# 雲端技術架構企業創新動能

VMware Steve
ssue@vmware.com

## 雲端技術為IT下一個技術發展的里程碑

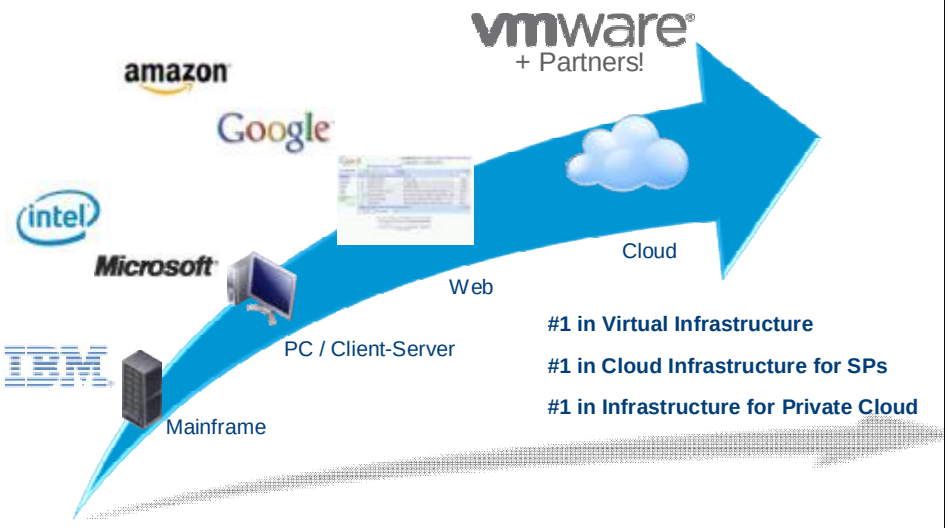


vmware

Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

## 雜亂的機房增加了管理的複雜性




vmware

Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

## 資料中心營運問題

§ IT常遇到……

- 效能使用率不平均或者過低
- 硬體維修造成停機時間拉長
- 沒有 HA 的建置，或 HA 建置與管理複雜度高
- 常常得解決硬體型號不同、驅動程式不同的問題
- 面臨效能瓶頸的舊硬體不易升級
- 過長、重複的系統部署時間
- 補丁、升級程式就頭痛，失敗時回退困難
- 災難回復計劃成本高、不好實踐與維護
- 機房空間、電力與空調不足
- 一直買新硬體

vmware

Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

# 虛擬化為雲端運算的基礎技術


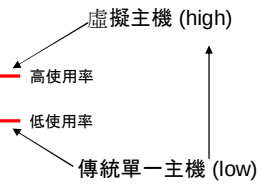


–現在的電腦效能相當強大，但卻有非常多的電腦平常只有使用不到20%的運算資源，甚至更少。但因為種種需求，又無法降低電腦數量，因此把數台電腦“虛擬化”塞進一台功能較強大的實體主機，榨乾實體主機的運算資源，成了一種受歡迎的方法。

vmware

Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

## 虛擬機的概念

§ 虛擬機可允許一台實體主機同時執行多個作業系統。

§ 一台強大電腦主機能做30台電腦的事，在佔用面積、耗電量、空調、維護人事成本上，都能有效節約。

§ 可在一台實體主機內執行多個虛擬主機，每一台虛擬主機既可以互相用線路連線、又可以獨立運作，互不干涉。

§ 可將多台實體主機合併成一台大的資源池(Resource Pool)，統籌管理、分配在上面的虛擬機。

vmware

Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

## 虛擬機的概念

- 只要實體主機的硬體能力許可，便能靈活地建立或移除多台虛擬機，可保持高度彈性。
- 虛擬主機往往在映像檔內直接執行、存取的檔案、作業系統，使得備份、搬移虛擬機較實體主機容易許多。
- 虛擬化可使實體主機擴增、維修、更換不須關閉虛擬機，只要把裡面的虛擬機搬移到別的虛擬機即可。
- 可使用虛擬機測試不穩定的軟體，再利用快照完全還原整個系統(僅需數秒鐘)，而不須在實體主機測試。

vmware

Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

# 何謂虛擬化

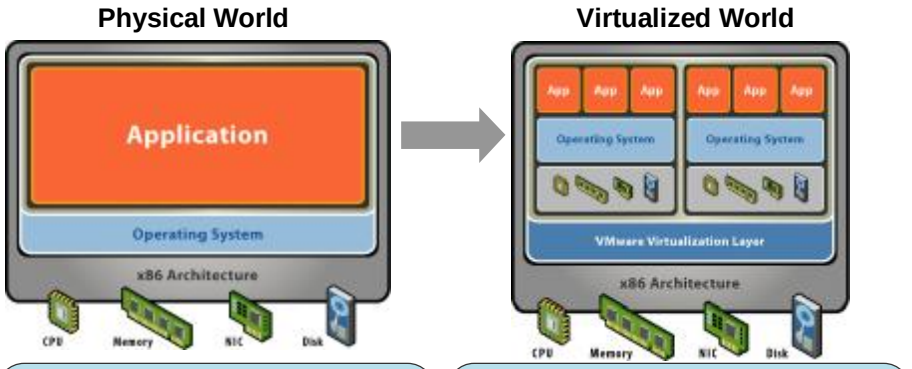


**Traditional x86 Architecture**

- 傳統是一台主機執行一套應用程式且整體使用率偏低平均約8%

**Virtualization:**

- 新一代資訊中心是在單一主機執行多種應用程式讓整體使用率提高到 80%

vmware

Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

## VMware平台支援六十種以上的作業系統

- Windows NT 4.0
- Windows 2000
- Windows Server 2003
- Windows Server 2008
- Windows Vista
- Windows XP
- Windows 7
- RHEL5
- RHEL4
- RHEL3
- RHEL2.1
- SLES 11
- SLES10
- SLES9
- SLES8
- Ubuntu 7.04
- Solaris 10 for x86
- NetWare 6.5
- NetWare 6.0
- NetWare 6.1
- Debian
- CentOS
- FreeBSD
- Asianux
- SCO OpenServer
- SCO Unixware
- …

vmware

Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

# 虛擬機的概念

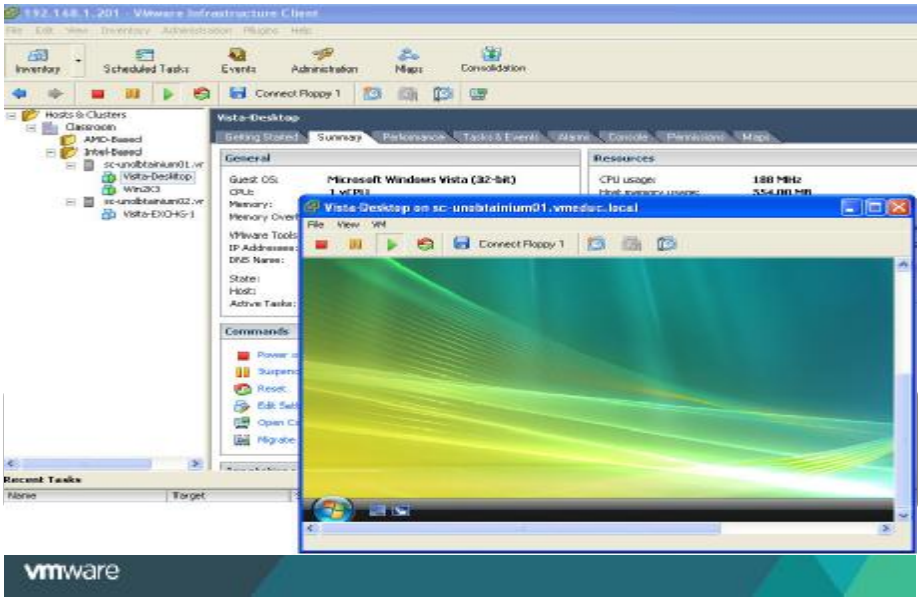


Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

## 虛擬機的概念

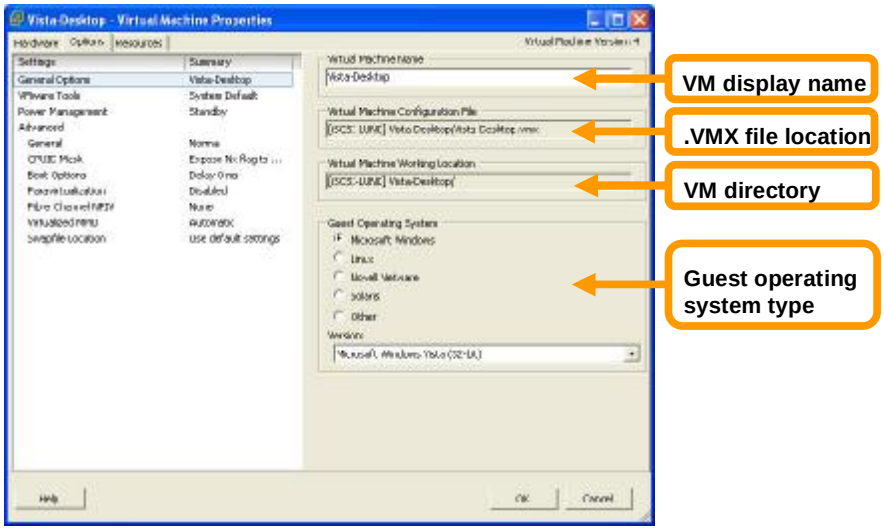


Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

- The General Options can be used to modify things like the display name used for the VM and the type of guest operating system installed. The location and name of the configuration file (.VMX file) is displayed and the location of the virtual machine's directory is also shown. You can select the text for the configuration file and working location if you need to cut and paste them into a document. But only the display name and the guest operating system type may be modified.

- Note: If you change the display name that is not going to change the names of all of the VM files or the directory the VM is stored in. When a VM is first created the file names and the directory name associated with the VM are based on its display name. But changing the display name later does not modify these file and directory names.

