

Database Deployment at RAL

Gordon D. Brown
Database Services Manager
e-Science, CCLRC

HEPiX - Spring 2006
Rome
3rd – 7th April

Overview

- Database Services
- Projects: 3D, FTS, SRB, CASTOR, NGS
- Monitoring, Back-ups
- Hardware & Structure
- Future

Database Services Structure

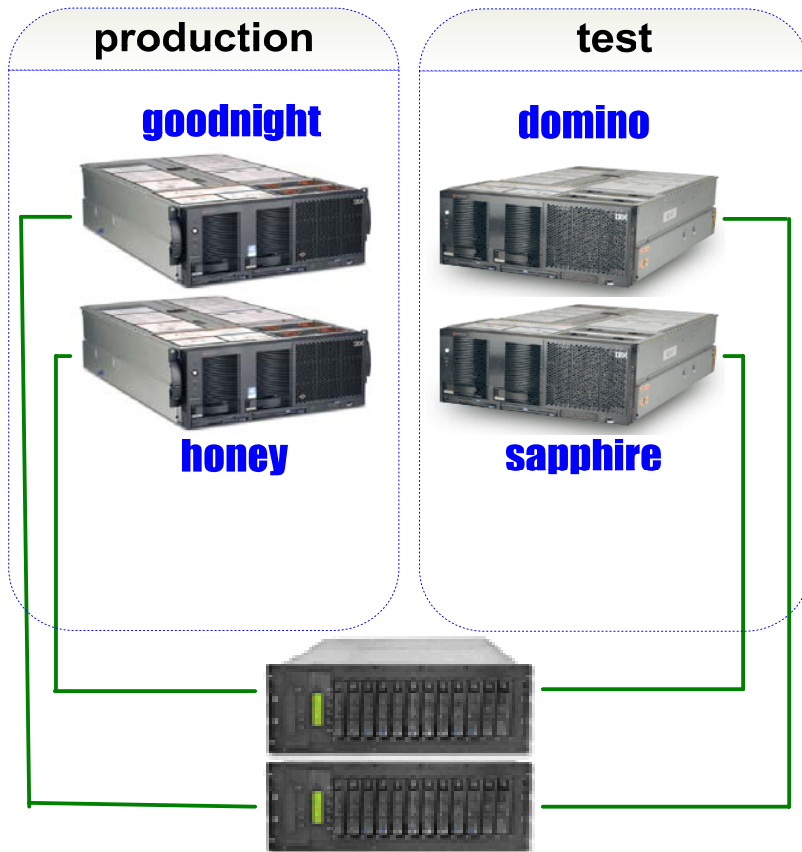
- Part of the Data Management Group in the e-Science Centre
- Structure
 - Database Services Manager
 - 4 DBA Staff Resource
 - 1 System Administration Staff Resources
 - More resources with other groups (including Tier-1)

