cloud STACK
COLLABORATION CONFERENCE

Automatic Workload Management in Clusters Managed by CloudStack

LINUX FOUNDATION

Problem Statement

In a cluster environment, we have a pool of server nodes with VMS running on them. Virtual Machines are launched in some nodes and released from other nodes periodically. After some amount of time, we will observe dispersion of VMs along the pool of the nodes.

- The terms “green computing” and “green IT” have been used frequently these days to refer to energy-efficient solutions in IT.

- Energy-efficiency issues of data centers are of major importance as costs of power and cooling make up a significant part of their operational costs.

- Energy-efficient initiatives done by data centers not only reduce their expenses, but overall, contribute to the Green IT policy that is targeted at minimizing carbon emissions.

How normal deployment works

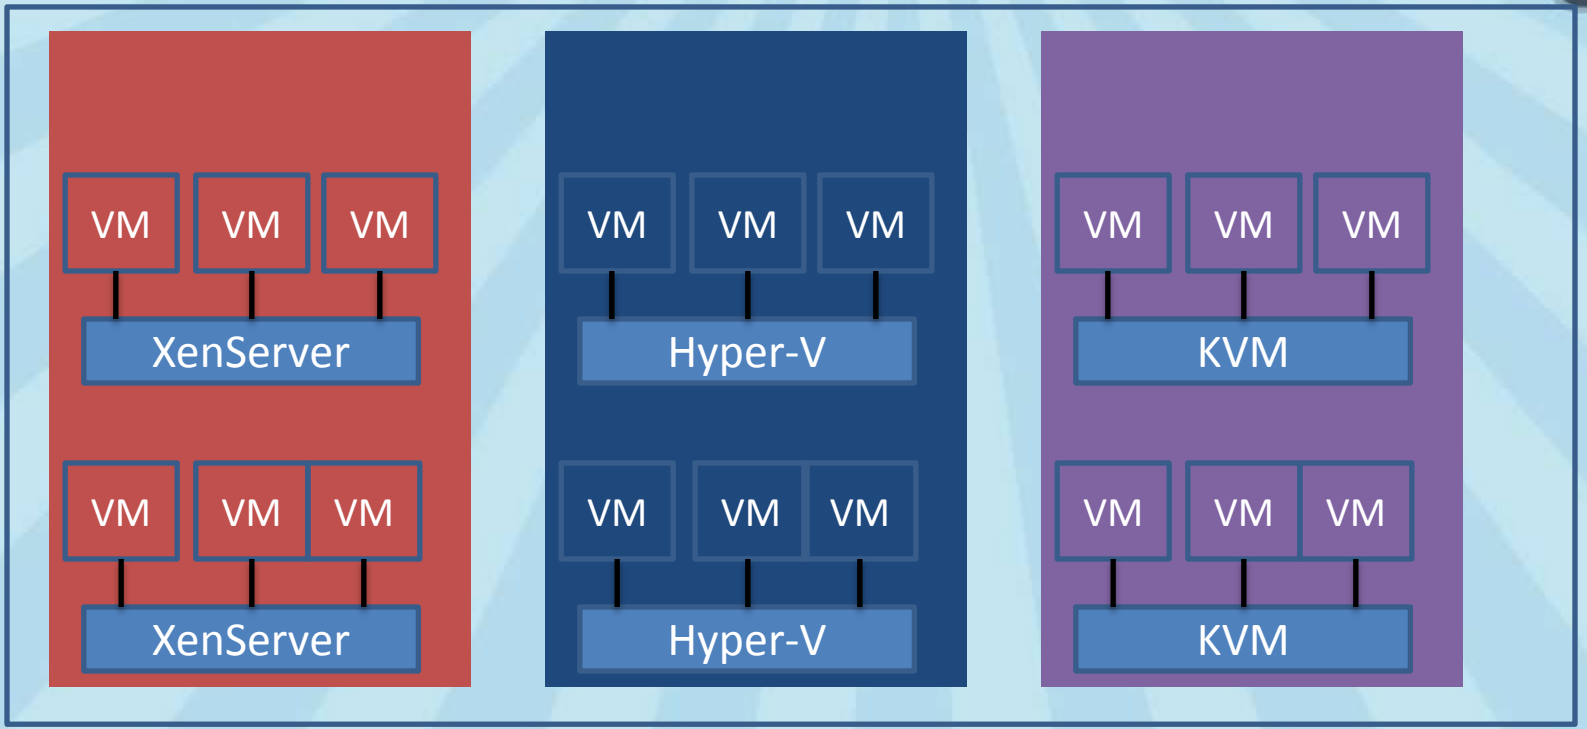

When VMs were destroyed/stopped/migrated

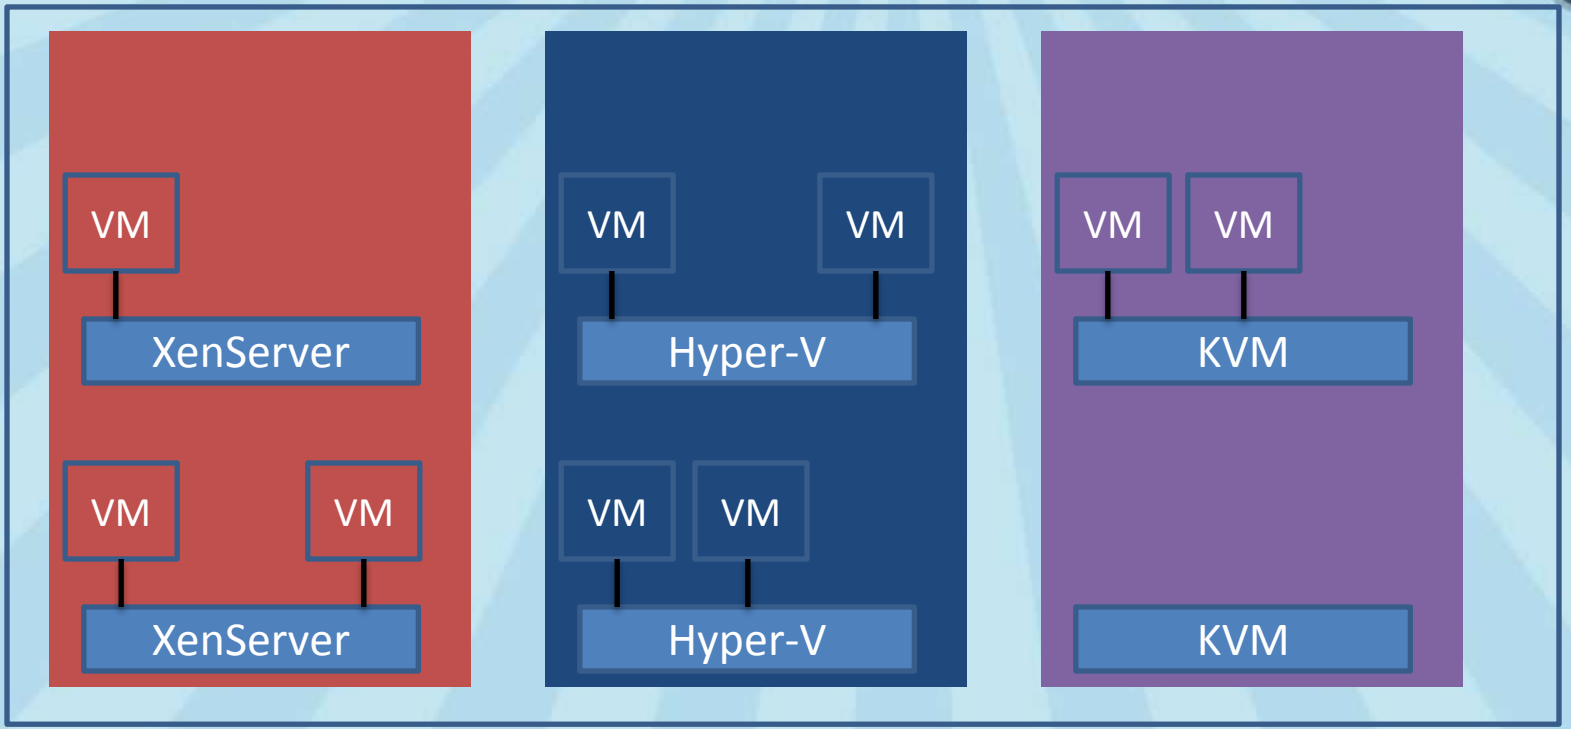

Disadvantages of the above Scenario

The above scenario requires Automatic workload management which will give benefits

Improve Utilization

Reduce Cost

Improved bandwidth availability and lot more etc.

How to resolve the use case

- Monitor Host Capacity and Usage
- Find the deployment model the host being used
- Like dedicated host, using local store, ZoneWidePrimary Store etc
- Live Migrate the VM's from less deployment host to the best possible host
- Choosing the best possible host depends on the characterises like
  - Cpu, memory, free capacity available to accommodate more
- Once the best hosts are chosen
  - Provide the info to the user to execute the action plan

