Host Thin Provisioning™ with Eco4Cloud

This document describes how to take advantage of Thin Provisioning with the physical hosts of a virtualized Datacenter through Eco4Cloud solutions, provisioning only the number of physical hosts that applications need.

Eco4Cloud — www.eco4cloud.com | Phone +39 0984494276 | E-mail info@eco4cloud.com

Thin Provisioning Storage

Thin Provisioning is a mature concept in storage management. It allows efficient disk space allocation, which lowers physical storage capacity needs and costs.

Using Thin Provisioning it is possible to create virtual disks, whose space gets allocated at run-time during write operations. This is achieved by virtualizing disk space, meaning that the space assigned to the virtual disk is different (larger) than the space allocated on the underlying physical datastore.

Storage Thin Provisioning only consumes capacity that is not empty (non-zero), empty areas of the virtual disk consume no physical capacity at all. This technology offers significant storage economics for the operators of a thin provisioned storage controller.

The opposite approach, called Thick Provisioning, or Fat Provisioning, is based on the immediate allocation of 100% of the assigned space of a virtual disk. This legacy approach leads to poor storage economics, the wasted capacity is clear in the following figure (allocated and unused).

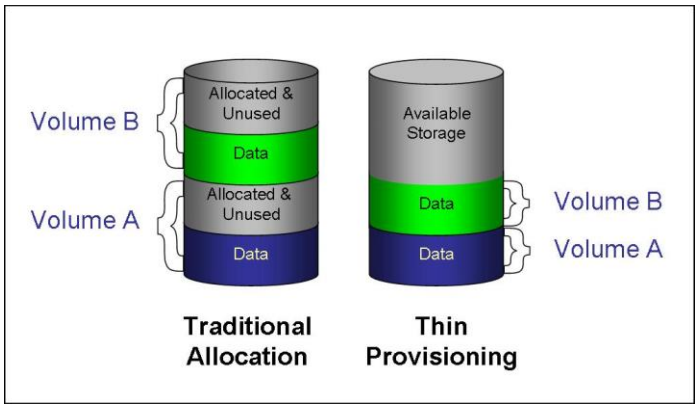

Memory Thin Provisioning

With memory over-commitment, the hypervisor ensures that host memory is consumed by active guest memory as much as possible, leading to a higher consolidation ratio.

The Thin Provisioning approach has also been used for many years in memory virtualization, and in memory over commitment.

Virtual memory is a well-known technique used in most operating systems, and almost all modern processors have hardware to support it. Virtual memory creates a uniform virtual address space for applications and allows the operating system and hardware to handle address translation between the virtual address space and the physical address space.

Hypervisors add a second level of memory virtualization. When running a virtual machine, the hypervisor creates a contiguous addressable memory space for the virtual machine. This allows the hypervisor to run multiple virtual machines simultaneously while protecting the memory of each virtual machine from being accessed by others. Therefore, from the view of the application running inside the virtual machine, the hypervisor adds an extra level of address translation that maps the guest physical address to the host physical address. As a result, there are three virtual memory layers: guest virtual memory, guest physical memory, and host physical memory.

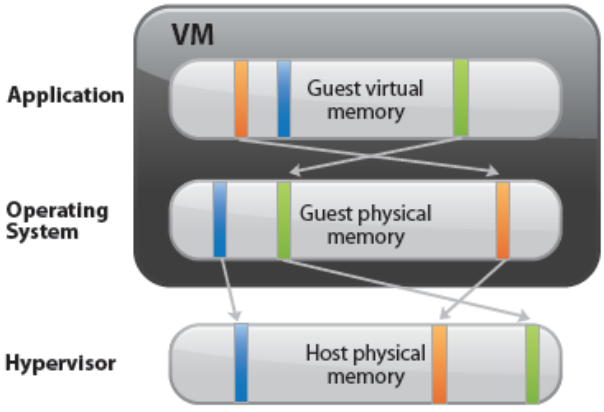

The final result is the memory over commitment: hypervisors virtualize memory space, because the memory assigned to a virtual machine is larger than the memory physically allocated on the virtual machine.

With memory over-commitment, the hypervisor ensures that host memory is consumed by active guest memory as much as possible, leading to a higher consolidation ratio. In fact, as shown in the next figure, it is possible to run three virtual machines with 2G guest physical memory each, in a host with 4G host physical memory. Without memory over-commitment, only one virtual machine can be run because the hypervisor cannot reserve host memory for more than one virtual machine, considering that each virtual machine has overhead memory.

