Datacenter Services
VIRTUALIZATION

Where, When & Why To/Not Virtualize

Introduction: Virtualization and Cloud Computing are two of the hottest topics within the IT industry today. However, as these subjects are growing rapidly there are still confusions being made in terminology, technology, implementation, licensing etc...

Virtualization: Allows the ability to run multiple virtual servers on one physical server. Today, advanced virtualized infrastructures consist of multiple physical servers attached to an enterprise storage device, implementing high availability configurations. Physical servers are accessing enterprise storage and run virtual servers; while actual data within a virtual server is stored on the enterprise storage. However, these abilities do not bring many benefits to complicated IT infrastructures; but if integrating this ability with enterprise storage and high availability clustering shifts the way IT organizations think about their infrastructures.

Such architecture allows:

- *Adding processing power by adding physical servers*
- *Adding storage by growing enterprise storage device*
- *High availability with virtual server migration from one physical node to another physical node*
- *Easier high availability as the whole infrastructure and inherently each virtual instance is high available*

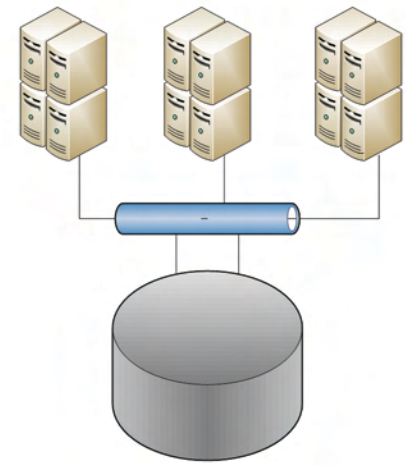

Benefits: There are 3 direct benefits from taking on a virtualization approach for almost any IT organization.

- *Increase of processing density*
- *Increase of ability to scale*
- *Increase availability*

Processing Density: Virtualization allows running multiple low processing virtual servers on one physical server. Total processing capacity can be dynamically allocated among virtual instances to make sure maximum CPU and RAM is utilized within hardware. Even though the goal is to use 100% of the CPU and RAM resources, there is a risk to overload a server because a small spike in one of the virtual servers can push the physical server to the edge of its processing capacity. Therefore, IT groups should measure processing profiles of servers and plan to load physical servers with somewhere around 80% of its maximum capacity.

***Digital Edge develops and successfully utilizes advanced density planning methodologies allowing clearly identified required capacity and planned resource allocation; providing enough processing power and loading of hardware proportionately.*

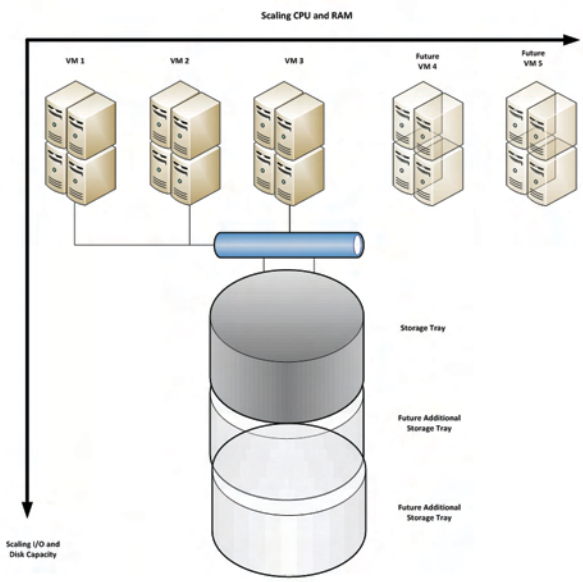

Scalability: In general there are 4 parameters that IT groups should be able to scale for their system processing:

1. *CPU power (Megacycles).*
2. *RAM (GB)*
3. *I/O (IOPS)*
4. *Disk capacity (GB)*

A virtualized infrastructure gives IT organizations the ability to linearly scale their infrastructures per processing requirement. CPU and RAM are scaled with adding extra physical servers, while I/O and Disk Capacity are scaled with adding more disks and disk trays.

Availability: Clustered physical servers talking to enterprise storage on the back end implements a high availability platform. Clustering software provides capabilities to re-locate Virtual Servers from one failed server to another working server. This can be done manually for maintenance purposes or automatically as a critical failover. In a pre-virtualized world, IT organizations needed to build clusters or other high availability configurations for each critical element of the infrastructure. Almost each hardware box had to have a redundant cold or warm standby clone. With Virtualization companies are given the ability to build an entire platform high available and have each Virtual server with inherit failover capabilities of the underlying platform; a virtualized failover server may play a role similar to a spare disk in the raid array.