# P2V 實體機轉虛擬機

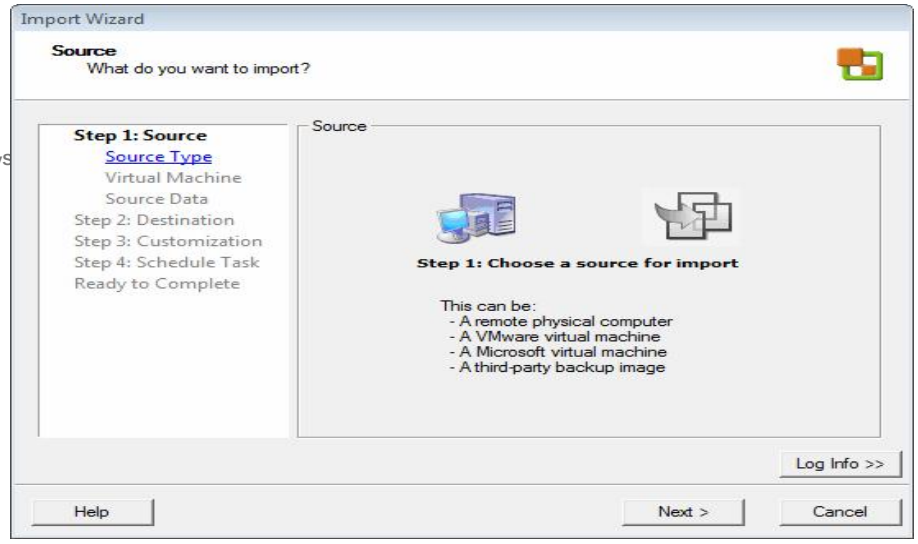


Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

# 映像檔轉虛擬機

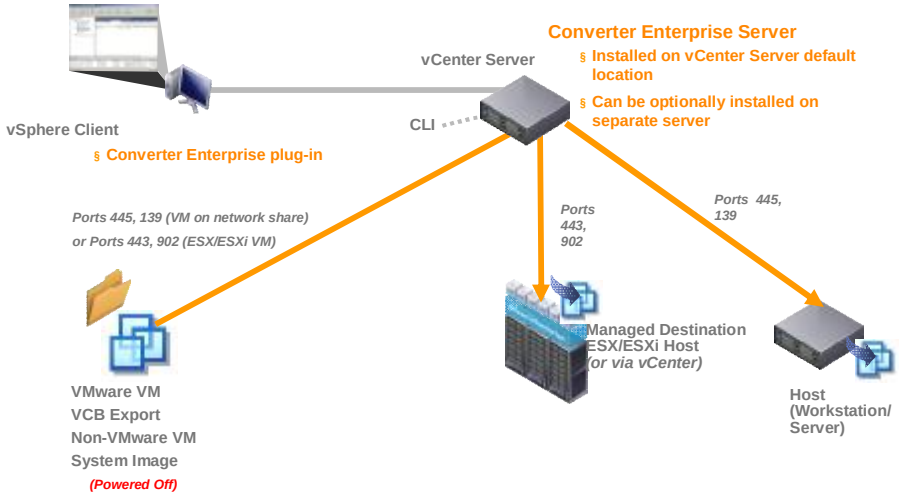


Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

## 虛擬機的成長首次超越實體機

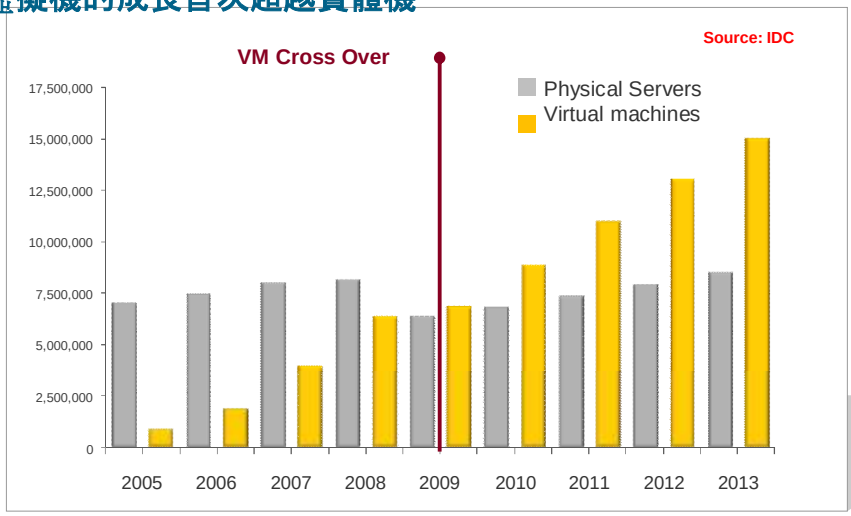


**We are past a virtual tipping point!**

vmware

Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

# Gartner: “VMware 是市場唯一領導廠商”

***“VMware stands alone as a leader in this Magic Quadrant”***


- “VMware is clearly ahead in”:
  - *Understanding the market*
  - *Product strategy*
  - *Business model*
  - *Technology innovation, Product capabilities*
  - *Sales execution*
- “VMware Strengths”:
  - *Far-reaching virtualization strategy enabling cloud computing, new application architectures and broader management*
  - *Technology leadership and innovation*
  - *High customer satisfaction*
  - *Large installed base (especially Global 2000), and rapid growth of service providers planning to use VMware (vCloud)*

vmware

Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

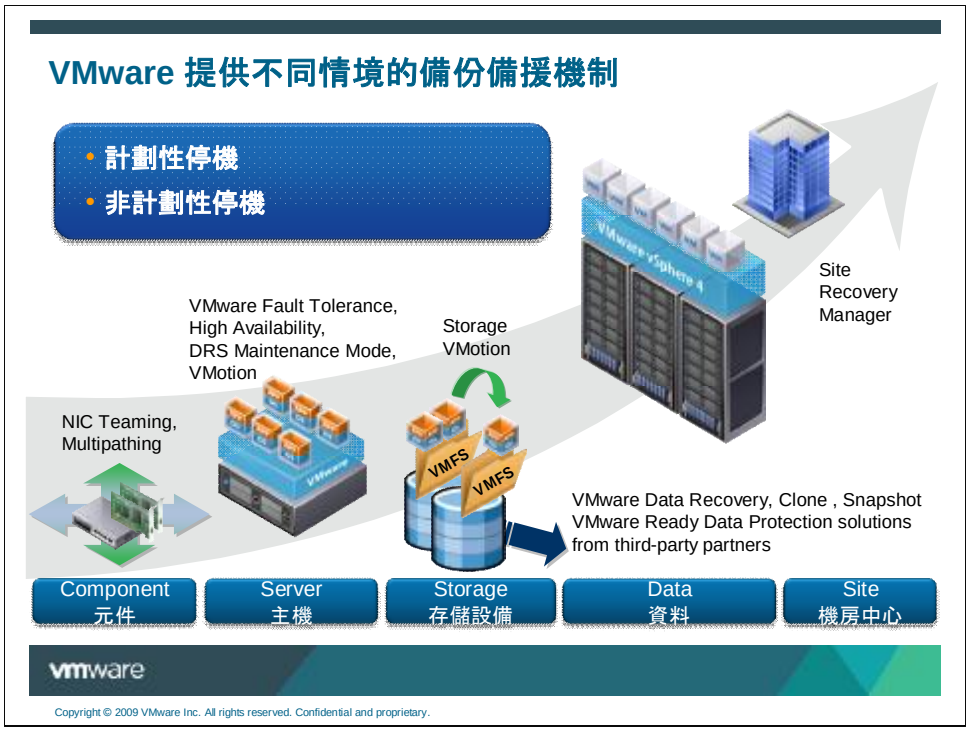


- **The power behind VMware's BC/DR solution is how we layer our protections. At every level of the datacenter, from individual components all the way up to the entire site, we provide protection against both planned and unplanned downtime.**

- **Many of the key properties of virtualization, such as encapsulation and hardware independence, already offer some inherent protections. From there, we provide additional protections throughout our platform to ensure your organization can meet its availability requirements.**

- **We'll be covering all of these features and products in today's presentation.**

## VMotion特性

- VMotion
  - 可線上將虛擬機器從其中一台ESX主機轉移到另一台
  - 此技術可應用在其他備援技術
    - Fault Tolerance
    - Storage VMotion
    - DRS and DPM


vmware

Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

# Enhanced VMotion Compatibility (EVC)特性

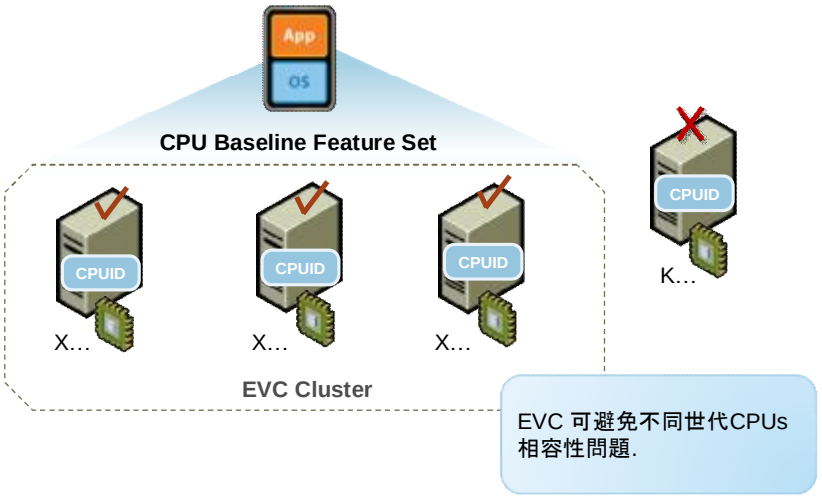


EVC 可避免不同世代CPUs相容性問題.

