

Attendees:

Local: Miguel (ALICE), Ulrich (IT/PES), Stefan, Maarten (ALICE),

Remote: Dirk (CMS), Antonio (CMS), Jan Just (NIKHEF), Manfred (FZK-LCG2)

Presentation (see slides attached to the agenda page)

===

Slide 2 (Status):

-----

Question: Is it OK for experiments to give feedback on the CERN deployment by end of February? The idea is to ask experiments to test the mjf API and the features information provided for completeness

Maarten: For Alice there is no imminent use case, but don't see a problem to test it by end Feb.

Dirk: For CMS this should be OK, for anything else concerning integration into the pilot system need to talk to CMS pilot developers

Stefan: For LHCb should be ok, will clarify with the LHCbDirac development team

Ulrich: How should the response be given?

Stefan: Let's discuss it on the MJF mailing list

Discussing the "problem" of concerns of overloading the OpenStack meta-service:

Jan to my understanding the meta service never leaves the hypervisor, its running local, so we should not be able to overload it.

Ulrich: to my knowledge we talk to a central service via nova.

-->>> We need to clarify this point.

Slide 4 (mjf example):

-----

Stefan: Is this sufficient for experiments to test mjf tool?

Maarten: Yes should be ok

Ulrich: In addition to the mjf python class also the /usr/bin/mjf command line tool has been deployed on all worker nodes

Slide 5 (Next Steps):

-----

- Deployment of mjf as command line tool @ CERN

Done, as discussed above

- Upload of different server side tools into the github repository

Ulrich: For CERN implementation could be of limited use b/c of very site specific implementation

Stefan: Please do it anyway, it may give hints to other LSF sites on how to implement it (pls also document what is site specific)

- Development of a key/value store for IaaS

Stefan: Idea is to have a simple nosql implementation (e.g. couchdb) with a web service in front, this then can be given to sites who wish to use it, otherwise they can also implement their own solution

Ulrich: might be good to go back to the "user stories", i.e. what do we need MJF on clouds for? I.e. power of the machine (can be provided by the VO) and shutdown time (needs to come from external source)

Stefan: Yes, good idea, but we need some communication site->VM at least for the shutdown time

Ulrich: In addition we should have the possibility to provide the deviation width for the power of a certain "VM type" on different sites. The mean value could be used for accounting, the "worst case" value could be used for calculating the job durations. We shall have a possibility to gather this information

Stefan/Maarten: good idea but on the boundary of the mandate of this task force.

Stefan: How long does it take to run a HepSpec benchmark?

Ulrich: Several hours

Stefan: We would need some easier way of calculating the machine power if we want to have many data points.

Jan: let's not forget with benchmarking that there are many more infos coming in.

E.g. disk, data access, threading which would not be seen in a HepSpec benchmark.

Jan: this power measurement should be kept as simple as possible, also to ease the work of the site admins.

Ulrich: for standard HepSpec benchmark such scripts exist.

Stefan: I regard this slide as our plan (also in priorities) for the next steps until the end of the task force. Did I forget any point? (no objection)

Stefan: Once we are at the bottom of this slide the TF can be closed :-)

Aob

===

- Need to get feedback from ATLAS concerning the testing of mjf at CERN by end of February

- Need to find a new slot for the TF meeting as Igor (in San Diego?) has joined, the earliest time possible is 5pm CET -> please fill in doodle <http://doodle.com/myifgsmddwdmqvr8>

- Clarify whether the meta-service in openstack can be run on a hypervisor level

The next meeting will be called latest by end of February or on request on the new slot decided from the doodle poll