Best practices for mixing workloads on servers

Michael Shaler
Product Manager
Microsoft Corporation

Agenda

- Requirements
- Solution strategy
- Architecture
- Advantages
- Licensing
- Support
- Availability

IT Spending – Opposing Forces

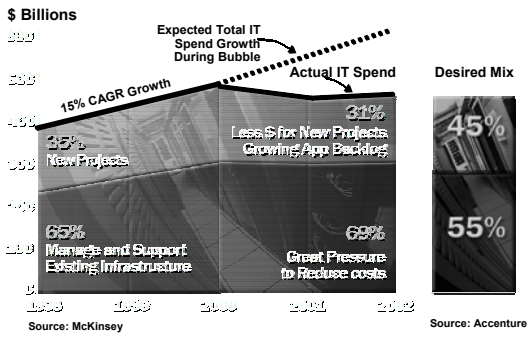

Today's infrastructure and operations are not aligned with business objectives

- Infrastructure is complex, brittle, inflexible
  - "One server, one app"
  - Server sprawl → underutilized servers
  - Diminishing hardware support for NT4
- Departmental line of business server-based applications drive operational costs
  - Care/feeding of application servers consumes disproportionately high system administrator time
  - High cost/risk to upgrade ISV/custom applications

Do More with Less

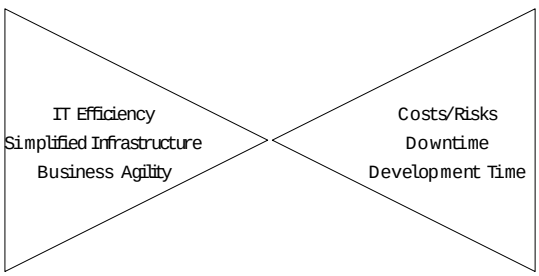

Customer scenarios

NT4 application migration

Departmental/branch server consolidation

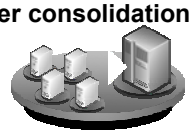

Test/development & training lab automation

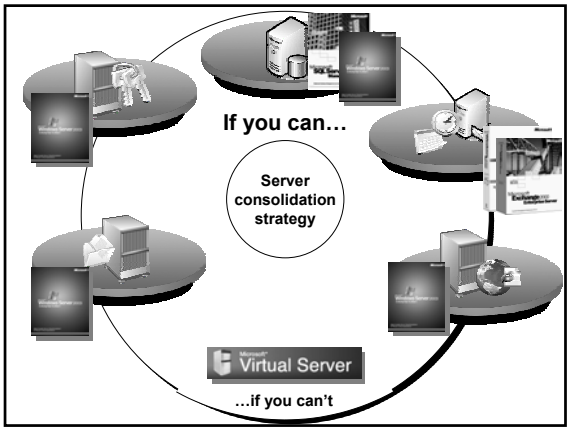

Server consolidation strategy

	Server role	Consolidation strategy	Consolidation scenario
If you can…	Database	Windows Server 2003 and SQL Server 2000	Enterprise customer consolidates enterprise databases on ES7000 running test and production WS03/SQL partitions
	Email	Windows Server 2000 and Exchange 2000	Medium-sized business consolidates email servers via Exchange on scalability cluster
	Web	Windows Server 2003 and IIS 6.0	Hosting service provider consolidates legacy web applications on IIS6 blade farm
	File/print	Windows Server 2003	Small-sized business consolidates on single file/print server WS03 using NAS
	Enterprise applications	Windows Server 2003 and WSRM	Medium-sized customer consolidates entire ERP application suite onto clustered 8-way systems running WS03 and WSRM
…if you can't	**Departmental/ legacy apps**	**Windows Server 2003 and Virtual Server**	**Enterprise customer migrates multiple NT4 applications on 1,000 stand-alone servers onto 50 4-way servers running VS**