vmware

Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

# 透過EVC讓不同世代的處理器可以相容

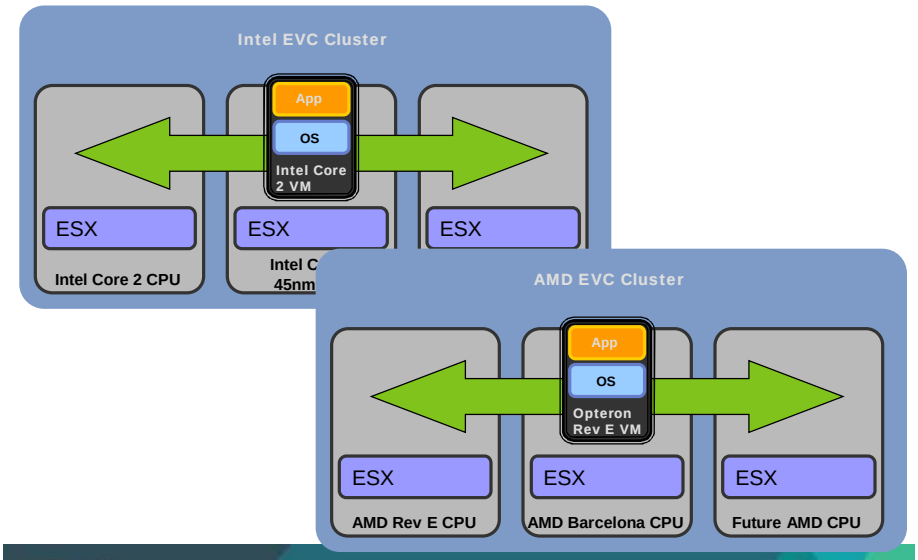


vmware

Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

# vSphere 重要高可用度功能

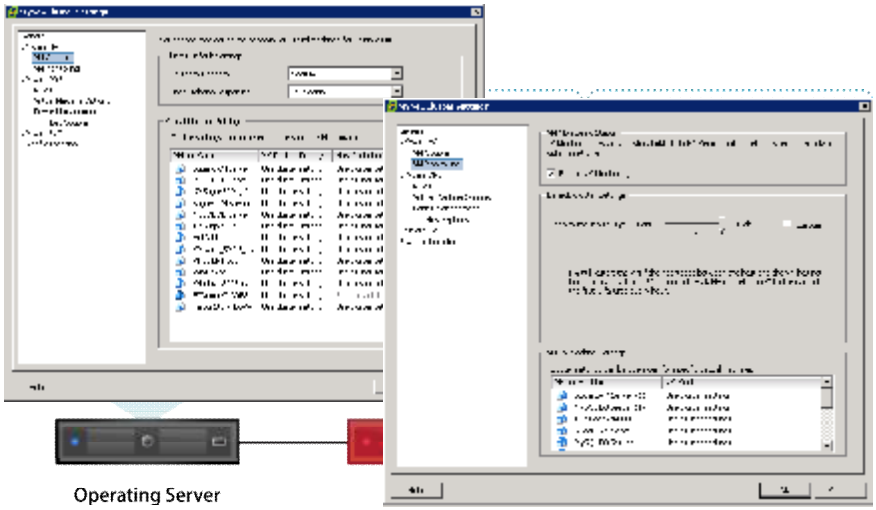


vmware

Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

## HA 診斷和可靠性改善

**HA 健康檢查狀態**

- HA 提供持續健康檢查功能，以確保能隨時符合所需的叢集組態。差異將導致叢集事件或警告。

**HA 運作狀態**

- 全新叢集運作狀態視窗會顯示更多目前 HA 運作狀態的資訊，包括 HA 叢集中每個主機的特定狀態和錯誤。

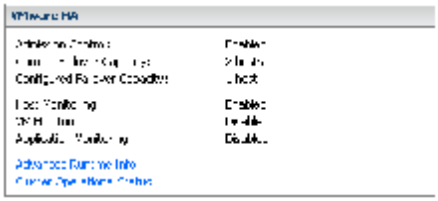

**改善 HA 容錯轉移時的 HA-DRS 互通性**

- DRS 會執行 vMotion 以釋放相鄰資源 (如，於某一主機上)，因此 HA 能夠將資源用於須重新啓動的虛擬機。

**HA 應用程式感知 – 提出 API 供協力廠商應用程式開發人員使用**

vmware

Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

- **VMware HA Healthcheck and Operational Status — The VMware HA dashboard in the vSphere Client provides a new detailed window called Cluster Operational Status, which displays more information about the current VMware HA operational status, including the specific status and errors for each host in the VMware HA cluster.**

# vSphere 重要高可用度功能

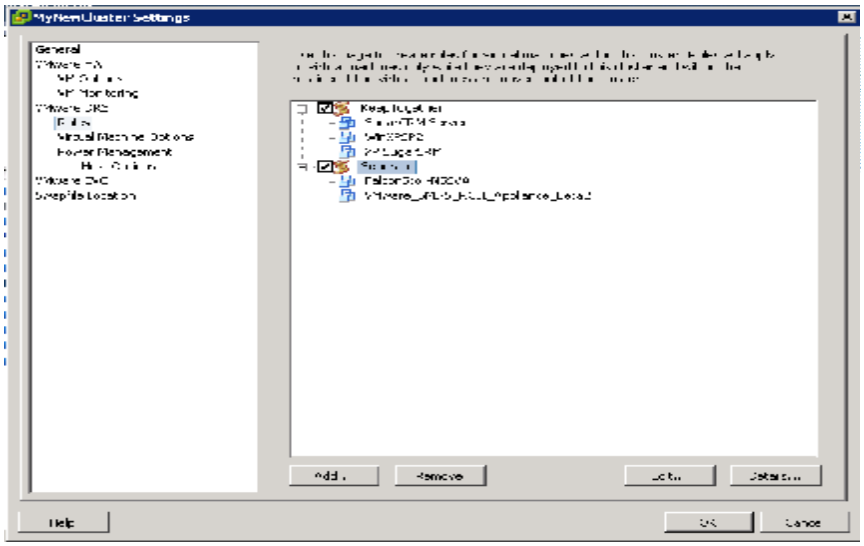


vmware

Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

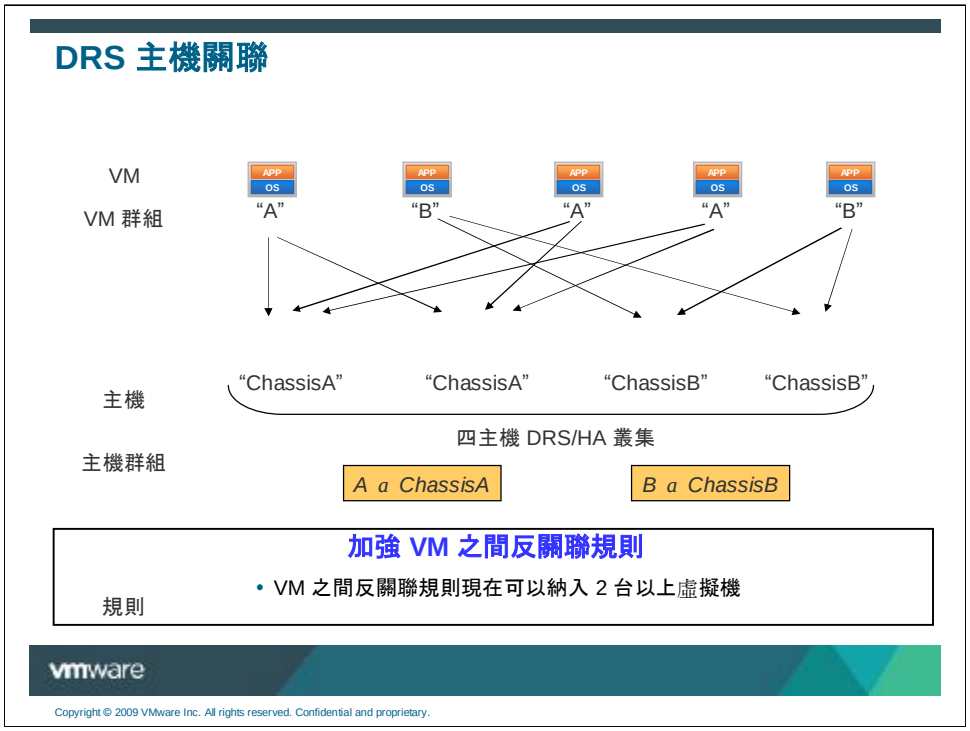


- **Strongly advise customers that these rules are not meant to be used often – the more constraints you put on VM mobility, the harder it is for DRS to balance load and to enforce resource allocation policies. You should only use them if you absolutely have to.**
- **Hard affinity rules are only advised to be used for enforcing host-based licensing of ISV apps. Soft affinity rules are meant for availability reasons – like keeping two VMs on different racks or blade chassis's.**

- **Preferential rules can be violated to allow the proper functioning of DRS, VMware HA, and VMware DPM.**

## DRS 主機關聯

- § DRS Groups
  - DRS Groups Manager
    - § Defines Groups
    - § VM groups
    - § Host groups

