PIC
port d'informació
científica

Service Challenge 3

PIC Status Report

Outline

- Introduction
- Dedicated hardware
- Services deployed
- Some issues/conclusions

PIC
port d'informació
científica

Introduction

- PIC is participating in SC3 as a T1 center with disk to disk file transfers using PhEDex and gLite FTS.
- The manpower commitment is around 2 effective FTE allocated to this activity (Gonzalo Merino, Marc Rodriguez, José Hernandez and Francisco Martinez) for the chain CERN (Tier-0) ▸ PIC (Tier-1) ▸ CIEMAT/IFCA (Tier-2)
- The hardware commitment is of 20 TB of disk and 5 dedicated servers connected to a shared gigabit external network to CERN.

PIC
port d'informació
científica

Hardware I

- Disk
  - Infortrend FC-to-SATA RAID (A16F-R1211)
  - 16 HDs SATA 400GB 7200rpm
- 2 Nodes
  - Dell PowerEdge 750
  - P4 3.4GHz, 2GB RAM

PIC
port d'informació
científica

Hardware II

- Disk
  - 4x (DELL PowerVault 220S)
    - 14 HDs 73GB SCSI
    - 3x (14 HDs 300GB SCSI)
- 2 Nodes
  - Dell PowerEdge 2650
  - 2x Xeon 2.8GHz, 1.5GB RAM
  - RAID SCSI controllers:
    - DELL PERC3 and PERC4

