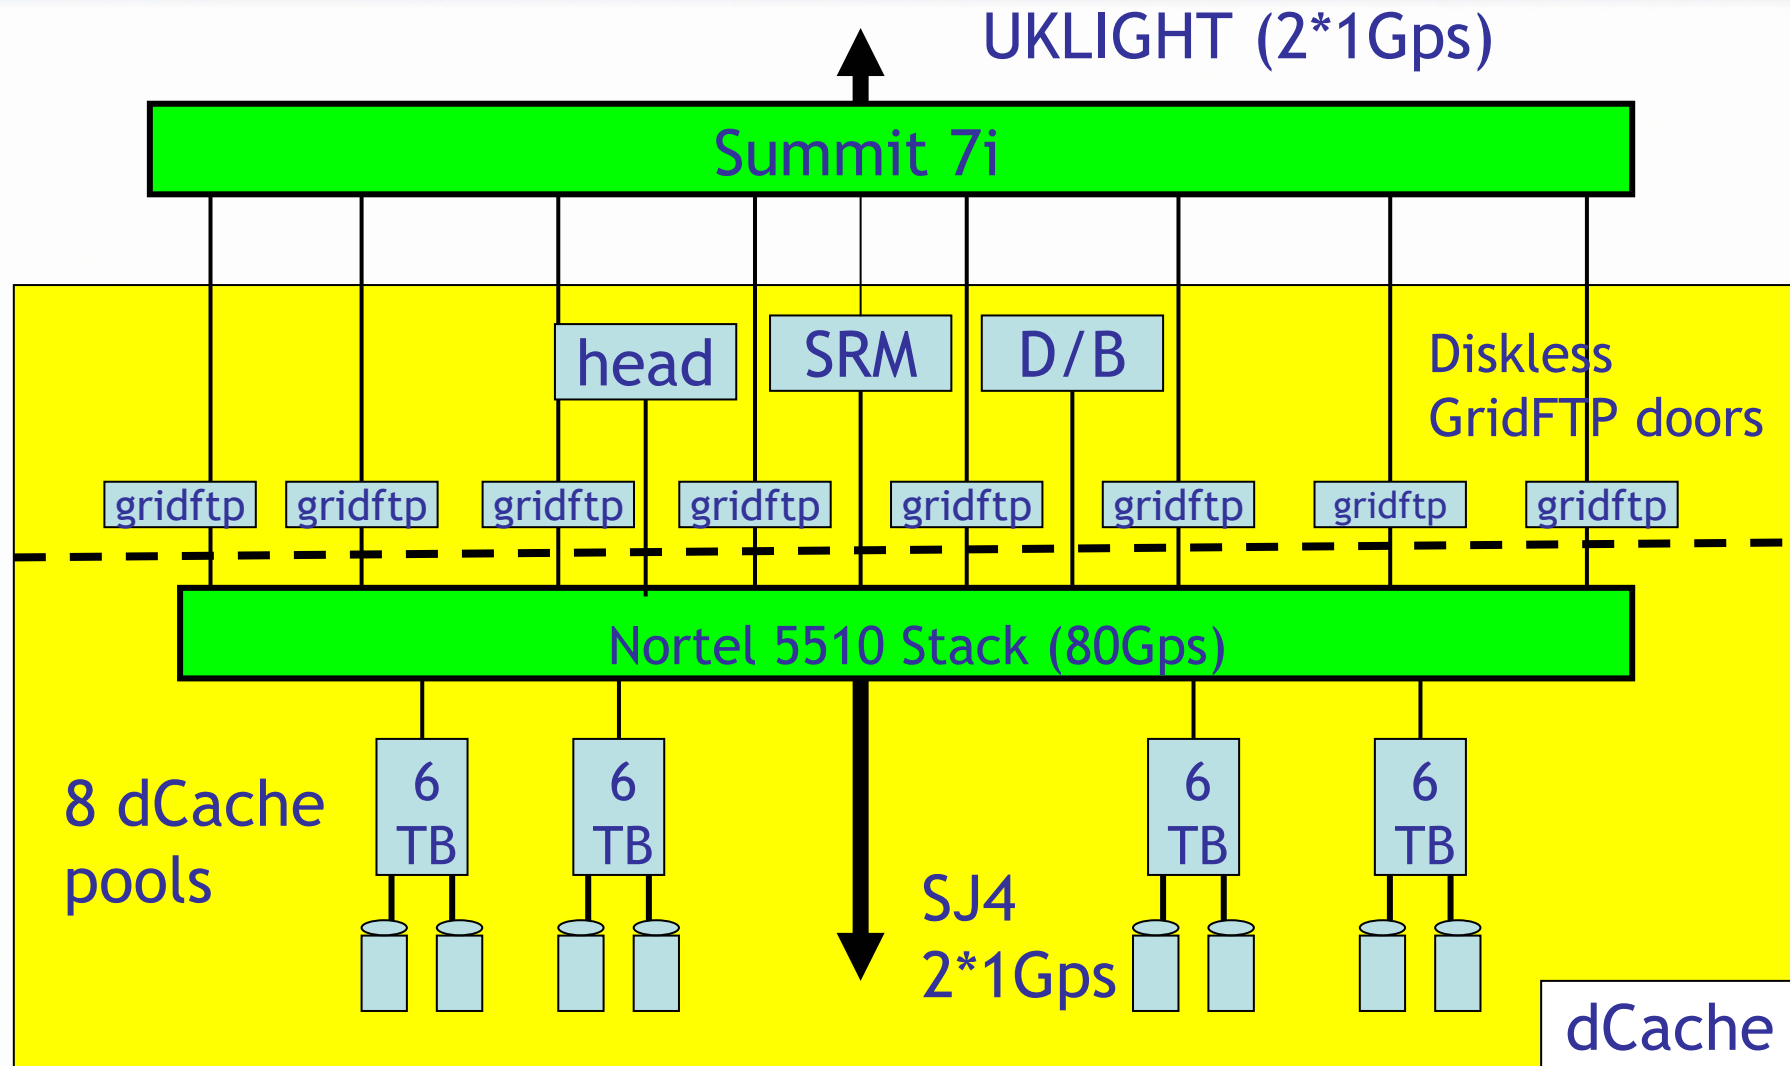




- Build another PRODUCTION dCache SRM
- Build from kickstart - therefore quick
- Build on PRODUCTION SRM experience
- Build in production monitoring
 - Exception monitoring (SURE/NAGIOS)
 - Ganglia Monitoring
 - helpdesk
 - Etc etc
- Build modular - avoid bottlenecks
- Decouple network transfers from disk servers - allows easier network/kernel tuning