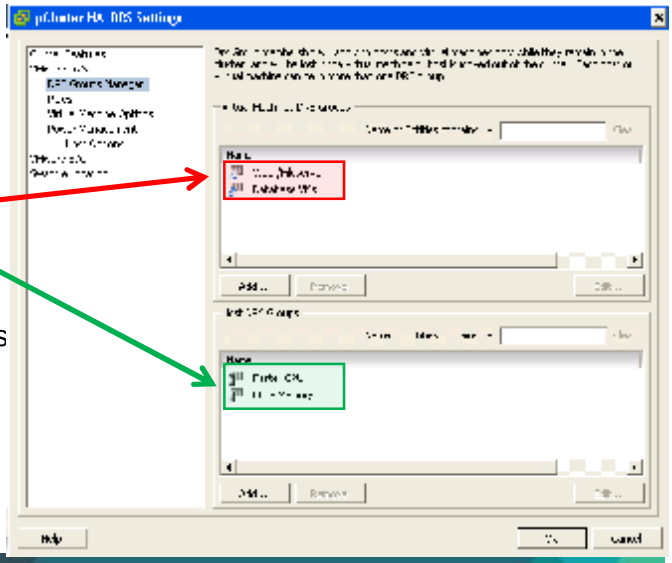


Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

## DPM電源管理加強功能

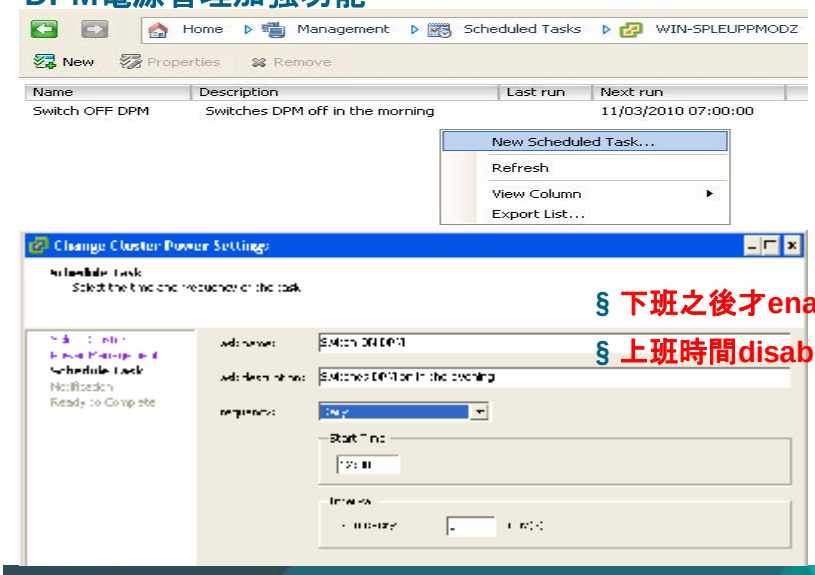


§ 下班之後才enable

§ 上班時間disable

Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

## vSphere 重要高可用度功能

§ Storage VMotion

- 在不同datastore 移轉
- 在不同的存儲設備移轉
- 可改變 VM 磁碟格式 (thick or thin)

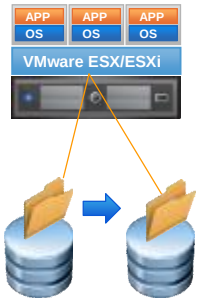


Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

# 線上複製Clone & 快照Snapshot

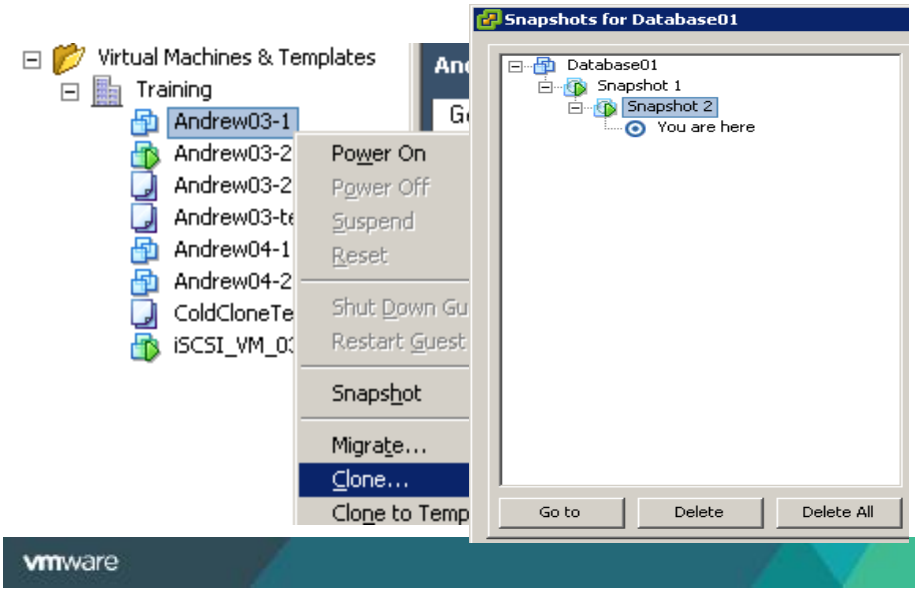


Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

# 範本檔使用

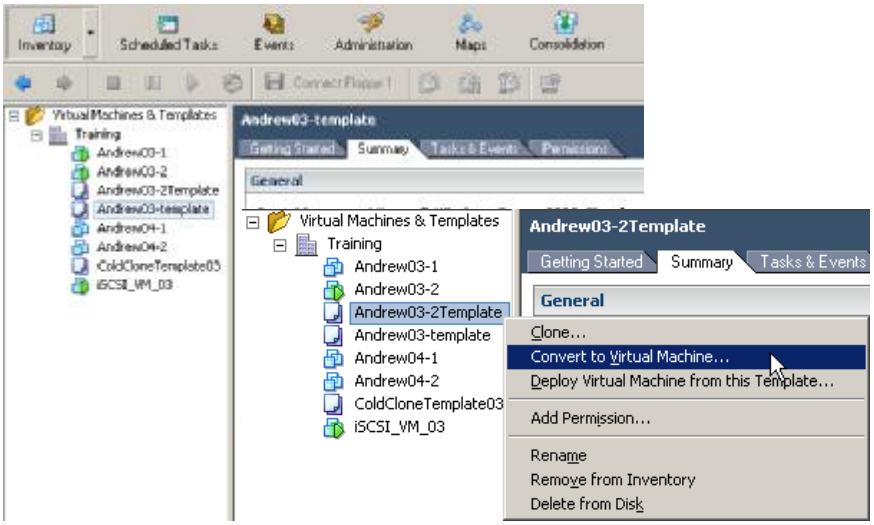


Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

# vSphere 重要高可用度功能

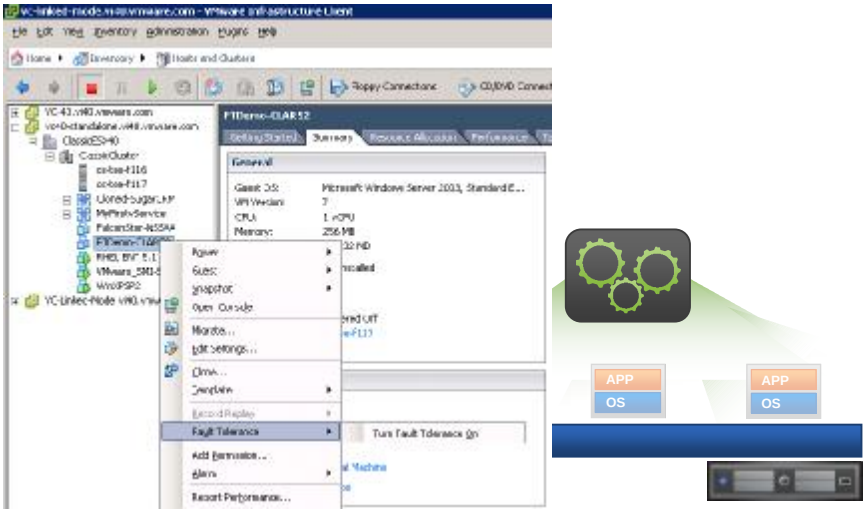


vmware

Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

# vSphere 重要高可用度功能

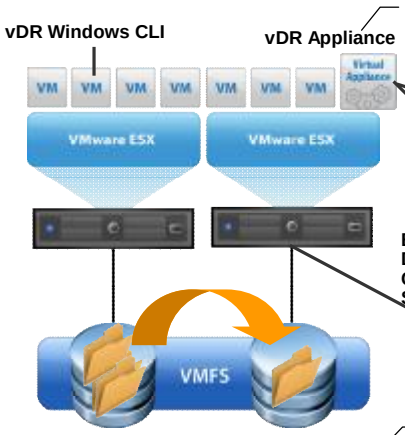


**備份與復原裝置**

- 每一的 vCenter Server 執行個體最多可支援 10 個裝置，因此可以保護達 1000 台虛擬機
- Windows Linux 虛擬機適用的檔案層級復原用戶端

