



Contribution ID: 93

Type: Poster

Virtual Expansion of a Departmental Cluster—Two MetaCentrum Use Cases

Monday, 11 April 2011 09:00 (8 hours)

Overview

In our previously published work, we described usage of virtualization technology in MetaCentrum, the Czech NGL. This poster describes a use case important for many of our user groups who perform most of their computations on their own local resources (e.g., a departmental cluster) and who want to transparently enlarge the resources by adding grid machines in cases of peak computation power requirements. We describe the advantages of the scenario and two prototype implementations.

Impact

Providing grid resources in the form of virtual nodes of existing clusters bring significant advantages for end users. It leverages existing know-how and eliminates a gap the users must cover when they want to use MetaCentrum by learning the MetaCentrum *modus operandi*.

Operating system type and version, batch system type and configuration, software installations, data storage and other conventions used by its local administrator can remain untouched. Users also prefer interaction with the local administrator (ideally dropping in a neighbouring office) compared to remote interaction with the MetaCentrum user support unit. The scientific expertise of the users is frequently firmly bound to local configuration and tools. Typical local administrator is a scientist doing this task as a part-time job, having the application expertise, and able to tune the environment to local specific needs.

The local cluster expansion setup allows to follow local security and networking policies. The cluster is still hidden behind the local firewall, users can still use software licenses bound to their IP segment or institution. We can also continue running technologies not suitable for large distributed environment, e.g., storage doesn't need to be scalable and security practices and protocols may be quite loose (e.g., ID based NFS authentication).

Description of the work

The solution is based on a MetaCentrum service called virtual cluster. Virtual cluster a set of virtual machines connected by a dedicated virtual network (VLAN). The VLAN is understood as an ordinary resource available to grid scheduler to be assigned to users. The VLANs can be deployed on a state-wide geographical area with a small overhead. We use advanced services of the CESNET2 (Czech NREN) network to connect MetaCentrum sites.

The idea of virtual expansion of a cluster is simple—we create a virtual cluster and connect it with the existing cluster on network layer. The additional virtual nodes run user-supplied OS image, behaving identically with the original departmental cluster.

We have implemented the first prototype on a departmental cluster managed by the New Technology Center (NTC), University of West Bohemia, Pilsen. The NTC uses a 12-node cluster connected on a dedicated local IP segment behind a firewall. The cluster runs local NFS, non-scalable security and system configuration and it is managed by one of its users. The node image was able to be directly transformed into a virtual cluster image. Proximity of a CESNET2 PoP allows direct connection of the IP segment of the cluster to the VLAN of the virtual cluster. New virtual nodes of the cluster (running elsewhere on the distributed MetaCentrum infrastructure) were added to the local batch system (Sun Grid Engine) and are transparently available to cluster users.

The other prototype implementation differed in two important aspects. First, it is based on collaboration with a non-academic subject (KitD, a video processing company). CESNET2 network is not therefore allowed to be used, so the network must have been tunnelled through NIX.CZ, the Czech peering node. Second, the cluster runs MS Windows. We demonstrated that MetaCentrum resources were transparently available regardless of their location and virtual node OS.

URL

<http://www.metacentrum.cz/>

Conclusions

The use cases described above represent a promising way to provide a significant part of MetaCentrum users with better services. It is intended for groups having their own resources and using grid services just on a time-to-time basis—for exceptionally resource demanding projects only. We can typically reach the state where using a MetaCentrum resources is completely hidden behind a local batch system and/or configuration tasks of the local administrator. This is not only a way to make access to grid resources easier but also how to increase efficiency of our users (who can focus on the scientific essence of a problem) and also MetaCentrum user support department who can provide services and support in synergy with the local administrator.

Primary authors: ANTOŠ, David (CESNET); SITERA, Jiri (CESNET); MATYSKA, Luděk (CESNET); RUDA, Miroslav (CESNET); HOLUB, Petr (CESNET)

Presenters: SITERA, Jiri (CESNET); RUDA, Miroslav (CESNET)

Session Classification: Posters

Track Classification: Poster