

GridPP Cloud Meeting

Agenda: <https://indico.cern.ch/conferenceDisplay.py?confId=218642>

Participants: David Colling (Chair), Daniela Bauer, Duncan Ran, Simon Fayer, Adam Huffmann, Ian Collier, Raja, Gareth Roy, Andrew Washbrook, John Green, Roger Jones, David Wallom, Peter Love, Matt Doidge, Chris Brew, anonymous?, Chris Walker, Kashif Mohammad, Pete Gronbech, Jeremy Coles (notes).

Introduction & Composition of Overview Body (David)

Slide 2. Welcome to new activity. GridPP cannot ignore clouds. In practice already ship good part of OS with job. Step in right direction with VM. Funding agencies expect us to have a perspective.

Slide 3: Commercial clouds not always appropriate. Scheduling. Data storage and transport – currently assumes xrootd. Wider academic world has important things to offer. *Modest resources and no new manpower!*

Slide 4: ToR attached to agenda.

Roger: Could say clouds are too expensive. Need to be clear what we mean. Using cloud technology is not the same as using (commercial) clouds.

Fortnightly <1hr meeting suggested. Meetings should be open.

Steering group suggested to make sure all areas remain represented.

ATLAS (Roger)

ATLAS has quite a few activities ongoing. Disparate projects. Agreed internally that it is time for more coherence. Will be a recognized service area. Will be pushed in jamboree next week. Some work with HelixNebula looking at MC on cloud. Explicit work has been in BNL/US using OpenStack and integrating with Panda. Experiments need to decide and align on technologies since we need to do this in common.

There is a discussion as next year the HLT farms at CERN are available. The default was to add them to T0, but the idea of a cloud implementation is gaining interest. Precise details after the meeting next week. There would be UK people who Alexi would like to see involved. Xrootd is also of interest.

ATLAS work has so far been side projects led by enthusiasts.

DW: Mentioned Nebula. What is the importance?

RJ: Operation does not depend on it. It is a proof of concept for a restricted use-case. CERN IT may be more keen than ATLAS at the moment.

PL: Getting ATLAS jobs running on Nebula had similar issues to getting stuff running on OpenStack so useful.

RJ: HC has also been imported. Adapted and can be used on the cloud. In EIS for HLT work need to find out contact (via Hans). Two weeks before I can speak to people directly will try to find out.

DC: CMS is doing something very similar.

CMS (Andrew)

Reasons to do it: quick access to resources, images in common which are more easily deployable... influence cloud directions. A goal needs to be sizeable increases in resources made possible.

A number of existing tests using StratusLab; LxCloud; Amazon; HLT

DC: UK is involved and getting more heavily involved.

LHCb (Raja)

There are intentions and plans in LHCb. Some work presented at CHEP.

DC: How does VMdirac work? Does it talk to EC2 frontend and start a VM and how then do they communicate with the scheduler and DIRAC?

RN: VMs send regular heartbeats back to DIRAC. Only done for simulation not processing.

DC: Please email slide. Who in the UK is involved?

RN: Nobody directly involved right now. Primarily France, Spain and Romania.

Relationship with others (David)

DW: Wide ranging discussion. Important points. Essentially there is limited funding and effort, so whatever we do we need to optimize. Leverage other activities or other relationships directly. Commercial providers engaged via HelixNebula with limited tests done by ATLAS. Commercial clouds also seen as very expensive. Work we are doing is looking at GridPP5 forward vision. Many providers have changed business models very quickly. Amazon halved costs through data transfer free at point of use change. So when documenting, be sure we are as up-to-date as possible.

Decent relationship within NGI with leader of Amazon effort in Europe.

Presented at GridPP meeting, a lot of effort already around general federated work that we can benefit from. Participating institutes are in both WLCG and EGI. Nationally need to link in as much as possible with those European activities.

Main concern, once up and running we should be engaging more as we can lead the activity.

DC: I agree. When up and running we should get more involved. Have had informal chats with those in EGI federated clouds and HelixNebula.

DW: There will be a stronger future relationship between HelixNebula and EGI. Boxes provided ... pushing providers to a common subset of interfaces. Open standards are the way forward. CERN presented at SWING meeting in Bern 2 weeks ago about the effort they will put into OpenStack. That has many people already bought in.. Pleased to see the CMS presentation looking at OpenStack. A lot of sharing can be done and there are deployment modules.

DC: We have OpenStack on HLT.

IC: When people say we should be involved we are. RAL is a StratusLab cloud contributor. Our community provided the HEPiX WG on sharing images (<http://w3.hepik.org/virtualization/>) , which is central to this work, adopted by EGI and now Federated CTF looking at it.

DW: We can be more active by supplying of use-cases. Currently 6 communities. Bio, musicology, space science.... Getting HEP in there would be useful especially with participation. CMS use-case with OpenStack is an example.

DC: Within CMS easy to get document together on use-cases. Interested to hear more on HEPiX work. (*Action?*)

IC: 3 years ago HEPiX did work on Virtualising WNs. Idea was VM images for WNs could be produced in one place and used at another, and there needed to be a trust model. That provided the basis for an EGI security policy on sharing images. In doing this work, we need to be aware of that framework for endorsing and revoking images. The work is done and dusted.

DC: Would be useful to set up a wiki with links to work already going on or completed. Will do this after the meeting (*Action*).