**儲存目的地**

- 擴大支援 DAS、NFS、iSCSI 或光纖通道儲存和 CIFS 權重以當做目的地
- 加強重複刪除效能

vmware

Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

# 改善傳統異地備援架構

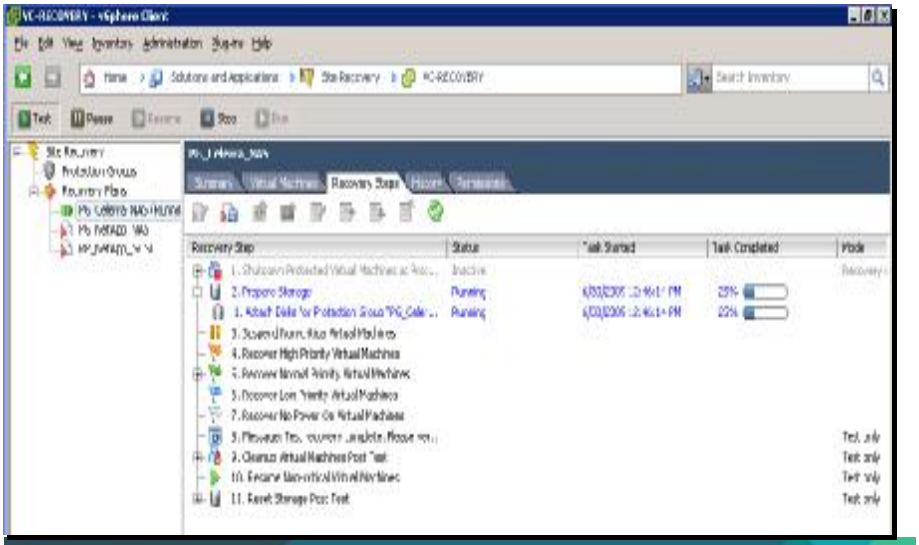

Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

# 改善傳統異地備援架構

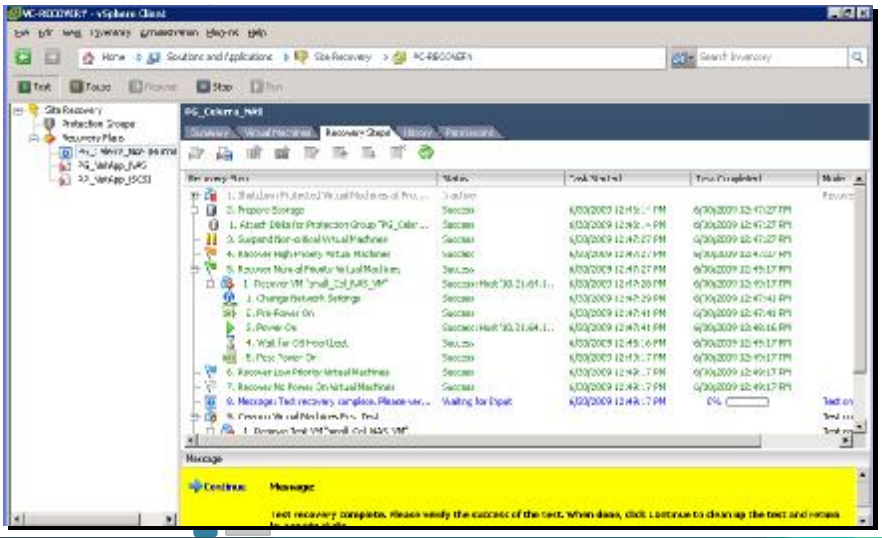


Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

# ESX & VM 主機規格

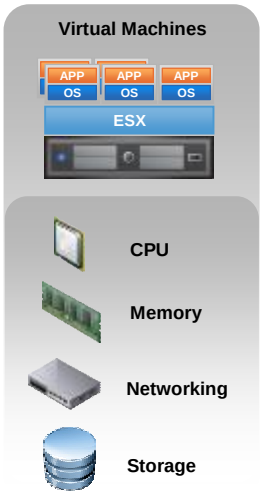


- **VM Scale Up**
  - 8-way vSMP and 255 GB of RAM per VM
- **Hardware Scale Up**
  - 64 cores and 1 TB physical RAM
- **Hardware Assist**
- **Purpose Built Scheduler**
  - Lowest CPU overhead
- **Hardware Assist**
- **Page Sharing**
- **Ballooning**
  - Maximum memory efficiency
- **VMXNET3**
- **VMDirectPath I/O**
  - Wirespeed network access
- **Storage stack optimization**
- **VMDirectPath I/O**
  - Greater than 30k iops per second Lower than 20 microsecond latency

**Current** **NEW**

vmware

Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

## vSphere 4.1 中的 NFS 和 HW iSCSI (儲存) 支援

更好的 NFS 效能

- § 讀取與寫入的 CPU 使用率皆下降達 15%
- § 讀取與寫入的資料流量皆改善達 15% 讀取增加12-40% 寫入增加32-124%

Broadcom iSCSI HW 卸載支援

- § 讀取時 CPU 使用率
  平均僅需原來的 89%！
- § 寫入時 CPU 使用率
  平均僅需原來的 83%

iSCSI SW卸載支援

- § 讀取時 增加原來的
  6-23%！
- § 寫入時 增加 原來的
  8-19%

!

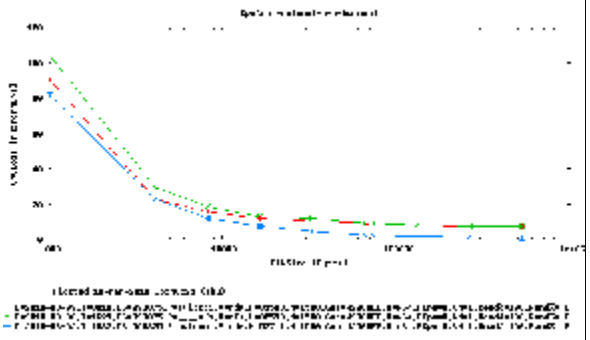

vmware

Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

- **NFS performance has also improved. NFS Read and write has improved by around 15% in CPU cost and throughput.**
- **Broadcom iSCSI HW offload support realizes an 80%+ improvement in CPU read and write cost.**

## 用於陣列整合的 vStorage API (VAAI)

藉由使用有效率的陣列操作，做為 VMware 主機式解決方案以外的另一項選擇，以改善效能

三個基礎功能包括：

1. 完整複製 – 類似 Xcopy 的功能，可卸載工作至陣列
2. 相同寫入 – 加快歸零區塊或寫入重複內容的速度
3. 最小測試與設定 – 鎖定整個 LUN 的替代方式

協助功能包括：

- Storage vMotion
- 從範本佈建虛擬機
- 改善精簡佈建磁碟的效能
- VMFS 權重儲存集區的延展性

- 必須向儲存裝置廠商取得韌體 (已有 6 家廠商參與，請見後續投影片)

vmware

Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

- **Our second major storage feature that is being added in 4.1 is the new vStorage API for Array Integration or VAAI. There are 6 storage partners which are working with VMware to release special firmware for their arrays which can support 3 primitives' that enable off loading certain tasks to the storage array. These primitives are 1) Full Copy, 2) Write Same and 3) Atomic test and set. The full copy enables arrays to make copies of certain virtualization objects with in the array without having to have the ESX server read and write those objects. The write same is useful in zeroing out a large number of blocks to speed up a eager zero thick vmdk. And the ATS provides an alternative means to protect the meta data for VMFS as a cluster file system and there by improving the scalability of large ESX server farms sharing a datastore.**
- **I the 4.1 release, the VAAI will be supported for block bases storage only.**

## Full Copy – Without VAAI

- Current VM cloning and migrations require data to be copied within ESX/ESXi host
- Copy process consumes host CPU, memory, and network bandwidth

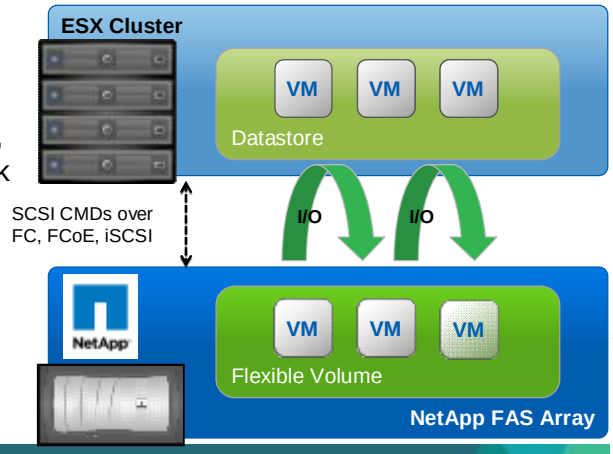


Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

## Full Copy – With VAAI

- With API VM cloning and migrations are instruction sets from ESX/ESXi hosts
- Copy process occurs on array
- Copy process is offloaded from host and network

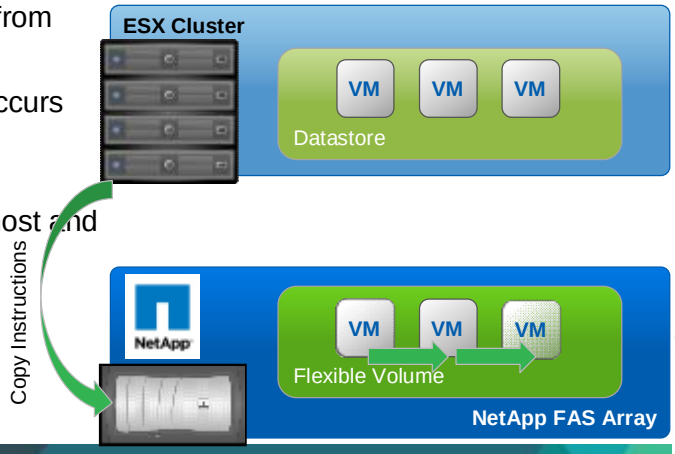


Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

## VAAI 改善 Storage vMotion 速度 – 範例

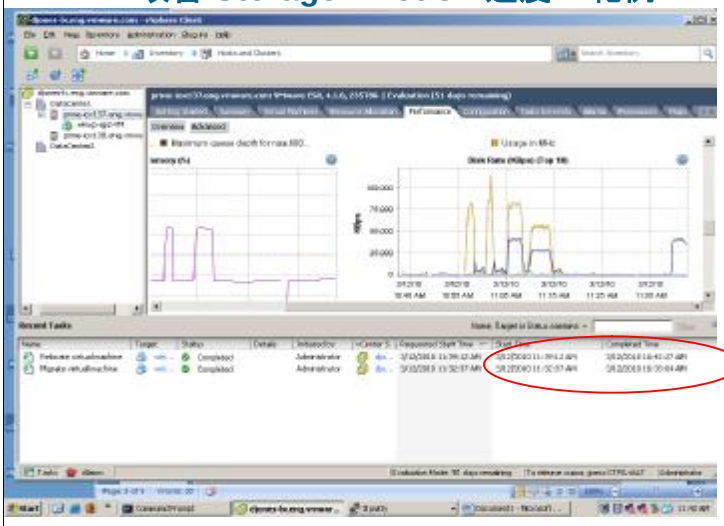

42:27 - 39:12 =
2 分 21 秒 (無 VAAI)
(141 秒)

33:04 - 32:37 =
27 秒 (含 VAAI)

141 秒與 27 秒相較

vmware

Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

- **One example of this was run to compare the time it took with the feature and without use of VAAI. The recent task shows without VAAI the Storage vMotion took 141 seconds and with the VAAI enabled, it took only 27 seconds. A considerable time savings.**

## Block Zeroing – Without VAAI

- Initializing unwritten blocks in various VMDK formats requires zeros to be written from ESX server to storage
- Writing process consumes host CPU and network bandwidth