Virtual machine strategy

- **Customers have been asking for VM solution from Microsoft**
- **Feb. 20 Microsoft acquired Virtual Server and Virtual PC from Connectix**
- **Transition plan**
  - **VM development team → Windows Core OS team, heads-down on security scrub**
  - **Connectix to sell/support VPC until 8/15**
  - **VPC RTMs Q3, VS RTMs Q4**

Virtual Server architecture

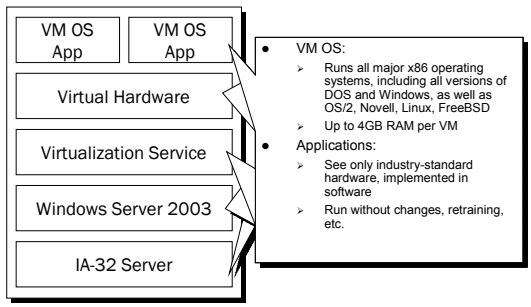

Virtual Server differentiators

- ***Flexibility***
  - **Deploy where needed—decoupled from hardware**
  - **Leverage existing infrastructure (storage/net)**
- ***Manageability***
  - **AD, MMC, perfmon/event log integration**
  - **MOM/SUS/SMS/ADS integration post-RTM**
- ***Automation***
  - **Scripted VM configuration/management**
  - **NT4 Application migration via P2V post-RTM**

Virtual Server features via customer scenario

- **Customer scenario**
- **Isolation**
- **Standardization**
- **Portability**
- **Automation**
- **Connectivity**
- **Manageability**
- **Integration**
- **Performance**

Customer scenario: before

Situation: 3-tier orphaned partner management application

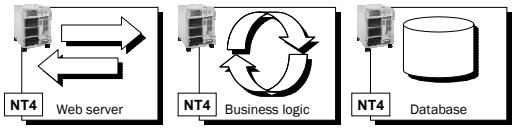

- Tightly coupled to business logic
- Application leaks memory, need rebooting daily
- Highly customized, undocumented code
- ISV out of business (consultants have given up as well)
- Frequent data-dependent crashes
- SLA: > x transactions/ hour with < y latency
- Central IT can't (re)host in datacenter

Virtual Server solution

Situation: 3-tier NT4 partner management application

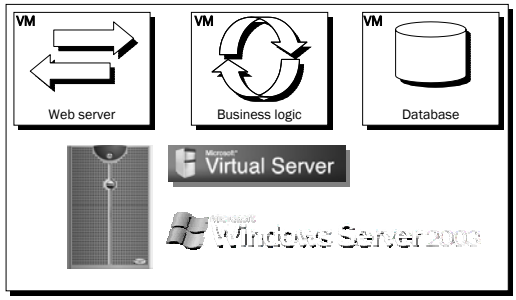

Isolation: Multiple OSes run concurrently on single server

Virtual Device Models

440BX chipset with PIIX4
System BIOS (AMI)
PCI Bus
ISA Bus
Power Management
SM Bus
8259 PIC
PIT
DMA Controller
CMOS
RTC
Memory Controller
RAM & VRAM
COM (Serial) Ports
LPT (Parallel) Ports
IDE/ATAPI Controllers
SCSI Adapters (Adaptec 2940)
SVGA Video Adapter (S3 Trio64)
VESA BIOS
2D Graphics Accelerator
Hardware Cursor
Ethernet Adapter (DEC 21140)
Keyboard
Mouse

- Virtual Server virtualizes:
  - CPU
  - Memory management sub-system
  - Hardware the VM OS sees, from synthetic motherboard on up
- Virtual Server emulates:
  - Device accesses are trapped and emulated in software through virtual device models
- VMs have no access to:
  - Host system physical memory
  - Other VMs' virtual memory or virtualized devices

Standardization: runs every major x86 operating system

- Support: Virtual Server only VM solution available supported by Microsoft
  - PSS supports Microsoft OS/apps in VM
  - Customer support for other OSes/apps comes from OSV/ISV support provider

