RED HAT ENTERPRISE VIRTUALIZATION: SYSTEM SCHEDULER

WHAT IS THE SYSTEM SCHEDULER

The Red Hat Enterprise Virtualization *System Scheduler* manages the allocation of physical resources within the virtual data center.

The System Scheduler continually monitors the utilization of host systems and virtual machines, dynamically managing the placement of virtual machines within the data center based on policies defined by the system administrator.

RESOURCE MANAGEMENT AND ALLOCATION

The physical resources in the data center such as servers, storage and networks can be assigned to logical pools for fine grain management and control.

Storage is added to the virtual data in the form of “storage domains” created from a Fiber channel LUN, iSCSI LUN or NFS export. Storage domains are grouped into logical pools and can be assigned to a virtual data center. Network resources such as physical networks or vLANs are defined at the virtual data center level and maybe allocated to one or more host clusters.

Physical servers are grouped into host clusters allowing an organization to break down their resources into multiple logical pools of resources that can be managed separately – for example assigning each business unit their own cluster to manage independently.

The System Scheduler is responsible for the placement of a virtual machine on a physical host system.

- **Starting a virtual machine**
  When a virtual machine is started the system scheduler automatically selects the host on which to run the virtual machine based on resource utilization and system policies.
- **Virtual Machine Affinity**
  An administrator may define a preferred host for a virtual machine for example to place a database server on a physical host with more I/O capacity or to collocate two virtual machines on the same host system to reduce their cross-network traffic and optimize their communication.

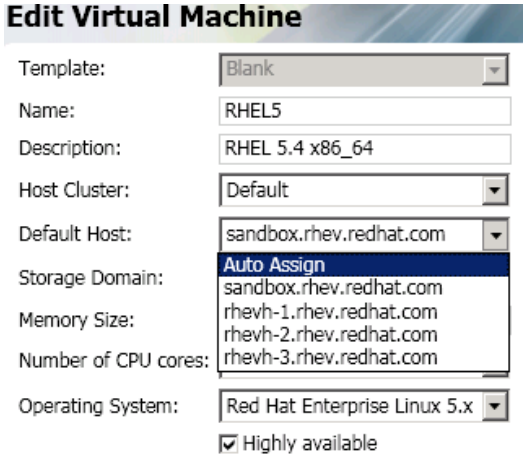

- **High Availability**
  In the event of a hardware failure any virtual machine configured to be “Highly Available” will be restarted on another host in the cluster. The system scheduler manages the selection of the host system based on resource utilization within the cluster.
- **Live Migration**
  When an administrator manually initiates a live migration the system scheduler will automatically select the destination host based on resource utilization. An administrator can override the system scheduler and manually select the destination host.

DYNAMIC RESOURCE SCHEDULING

The System Scheduler continually monitors the utilization of host systems and virtual machines. The administrator can define policies to govern the resource allocation within a cluster of host systems. Based on these policies the System Scheduler will automatically balance the workload between physical hosts, using live migration to move virtual machines without any down time or service interruption.

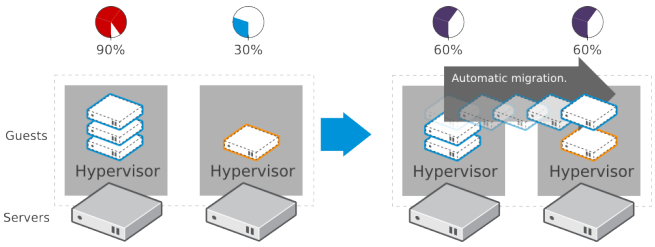

The administrator can define the thresholds at which the System scheduler will automatically use live migration to relocate virtual machines, defining both the utilization threshold the physical host should reach and the duration for which this threshold must be exceeded. This ensures that a momentary spike in resource usage does not trigger reallocation of resources.

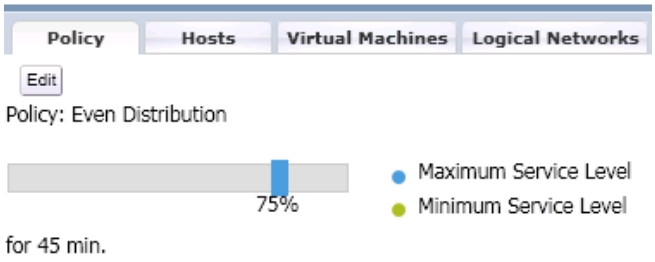

With System Scheduler organizations can rapidly adapt to the changing needs of the business. As server utilization increases the System Scheduler will balance workload across available resources to improve performance by ensuring that virtual machines can take advantage of all available resources. If the cluster is operating at maximum capacity an administrator can dynamically add another server to the cluster, either from another cluster or a new physical server. The System Scheduler will automatically leverage the extra capacity and re-balance workload within the cluster.

Different System Scheduler polices can be defined for each cluster in the virtual data center allowing organizations to delegate system utilization policies to individual business units.

REDUCE POWER CONSUMPTION

"Power Saver" extends the System Scheduler to add policies to reduce power consumption by consolidating more virtual machines onto a smaller number of physical hosts.
Since most organizations size their infrastructure to cope with their peak capacity requirements, during off peak hours, such as nights and weekends the extra physical capacity is consuming expensive physical resources such as power and cooling. The administrator configures the minimum service level at which the Power Saver policy is triggered. For example if the utilization of a single host goes down to 10% for 20 minutes or more then System Scheduler with use live migration to relocate the virtual machines running on this host to other hosts in the cluster.

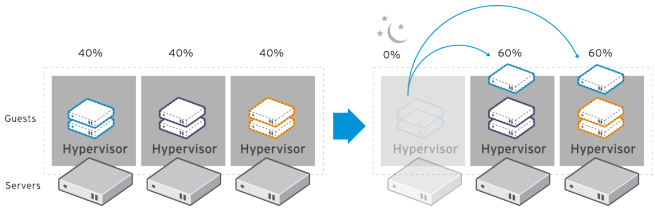

The host will now run idle, consuming significantly less resources - typically 10 to 15% of the power of an active server. If other hosts in the cluster reach the maximum service level the System Scheduler will live migrate virtual machines onto the host and re-balance workload within the cluster.

SOFTWARE REQUIREMENTS

System Scheduler and Power Saver are included in all editions of Red Hat Enterprise Virtualization.