- According to the action plan the VMs will be migrated.
- Once the VM's are migrated the some of the hosts will be empty.
- Mark the hosts in Maintenance mode and shutdown those hosts.
- When there is a peak of demand or on user request, WOL packet will be sent to the host which will bring up the server.
- When the server comes up the agent will be created for the same and it will be moved to "up" state which can be used for deployment plan

- This will help the IT to reduce costs and manage their hosts efficiently and effectively

How to choose Best hosts or Suitable hosts

- We need to select Best Server Consolidation Algorithm to find the best suitable hosts for consolidation
- We need to take into account several constraints to select consolidation model. The first set of constraints is related to live migration technology. As we are going to re-allocate VMs by migrating the nodes, we need to consider the constraints of live migration technology. Currently, live migration requires having compatible virtualization software, comparable CPU types, similar network connectivity, and usage of shared storage on both source and destination nodes

How to choose Best hosts or Suitable hosts

- Second, it is necessary to limit the upper bound of CPU utilization of a single node by some threshold value. This is done to prevent a node's CPU from reaching 100% utilization. The motive behind this is that 100% utilization can lead to performance degradation.
- Furthermore, live migration technology also consumes some CPU cycles. Keeping CPU utilization below that threshold value allows retaining a certain level of CPU throughput.

Server Consolidation Algorithms

