

Table of Contents

Virtualization technologies.....	1
Introduction.....	1
Virtualization of service nodes.....	1
Virtualization of non-reliable resources.....	1
Relationship with clouds.....	2
Recommendations:.....	2

Virtualization technologies

[draft version 0.1]

Introduction

Virtualization was used in the past for expensive main frames. With upcoming multi-core CPUs in the PC world it is becoming interesting again. Several solutions exist, and are mature enough to be interesting for sites.

Virtualization is used at many sites already, partly for many years. The advantages of due to migration and life-migration, pay load encapsulation, and resource usage optimization often pay off the (small) penalties due to the virtualization overhead.

Virtualization of service nodes

Many services which sites have to provide to the experiments have little need for CPU power nor large requirements for disk space and memory. On the other hand, physical CPUs come with more and more cores, along with more memory and often disk space. For the sake of good resource usage virtualization of such resources is a good practice. If combined with life migration capabilities, operations of the resources become easy because intrusive interventions on the hardware (disk exchanges, memory failures etc) can be dealt with often without the users noticing.

Virtualization of services is used at many sites for a long time.

Virtualization of non-reliable resources

Less obvious is how much sense it makes to virtualize unreliable resources, such as cheap batch worker nodes. It has been tried by some sites though.

- while virtualization penalties for services often don't really matter due to the low requirements of such services, this is different for "number crunchers". A virtualized batch farm must therefore be optimized for efficiency.
- as in the service consolidation case, life migration of batch nodes can make the the life of the operations team easier. Possible caveats include:
 - ◆ it may not be practible in all cases though. In case batch worker nodes are only running short payload, it may be faster to drain the batch node and reinstantiate it somewhere else.
 - ◆ the scratch space needs to be either shared or has to be copied over to the new location in case of a life migration request. This can take a while, and will use a significant amount of resources eg for network bandwidth
 - ◆ there must be enough "free" resources to which batch nodes and be migrated.
 - ◆ it is not clear if all jobs would survive such a migration. This needs testing
- Further investigations are required to judge if a static virtualization of batch resources is useful or not. If the setup is made in a

more dynamic way though, additional benefits are possible for the site:

- - ◆ automate the deployment of intrusive updates
 - ◆ encapsulation of user jobs. A job which triggers a kernel bug will only kill the virtual machine on which this actual job is running.

- ◆ chop down of large nodes with many cores so that they can be used by multithreaded jobs which only scale to a lower number of CPUs.
 - ◆ easier draining of worker nodes if the number of simultaneously running jobs per VM is limited
- a clever virtual machine provisioning system will be able to adapt the batch farm to incoming requests, with implications on the job turn around time, and resource usage patterns

Relationship with clouds

Virtualization and clouds are two different things, which must not be mixed up.

Recommendations:

- virtualization is a decision of the site
- sites should ensure that the virtualization penalty is lower than the gain in TCO

-- UlrichSchwickerath - 24-Jan-2012

This topic: LCG > WLCGWIMgmtTEGVirtualization

Topic revision: r1 - 2012-01-24 - UlrichSchwickerath



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](#)