

Top level BDII Deployment Proposal

Flavia Donno (IT-ES), Laurence Field (IT-GT)
CERN

- The WLCG Information System is to discover services and get status information.
- Multi-component system: Information providers, resource BDII, site-level BDII, top level BDII
- It is important to ensure reliability of the WLCG Information System Service.
- The service must be treated as critical.
- New proposal for a deployment strategy for top level BDIIs.

- Criteria to take into consideration:
 - Query load
 - Network Latency
 - Available support effort
- Deploy a well managed number of top level BDII instances in the continents: North America, Europe, Asia
- The topology could take into account Monarc hierarchy of T1/T2 per experiment

- At least 3 4-core physical (avoid virtual) machines with a minimum of 4GB physical memory in load-balancing (round robin) configuration
- The machines must be strictly monitored so that the average CPU and Network utilizations do not exceed 40%.
- In case that the average exceeds the indicated threshold, more machines must be added to the pool.

- The service must be 99% available (MoU)
 - Average availability measured on an annual basis
 - $(\text{time running})/(\text{total time})$
- Availability means:
 - All BDII processes must be running
 - DB query response time less than 10 seconds
 - Data freshness less than 1 hour
- Sites are encouraged to instrument their own monitoring probes to check data freshness or use Nagios probes
 - `ldapsearch -LLL -x -h lcg-bdii -p 2170 -b o=infosys modifyTimestamp`
- It is recommended to install a watchdog to automatically restart the BDII in case of problems.

- The service must be strictly monitored
- The metrics to monitor for are the following:
 - CPU load and Network utilization
 - Host uptime
 - Memory usage (should not exceed 1GB for openldap processes)
 - Freshness of data (should not exceed 1hour)
 - Uptime for BDII processes
 - DB query response time (less than 10 seconds locally for all queries)

- The service must be managed as a critical service for WLCG.
- Releases must be carefully followed.
 - Upgrades are expected to happen within 1 month after the tested candidate release is approved by the information officer and announced at the T1 Service Coordination meeting.
- Service administrators should include rolling back to previous release as part of the upgrade plan.
 - EMI developers commit to always produce backward-compatible releases.
- The service must be monitored to identify and fix failures within hours and to measure the query load in order to provide additional resources when required.

- Finalize this proposal with the involvement of the WLCG experiments, EMI, EGI, NDGF, OSG
- Identify hosting Tier-1s
- Agree on a fail-over strategy for services depending on the WLCG Information System.

- The need to increase the robustness of the WLCG Information System service has been identified
- A new top level BDII deployment strategy has been put in place
- The presented proposal lists the technical specifications for top level BDII services hosted at Tier-1s
- The fail-over strategy foresees the identification of candidate hosting Tier-1s and configuration details for BDII depending services.