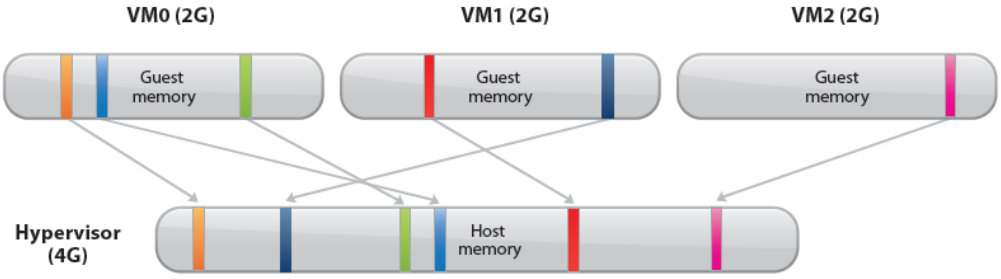

Host Thin Provisioning™

If you are operating a medium to large-scale virtual environment either as an enterprise or MSP, why waste capital, energy and data center capacity? Be thin; use E4C solutions for Datacenter Maximization.

Thin Provisioning approach can be applied to host cores, too.

Having all physical hosts utilized at only at 30% - 50% is a very common scenario in the data center. The vast majority of hosts are hence operating with low efficiency, because:

- Energy consumption of a fully utilized host is quite similar to that of the same host in idle condition.
- The data center is likely to have too many physical servers for the applied workload, wastefully consuming capital and data center capacity

Eco4Cloud Workload Consolidation is able to consolidate any number of virtual machines onto the minimum number of physical server hosts, hence delivering autonomic ***Host Thin Provisioning™***, provisioning only the number of physical servers and cores that applications need, rather than the number provisioned. Host Thin Provisioning™ works by switching on and off physical cores as workload varies in real-time.

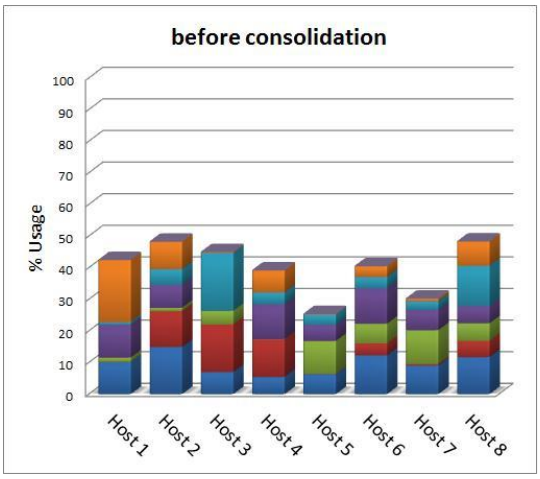

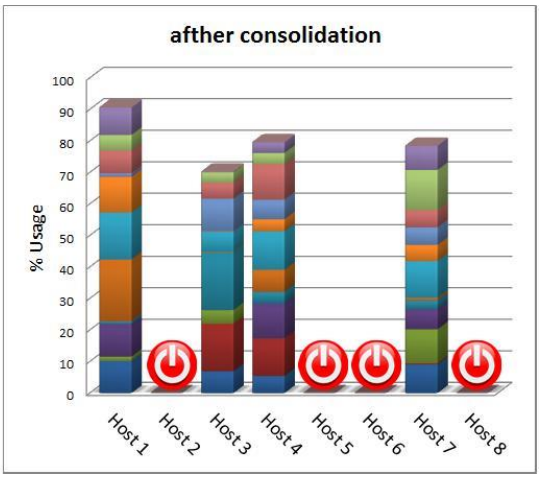

If you are operating a medium to large-scale virtual environment either as an enterprise or MSP, why waste capital, energy and data center capacity? Be thin; use E4C solutions for Datacenter Maximization.

Useful Links

- Eco4Cloud – www.eco4cloud.com

- Architecture and Requirements - www.eco4cloud.com/download/E4C-Architecture-and-Requirements.pdf
- Saving energy in datacenters through workload consolidation - www.eco4cloud.com/download/e4c-white-paper.pdf