

# Leveraging Local Image Redundancy for Efficient Virtual Machine Provisioning

May 2012 [Network Operations and Management Symposium 2012](#)

Image-based provisioning provides a fast and reliable mechanism for handling the demands of Cloud Computing. Typically, a Cloud data center contains a catalog of images in the image library, multiple hypervisors with inexpensive direct attached storage (where the instances are created), and a placement mechanism that allocates and reserves resources. Image-based provisioning is a deployment and activation mechanism that clones a "golden" read-only virtual machine (VM) image residing in the image library to create a new virtual machine instance on a hypervisor. The main steps of the provisioning process are: 1) selection of hypervisor based on a placement policy; 2) copying VM image from a storage server to the direct attached storage of the hypervisor, and 3) image activation to create an instance. The image copy from the storage server to the direct attached storage of the hypervisor is time consuming and network intensive, directly contributing to user perceived provisioning latency.

This article proposes a mechanism that can reduce the network bandwidth requirements by efficiently selecting a hypervisor for placement and reconstituting the required image based on content already available on the local storage. The proposed system leverages virtual machine image similarity and provisioning frequencies to reduce the data volume transferred from the storage server to the hypervisor on which the virtual machine is being instantiated. There are multiple situations leading to significant degree of image similarity, such as: when two different images are created from the same base image, two different images have the same middleware or applications installed, or when users modify configurations and recapture images for later provisioning. In such situations, there are clusters of blocks that are the same across the images in the library. These blocks within the cluster could even be non-contiguous. The image redundancy information is used to supplement capacity based placement by utilizing overlap across images present on direct attached storage for reconstituting a virtual image. The algorithm is implemented in a testbed and also validated using extensive discrete event simulations based on a library representative of typical Cloud provider's catalog. Analytical model and simulations measure impact of degree of image similarity, system utilization, hypervisor capacity and image provisioning frequencies on expected gain. The system achieves up to 80% reduction in amount of data transferred from the storage server to hypervisors. It is especially effective for large and highly utilized hypervisor clusters.

**Source URL:** <http://www.comsoc.org/ctn/leveraging-local-image-redundancy-efficient-virtual-machine-provisioning>