

Understanding Dashboard

There is a subtle yet important difference between **real time usage** vs the **allocated view** of cloudstack. Allocated capacity refers to the resources reserved by MS whereas real time usage as the name refers is the actual usage (So for storage it is analogous to "df -h" (Linux command)).

Example MS has the following configuration : 2GB RAM, 2GHZ CPU, 10GB storage.

Lets say you choose a service offering of 1GB RAM, 1GHZ CPU and 5gb disk for deploying a vm, MS then reserves these capacities from its physical resources exclusively for this vm. While the vm's "actual usage" might be significantly less (say .1gb RAM , 128MHZ CPU and 1 gb storage) than the allocated(reserved) capacity. Assuming this to be the only vm in the system, the dashboard would display the following values.

Allocated RAM - 1/2 (50%)

Allocated CPU - 1/2 (50%)

Allocated Storage - 5/10 (50%)

Storage Used - 1/10 (10%) (analogous to "df -h" (Linux command))

- Ideally we wouldn't want to have a huge difference between allocated and real time usage and hence the notion of **overprovisioning**.
- While deploying vms we always consider the total allocated capacity and never the real usage of the resources (except when the real usage for storage > threshold).
- The current granularity of dashboard is shows it till the cluster level. So though the aggregate view might show some vacant space there might be no physical resource in the cluster/pod/zone with enough space to accommodate the vm in terms of cpu, ram or storage.
- Dashboard view is updated real time for cpu and ram and every 5 mins (capacity.check.period) for allocated storage. For real time storage we update it every minute.
- When you increase the factor, the allocated percentage decreases and hence **you can fit in more vms**.

Say allocated = 100% = 100gb out of 100gb total space and overprovisioning

factor = 2 and at this moment you can't deploy more vms.

Say you change overprovisioning to 3 and allocated now becomes = $100/150 =$

67%, so you can fit more vms.

- The dashboard calculation is for the entire cloud and that means the systemvms (ssvm, cpvm, routervm) as well.
- Once the instances are stopped their capacity is released after a configurable time (1 hr default).
- If they change over provisioning when their instances are already installed the calculation is not straightforward. Refer the [wiki](#) for understanding the calculation with over provisioning.