3.8.1 PHYSICAL VERSUS VIRTUAL CLUSTERS

Virtual clusters are built with VMs installed at distributed servers from one or more physical clusters. The VMs in a virtual cluster are interconnected logically by a virtual network across several physical networks. Figure 3.8 illustrates the concepts of virtual clusters and physical clusters. Each virtual cluster is formed with physical machines or a VM hosted by multiple physical clusters. The virtual cluster boundaries are shown as distinct boundaries.

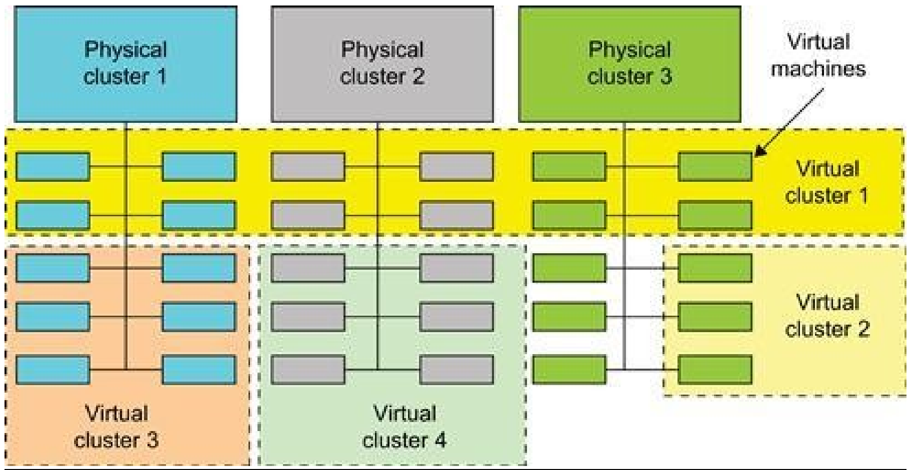

Figure 3.8 A cloud platform with four virtual clusters over three physical clusters

3.8.2 Properties of Virtual Cluster:

- The virtual cluster nodes can be either physical or virtual machines. Multiple VMs running with different OS can be deployed on the same physical node.
- A VM runs with a guest OS, which is often different from the host OS, that manages the resources in the physical machine, where the VM is implemented.
- The purpose of using VMs is to consolidate multiple functionalities on the same server. This will greatly enhance server utilization and application flexibility.
- VMs can be colonized (replicated) in multiple servers for the purpose of promoting distributed parallelism, fault tolerance, and disaster recovery.

Fast Deployment and Effective Scheduling

- The system should have the capability of fast deployment.
- The concept of “green computing” has attracted much attention recently.
- However, previous approaches have focused on saving the energy cost of components in a single workstation without a global vision.
- The live migration of VMs allows workloads of one node to transfer to another node. However, it does not guarantee that VMs can randomly migrate among themselves.
- In fact, the potential overhead caused by live migrations of VMs cannot be ignored.
- The overhead may have serious negative effects on cluster utilization, throughput, and QoS issues.
- Therefore, the challenge is to determine how to design migration strategiesto implement green computing without influencing the performance of clusters.
- Another advantage of virtualization is load balancing of applications in a virtual cluster. Load balancing can be achieved using the load index and frequency of user logins.

- The automatic scale-up and scale-down mechanism of a virtual cluster can be implemented based on this model.
- Consequently, we can increase the resource utilization of nodes and shorten the response time of systems. Mapping VMs onto the most appropriate physical node should promote performance.

High-Performance Virtual Storage

- The template VM can be distributed to several physical hosts in the cluster to customize the VMs.
- In addition, existing software packages reduce the time for customization as well as switching virtual environments.
- It is important to efficiently manage the disk spaces occupied by template software packages.
- Some storage architecture design can be applied to reduce duplicated blocks in a distributed file system of virtual clusters.
- Hash values are used to compare the contents of data blocks.
- Users have their own profiles which store the identification of the data blocks for corresponding VMs in a user-specific virtual cluster. New blocks are created when users modify the corresponding data.
- Newly created blocks are identified in the users' profiles.
- Basically, there are four steps to deploy a group of VMs onto a target cluster:
- Preparing the disk image, configuring the VMs, choosing the destination nodes, and executing the VM deployment command on every host.
- Many systems use templates to simplify the disk image preparation process.
- A template is a disk image that includes a preinstalled operating system with or without certain application software.