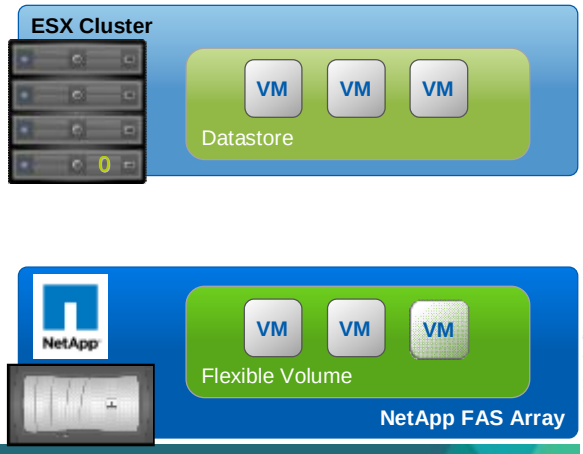


Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

## Block Zeroing – With VAAI

- With the VAAI Block Zeroing API blocks can be initialized by the storage device
- Initialization process is faster and consumes negligible CPU and network bandwidth

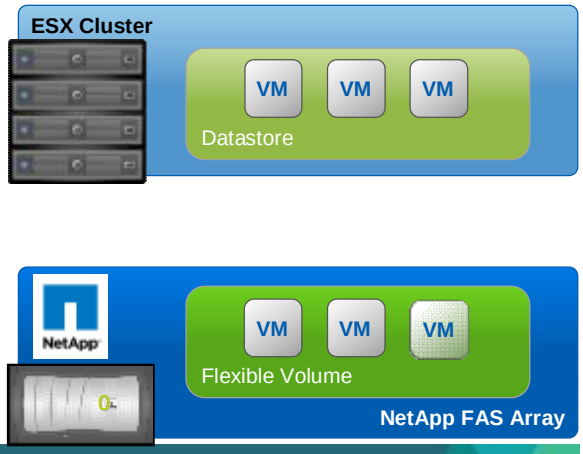


Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

## VMFS Locking – Without VAAI

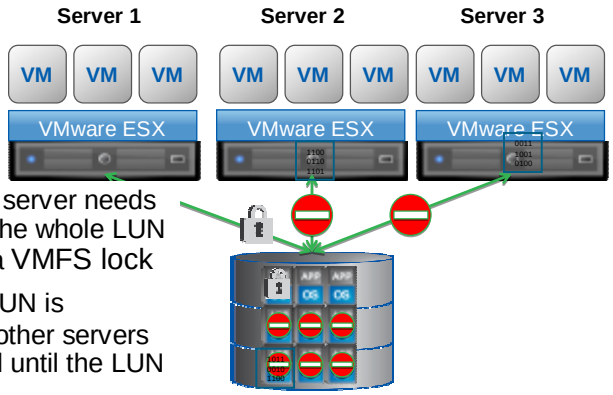


- Currently a server needs to reserve the whole LUN to acquire a VMFS lock
- While the LUN is reserved other servers are blocked until the LUN is released
- This limits performance and scalability

Server 1 wants to start a VM and needs to lock the LUN

vmware

Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

## HW Assisted Locking – With VAAI

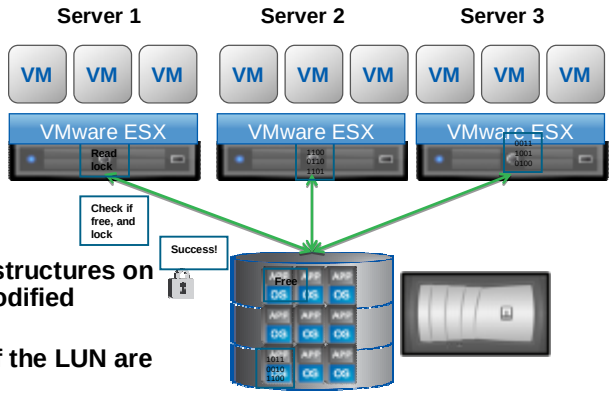


- Lock data structures on disk are modified atomically
- All areas of the LUN are accessible
- Performance and Scalability increase

Server 1 tells storage "If lock still free, lock it for me"
Servers 2 & 3 can still access the LUN

Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

## Memory Compression：記憶體壓縮

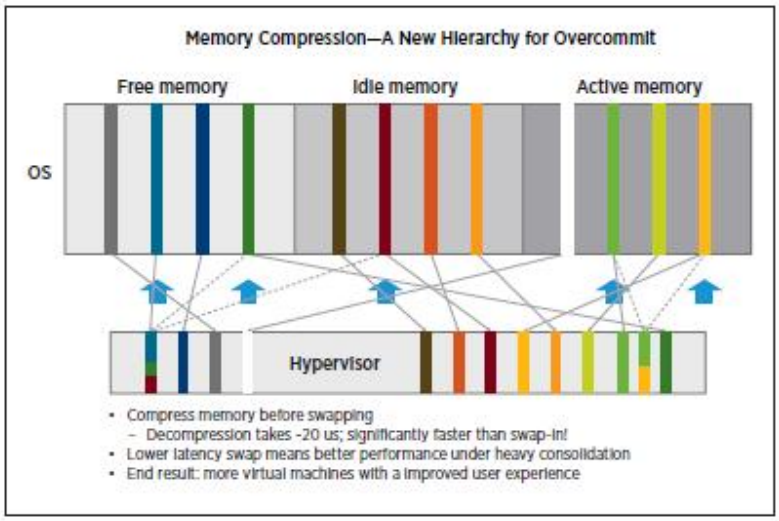


vmware

Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

## 更強大的 vCenter 延展性 – 雲端運算規模

| | vSphere 4 | vSphere 4.1 | 比例 |
|---|---|---|---|
| 每一主機的 VM 數目 | 320 | 320 | 1x |
| 每一叢集的主機數目 | 32 | 32 | 1x |
| 每一叢集的 VM 數目 | 1280 | 3000 | **3x** |
| 每一 VC 的主機數目 | 300 | 1000 | **3x** |
| 每一 VC 的註冊 VM 數目 | 4500 | 15000 | **3x+** |
| 每一 VC 的開啓 VM 數目 | 3000 | 10000 | **3x** |
| 同時進行的 VI 用戶端 | 30 | 120 | **4x** |
| 每一 DC 的主機數目 | 100 | 500 | **5x** |
| 每一 DC 的 VM 數目 | 2500 | 5000 | **2x** |

vmware

Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

44

- 64-bit vCenter removes 4GB virtual address space limitations and improves application scalability. This is evident in the number of VMs and Hosts that vCenter 4.1 can support per VC instance or per data center. You can see that we scale 2x to 5x better than vCenter 4.0. Note that a single vCenter can now support up to 10,000 powered-on VMs and up to 15,000 registered VMs.
- This is important as we plan to match cloud scale deployments with vSphere private clouds.

# 權限控管功能

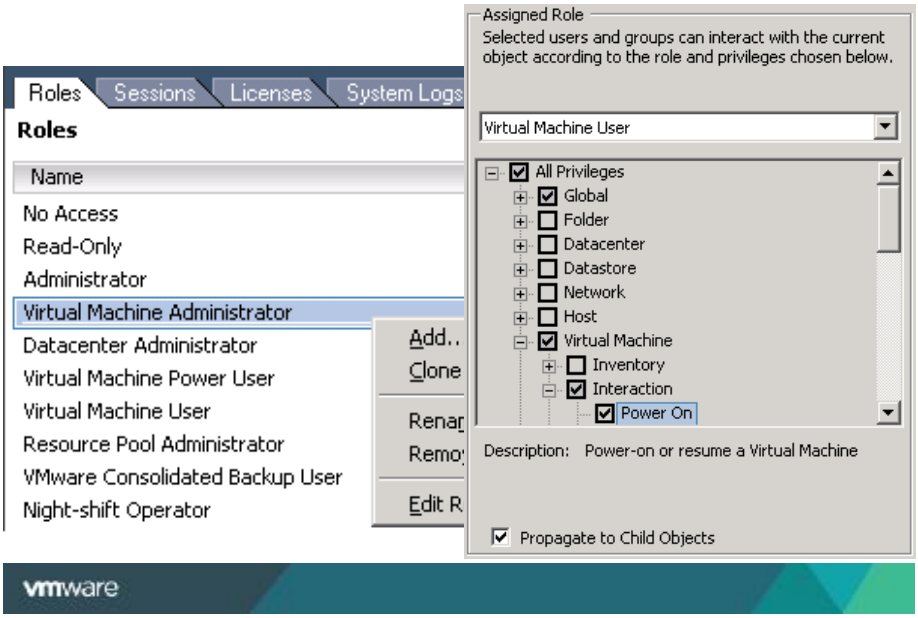


Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

## 安全管理功能

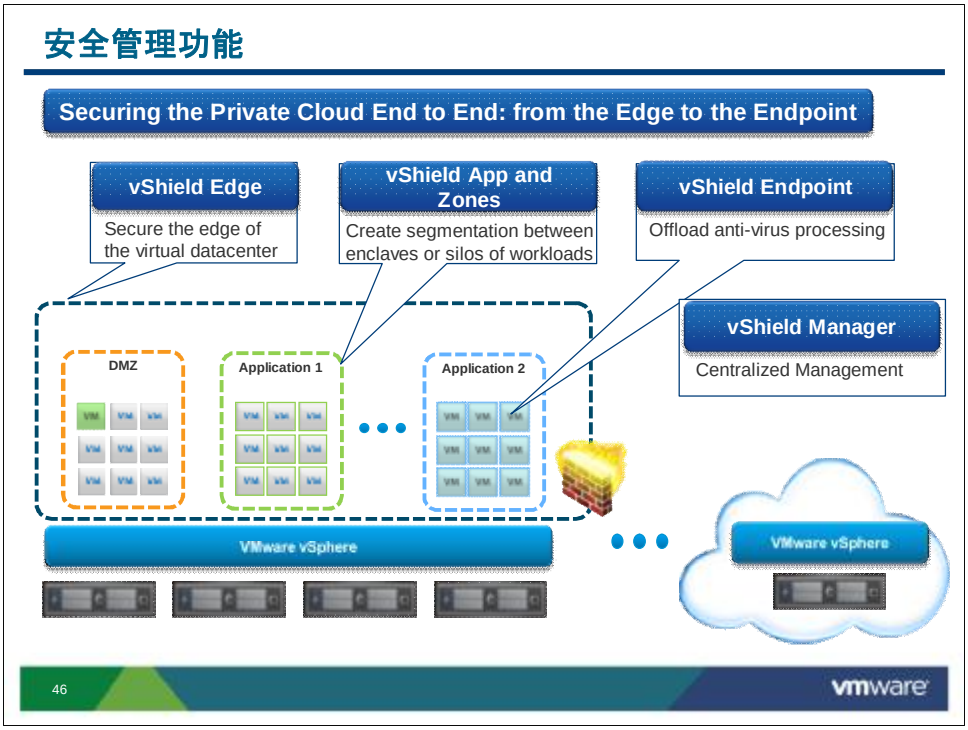
Securing the Private Cloud End to End: from the Edge to the Endpoint
vShield Edge
Secure the edge of the virtual datacenter
vShield App and Zones
Create segmentation between enclaves or silos of workloads
vShield Endpoint
Offload anti-virus processing
vShield Manager
Centralized Management
DMZ
Application 1
Application 2
VMware vSphere
VMware vSphere
46
vmware