VM Operating Systems

MS-DOS	OS/2 Warp
PC DOS	OS/2 LANManager
Windows 1.0	NetWare 5
Windows 2.0	NetWare 6
Windows 3.11	NetWare 7
Windows 95	RedHat Linux
Windows 98	SuSE Linux
Windows Me	Turbo Linux
Windows NT	Slackware Linux
Windows 2000	Mandrake Linux
Windows XP	FreeBSD
Embedded XP	NetBSD
Windows Server 2003	OpenBSD
Windows CE	Solaris 7
	Solaris 8
	OpenStep
	Darwin

Memory features in action

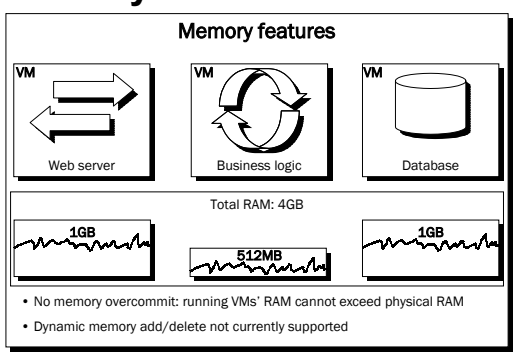

Workload management features

- Multithreaded service
  - 1 VM ≤ 1 physical CPU
  - Each VM operates in own thread of execution, I/O occurs in child threads
- CPU resource allocation policies
  - Weight-based (priority)
  - Constraint-based (min/max)

CPU features in action

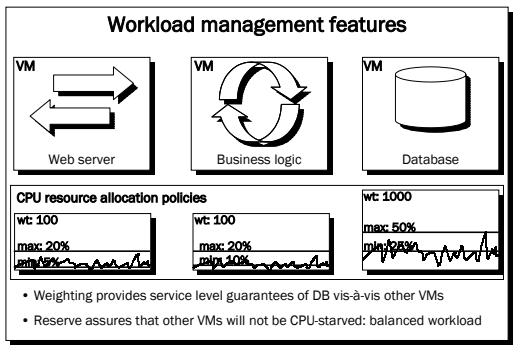

- Weighting provides service level guarantees of DB vis-à-vis other VMs
- Reserve assures that other VMs will not be CPU-starved: balanced workload

Portability → flexibility

- **VHD contains VM data, and XML file contains all configuration metadata**
  - **Dynamic drives**
  - **Fixed-size drives**
  - **Differencing drives**
  - **Undo drives**
- **Benefits: fast, flexible, economical deployment**
  - **Leverages existing storage, networking, security and management infrastructures**
  - **VHDs on SAN enable effective VM-based disaster recovery**

VHD usage scenarios

Drive type	VM sees	System OS sees	Differencing drive?	Undo?	Benefit
Dynamic drives	16GB (default)	Sparse file: 32kb initially, growing as needed. COM API fires alert on low disk warning.	Yes, as parent.	Yes.	Flexible, portable, simple
Fixed-size drives	16GB (default)	16GB. COM API fires alert on low disk warning.	Yes, as parent.	Yes.	Quotas/capacity planning policies
Differencing drives	Single drive (read-only parent/read-write child)	COM API fires alert on low disk warning.	N/A	No.	Powerful DR capabilities (child back up/restore)
Undo drives	Single drive (read-write parent/read-write child)	Parent drive plus saved state file	No.	N/A	Development/test/ demonstration
Saved state	"RAM-to-disk"	N/A	No.	N/A	Hibernation→ failover

Virtual Hard Drives in action

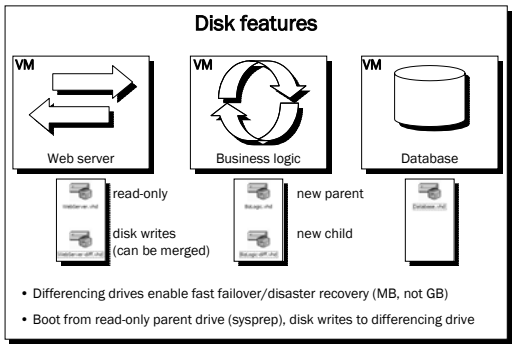

- Differencing drives enable fast failover/disaster recovery (MB, not GB)
- Boot from read-only parent drive (sysprep), disk writes to differencing drive