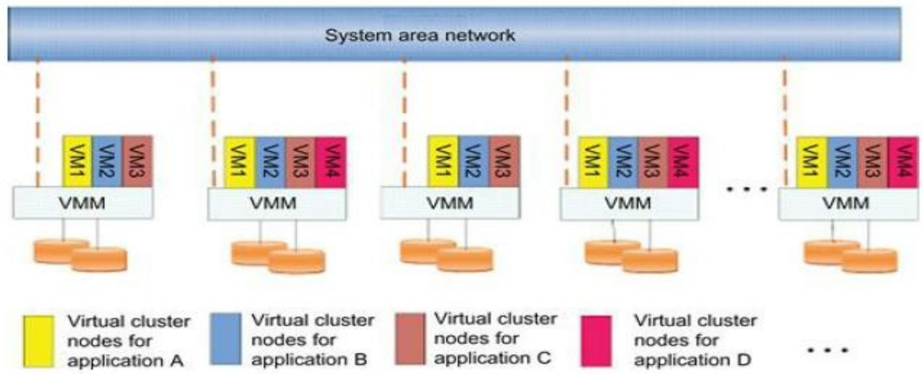

FIGURE 3.19 The concept of a virtual cluster based on application partitioning. Courtesy of Kang Chen, Tsinghua University 2008

Live VM Migration Steps and Performance Effects

- In a cluster built with mixed nodes of host and guest systems, the normal method of operation is to run everything on the physical machine.
- When a VM fails, its role could be replaced by another VM on a different node, as long as they both run with the same guestOS.
- In other words, a physical node can fail over to a VM on another host.
- This is different from physical-to-physical failover in a traditional physical cluster.
- The advantage is enhanced failover flexibility.
- The potential drawback is that a VM must stop playing its role if its residing host node fails. However, this problem can be mitigated with VM life migration.
- The migration copies the VM state file from the storage area to the host machine.

