



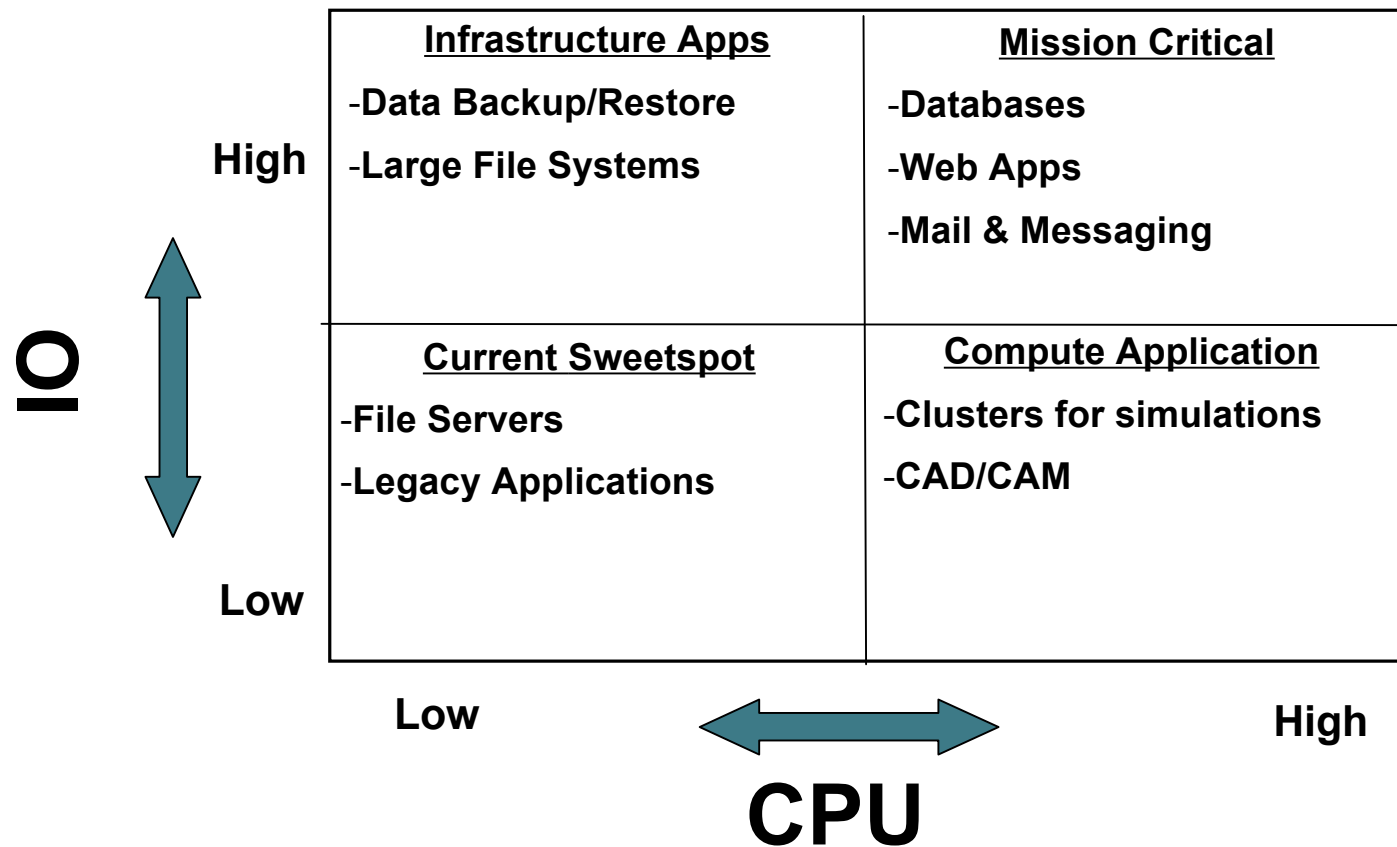
Network Resources and VMs

8nd September, 2006

Tony Vaidya

Cisco Systems, Inc.

Where is virtualization today ?



Myth: Server virtualization technologies are not for applications with high CPU and I/O bandwidth requirements.

Demystifying

- **CPU cores per socket and I/O bandwidth are both increasing, so what are the other issues ?**
 - **There are network mappings that control how a server can consume network resources. Those mappings need to be extended to Xen domains.**
 - **Networking**
 - **VLANs can be extended through Linux bridge**
 - **Storage access**
 - **NPIV, VSANs**
 - **Security**
 - **802.1x**

Storage Access

- **Allow a domain to have access to multiple tiers of storage.**

Possible solutions:

- **Dedicate an HBA for each tier of storage.**
 - **Limited to the number of PCI slots**
 - **Does not allow lun masking and zoning based on domain.**
- **NPIV**
 - **Allows unique mapping between domain and WWPN.**
 - **Allows for FC zoning and storage side lun masking.**
- **VSANs / Virtual Fabrics**
 - **Allows for unique mappings between WWPN and VSAN.**
 - **Allows segmentation of FC fabric and services.**

...but still need more

SANs and live migration

- **Live migration requires that I/Os in the backend driver are completed before restore operation can start.**
- **Link failures, failures in the storage array, or congestion cause undesirable delays during domain save operation.**
 - **This is the time to migrate !**
 - **HBA's with NPIV or VSAN support needs semantics to quiesce a low-level driver instance.**
 - **Force a low level FC driver to log out from targets and then FC fabric. - implicit aborts**

Security

- **Desire is to authenticate a VM before it is allowed to use a network resource (i.e. VLAN)**
- **Use 802.1x to provide network access control.**

**"Port-based network access control makes use of the physical access characteristics of IEEE 802 LAN infrastructures in order to provide a means of *authenticating* and *authorizing* devices attached to a LAN port that has point-to-point connection characteristics, and of *preventing access* to that port in cases which the authentication and authorization fails. A port in this context is a single point of attachment to the LAN infrastructure."
--- 802.1X-2001, page 1.**

Security

- **Need a linux bridge and a 802.1x authenticator**
 - **Combine a linux bridge and 802.1x authenticator in a dom0**
 - **Map DomUs uniquely to physical NICs and let the upstream switch provide authentication.**
 - **Cannot scale number of VMs beyond number physical NICs**
- **Need to have all domains running a supplicant**
 - **Use xsupplicant for Linux guests**

Even If this all works, still need more

Security

- **What if the domU is running a legacy OS, for which there are no supplicants ?**
 - **Need a mechanism to create and securely store a credential for a DomU.**
 - **Need a mechanism to pass the credential to the authenticator.**
 - **Allow the domain running linux bridge + 801.x authenticator to obtain that credential.**

Next set of questions

- **How do you migrate domains across different smart hardware that may have different network/storage attachment characteristics?**
- **In cases where you have similar capable hardware, how do you ensure that the devices attach in the same manner to the network?**
- **How do you ensure that smart hardware synchronizes I/O completion during migration?**

Q & A