

# Migrating INFN-T1 to HTCondor

Stefano Dal Pra

INFN-T1

[stefano.dalpra@cnaf.infn.it](mailto:stefano.dalpra@cnaf.infn.it)

HTC Week, 23/05/2018



# Summary

- 1 **Current status**
- 2 **Tools**
- 3 **batch special configurations**
- 4 **Planning migration**
- 5 **Acknowledgments**

## The INFN Tier-1

- main italian computing centre for HEP, located in Bologna
- delivers ~ 370 KHS06 (~ 32000 slots, 1000 WN)
- two remote sub-sites (Bari-Recas, 800Km; Cineca, 16Km)
- opportunistic (cloud) resources
- All WN managed by IBM LSF 9.1.3 batch system.
- 3 to 4 GB RSS/slot
- Grid access through Cream CEs

## Users

- **Grid:** The four LHC experiments, astronomy (Virgo/Ligo, AMS et al)
- **Local:** about 25 more minor experiments
- Access arbitrated by LSF using fairshare policy: priority of a pending job depends on the dynamic priority of its submitter, which depends on his overall activity

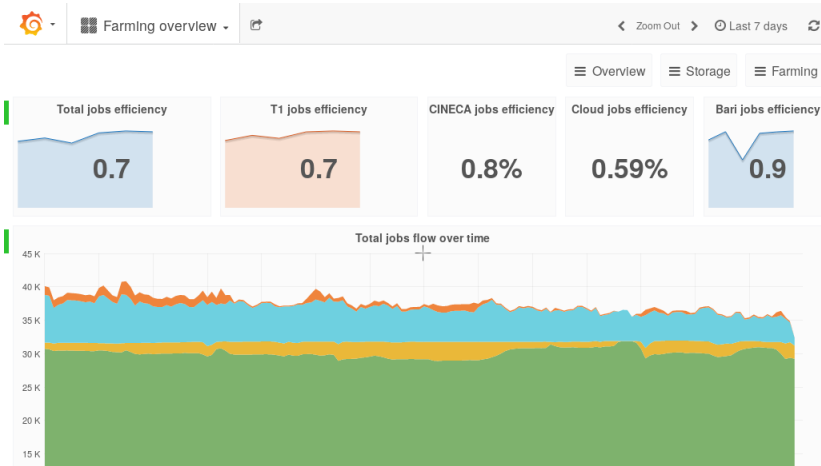
## Job types

- **singlecore**
- **multicore** (2,4,8 slots, dynamic provisioning)
- **special need** (more ram, longer lifetime, direct readonly MySQL or PostgreSQL access)

## Monitoring, cluster management, accounting

- **raw data** mostly collected from LSF
- used by several custom management tools (MachineJobFeatures, multicore provisioning, host failure management . . . )
- **custom accounting**: batch and blah records of done jobs are `INSERTED INTO` a PostgreSQL DB. SQL join to create grid records. Custom script to generate apel records, delivery through ssm send
- **concern**: get equivalent data from HTCondor instead

# Monitoring



# Accounting

Subject **[CNAF farming] HS06 Farm power: 360828.10 HS06 Wall, 297111.99 HS06 cpt used from 2018-04-18 to 2018-05-18**

To [farming@lists.cnaf.infn.it](mailto:farming@lists.cnaf.infn.it)★

| Queue         | NJobs   | AVG HS06 Wall | AVG HS06 CPT | Efficiency |
|---------------|---------|---------------|--------------|------------|
| lhcb          | 390916  | 89879.53      | 87544.11     | 0.97       |
| alice         | 230894  | 64319.97      | 56902.25     | 0.88       |
| cms_mcn       | 7401321 | 46069.46      | 32335.02     | 0.70       |
| atlas7        | 233686  | 32204.41      | 27630.10     | 0.86       |
| cms_mcore     | 7021983 | 23204.19      | 18513.52     | 0.80       |
| mcore7        | 251840  | 15643.79      | 10761.71     | 0.69       |
| T3_B0         | 417     | 15623.09      | 0.00         | 0.00       |
| atlas         | 480324  | 12994.98      | 13601.41     | 1.05       |
| virgo         | 85889   | 11591.92      | 5837.43      | 0.50       |
| atlas_bari    | 72929   | 10093.98      | 9584.10      | 0.95       |
| mcore         | 190024  | 8711.92       | 6788.71      | 0.78       |
| borexino_prod | 238743  | 7673.04       | 7397.39      | 0.96       |
| cta           | 63194   | 5115.39       | 5115.51      | 1.00       |

# Reporting

Subject [CNAF farming] LSF REPORT: 11/206 new down WN, 5890/31712 slots, 342580 = 403708 - 61128  
HS06, 4248/29928 CORES

