

Collaborative Project in Cloud Computing

Project Title: Virtual Cloud Laboratory (VCL) Services for Education

Transfer of Results: from North Carolina State University (NCSU) to IBM

Technical Description:

1. Types of services:

From the end-user perspective, VCL offers a series of services ranging from: single-seat desktop-type offerings with access to either routine or specialized computational resources and applications, to:

- groups of seats that can be reserved for a particular timeslot, to
- long-term reservation of one or more servers, to
- reservation of homogeneous or heterogeneous aggregates (or bundles) of resources called **environments**, which are the building blocks of virtual clouds that VCL supports, to
- long-term reservation of research clusters, to
- HPC cluster and facilities, and so on.

A **seat** can represent either access to a sole-use virtual or bare-metal resource (e.g., a Windows or Linux server) that in turn can be used on its own, or as a management node for a group of services or resources, or it can represent access/portal to an already preconfigured shared service or resource (e.g., access to an account on an IBM System z machine, or access to HPC clusters via a sole-use or shared-login node).

One of the primary characteristics of VCL is that it can dynamically change its configuration and move resources from one type of service to another. One can distinguish two major groups of service categories:

1. HPC services: consists of access to cluster-based resources controlled through an HPC scheduler such as Platform LSF (load-sharing facility), to shared memory resources, and in special cases to supercomputers.
2. Other services include:

- Access to single-node bare-metal or virtual computers; typically in the desktop mode, these resources come to an end-user with administrative privileges (root access)
- Ability to make a reservation for a group of desktops for a particular timeslot (e.g., for use in a class during a class period or a bank or office during working hours)
- Ability to reserve one or more (bare-metal or virtual) servers
- Ability to reserve aggregations of tightly coupled or loosely coupled computers and servers either of the same type (e.g., blades for a computational test cluster) or an integrated set of diverse components (e.g., a heterogeneous aggregate might consist of a Web server, a database server, a System z resource running applications, and a cell cluster to perform some related analytics).
- Ability to use VCL to request a portal into HPC cluster resources and submit batch or real-time jobs to such an environment.

2. Types of resources:

We distinguish between two types of resources, **undifferentiated** and **differentiated**.

1. **Undifferentiated resources** are those that can be reconfigured and reloaded at will with whatever suits the end-user:
 - A group of blades that can be loaded from scratch (bare metal) with Linux, Windows, or some other OS and applications on short notice represent undifferentiated resources.
 - A group of servers that is already loaded with a hypervisor (e.g., VMware ESX) and can receive any virtual image of choice represents an undifferentiated resource for that type of virtual image. When a user is finished using the resource, it is again returned to the pool of undifferentiated resources.
2. **Differentiated resources** are preconfigured but can be made available to the end-user at will or on schedule.
 - A group of machines that may be located in a university computing laboratory and that are made available to users over the network when the laboratory is closed may be classified as differentiated if the end-user does

not have the right to reload them at will, that is, can use them only in the already configured (differentiated) state.

- Similarly, access to a logical partition (LPAR) or an account on an IBM System z resource may be considered as access to an already differentiated resource. In this case, when the user is finished with the resource, it is returned to the pool but remains differentiated.
- Any Web services offered through VCL-affiliated resources fall into the same category.

VCL offers the ability to incorporate new hardware with ease, to administer metadata about the hardware, and the hardware itself (reload, default image, etc.). Hardware can be grouped into logical groups that may reflect its properties (interconnectivity, processor type, and ownership), and images can be mapped onto any hardware component, some hardware groups (e.g., those tightly coupled through a high-speed interconnect), & assigned to specific user groups.

3. VCL architecture: basic VCL network configurations

One of the key features of the undifferentiated VCL resources is their **networking setup**. It allows for secure dynamic reconfiguration, loading of images, and isolation of individual images and groups of images. Every undifferentiated resource is required to have at least two networking interfaces:

- one on a private network,
- the other on either a public or a private network depending on the mode in which the resource operates.

For full functionality, undifferentiated resources must have a way of managing the hardware state through an external channel, e.g., through the IBM BladeCenter chassis **management module** (MM).

3.1. The dual-use physical network configuration

Typically, the eth0 interface of a blade is connected to a private network (10.1 subnet in the example), which is used to load images. Out-of-band management of the blades (e.g., power cycling) is effected through the management network (172.30.1 in the example) connected to the MM interface. The public network is typically connected to the eth1 interface. The VCL node manager (which could be one of the blades in the cluster or an external computer) at the VCL site has

access to all three links; that is, it needs to have three network (public, management and image-loading) interfaces (Fig. 1).




Fig. 1 VCL dual-use physical network configuration

This VCL configuration can be used for environments where seats and services are assigned individually or in synchronized groups, or when assigning or constructing an end-user image aggregate or environment where every node can be accessed from a public network.

3.2. VCL high-performance computing physical network setup

Now, the node manager is still connected to all three networks - *public*, *management* and *image-loading*, and preparation private network-but now eth1 is connected (through VLAN manipulation, VLAN 5) to what has now become a **Message Passing Interface** (MPI) network switch. This network now carries intra-cluster communications needed to result in tightly coupled computing usually given to the HPC cloud services. Switching between non-HPC mode and HPC mode takes place electronically, through VLAN manipulation and table entries; the actual physical setup does not change. We use two different VLANs to eth1 to separate public network (external) access to individual blades when those are in individual-seat mode (VLAN 3), from the MPI communications network to the same blade when it is in HPC mode (VLAN 5 in configuration 2), see Fig. 2.




Fig. 2 VCL high-performance computing physical network setup

This second VCL architecture has a configuration with blades assigned to a tightly coupled HPC cluster environment or to a large overlay (virtual) cloud that has relatively centralized public access and computational management.

Benefits for NCSU

NC State's intelligent provisioning system also improves its ability to flexibly allocate resources between instructional, administrative and research activities, each of which has its variations of resource requirements.

A key example is the near complete drop-off of student computing activity that is typical during breaks between semesters. The new system gives NC State the ability to quickly and easily switch the bulk of its roughly 1,000 IBM BladeCenter server blades to the computationally intense requirements of researchers—such as running complex models and simulations—and in the process leveraging what would have been idle server capacity to advance the goals of the university's researchers. The same capability would enable NC State to shift capacity to administrative functions like class registration that produce a surge in processing activity before each academic semester.

All of these benefits point to how intelligent provisioning and similar "cloud" initiatives effect the more granular optimization of computing resources, which

enables NC State to handle the academic computing requirements of a growing student population while minimizing the growth of its infrastructure.

NCSU's goal was to rethink the way it met the academic computing needs of students, instructors and the other populations the university serves. By collaborating with IBM and transferring R&D results, NCSU was better able to deliver on that mission.

VCL Project Partners:

- NCSU: North Carolina State University, US
- IBM US