

Monitoring the ATLAS Data Grid

Patrick McGuigan

ATLAS Data Grid

- ATLAS experiment to generate 2 petabytes of data per year
- Researchers are geographically dispersed
- Grid computing provides mechanisms to aggregate computing resources
 - Creation of Virtual Organizations
 - Data storage and retrieval
 - Distributed analyses
- ATLAS computing to be organized in tier structure.
 - Tier 0 – CERN
 - Tier 1 – National computing center (BNL)
 - Tier 2 – Regional computing centers (TBD)
 - Tier 3 – Organizational centers
 - Tier 4 – Researcher's desktop

Monitoring

- Types of monitoring:
 - Site monitoring (state of collection of computers)
 - Resource monitoring (state of computer)
 - Internal application monitoring (trace, debugger)
- Monitoring in grid should use available grid tools (MDS)
- Two scenarios for monitoring
 - Monitoring sites to determine ATLAS induced load on site's resources.
 - Monitoring computational resources to allow researchers to make optimal choice of hosts.

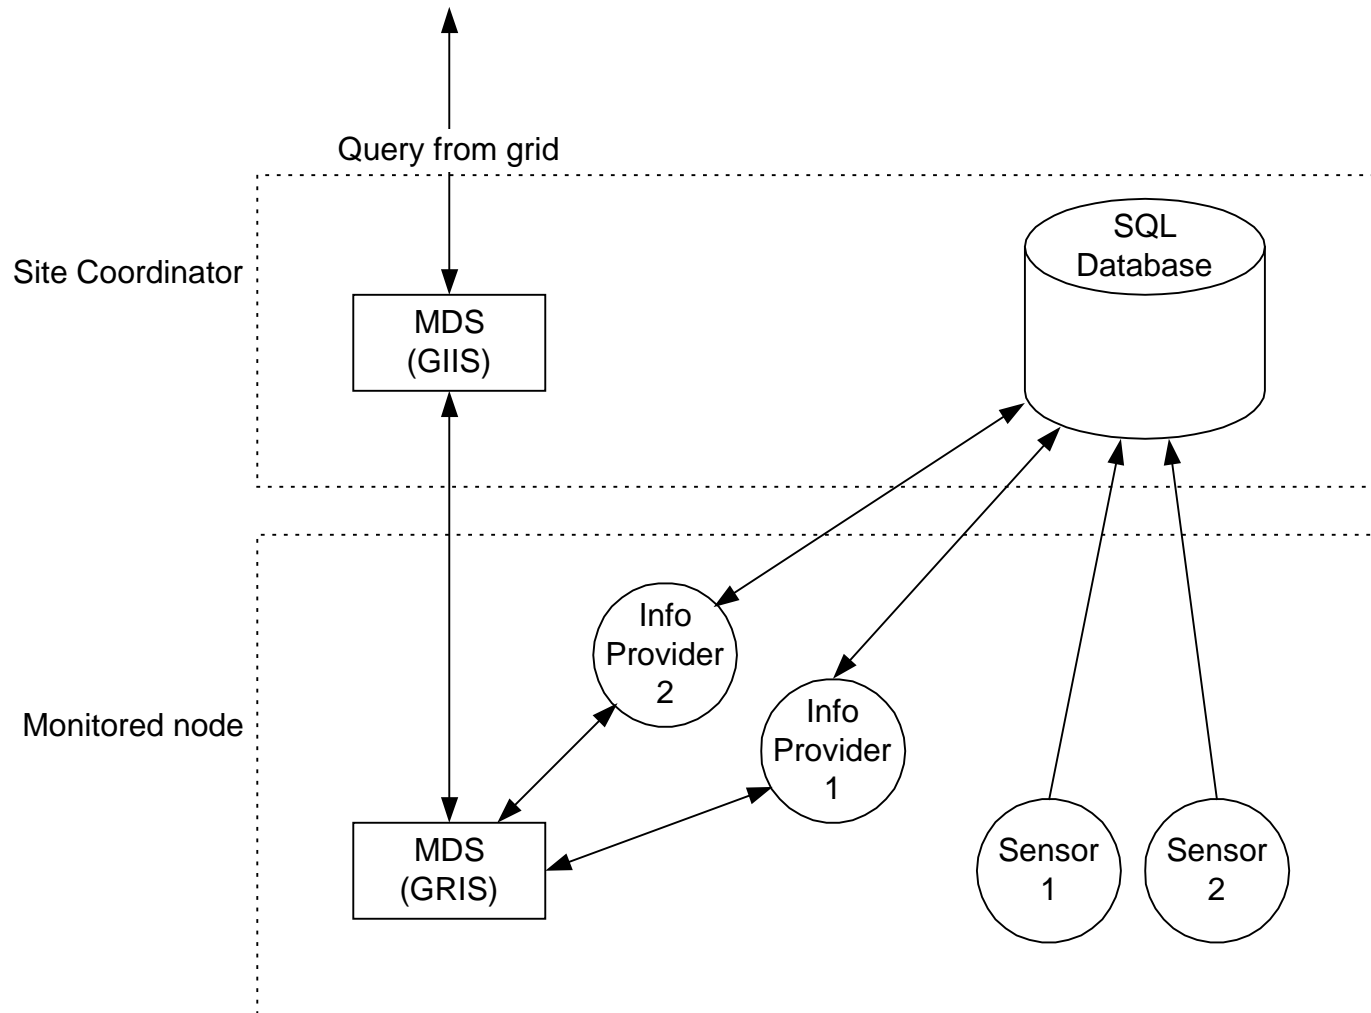
Site Monitoring Scenario

- Each node will have sensors for
 - CPU time consumed by grid users
 - Network utilization (bytes sent/received) by grid users
 - Secondary storage utilization by grid users
- Provide cumulative statistics for both node and site.
- Provide archive of measurements to determine total usage over time.
- Allows one to discover underutilized resources that may be candidates for upgrades and/or retirement.

Computational Resource Monitoring Scenario

- Each node will have sensors for dynamic information
 - Processor load
 - Network load
 - Available free memory
 - Available free storage
- Provide archive of measurements so that prior history can be used as predictor
- Allows one to determine “best” host based not only on current measurements but can include historical data to make decisions.

Monitoring Architecture



Monitoring Research Issues

- Impact of monitoring on resource performance
- Scalability of design
- Data format design
- Choice of sensors
- Frequency of measurement
- Timespan contained in archive