# vShield Zones 防火牆及流量管理功能

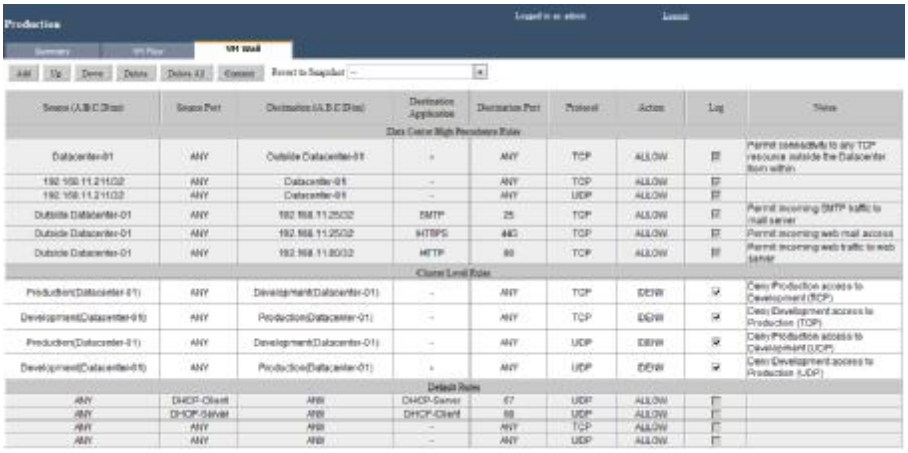




vmware

# vShield Zones 防火牆及流量管理功能

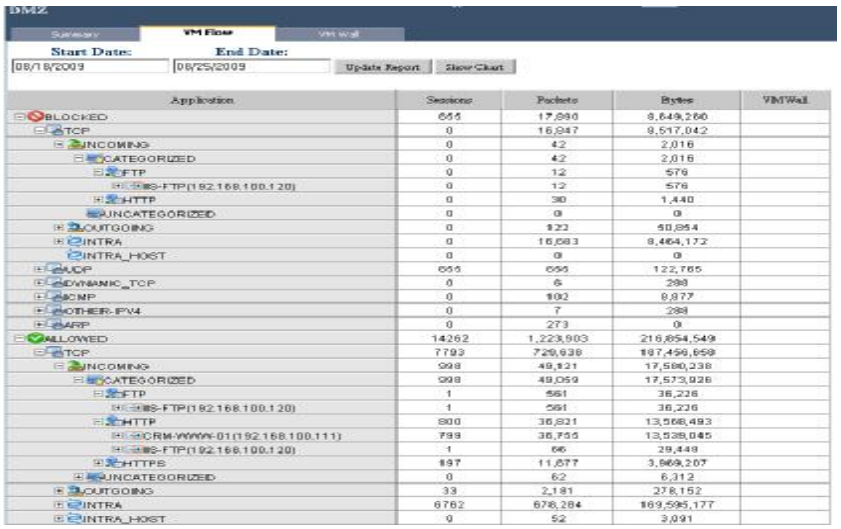


| Application | Sessions | Packets | Bytes | VMWall |
|---|---|---|---|---|
| BLOCKED | 655 | 17,880 | 8,649,280 | |
| TCP | 0 | 16,947 | 8,517,042 | |
| INCOMING | 0 | 42 | 2,016 | |
| CATEGORIZED | 0 | 42 | 2,016 | |
| FTP | 0 | 12 | 576 | |
| FTP(192.168.100.120) | 0 | 12 | 576 | |
| HTTP | 0 | 30 | 1,440 | |
| UNCATEGORIZED | 0 | 0 | 0 | |
| OUTGOING | 0 | 122 | 50,854 | |
| INTRA | 0 | 16,683 | 8,464,172 | |
| INTRA_HOST | 0 | 0 | 0 | |
| UDP | 655 | 655 | 122,765 | |
| DYNAMIC_TCP | 0 | 6 | 288 | |
| ICMP | 0 | 102 | 8,977 | |
| OTHER-IPV4 | 0 | 7 | 288 | |
| ARP | 0 | 273 | 0 | |
| ALLOWED | 14262 | 1,223,903 | 216,854,549 | |
| TCP | 7793 | 729,938 | 197,456,658 | |
| INCOMING | 998 | 49,821 | 17,580,238 | |
| CATEGORIZED | 998 | 49,059 | 17,573,926 | |
| FTP | 1 | 561 | 36,226 | |
| FTP(192.168.100.120) | 1 | 561 | 36,226 | |
| HTTP | 800 | 36,821 | 13,568,493 | |
| CRM-WWW-01(192.168.100.111) | 799 | 36,755 | 13,539,045 | |
| FTP(192.168.100.120) | 1 | 66 | 29,448 | |
| HTTPS | 197 | 11,677 | 3,969,207 | |
| UNCATEGORIZED | 0 | 62 | 6,312 | |
| OUTGOING | 33 | 2,181 | 278,152 | |
| INTRA | 6762 | 678,284 | 169,595,177 | |
| INTRA_HOST | 0 | 52 | 3,091 | |

