Virtualization Technology Outlook and IBM Directions
Analyst Briefings, IBM Bedfont, UK

Date: July 5th, 2006
Revision: 1.0

PUBLIC

Mark Cathcart, IBM Distinguished Engineer,
Systems and Technology Group, IBM Somers, NY
m_cathcart@us.ibm.com

The following material regarding IBM's future direction and intent is subject to change or withdrawal without notice, and conveys goals and objectives only.

Virtualization Concept

Virtual Resources

Proxies for real resources: same interfaces/functions, different attributes.

May be part of a physical resource or multiple physical resources.

Virtualization

Creates virtual resources and "maps" them to real resources.

Primarily accomplished with software and/or firmware.

Resources

Components with architected interfaces/functions.

May be centralized or distributed. Usually physical.

Examples: memory, disk drives, networks, servers.

Separates presentation of resources to users from actual resources

Aggregates pools of resources for allocation to users as virtual resources

Virtualization Functions and Benefits

Sharing

Examples: LPARs, VMs, virtual disks, VLANs

Benefits: Resource utilization, workload manageability, flexibility, isolation

Aggregation

Examples: Virtual disks, IP routing to clones

Benefits: Management simplification, investment protection, scalability

Emulation

Examples: Architecture emulators, iSCSI, virtual tape

Benefits: Compatibility, investment protection, interoperability, flexibility

Insulation

Examples: Spare CPU subst., Capacity Upgrade on Demand(CUoD), migration

Benefits: Continuous availability, flexibility, investment protection

Virtualization Management Services

Manual and automated control of virtualization technologies to optimize resource utilization and workload service level achievement

Dynamic provisioning, resource level adjustment, and management of virtual resources

Management services are critical to achieving the full benefits of virtualization.

Server Virtualization

Roles:

* Consolidations
* Dynamic provisioning / hosting
* Workload management
* Workload isolation
* Software release migration
* Mixed production and test
* Mixed OS types/releases
* Reconfigurable clusters
* Low-cost backup servers

Benefits:

* Higher resource utilization
* Greater usage flexibility
* Improved workload QoS

* Higher availability / security
* Lower cost of availability
* Lower management costs
* Improved interoperability
* Legacy compatibility
* Investment protection

Virtualization benefits take three forms:

* Reduced hardware costs
* Higher physical resource utilization
* Smaller footprints
* Improved flexibility and responsiveness
* Virtual resources can be adjusted dynamically to meet new or changing needs and to optimize service level achievement
* Virtualization is a key enabler of on demand operating environments
* Reduced management costs
* Fewer physical servers to manage
* Many common management tasks become much easier

Server Virtualization Approaches

Hardware Partitioning

Physical partitioning
Sun Domains, HP nPartitions

Logical Partitioning
pSeries LPAR, HP vPartitions

Hypervisor: Type 1

Hypervisor software/firmware runs directly on server.

zSeries PR/SM and zVM

POWER Hypervisor

VMware ESX Server

Open Source Hypervisor (Xen)

Virtual Iron, ScaleMP

Hypervisor: Type 2

Hypervisor software runs on a host operating system.

VMware GSX

Microsoft Virtual Server

Win4Lin

User Mode Linux

Hardware partitioning and hypervisors are the two main implementation approaches.

* Hardware partitioning subdivides a server into fractions, each of which can run an OS.
* Hypervisors use a thin layer of code to achieve fine-grained, dynamic resource sharing.
* Type 1 hypervisors with high efficiency and availability will become dominant for servers.
* Type 2 hypervisors will be mainly for clients where host OS integration is desirable.

System Virtualization – Hierarchy of Methods

Application Containers

Virtual Runtimes

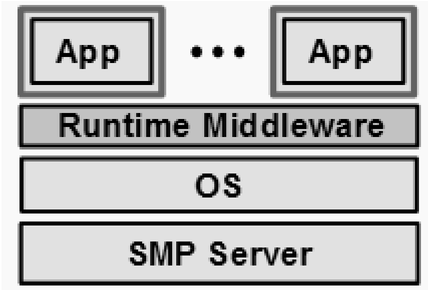

Middleware provides JVM, J2EE, or CLR container per application

Virtual runtimes can be OS independent

Examples:

WebSphere XD

….

Virtual OS Environments

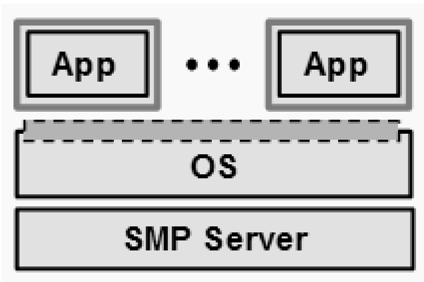

OS and middleware creates virtual OS environment per application

Each container has its own name space, files, root, etc.

