

Understanding The Virtualization Lifecycle In The Context Of Cloud Computing: A Beginner's Approach

The building blocks of Cloud Computing are Operating Systems, Virtualization, Principles of Networking, Information & Network Security and Storage. All of these areas combined, form the basis of the Infrastructure as a Service (IaaS) model of Cloud Computing. In this article we shall take a brief look at the Virtualization lifecycle in the context of Cloud IaaS.

Virtualization is one of the key enablers of cloud computing. Virtualizing hardware, applications and consolidating them aims to reduce IT Infrastructure costs (which includes purchasing hardware and maintaining it), and allowing easier management of resources in the data center. The IaaS model of the cloud deals with provisioning of compute capacity, storage and networking resources. Infrastructure costs are reduced by virtualizing hardware, thereby avoiding under-utilization of resources. As a result, multiple applications can be virtualized over the same physical hardware ensuring optimal usage of resources. The other terms that we often hear w.r.t cloud computing are 'on-demand' and 'elasticity'. These two terms go hand-in-hand; on-demand refers to a 'pay-as-you-go' model. You need to pay only for the resources you use, and elasticity refers to scaling up/down the resources at will. In the context of cloud computing, the virtualization lifecycle comprises a set of technical assessment activities which are governed by business and operational decisions. Technical assessment for virtualizing candidates revolves around meeting end-user Service Level Agreements (SLA's), reducing IT costs, and designing an optimized data center. Every phase in the virtualization lifecycle for cloud computing is highly challenging with a wide variety of complex and open problems which are currently being tackled.

Analysis & Discovery: When migration from Physical environments to Virtualized environments (P2V), in-depth analysis of the virtualization candidates must be performed. This stage involves discovering the data center entities (servers, networks,

storage devices), collecting utilization profile data of applications along the different dimensions (CPU, memory, network i/o, disk i/o). The main theme of P2V is to move applications from an under-utilized bare metal environment to a virtualized / hypervisor environment to enable optimal utilization of hardware. In addition to discovering the heavy artillery in the physical environment, it is important to assess the applications deployed on them. The OS characteristics, application performance footprints play a vital role in determining capacity in a virtualized environment. On completion of these assessments, capacity management models need to be developed for the virtual environments.

Implementing Capacity Models: Developing capacity models for a virtual environment is a tricky task since it is governed by other business and operational factors. Target SLA's (performance, availability), power consumption levels are to be considered along with the possible impacts (side effects) of virtualization (hardware normalization, hypervisor overheads, IO interference etc). The idea is to come up with a 'pre-VM placement' strategy which describes the 'footprints' of VM's. Capacity planning for virtualized data centers in the light of cloud computing has become a highly sought after topic. Determining capacity size of virtual machines a priori to migration becomes an extremely critical step. If done accurately it can result in optimal allocation and usage of resources, if over-cooked, can lead to resource wastage. Similarly if under-cooked, can result in poor performance and violation of SLAs. There are many useful P2V, V2V and capacity analysis tools that can help you achieve this, viz. Platespin Recon, Microsoft SCVMM, VMware P2V Assistant, to name a few. Also researchers are exploring intelligent ways of doing capacity sizing in Virtual environments.

VM Placement & Management: This is the most critical process from a data center administrator point of view. Academic and Industrial groups are grappling to identify '*best-fit*' placement strategies to enable highly optimized virtual environments. This refers to the concept of 'packing' VMs appropriately. VMs need to be packed in such a way that performance of isolated (individual) VMs is not hampered due to interference, and to avoid fragmentation in the data center. The on-demand

provisioning of virtual servers will eventually lead to a server sprawl, thus complicating management of virtual servers. Hence efficient techniques for placement and management hold key in having a greener and well maintained data center. Other issues may also involve cross data center migration, synchronization between different servers, enabling and managing hybrid clouds (a combination of public and private / in-house environments).

Thus, the virtualization life cycle poses many challenges in different areas. Leading Cloud providers and academicians are busy solving these problems, and we hope to see greener data centers soon !!

Source:<http://thesemanticway.wordpress.com/2011/06/09/understanding-virtualization-in-the-context-of-cloud-computing/>