Database Services Set-up

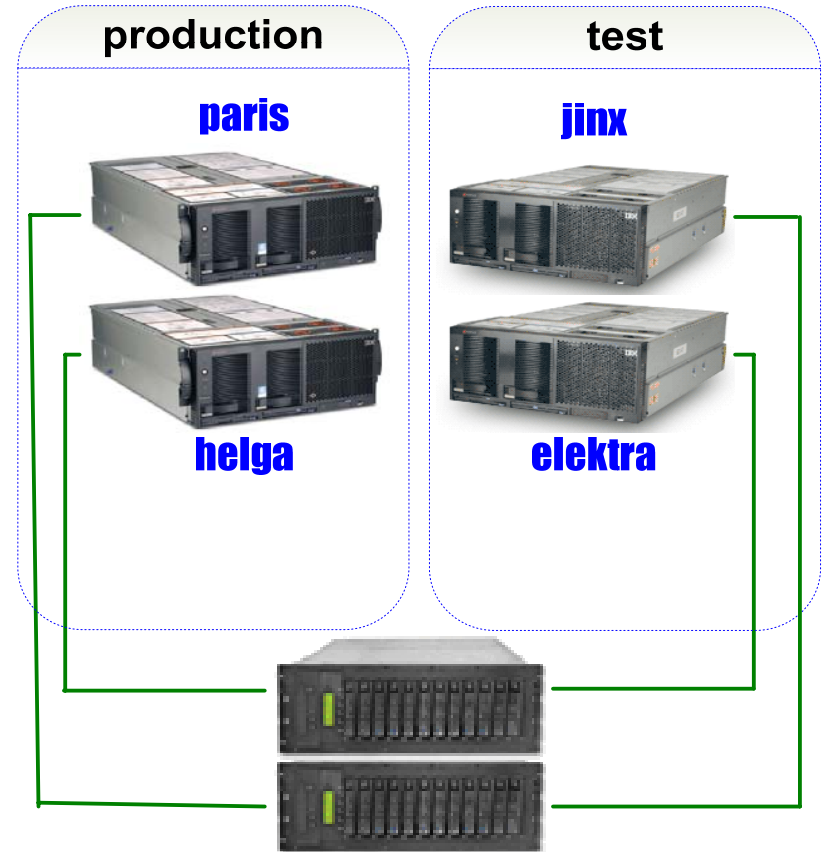
- Running Oracle 10g Enterprise Server with Real Application Clusters
- Using RedHat Linux (RHEL 4.0 AS) with Oracle Automatic Storage Management (ASM)
- Oracle Data Guard
- Recovery Manager (RMAN)
- Grid Control (Oracle Enterprise Manager)

Database Services Hardware Structure

Rutherford Appleton Laboratory



Daresbury Laboratory



RAL 3D Status

- Test implementation early last year
- Still in test phase, using test machine
 - 2 x 2.8GHz xeon processors
 - 2GB RAM
 - 120GB HDD
 - RedHat Enterprise Linux 3

RAL 3D Status


- Oracle 10.1.0.4 (single instance) installed
- Back-ups configured for RMAN
- Successful set-up of:
 - Oracle Streams (3D)
 - File Transfer Service (FTS)
 - Oracle Enterprise Manager (OEM)
 - COOL

RAL 3D Plans

- Hardware arrived Feb 2006
 - 2 x 2 node Oracle cluster
 - Dual Opteron CPU, 4GB Memory
 - Shared storage of 4TB
- Oracle installation and pre-production testing April 2006
- Production in autumn 2006