Examples:

- z/OS Application Servers
- i5/OS Subsys.
- Solaris Containers

- AIX Corrals (2006-07)
- HP-UX Secure Res. Part.
- Windows & Linux offerings from Softricity, SWsoft, VERITAS, Trigence, …

Benefits:

* Extremely fine granularity
* Best efficiency through single shared OS
* Reduces number of OS kernels and physical servers

Issues:

* Container management(n OS's n AC's + m OS's)
* All apps must support same OS release level
* OS service fixes affect all application containers
* Can't migrate to new OS release until all apps ready
* OS kernel outages limit use, but are infrequent
* Isolation of user spaces only, not kernel level
* Software license cost

Virtual Servers

Virtual Machines, LPAR's

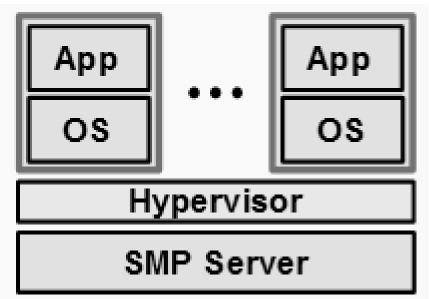

Hypervisor provides virtual server per OS

Virtual servers are isolated but can share hardware resources

Examples:

- zSeries PR/SM and z/VM
- POWER Hypervisor
- Itanium: HP IVM
- x86: VMware, MS VS, Xen, Virtual Iron, ScaleMP, UML, Win4Lin, Jaluna OSware, …

Benefits:

- Very fine granularity
- High efficiency
- Heterogeneous OS types and releases
- Reduces number of physical servers

Issues:

- Availability depends on hypervisor design and integration
- Does not reduce number of systems to manage
- Software license cost
- Hardware Partitions, PPAR's

Virtual Machines, PPAR's

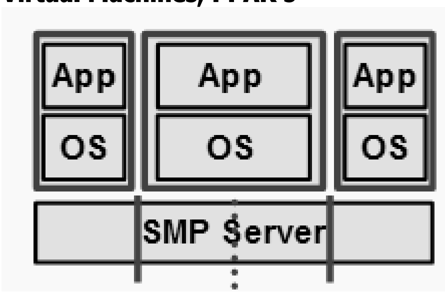

Physical server is divided into fractions to form virtual server per OS

Physical partitions are made of whole boards to achieve electrical isolation

Examples:

- IBM 3033(in 1978!)
- Sun Dynamic Domains,
- HP nPars

Benefits:

- Electrical isolation of partitions isolates most hardware failures
- Independent power on/off per partition

Issues:

- Very limited number of partitions (V=F)
- Very course granularity
- High-end servers only
- Not true HA cluster
- Move workloads instead

Consolidations: Virtual Servers or Application Containers?

Goal: Consolidation - Each Application Has Its Own AIX LPAR, Multiple LPARs Per Server

Fault Isolation: Full isolation: OS faults will only impact the specific LPAR

Security Isolation: Complete isolation between OSes – shared nothing

OS levels: Different OS levels in LPARs possible

OS Service/Fix Level: OS service fixes can be applied to individual LPAR OS images

ISV Software Cost: ISVs will count only the number of processors in LPAR

Efficiency: Good, but little sharing of OS resources leads to lower efficiency than with container approach

System Admin costs: System admin costs don't go up linearly per OS image as cluster sys mgmt tools lower costs for managing multiple OS images

Resource Granularity: Fine grained

Live Migration: Live LPAR migration in 2006

Goal: Consolidation - Each Application Is In Its Own AIX Container, Multiple Containers Per AIX

Fault Isolation: Limited isolation: OS faults will bring down all containers, but happen infrequently

Security Isolation: Only user space level isolation; kernel level is still exposed

OS levels: Has to be the same across all containers

OS Service/Fix Level: OS service fixes affect all containers, negatively impacting multi application - OS dependency

ISV Software Cost: Many ISVs price based on number of CPUs, they will count all processors in OS image

Efficiency: Very good, due to extensive sharing of OS resources such as code or text

System Admin costs: TCO analysis based on number of OS images would favor containers, although in reality sys admin costs will shift to per container level

Resource Granularity: TCO analysis based on number of OS images would favor containers, although in reality sys admin costs will shift to per container level

Live Migration: Live container migration in 2007

Each approach comes with pros and cons, so the choice depends on customer needs and preferences.

POWER Hypervisor With Virtual I/O Servers

Virtualization of POWER5 servers is accomplished using two layers of firmware

- Thin core Hypervisor that virtualizes processors, memory and local networks
- One or more virtual I/O Server partitions that virtualize I/O adapters and devices

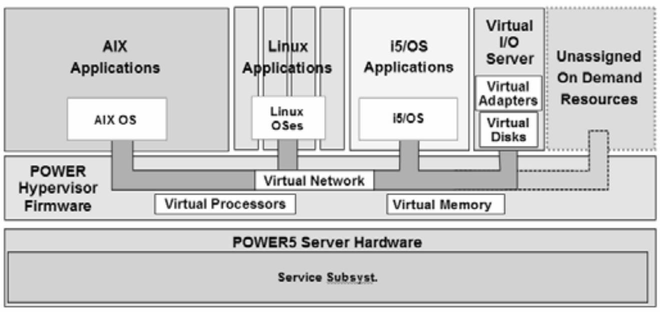

IBM System z offers similar and in some cases more extensive capabilities than POWER5, the implementations though are different.

