

Table of Contents

CE Notes	1
Introduction.....	1
Components.....	1
Data.....	1
Configuration.....	1
High Availability.....	1
Equipment required.....	2
Engineering required.....	2
Questions.....	2

CE Notes

Introduction

The Computing Element (CE) is the service representing a computing resource. Its main functionality is job management (job submission, job control, etc.). The CE may be used by a generic client: an end-user interacting directly with the Computing Element, or the Workload Manager, which submits a given job to an appropriate CE found by a matchmaking process. For job submission, the CE can work in push model (where the job is pushed to a CE for its execution) or pull model (where the CE is asking the Workload Management Service for jobs). Besides job management capabilities, a CE must also provide information describing itself. In the push model this information is published in the information Service, and it is used by the match-making engine which matches available resources to queued jobs. In the pull model the CE information is embedded in a "CE availability" message, which is sent by the CE to a Workload Management Service. The matchmaker then uses this information to find a suitable job for the CE.

In the production cluster for CERN, the local batch system is based on LSF.

The CE participates in the following flows

- WmsFlows

Components

The CE consists of

- An LDAP server which presents the information on the CE to the BDII
- A gatekeeper to receive jobs from the RB and submit them to the local batch system
- A gridftp server to receive status information for workers

Data

The CE data storage is as follows

Location	Purpose
/pool/	Storage of GASS data (?). There are a LARGE number of files for each user. Purpose unknown
/var/log	Log data currently at 1.1GB. Main user is the gatekeeper log which does not seem to be rotated

Configuration

Configuration is performed via YAIM.

High Availability

If the CE is down,

- New jobs cannot be submitted to the site
- Completed jobs will not be able to report their status
- Accounting data is not reported

An IP alias ce001.cern.ch will be defined which allows the service to be switched between machines if required.

All state data will be stored 'off the box'. The state data consists of several directories.

..

Thus, in the event of failure of the master, the slave would take over the external disks. The state data stored on file systems would be 'rolled back' using ext3 functions. The MySQL database would be restarted and would play its redo log to arrive at a consistent state.

Equipment required

Assuming 1 master CE and 1 spare, the hardware required is

Component	Number	Purpose
Midrange Server	2	CE masters and standby machines
FC HBA	4	Fibre channel connectivity
FC Switch Ports	4	Connectivity for the two servers
FC Disk space	20	Storage for job information (2x10GB on different disk subsystems)

Engineering required

Development	Purpose
Start/Stop/Status procedure	Scripts for operations
Lemon GridFTP availability test	A lemon aware sensor for GridFTP
Lemon CE availability test	A lemon aware sensor which can be used for reporting availability CeMonitoringNotes
Linux Heartbeat availability test	A Linux-HA aware sensor which would activate the procedure for automatic switch from master to slave
Switch procedure	Automatic switch from master to slave changing the DNS alias, disabling the master, enabling the slave in its new master role
Capacity Metric	Capacity metrics defined for Number of jobs/second Active accounts / VO for tracking cmsNNN usage
Quattor configuration for Linux-HA	NCM component to configure Linux-HA/Heartbeat

Questions

Nr	Description	Status	Open Date	Who	Log
1	What is the purpose of the MySQL database	open	2005/09/21	Tim	
2	Where does the CEMON fit in to the CERN installation ?	open	2005/09/23	Tim	

-- TimBell - 15 Sep 2005

This topic: LCG > CeNotes

Topic revision: r8 - 2006-11-28 - LaurenceField



Copyright &© 2008-2021 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.
or Ideas, requests, problems regarding TWiki? use [Discourse](#) or [Send feedback](#)