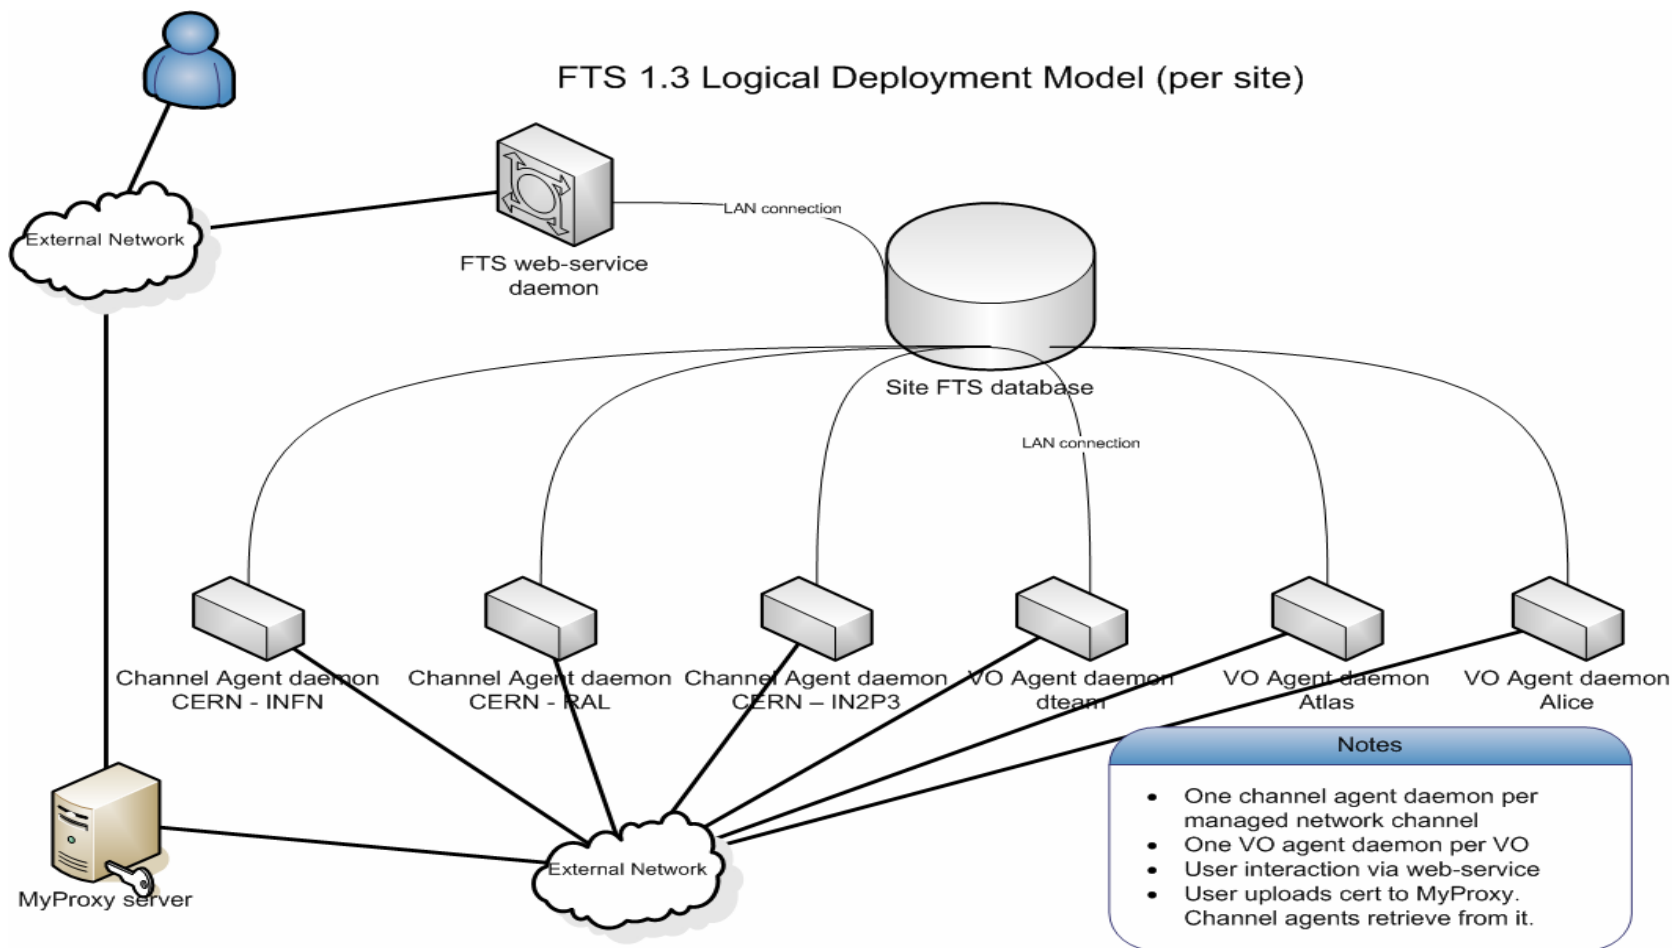
3D Database Hardware Structure

Rutherford Appleton Laboratory



LCG Grid Applications

- FTS
 - The File Transfer Service provides a controlled channel for reliable delivery of files from one SE to another. It maximises utilisation of lines and aims to avoid overloading servers with excessive parallel requests.
- LFC
 - Can run using Oracle



Storage Resource Broker

- Developed at San Diego Supercomputer Center
- A distributed file system (Data Grid), based on a client-server architecture
- Allows users to access files seamlessly across a distributed environment, based upon their attributes rather than just their names or physical locations
- It replicates, syncs, archives, and connects heterogeneous resources in a logical and abstracted manner