IBM Servers exploit virtualization to overcome PCI Architecture deficiency – System Z, P, I

Exploiting Virtualization

High Availability Hypervisors

The System z and POWER hypervisors are designed to provide business-critical availability

- Integrated server firmware is optimized for hypervisor role - "lean and mean"
- Restricted function is well tested - not a significant source of failures
- Intelligent machine check processing
- Machine check interrupts are used to localize hardware failures to individual partitions
- CPU Gard™

- If the rate of soft failures on a given processor crosses a threshold, then a spare (COD) processor is substituted transparently or the failing processor is taken offline
- Mature and field-proven technology
- PR/SM widely used since 1988; POWER hypervisor based on iSeries hypervisor, since 1999

System z - The Virtualization and Availability Gold Standard

Continuous Availability via Clustering and Hypervisors

High-availability clustering is required to achieve continuous availability.

Hypervisors and HA clustering may be combined to achieve the benefits of both.

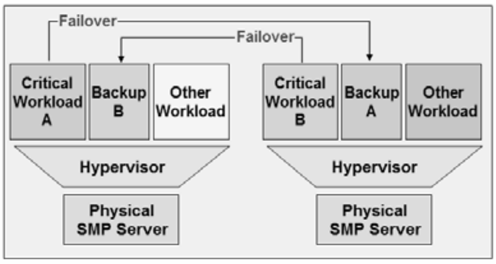

- Each critical virtual server has a backup in a different physical server.
- Backup virtual servers can consume little resource while idle.
- At failover, backup virtual servers can grow automatically by taking resources from lower priority virtual servers.

Security Role of Hypervisors

General purpose operating systems are a weak foundation for secure computing.

- Large and complex with many latent bugs
- Constantly changing – security certifications are quickly made irrelevant
- TCP/IP stacks make them vulnerable to attack by viruses, worms, and hackers

Hypervisors will be used to establish a solid foundation for secure computing.

- Small enough to be fully inspected and certified
- Limited functionality – may rarely need to be changed – suitable as a BIOS extension.

Server Virtualization: Virtual Storage Servers

Storage servers will leverage hypervisors for improved integration, resource utilization, usage flexibility, availability, security, serviceability, and manageability.

Virtual servers (LPARs) will be used within storage servers for multiple roles:

- Virtual storage servers – well isolated for mixed production, development, and test roles
- Storage management servers – Tivoli Storage Manager and others
- File / tape servers – CIFS/NFS gateway, SAN File System, Virtual Tape Server, …
- Virtual I/O servers – provide virtual I/O adapters and networks
- Security services – rooted in trusted platform modules
- Cluster services – for storage server scale out of DS-8000 family servers
- Database applications – leveraging high performance local DMA

Virtual storage servers will be improve storage server availability and serviceability.

- They will enable new storage server failover capabilities, the addition of new controller functions without destabilizing the core code base, and less disruptive service procedures.

Virtualization Management

Emergence of Multi-Server Virtualization

- Server virtualization will be extended in scope from single servers to aggregations of servers, storage, and network components.
- Make a large system look like many – partitioning technology
- Make many small systems look like one from a management perspective

Extends scope of virtualization to aggregations of servers

Shared resource pools will be extended from a single SMP server to multiple servers.

- Movable virtual servers – from one physical server to another dynamically with ease.
- End-to-end virtual networks – span physical networks and hypervisor-based networks.
- Distributed workload management, provisioning, resource optimization – spans multiple servers.
- Applications will no longer be tied to a single physical server

Summary

Virtualization of servers, storage, and networks will become pervasive

- It provides high customer value by reducing costs and improving responsiveness
- "Companies that don't leverage virtualization technologies will pay 40% more in acquisition and 20% more in administration" - Gartner Group, 2004 Fall Symposium

Hypervisors will be the most important server virtualization approach

- They provide fine-grained resource sharing and are easy to apply within scale-out datacenter environments

The greatest ultimate benefit of virtualization will be simplified IT management

- Management labor costs generally far exceed server hardware costs
- Virtualization will reduce management costs by simplifying management tasks

IBM will extend its leadership in virtualization

- The zSeries and Power hypervisors will continue to lead the industry.

Virtualization will simplify IT Management

- Automation software will optimize QoS achievement and resource utilization
- Virtual servers are fully self-describing including processor and I/O architecture => durable entities
- Integration of these technologies within servers will make them easier to use
- Application containers, will provide functions above the OS
- Pool of physical servers (blades, drawers, frames)
- Several generations with evolving architecture
- Same general architecture family (all Power, all x86, …)
- Hypervisor and "VACs" integrated with each server

The Virtualization Engine™ will provide a full, unified virtual systems environment in support of automated, responsive, and efficient On Demand data centers.