

The **Astro-Wise** Federation

... staying connected

Lorentz center, Monday, 31 March 2008

Components

- Databases
- Dataservers
- Distributed Processing Units
- Users
- Networks

The components of the federation

- Databases
 - Identical copies at all AstroWise nodes
- Dataservers
 - Multiple dataservers per node
 - Cache frequently accessed data from other nodes
- DPUs – Distributed Processing Units
 - Process data using database and dataservers at the “local” node
- Networks
- Users

Databases

- Hub-and-spoke replication using Oracle Streams
- Groningen as hub
- Bonn, München, Napoli as spokes
- Databases can be used locally if network to hub is down
- Active since about one year

Databases

- 15 GByte metadata
- 60 GByte sources and associations
 - excluding USNO and SDSS
- 50 GByte indexes
- 600 GByte filespace
 - including USNO and SDSS

Databases

- Groningen 581 GByte
 - Includes SDSS, USNO
- Bonn 1 GByte
 - Copying from scratch after crash
- München 70 GByte
 - Raw FITS headers missing
- Napoli 7 GByte
 - Sources and Associations still missing

Databases – what next

- Upgrade database hardware
 - Groningen
- Upgrade to Oracle 11g
 - Groningen, München
- Increase network bandwidth
 - Napoli
- Copy remaining data
- Local accounts for federated users

Dataservers

- Peer-to-peer
- Write locally, read anywhere
- Local copy in cache of remote files
- Connected for more than a year
- Today connected between Bonn, Groningen, München, Napoli and Nijmegen

Dataservers

- 25 Dataservers
- 363 TByte total online storage
- 23 TByte in use
- 2 million files

DPU's

- Distributed Processing Unit
- Clusters with 8 nodes up to 100's of nodes
- Front end to GRID processing
- Active for more than a year
- Available in Bonn, Groningen, München, Napoli and soon in Nijmegen

Networks

- Speed from or to Groningen
 - Bonn 2~5 MByte/s
 - München 2~5 MByte/s
 - Napoli <0.5 MByte/s
 - 5MByte/s radio link to be installed
 - Nijmegen >5 MByte/s ?
- Databases and dataservers share the *same* bandwidth

Discussion

- Questions?
- Comments?
- Suggestions?

The components of the federation

- Databases
 - Identical copies at all AstroWise nodes
- Dataservers
 - Multiple dataservers per node
 - Cache frequently accessed data from other nodes
- DPUs – Distributed Processing Units
 - Process data using database and dataservers at the “local” node
- Networks
- Users

Today's federation

- 5 active nodes (Bonn, Groningen, München, Napoli, Nijmegen)
- 4 databases (Bonn, Groningen, München, Napoli)
 - 150GB of sources and associations
 - Includes USNO and SDSS
 - 5GB of metadata
 - 50GB of indexes
- 10 dataservers
 - 22TB of science and calibration images
- 5 DPU's
- Number of users 20-30

Tomorrow's federation

- 6 active nodes (+Leiden)
- 6 DPU's (+Nijmegen)
- 4 databases
 - 150GB of sources and associations
 - Includes USNO and SDSS
 - 5GB of metadata
 - 50GB of indexes
- 18 dataservers
 - 22TB of science and calibration images

Databases

- Query locally and add data locally
- Peer to peer copies using Streams
- Metadata is almost instantly available
- 100,000 sources with 300 double precision parameters take less than a minute to appear on a different node (roughly 4MByte/s depending on the network bandwidth)
- Each database knows each filename in the federation

Remote database not accessible

- Remote database can be down
or
The network in between can be down
- Local queries are not affected
- Data that is added locally will be copied as soon as the remote database is up again
- Users will not be aware that a remote database is down, but administrators will

Dataservers

- Peer to peer access without copies
- Dataservers know of each other
- Store locally and retrieve from anywhere
- Cache frequently used remote data because of bandwidth
 - A small fraction of the data is accessed very frequently when processing
 - Local network speed is about ten times or more the network speed between nodes

Dataserver not accessible

- Any dataserver can be down
or
The network to it can be down
- If a copy in the local cache is available it will be used
- If data is not available in a local cache processing can only succeed if the dataserver can be accessed again

DPU

- DPU is independent of other DPU's
- DPU functions independently of which database it is told to use
- DPU will retrieve data from the datasevers that are federated