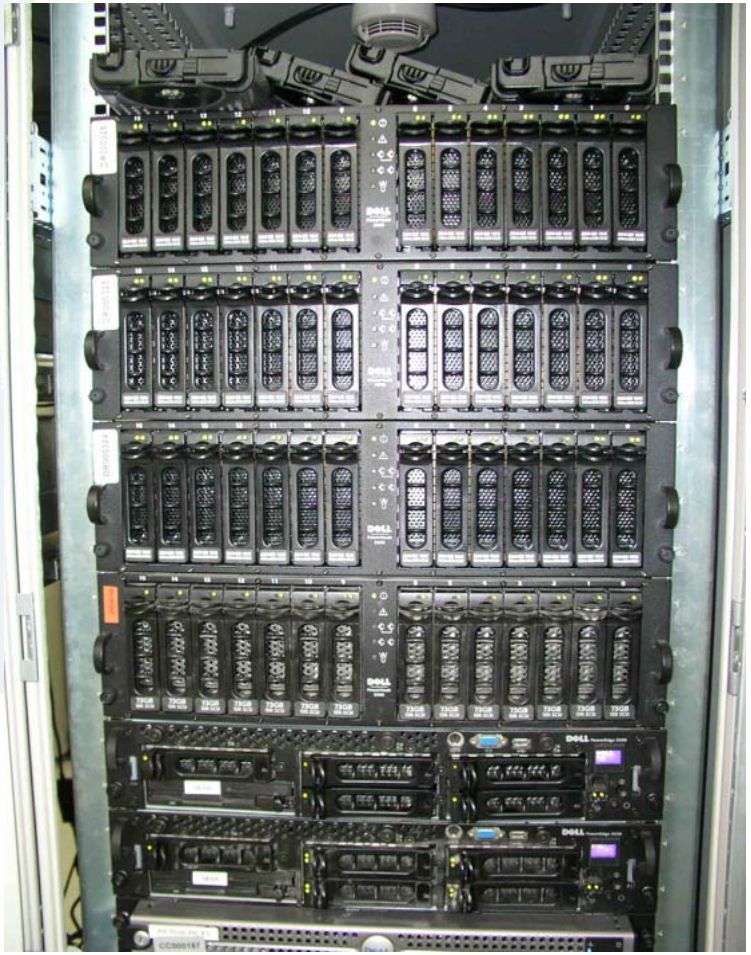

SC3 dedicated servers

Network connectivity: **1 Gbs shared**

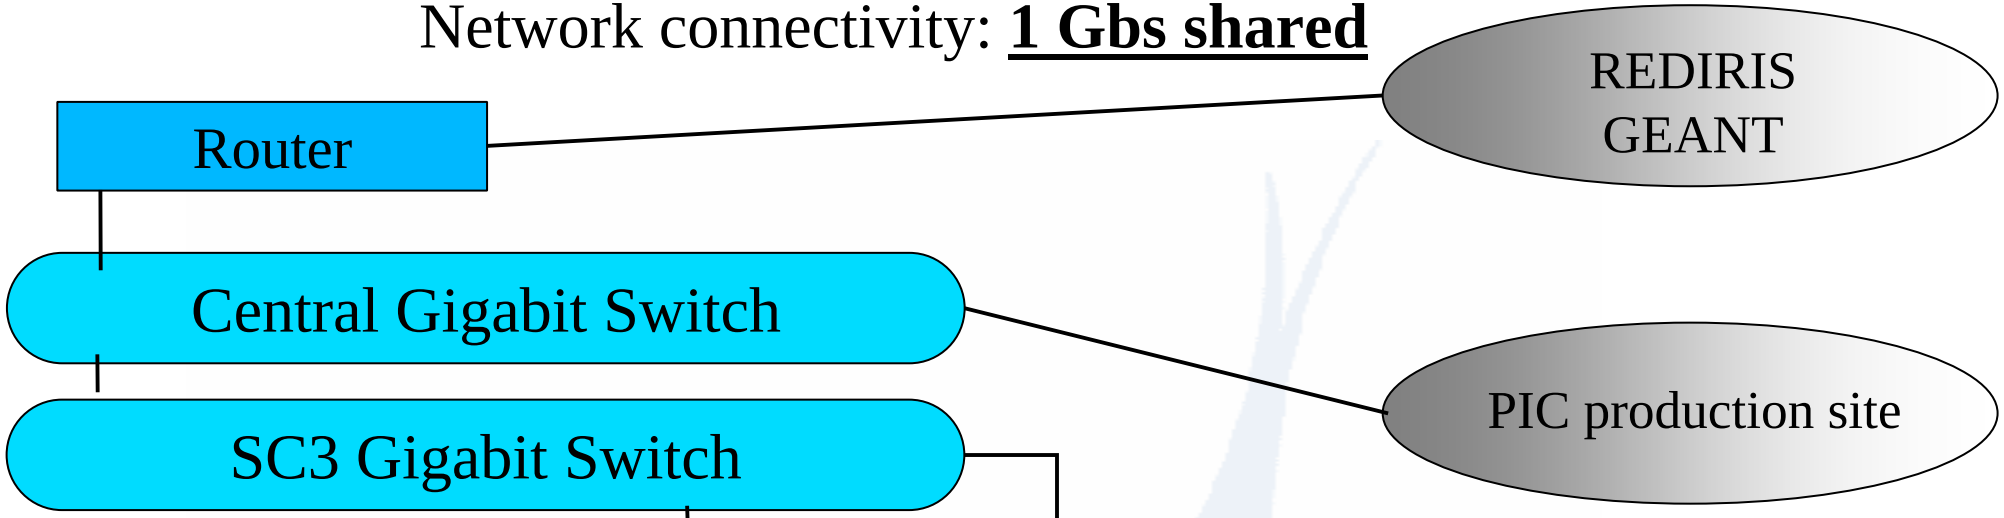

All internal connections are at 1 Gbs

castorsrmsc.pic.es

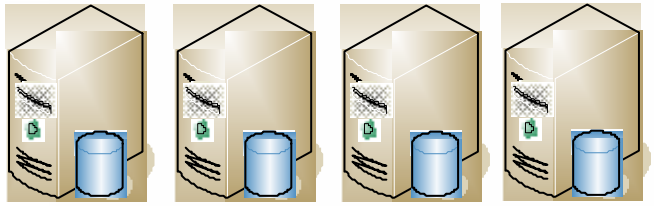

Each node has:
- SRM-Castor (SRM-1.4.3-1)
- GridFTP server
- Castor Disk Server

Castor Central Services
- Stager, NS, vdqm, vmgr...
- Dedicated to SC3
  /castor/pic.es/**sc3/**...

PhEDex catalogue

Issues/conclusions

- With the current 4 servers we ~ fill the available bandwidth

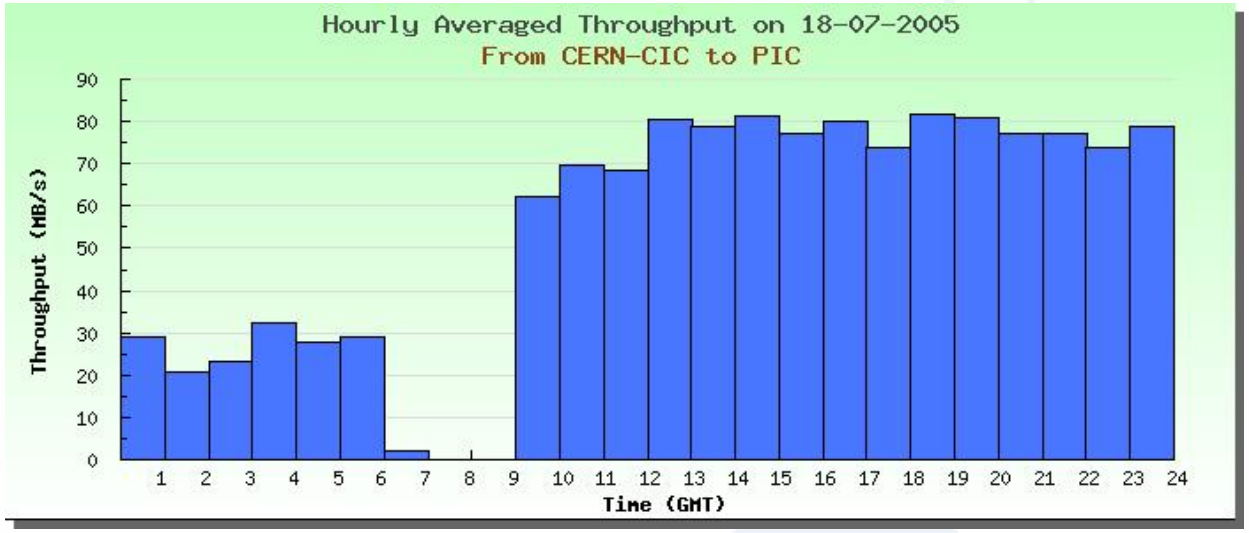

- Two main parameters for the FTS transfers:
  - Nr. of files being transferred in parallel
  - Nr. of streams per file
- Had no time yet to tune them carefully, but currently using Nfiles~20 and Nstreams~20

PIC
port d'informació
científica

Issues/conclusions

Number of streams issue:

- TCP window size set to 2MB, as on CERN side:

  net.ipv4.tcp_rmem = 262144 2097152 4194304
  net.ipv4.tcp_wmem = 262144 2097152 4194304

- But still we were not able to get more than ~2MB/s from a single stream...

- Maybe the max. bandwidth per stream is limited at some network element in the CERN-PIC path?

PIC
port d'informació
científica

Issues/conclusions

Writing to tape

- Currently we have five 9940B tape drives at PIC.
- Pushing up the throughput in writing to tape to try to reach the SC3 goal, would require dedicating large fraction of our resources.
- For this reason, for the moment PIC is not committed to contribute to the disk-to-tape exercise in the SC3-throughput phase.
- Depending on the load on the production system, we might try and allocate some tape resources online.