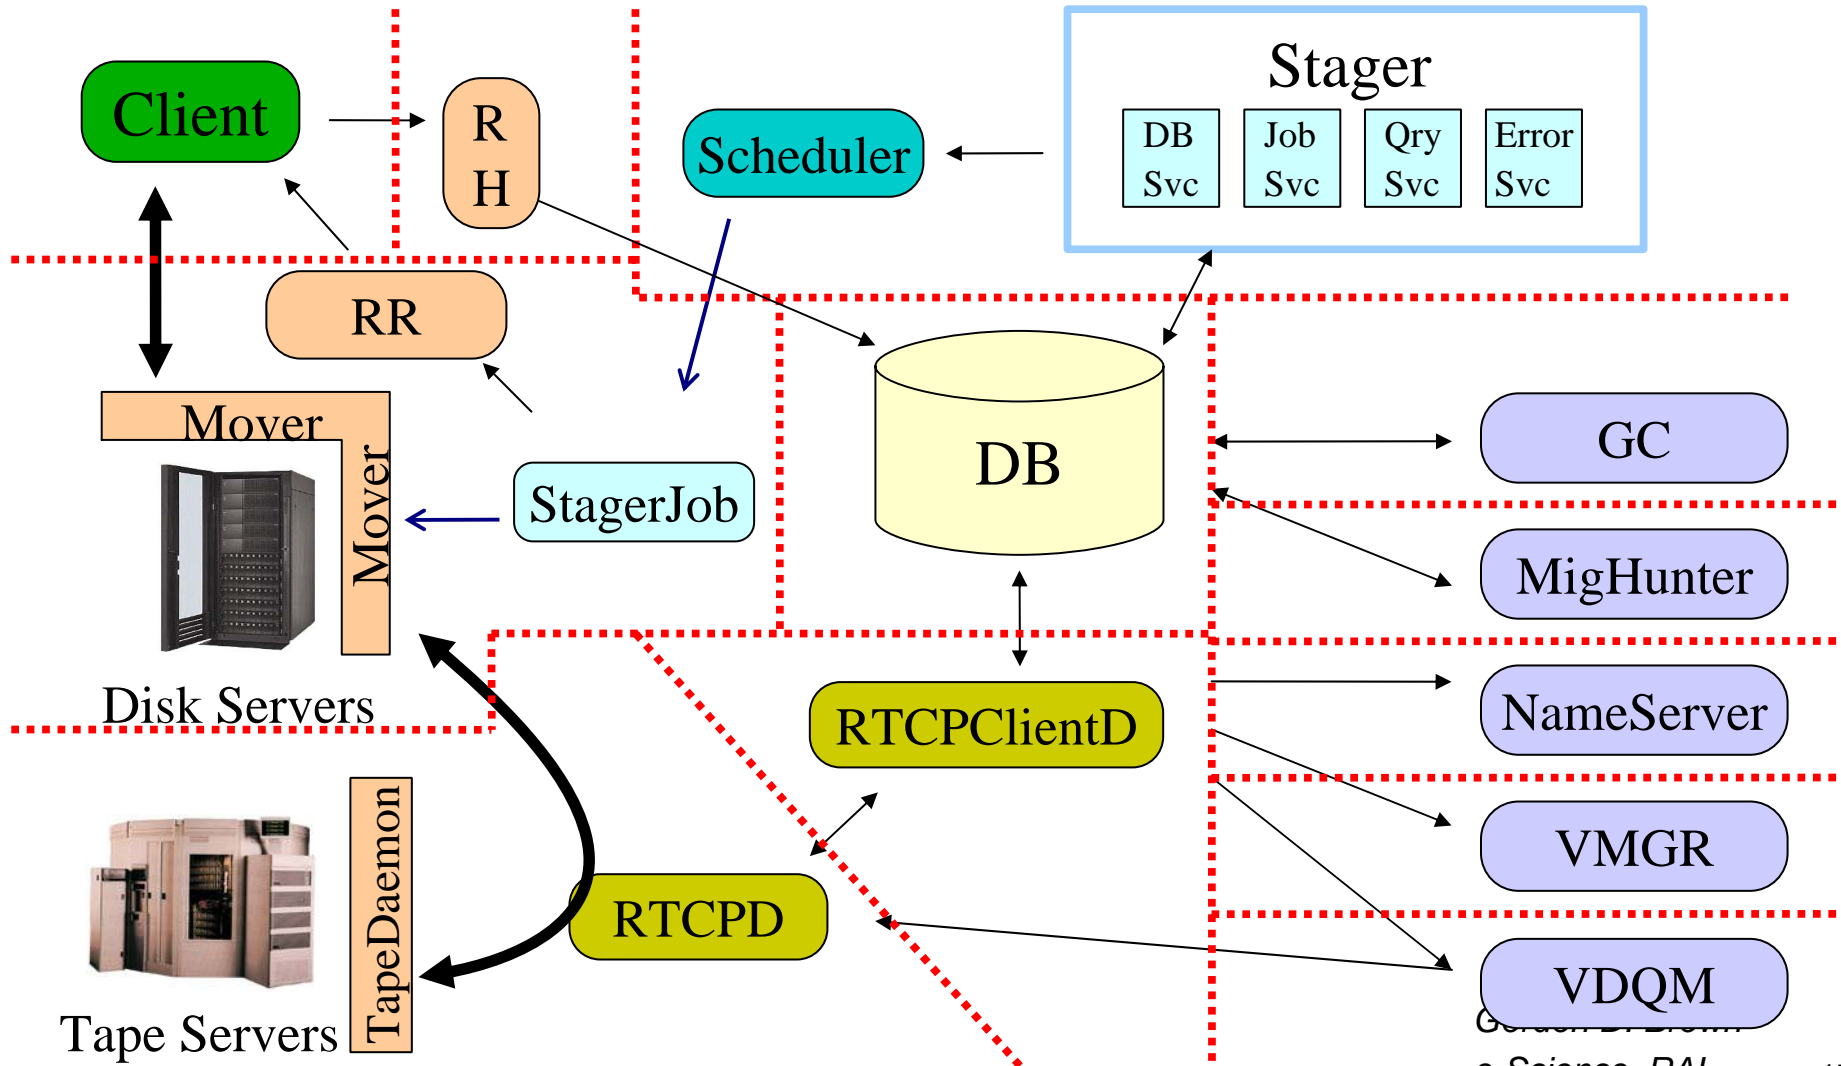
Storage Resource Broker

- The MCAT database is a metadata repository that provides a mechanism for storing information used by the SRB system
- Includes:
 - Internal system data required for running the system
 - Application (user) metadata regarding data sets being brokered by SRB.

CASTOR

- Cern Advanced STORage system
- Hierarchical Storage System developed by CERN for large scale storage of particle physics data
- Data access to/from scalable farm of disk servers
- Data transparently migrated from disk servers to tape according to defined policies
- Database centric architecture (Oracle)

CASTOR Architecture




National Grid Service

- Facilities for grid computing and data analysis
- To provide the broad user community with access to large scale resources through the grid interface
- To meet users' evolving requirements for a production grid service by working with the user community, through the User Group established by the Core Programme