Automation features

- **COM API enables extensible platform for automation**
  - **Web console as "reference implementation"**
  - **Fully-documented: 28 classes and 363 calls**
  - **Support for all COM-capable languages**
  - **Partners are leveraging in management solutions**
- **XML configuration files**
  - **Extensible VM descriptors**
  - **Management solutions can leverage/enrich metadata**

Automation in action

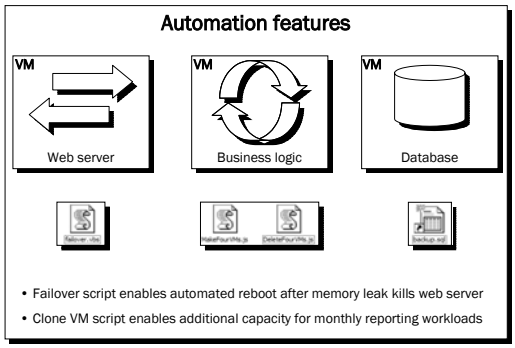

- Failover script enables automated reboot after memory leak kills web server
- Clone VM script enables additional capacity for monthly reporting workloads

Virtual Networking features

- Virtual Networks (VM to any)
  - Up to 4 Virtual NICs per VM
  - Each NIC connects to any virtual network
    - Can be bridged to a host Ethernet adapter
    - No custom drivers needed for VM OS
    - Includes support for teamed NICs
  - Virtual switch performs local and external routing
- Local-only Networks (VM to VM)
  - Uses Virtual DHCP server
  - No host NIC connection—no packets on wire
  - All routing local to VS

Virtual Networks in action

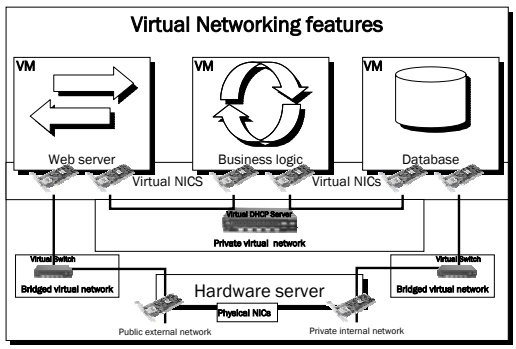

Manageability features

- Event Logging
  - Integrated with host Windows event log
  - Integrates with management solutions
- Performance Monitoring
  - Integrated with host PerfMon
  - Multiple counters available: CPU, RAM, heartbeat, etc.
  - Integrates with management solutions

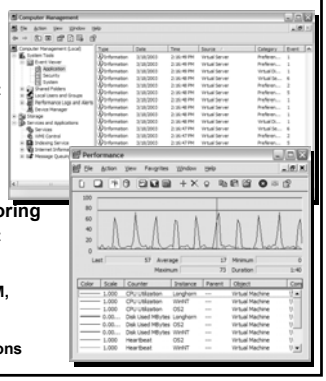

Virtual Server demo

Virtual Server benefits

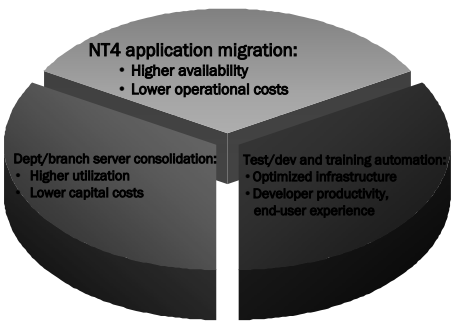

Licensing and support

- Licensing
  - VS is licensed on per-server basis, pricing not yet announced (*strongly* value-based)
  - Same requirements as standalone servers, licensed per installed OS and application
  - Incremental: legacy CALs upgrade required
- Support
  - PSS supports Microsoft OS/apps in VM
  - Support for other OSes comes from current support provider