Bin Packing:
- best fit decreasing
- first fit decreasing

Sercon Algorithm:
This algorithm inherits some properties of First- and Best-Fit, which try to minimize the number of nodes used. However, we not only aim to minimize the number of nodes used, but also to minimize the number of migrations. Hence, this algorithm behaves differently from the ones related to bin-packing.

What is WOL

Wake-on-LAN (**WOL**) is an Ethernet computer networking standard that allows a computer to be turned on or *awakened* by a network message.
Wake-on-LAN ("WOL") is implemented using a specially designed packet called a magic packet, which is sent to all computers in a network, among them the computer to be woken up.

The magic packet contains the MAC address of the destination computer, an identifying number built into each network interface card ("NIC") or other ethernet device in a computer, that enables it to be uniquely recognized and addressed on a network. Powered-down or turned off computers capable of Wake-on-LAN will contain network devices able to "listen" to incoming packets in low-power mode while the system is powered down.

If a magic packet is received that is directed to the device's MAC address, the NIC signals the computer's power supply or motherboard to initiate system wake-up, much in the same way as pressing the power button would do.

What is Magic Packet?

The *magic packet* is a broadcast frame containing anywhere within its payload 6 bytes of all 255 (FF FF FF FF FF FF in hexadecimal), followed by sixteen repetitions of the target computer's 48-bit MAC address, for a total of 102 bytes.

Q & A

Thanks

cloud STACK
COLLABORATION CONFERENCE

LINUX FOUNDATION