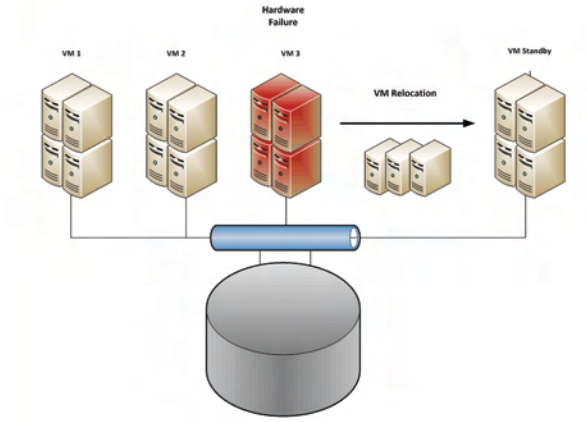

A Single standby physical server can be a failover resource for multiple active physical servers. The ratio of active/standby and failover logic depends on the size and complexity of the infrastructure. Enterprise storage, SAN or iSCSI provides its own built in redundancy so no additional provisioning except multi-pathing to the storage is required

When Not to Virtualize: A main reason not to virtualize is when a system within an infrastructure has a very high load. When load balancing or any scale up / scale out methodologies are implemented. Consider a very active web site that has 4 load balanced web servers and a database cluster. The system was stress tested and the IT group knows that it sustained its targeted load. There is nothing much to virtualize in such an environment.

1. Web tier: if the application works on 4 physical servers it may not be able to run with the same performance characteristics on 4 physical servers running multiple Virtual Servers. It will not run the same even if only one Virtual Server per Physical Server will be created. Hypervisor, the software allowing single physical hardware running multiple Virtual Servers, takes hardware resources. Hypervisor also abstracts and virtualizes hardware specifics for Virtual Servers which also require processing resources decrease the ability of end systems to use all the hardware specifics. To achieve the same performance, the same number of servers or even more servers needs to be used. There is no opportunity to increase processing density in this case.

2. Database tier: Database in high performance, high I/O, high transactions systems need to utilize every possible hardware and OS feature to gain required performance. Very precise memory and I/O tuning and mechanisms are available for SQL performance optimization. When SQL Server is running inside Virtual Server it automatically inherits high availability. However such configuration will prevent DBAs from using hardware and OS specifics as hypervisor will abstract access to RAM and I/O. Heavy transactional and I/O "hungry" systems will suffer the most.

Cloud: Cloud Computing should be defined as a virtualized infrastructure built by Data Centers in order to sell Virtual Private Servers or VPSs to their clients. In addition to virtualizing servers and storage for their clients, Cloud infrastructure includes automatic billing, fast creation of VPSs based on pre-built libraries of templates, GUI etc...
In today's IT some people call Cloud – "Software as a Service" infrastructures. "Software as a Service" or SaaS are web-based systems that provide business functions that could be used as building blocks, performing specific tasks. A good example of SaaS would be a NetSuite. NetSuite implements product catalog functionality that could be accessed through SOAP interfaces. Businesses, instead of programming their own product catalogs can use NetSuite GUI to upload products and manage them and programmatic interfaces (SOAP APIs) to power E-Commerce web sites. In such situations, the business using NetSuite would have to draw only E-Commerce pages but the whole product database would be hosted and managed by NetSuite. There are other products that are SaaS and people call them "Cloud" because they don't really care how the information is stored at the service provider and how it is manipulated.

When to Cloud: In order to go to Cloud, you need to understand your processing requirements, growth needs and strategies; IT infrastructures that are fragmented require multiple independent servers / roles and when each server is low resource consumption.

When Not to Cloud: Consider this, if your infrastructure cannot be virtualized, can it be in a cloud? Besides Virtualization assessments and qualifications, the following potential drawbacks of Clouds should be considered:

1. *PCI Requirements*

2. *Virtual Server to Virtual Server security and Virtual Server to Physical Server security. Cloud security is a very tricky aspect. Depending on Virtualization implementations in Clouds, there are possibilities of security weaknesses in communication between Virtual Servers or Virtual Servers and a Physical Server. The access through the storage network is a potential weakness as well. Some separation relies on VLAN protocols that are not strong security mechanisms considering then a client has root access to virtual instance. A not friendly neighbor can try to spoof your network, tag VLAN packets, assign duplicate IP addresses and send malicious network traffic destabilizing the whole Cloud infrastructure. Most Clouds today keep its security and architecture in secret. This makes IT organizations impossible to assess the risk of being hacked in a Cloud.*

3. *Data security of Virtual Servers is another consideration. In any configuration, Cloud operator has FULL access to ALL data that resides in the Cloud. Without having root access to client's virtual access, the Cloud operator can gain access to your own data.*