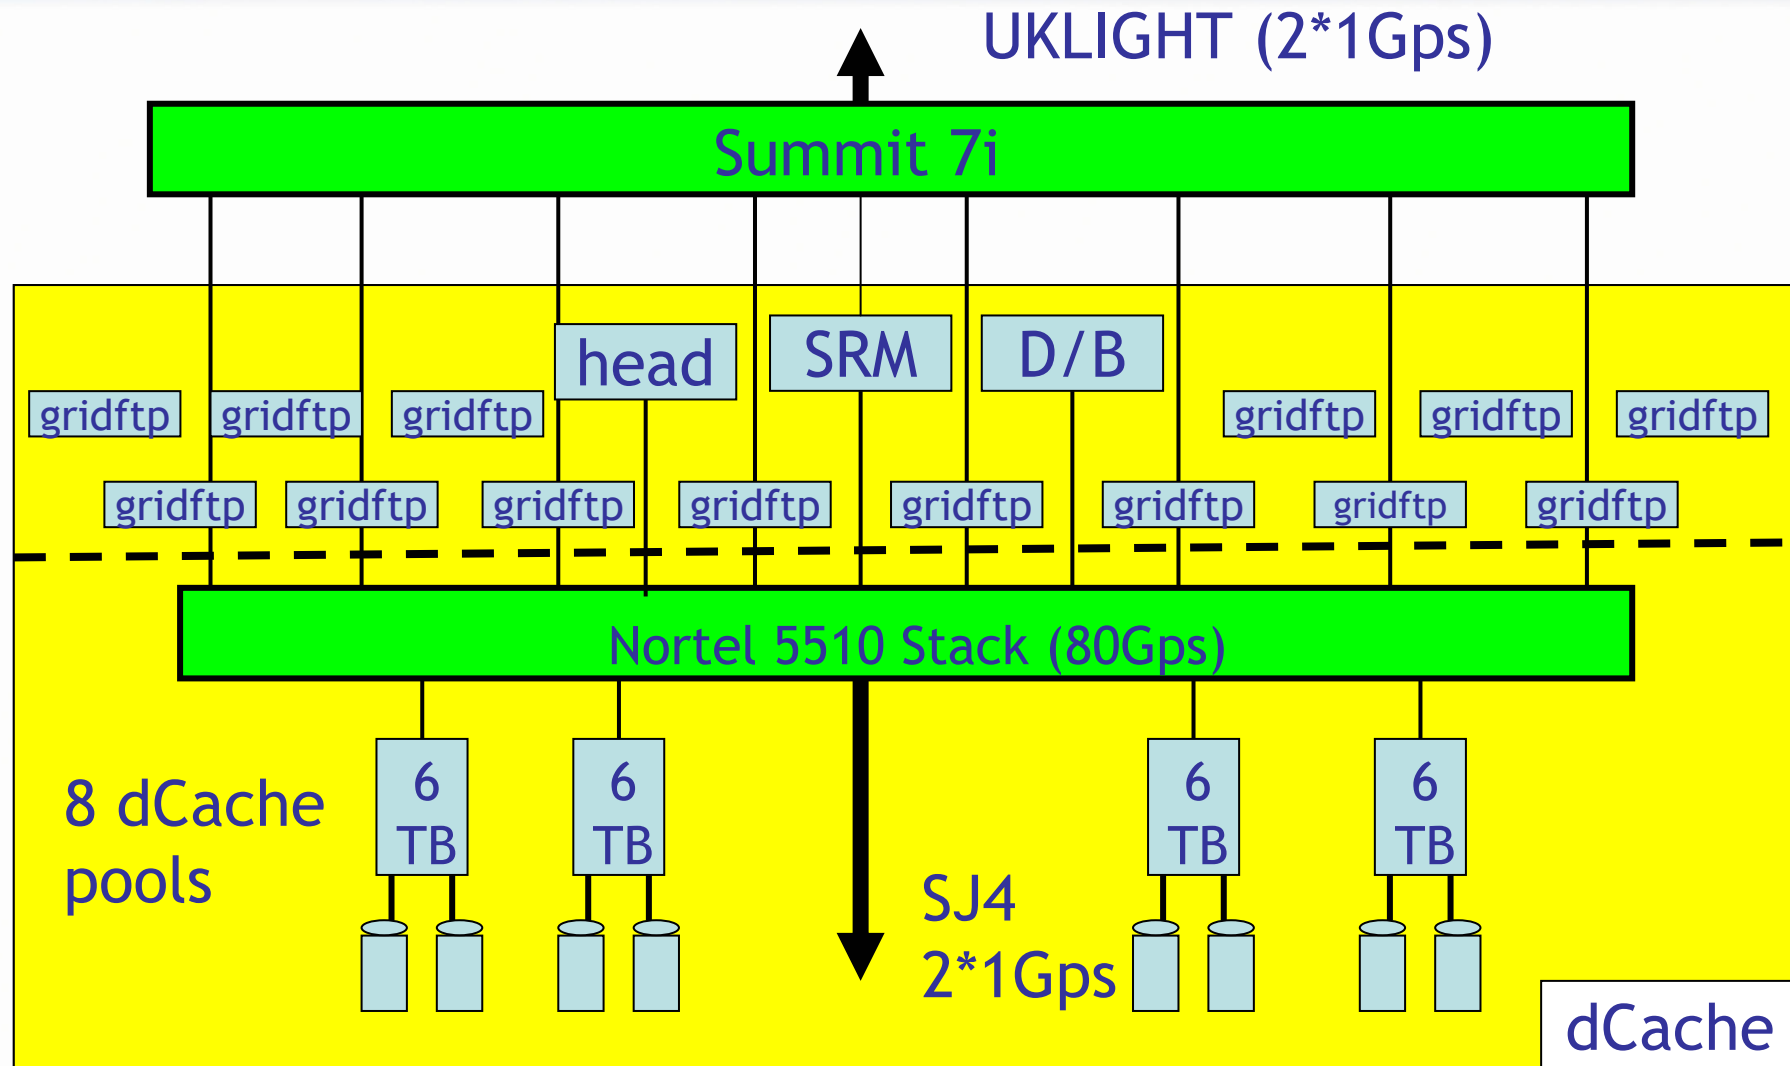


SC2 dCache SRM



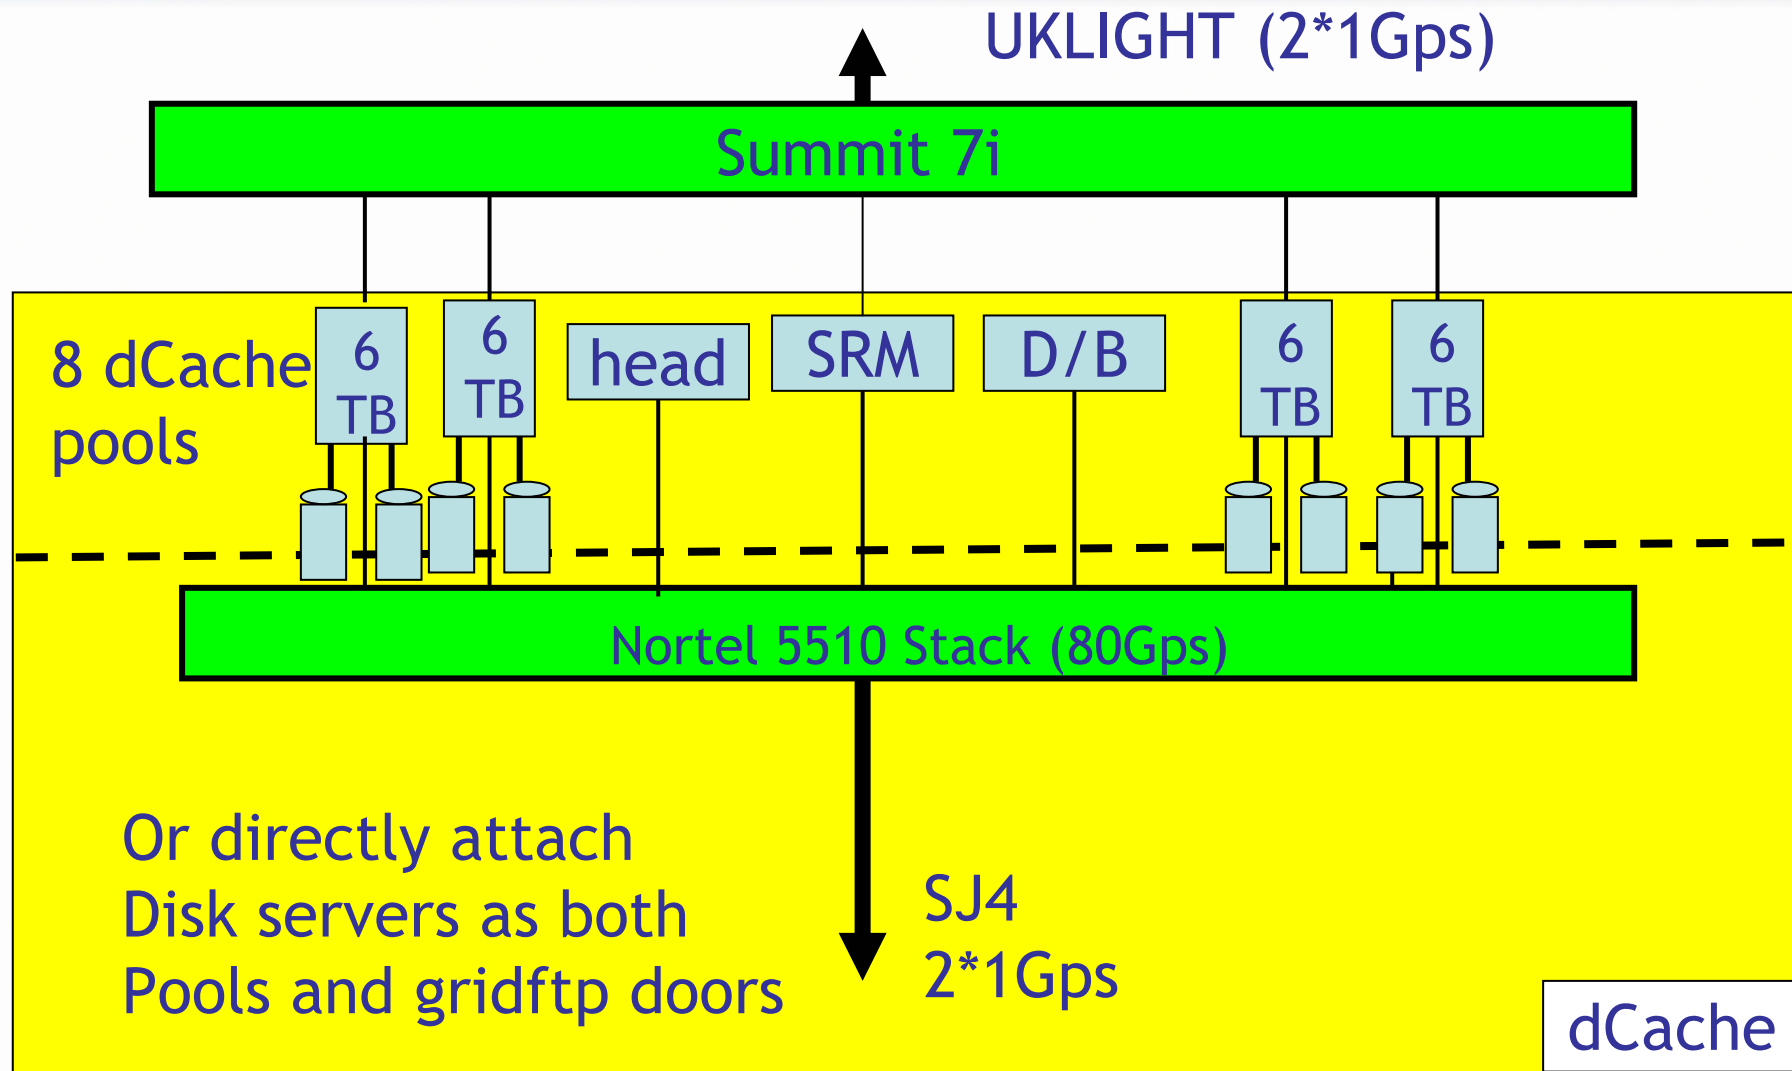


SC2 dCache SRM






SC2 dCache SRM



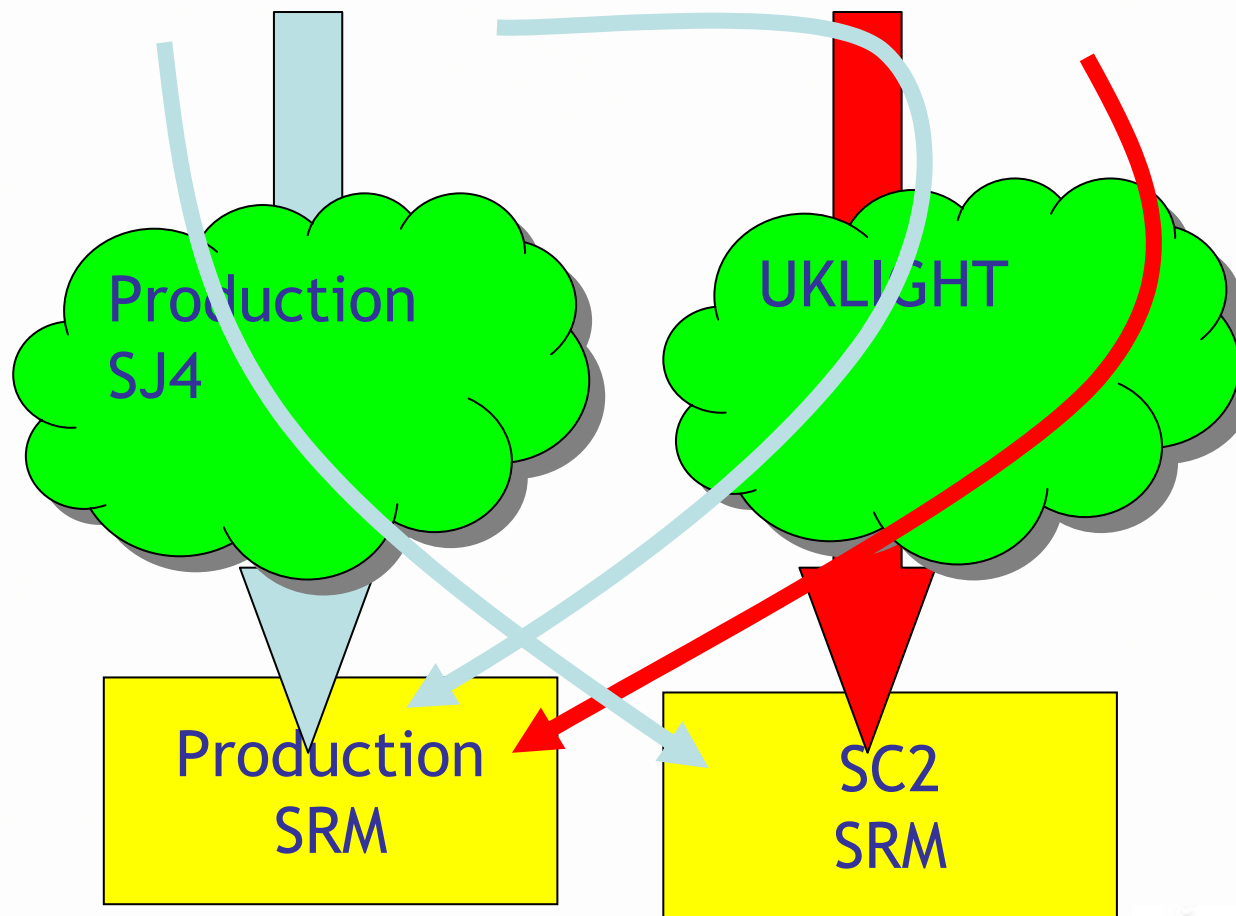


SC2 dCache SRM





Experiment Production Service Challenges





MSS Integration

- Back end from dCache to MSS (STK plus RAL ADS) written and tested. Just beginning experiment testing (CMS initially).
- MSS still bottlenecked at MSS cache disk - probably will not exceed 60MB/s (writing) until new hardware delivered (in time for July).
- Will not commit to MSS testing for SC2, but may turn dCache MSS back end on at some stage in SC2 and see what we can achieve.



- Network
 - Access agreement for UK Light
 - Peering agreed with NetherLight
 - Attach to UKLIGHT 25 February
 - Start throughput tests to CERN, W/B 28 Feb?
- Hardware
 - All hardware now drained from production
 - Extra NICs purchased
 - Disk Stress tests complete H/W ready to deploy
- dCache SRM
 - Just starting deployment (2-3 days)
 - Ready for SRM tests from CERN to RAL W/B 7 March
 - Peak load test 2Gb/s sometime in W/B 14 March