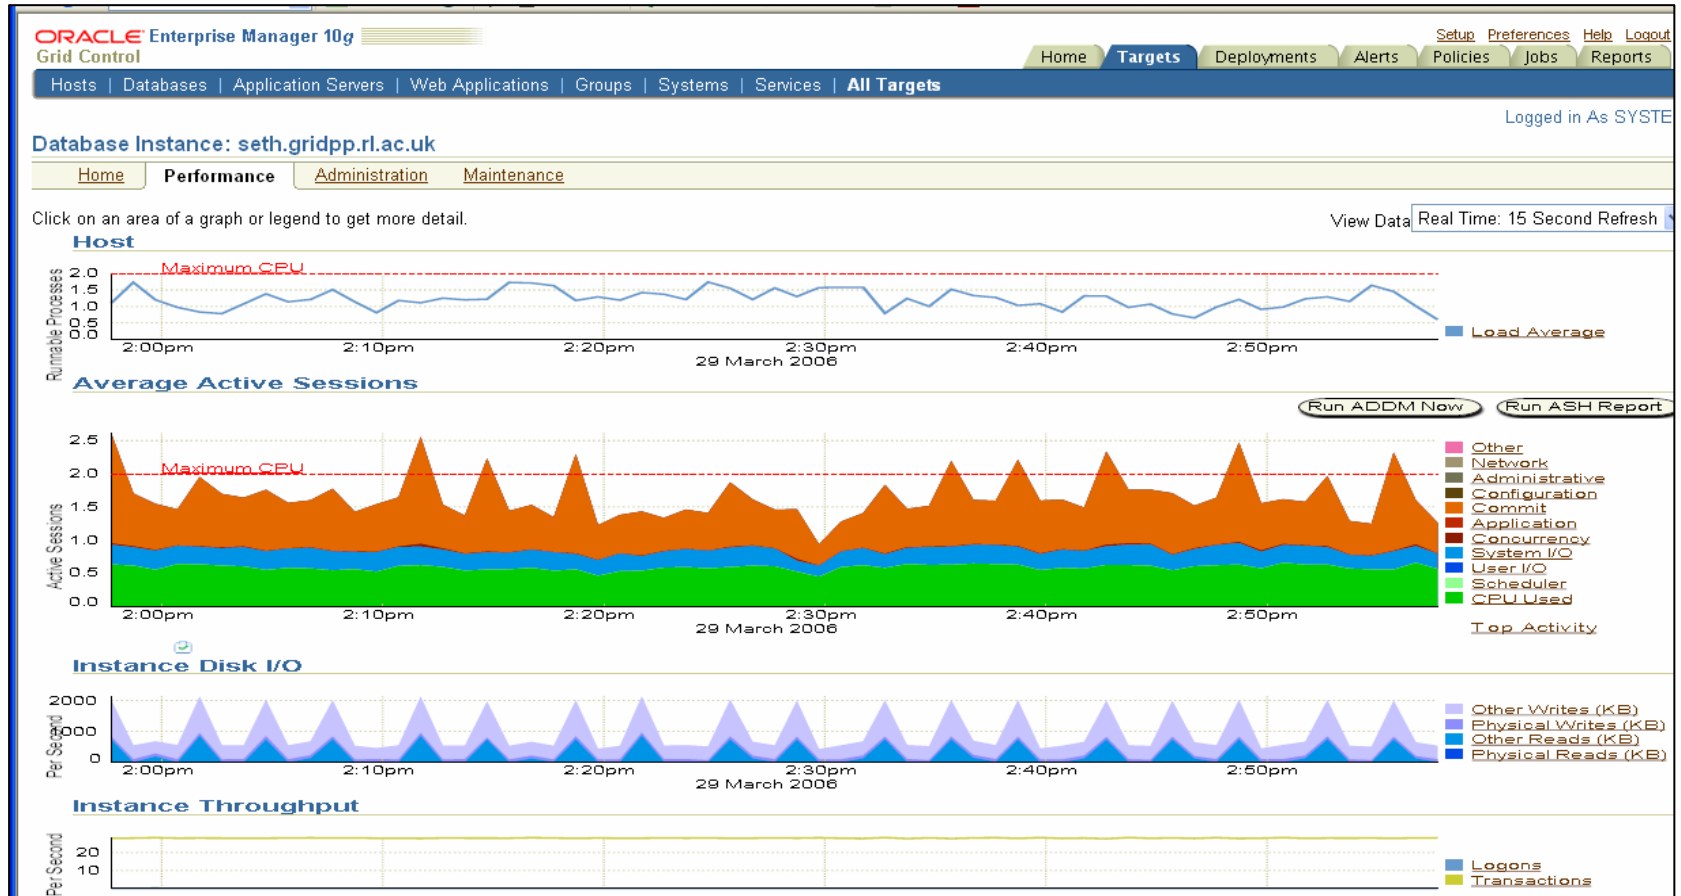
National Grid Service

- Configuration



Back-ups

- Current configuration
 - RMAN
 - BMAN
 - Atlas Data Store
- RMAN Recovery Catalog
- Efficiency Improvements
- Improvements to CCLRC Systems
- Extension to NGS



Questions & (hopefully) Answers