To farming@lists.cnaf.infn.it

## Problem WN LIST

```
wn-201-04-39-02-a: 16,0x20004: closed_Admin,eadmin: JOB EXIT RATE THRESHOLD EXCEEDED
wn-201-04-07-04-a: 16,0x20004: closed_Admin,eadmin: JOB EXIT RATE THRESHOLD EXCEEDED
cn-610-01-01: 72,0x20144: Unknown,SDP: NO IPMI
wn-200-10-31-03-a: 32,0x20004: closed_Admin,eadmin: JOB EXIT RATE THRESHOLD EXCEEDED
wn-201-04-13-02-a: 16,0x20004: closed_Admin,eadmin: JOB EXIT RATE THRESHOLD EXCEEDED
wn-200-08-31-02-a: 16,0x20014: Unknown,eadmin: JOB EXIT RATE THRESHOLD EXCEEDED
wn-200-08-31-04-a: 16,0x20004: closed_Admin,eadmin: JOB EXIT RATE THRESHOLD EXCEEDED
wn-200-13-21-13-a: 32,0x20014: Unknown,eadmin: JOB EXIT RATE THRESHOLD EXCEEDED
```

```
+-----+-----+
| 2018-05-20 | Summary |
+-----+-----+
| okjobs      | 84727   |
| failjobs    | 3038    |
| ok_cpu      | 22036 days |
| fail_cpu    | 1838 days |
| ok_wall     | 24122 days |
| fail_wall   | 1110 days |
| cpuratio    | 92.50%  |
| wallratio   | 93.46%  |
| jobratio    | 94.38%  |
+-----+-----+
```



# Inventory database

The screenshot displays the INFN Doce inventory database interface, divided into several sections:

- Computers & Others:** A table listing system names and hostnames.
 

| NAME         | HOSTNAME     |
|--------------|--------------|
| cn-610-01-08 | cn-610-01-08 |
| cn-610-01-09 | cn-610-01-09 |
| cn-610-01-04 | cn-610-01-04 |
| cn-610-01-05 | cn-610-01-05 |
| cn-610-01-06 | cn-610-01-06 |
| cn-610-01-07 | cn-610-01-07 |
| cn-610-01-02 | cn-610-01-02 |
- Network Ports:** A table listing network interface names and MAC addresses.
 

| NAME              | MAC               |
|-------------------|-------------------|
| cn-610-01-04      | A0:36:9F:E2:C1:AE |
| mgng-cn-610-01-04 | 40:12:e9:c6:bd:e0 |
- Contained Devices:** A table listing devices contained within the system.
 

| POS. | DEVICE CONTAINED |
|------|------------------|
|      |                  |
- Hardware Configurations:** A detailed view for a specific device, identified as a **Lenovo NeXTScale nx360 M5** blade server.
 


| HARDWARE TYPE | NAME                      | LOCATION |
|---------------|---------------------------|----------|
| Blade         | Lenovo NeXTScale nx360 M5 | 0        |
- Components Part:** A table listing the hardware components and their quantities.
 

| COMPONENT                | QUANTITY |
|--------------------------|----------|
| Xeon CPU E5-2697 v4      | 2        |
| DDR4 SDRAM 16384MB       | 16       |
| HD 2TB SATA ST2000NX0253 | 2        |
| RAM                      | 16       |
| HS06_HTON                | 832      |
- Image Upload:** A sidebar on the right shows an "Upload Image" button and a thumbnail of a server blade with red arrows pointing to a "USB hypervisor BIOS socket" and a "PCIe slot".


## Multicore provisioning

- At any time, WN accepts single core only or multicore only.
- A multicore only node move back to singlecore when free for more than  $\sim 15$  min.
- Used to be good with older machines (16, 24 or 32 slots)
- Current WNs are up to 72 slots. It would be better having multi and single core jobs on the same node.

### July 2016, 7 days




### May 2018, 20 days



## Issues with short jobs

### Short job flooding



### LSF fairshare and short jobs

- User x start a flow of short (few seconds) jobs.
- The submission tool wants to have, say, 300 pending
- User x has high priority → LSF dispatches many jobs
- jobs exit immediately → priority remain high

# Preparing to move

## Planning a migration



## Early considerations

### Migrating to HTCondor

- **PRO:** all of our use cases are covered already with HTCondor
- **PRO:** several HEP sites already there (wider knowledge base)
- **Effort needed** to adapt our management tools to accept data retrieved from HTCondor instead of LSF
- **Effort needed** to learn and master HTCondor–CE
- **Effort needed** to support our local users to adapt their submission tools

## Hatching a HTC instance

### Current HTC test instance

A small HTC test instance is currently in use

- One CM + Collector
- two submit nodes
- a few WNs
- One HTCondor-CE
- mixed SL6, Cos7

### Envisioned future



- Test Grid job running through HTC-CE (dteam VO)
- next: add real VOs (CMS, ATLAS, ...)
- hardening (HA setup)

## Hatching a HTC instance (2)

### Problems to prevent

- Wrong strategy/instrument for cluster administration
- being unable to troubleshoot problems
- slow cluster, hidden bottlenecks

- 1 Follow good practices!
- 2 when in doubt, ask
- 3 destroy test instance first

### Unwanted future



## L\$F migration strategy

### At first

- stay small (few Exec Nodes only)
- begin with Grid users (more standard requirements)
- support local users converting `bsub` → `condor_submit`
- adapt (or drop) custom management tools

### Second stage

- HTCondor hardening (HA setup)
- add more schedd/CEs
- move more WN to HTCondor
- Production



## Concerns

### re-configurations

- One LSF master configured to “rule them all” vs a puppet master to “re-configure them all”.

### Getting confident

- still need to go on with training
  - need to reach a “know what to do” state in case of incident
  - fallback: “know who to ask”
- 
- Disclaimer: several (maybe all) of my concern and worries already found answers during this event. I just need to work as advised.

## Credits

Several people have provided useful advices, examples and pointers.

To make a long list short:

Thanks to All!