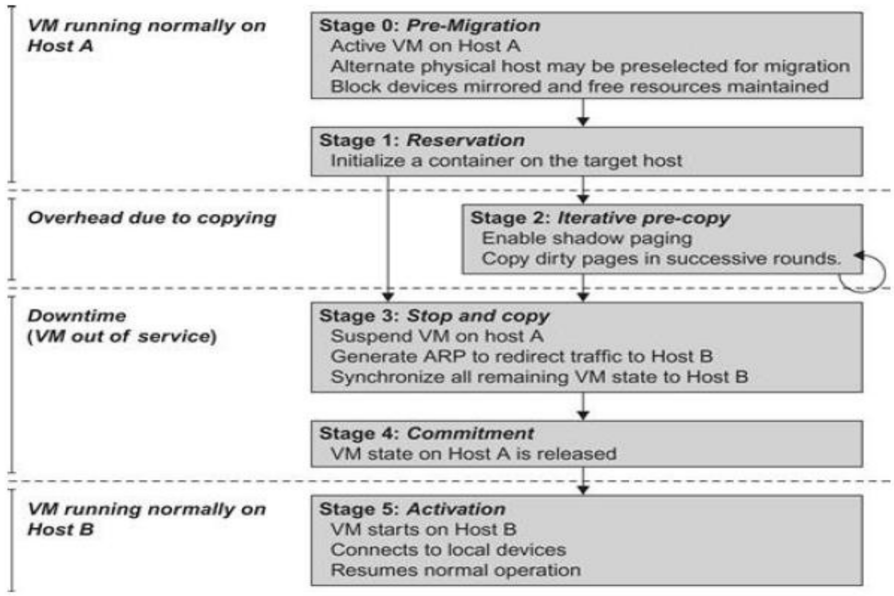

- There are four ways to manage a virtual cluster.
- First, use a guest-based manager, by which the cluster manager resides on a guest system. In this case, multipleVMs form a virtual cluster.
- For example, openMosix is an open source Linux cluster running different guest systems on top of the Xen hypervisor.
- Second, you can build a cluster manager on the host systems.
- The host-based manager supervises the guest systems and can restart the guest system on another physical machine. These two cluster management systems are either guest-only or host-only, but they do not mix.
- A third way to manage a virtual cluster is to use an independent cluster manager on both the host and guest systems.
- Finally, you can use an integrated cluster on the guest and host systems.This means the manager must be designed to distinguish between virtualized resources and physical resources.
- Various cluster management schemes can be greatly enhanced when VM life migration is enabled with minimal overhead.
- A VM enters the suspended state if its machine file and virtual resources are stored back to the disk. As shown in live migration of a VM consists of the following six steps:
- **Steps 0 and 1: Start migration.** This step makes preparations for the migration, including determining the migrating VM and the destination host.
- **Steps 2: Transfer memory.** Since the whole execution state of the VM is stored in memory, sending the VM's memory to the destination node ensures continuity of the service provided by the VM. All of the memory data is transferred in the first round, and then the migration controller recopies the memory data which is changed in the last round. These steps keep iterating until the dirty portion of the memory is small enough to handle the final copy. Although precopying memory is performed iteratively, the execution of programs is not obviously interrupted.
- **Step 3: Suspend the VM and copy the last portion of the data.** The migrating VM's execution is suspended when the last round's memory data is transferred. Other non memory data such as CPU and network states should be sent as well. During this step, the VM is stopped and its applications will no longer run. This "service unavailable" time is called the "downtime" of migration, which should be as short as possible so that it can be negligible to users.

- **Steps 4 and 5: Commit and activate the new host.** After all the needed data is copied, onthe destination host, the VM reloads the states and recovers the execution of programs in it, and the service provided by this VM continues. Then the network connection is redirected to the new VM and the dependency to the source host isb cleared. The whole migration process finishes by removing the original VM from the source host.
- Before copying the VM with 512 KB files for100 clients, the data throughput was 870 MB/second. The first pre copy takes 63 seconds, during which the rate is reduced to 765 MB/second. Then the data rate reduces to 694MB/second in 9.8 seconds for more iterations of the copying process.
- The system experiences only 165 ms of downtime, before the VM is restored at the destination host.

3.9 VIRTUALIZATION FOR DATA-CENTER AUTOMATION

Part –B	**1. What do you mean by Data Center Automation using Virtualization? (AU/April/May 2017)**

- The dynamic nature of cloud computing has pushed data center workload, server, and even hardware automation to whole new levels.
- Now, any data center provider looking to get into cloud computing must look at some form of automation to help them be as agile as possible in the cloud world.
- New technologies are forcing data center providers to adopt new methods to increase efficiency, scalability and redundancy.
- Let's face facts; there are numerous big trends which have emphasized the increased use of data center facilities. These trends include:
    - More users
    - More devices
    - More cloud
    - More workloads
    - A lot more data
- As infrastructure improves, more companies have looked towards the data center provider to offload a big part of their IT infrastructure. With better cost structures and even better incentives in moving towards a data center environment, organizations of all sizes are looking at colocation as an option for their IT environment.
- With that, data center administrators are teaming with networking, infrastructure and cloud architects to create an even more efficient environment. This means creating intelligent systems from the hardware to the software layer. This growth in data center dependency has resulted in direct growth around automation and orchestration technologies.
- Now, organizations can granularly control resources, both internally and in the cloud. This type of automation can be seen at both the software layer as well as the hardware layer.
- Vendors like BMC, ServiceNow, and Microsoft SCCM/SCOM are working towards unifying massive systems under one management engine to provide a single pain of glass into the data center workload environment