Sites (Ian)

Sites have drivers and reasons to investigate cloud technologies (e.g. CERN's agile infrastructure). Within GridPP we should be looking at this. At RAL developments with config management and Virtualisation of services shows cloud technologies may have a role in managing infrastructure. Also other use-cases within STFC community. Scientific community... Andrew has been using cluster to do tests for CMS. Could look at capacity provision being rolled into this if working efficiently. We are close to being in a position where we can provide

capacity provision. That makes it easier to use other resources that become available to us. That's the driver in other places. If we are in the Federated Cloud we get potential access to other resources. Our community is organized enough to allow it.

One of the site perspectives, we may have our own infrastructure reasons to investigate this. At other sites cloud activities are taking place – e.g. Oxford (OpenStack not yet in T2), don't know what other sites are actively doing but we may produce blueprints for overlays on present infrastructure.

DC: We are coming up to LS1. Much will change in computing models. If we show this is a reliable solution then many sites will go in this direction.

IC: Not sure how easy it will be at sites to make major changes. During LS1 we don't expect things to stop. What it does mean is that the HLT work will be a significant change.

DC: LS1 is a 'post paper writing phase'; it gives an opportunity to experiment a little more. Will look at parked data and reanalyze things but it will be less intense. Easier than post LS1.

KM: More information on Oxford setup. We have small OpenStack setup with 20 Dell machines. Does not belong to GridPP - it was done by OeRC and Oxford SC. Interested as CMS have setup to send jobs to Cloud with appropriate API. Would be easy for us to participate. We have an EC2 and No? API.

DC: Andrew, have you been submitting jobs with glideIN WMS?

AL: No. Experimented with CREAM CE with Condor .. manager to create WNs when needed. Eg. Submit 4000 jobs and that results in the creation of WNs to accommodate jobs. Other batch systems have had to do clever things to get this to work but Condor makes it easier.

DC: Are you starting the machines by hand or the CE?

AL: Script checks status and if it sees jobs waiting creates machines.

AW: ECDF have pilot service on OpenStack. Would like to use this in GridPP work.

PL: These instances, are they open?

AW: Not at the moment but it is part of the work to enable it.

DW: Oxford one will not be open to production work, system is not enabled for production services.

PL: Testing various aspects would be the main aim.

IC: At RAL keen to get people involved in testing. Will add resources if use-case is seen to be a real need. For testing work, a connection is needed with someone internal. But intention is to have a user-automated route.

DC: Access to resources and security model is an important aspect of work.

DW: We would look at ours as being a cloud. Exists outside of University firewall. Where we give people accounts we would expect them to do as they can on a commercial cloud. Those doing this already will have come across the issue. NGS pilot Ox-ECDF, the conditions of connections ... arbitrary access to people at root level is not allowed with current policies.

DC: A lot of negotiation is needed there.

DW: Steve Thorn's experiences in this are worth following up.

Equipment at IC (Adam)

GridPP have invested around £10,000 for testbed resources. Adam gave an overview. Specifics in document <https://indico.cern.ch/materialDisplay.py?contribId=0&materialId=0&confId=218642>.

It will be another test infrastructure. Will be configured as people need for tests.

IC: DW was talking about OpenStack instance at Oxford. At RAL for now we will run two clouds, the internal one Andrew is using and a public one which is being put in to work as part of federated infrastructures. Need to think about whether resource connects to site or has public connections. (i.e. in DMZ).

Free Range discussion:

CW: Questions and comments to feed in. The VO motivation seems to be to expand to accommodate peak demand, and the site motivation is ease of deployment. There is a lot of potential in this. There have been some GDB talks about a grid of clouds. If firing up many nodes that is what you would need to do. This would allow the leveraging of additional resources and access to other communities. If we (QMUL) try it, do we go with Nebula or OpenStack? In Nova community talk there was a machine that plugs into a rack and provides a pluggable option.

DW: One of the things this community can do is learn what others are doing. E.g. look at work done at CERN for automating openStack implementation.... Many opportunities. Don't put all eggs in one basket. Agree different groups will try different things.

IC: Where do you try running jobs.... A user will try anywhere they can get at. We should experiment with different infrastructures. Until recently OS was harder to

get working, but it is at a tipping point. StratusLab on top of ON is almost working out of the box.

DC: Try OpenNebula in QMUL. IC tries OpenStack – matches HLT work.

CW: Part of me says, if going to make an impact, then if we all go different ways then there will be a lot of wasted effort.

DC: I don't think we are at the stage where that is actually the case. I think we benefit more from diverse experience at this stage.

IC: The other thing, we must benefit from the work others have done. A resource agnostic approach is being pursued by the federated cloud group. Job is to build shared interfaces. Work to make sure images contextualise properly. HEPiX WG had early on could do images on Xen etc. They are solved problems. We do not have to and should not reproduce that work.

DW: The way to install a cloud is no longer as important as how you use your cloud. The important thing is to get 100% out of that cloud.

DC: What next?

CW: xrootd was mentioned a few times. I feel I should mention WebDav. Missed WebDav for http meeting on Tuesday, but it holds promise.

DC: There is a lot of activity going on at a low level.

DW: Those sites that have clouds or are thinking about it, should join federated cloud work as providers.

DC: We cannot force sites to do it. But sites can do it and we should encourage it.

DW: We/they will get a lot of expertise back quickly on how to set things up.

DC: I have some CMS quite specific ideas. Adam to setup OpenStack here at IC. What tests or things would people like to do on the resources here?

RJ: For ATLAS I will email ideas in a week or two, but it is true that we are not queuing up at the moment.

DC: We will have another meeting before Christmas. Before then we will setup a twiki links to other projects). Setup resources at IC with openStack with CMS tests started. Then in a few weeks meet to discuss ways forward.

JC: It will be useful then to get the ATLAS activities input from Roger.

CW: Please post the email list details.

DC: Will start doodle poll for next meeting (*action*)

AOB: None.