vmware

Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

# 運算能力可隨企業業務需求作彈性擴充

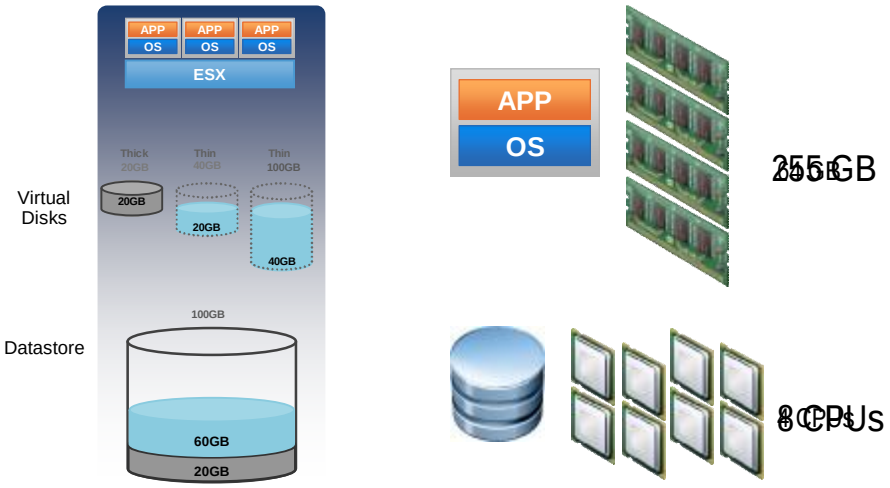


Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

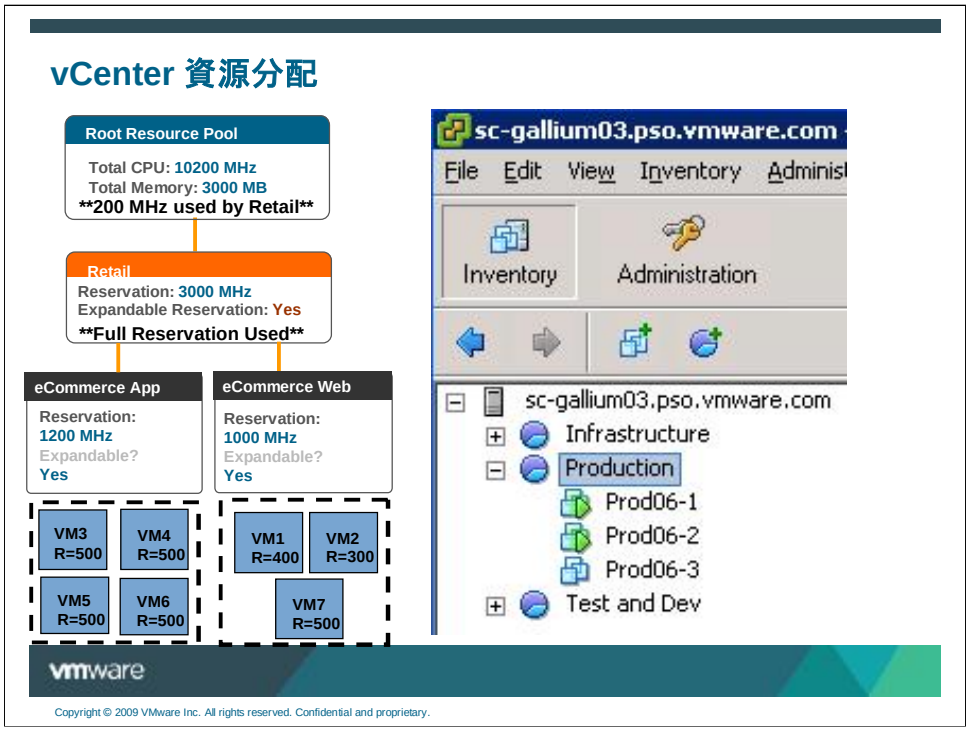


- To continue the example, enable expandable reservation on the eCommerce Web resource pool. This will then allow VM7 to power on because its CPU reservation of 500 MHz can be satisfied by the eCommerce Web resource pool and its parent resource pool, Retail.
- Let's take this example one step further. The eCommerce Apps resource pool contains four virtual machines: VM3, VM4, VM5 and VM6. Each of these VMs has a CPU reservation of 500 MHz. The eCommerce Apps resource pool has a CPU reservaton of 1200 MHz and has expandable reservation set. In order to satisfy the CPU reservations of all the VMs in both resource pools, CPU resources must be taken from as high as the root resource pool.
- Let's explain what's happening, starting from the top:
- The root resource pool has a total of 10200 MHz availabe for its child resource pools to use.
- The Retail resource pool has a total of 3000 MHz available for its child resource pools to use. It has expandable reservation set.
- eCommerce App and eCommerce Web are child resource pools of the Retail pool. They both have expandable reservation set. Together, they have reserved a total of 2200 MHz in the Retail pool. Therefore, the Retail pool has 800 MHz left of its reservation for others to use.
- The total amount of VM CPU reservation in the eCommerce App resource pool is 2000 MHz. Since eCommerce App only has 1200 MHz reserved, the remaining 800 MHz needed to satisfy the VMs' reservations is taken from the Retail resource pool, which has 800 MHz to give. At this point, the Retail pool's full reservation is used.
- The total amount of VM CPU reservation in the eCommerce Web resource pool is 1200 MHz. Since eCommerce Web only has 1000 MHz reserved, the remaining 200 MHz needed to satisfy the VMs' reservations is taken from the parent resource pool, Retail. But since the Retail pool has no more reservation to give, the 200 Mhz is taken instead from Retail's parent, the root resource pool.
- Modify the eCommerce Web resource pool to use expandable reservation and VM7 will be allowed to start
- Borrowing resources occurs recursively from the ancestors of the current resource pool
  - **Expandable Reservation** option must be selected.
  - Offers more flexibility, but less protection

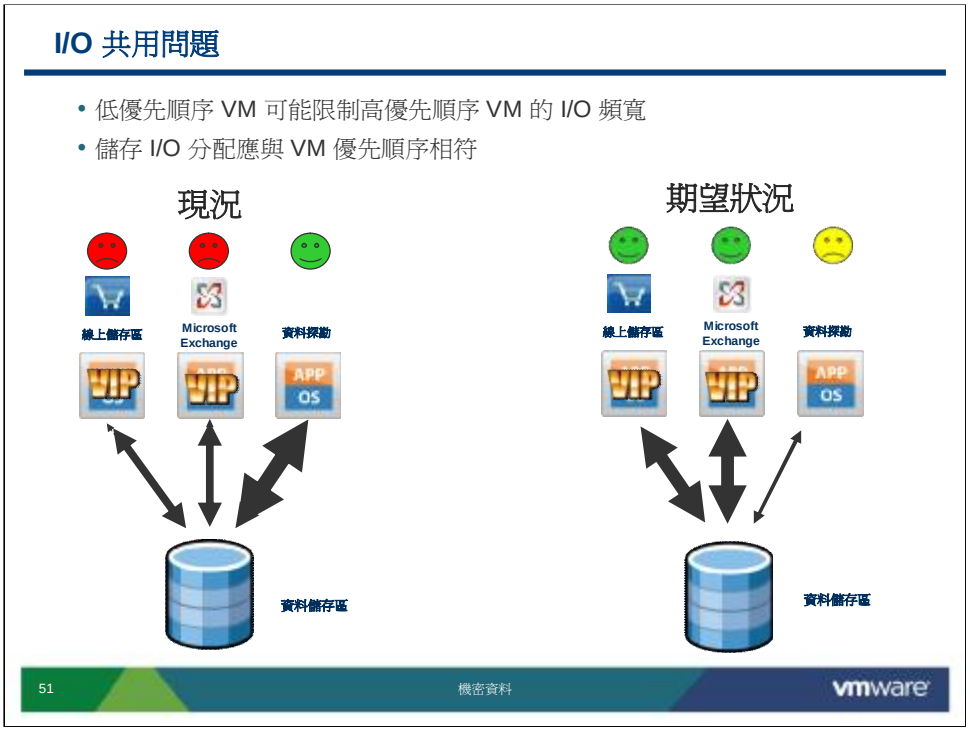


> **The problem Storage I/O control is addressing is the situation where some less important workloads are taking the majority of I/O bandwidth from more important applications. In the case of the three applications shown here, the data mining is hogging a majority of the storage I/O resource. And the two more important to the business operations are getting less performance than needed.**

> **<Click> what one wants to see is a distribution of I/O that is aligned with the importance of each virtual machine. Where the most important business critical applications are getting the I/O bandwidth needed for them to be responsive and the less critical data mining application is taking less I/O bandwidth.**

## 分配 I/O 資源

### 權重轉換爲 ESX I/O 佇列插槽

- 權重較多的 VM 能一次傳送較多 I/O
- 插槽指定是動態的，並依據 VM 權重和現有負載來決定
- 可用插槽總數也是動態的，依據擁塞狀況來決定

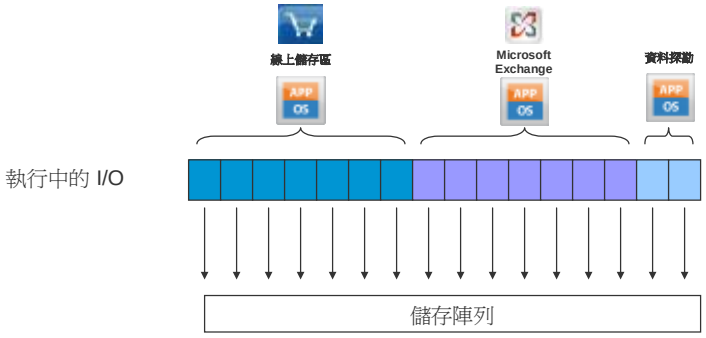


52 機密資料 vmware

> **The way in which these I/O shares are used to effect performance is that queue depth for each ESX server can be assigned and throttled to align the specific shares assigned for each VM running on the collective pool of storage. In the case of our 3 VMs displayed earlier, we have the data mining vm getting the least number of queues assigned while the other two VMs are getting many more queuing slots enabled for their I/O.**

## 儲存 I/O 控制

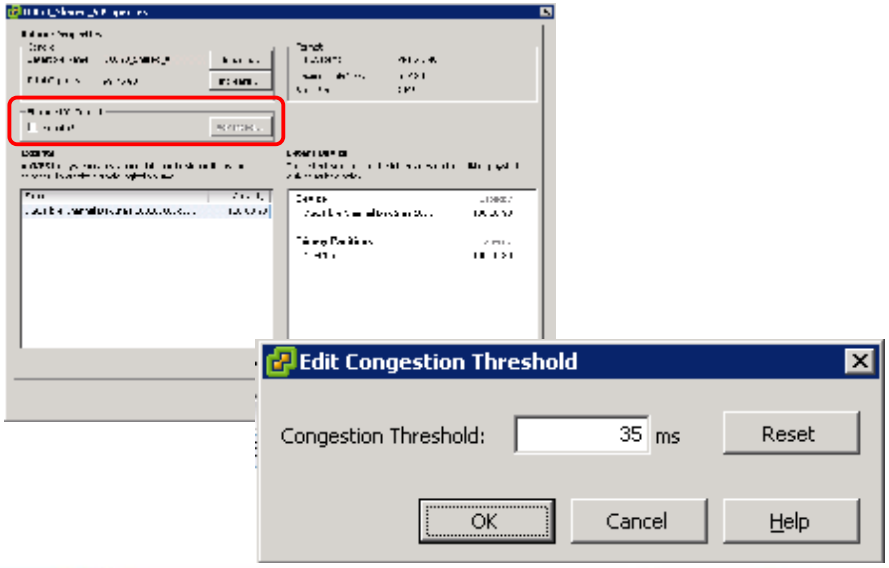
Edit Congestion Threshold
Congestion Threshold:
35
ms
Reset
OK
Cancel
Help

## 儲存 I/O 控制實例：範例 1

兩台 VM 在兩台主機上執行 IOmeter*

- 16 KB 隨機讀取

**VM1：1000 權重**

**VM2：2000 權重**

*結果：這兩個 VM 以 1:2 的比例取得 IOPS*

| | 無儲存 I/O 控制 | | 有儲存 I/O 控制 | |
|---|---|---|---|---|
| | IOPS | IOmeter 延遲 (毫秒) | IOPS | IOmeter 延遲 (毫秒) |
| VM1 | 1500 | 20 | 1080 | 31 |
| VM2 | 1500 | 21 | 1900 | 16 |

* http://www.iometer.org



> **One example of performance impact compares using SIOC with two VMs that have different shares assigned. In this case VM2 has 2x the number of shares as VM1. Without storage I/o control the IOPs and latency are similar. With a background load generation to kick in latency for both the SIOC and without SIOC running to generate latency for the datastore, we see use of SIOC delivers almost 2x the through put and half the latency for the VM with higher shares.**

## 儲存 I/O 控制實例：範例 2

**兩台 Windows VM 在兩台主機上執行 SQL Server**

- 250 GB 資料磁碟、50 GB 記錄磁碟

**VM1：500 權重**

**VM2：2000 權重**

***結果：權重較多的 VM2 每分鐘處理的程序較多，而且較少延遲！***

| | 無儲存 I/O 控制 | | 有儲存 I/O 控制 | |
|---|---|---|---|---|
| | 程序/每分鐘 | 處理時間 (毫秒) | 程序/每分鐘 | 處理時間 (毫秒) |
| VM1 (500 權重) | 8800 | 213 | 7000 | 275 |
| VM2 | 8500 | 220 | 12400 | 150 |
| 綜合 | **17300** | | **19400** | |



> **In this second example, we see that two workloads with one having a 4x number of shares assigned will process a much higher number of orders per minute. With lest time taken. And <click> the combined/aggregate through put of these two work loads using SIOC is greater than the two running without SIOC.**

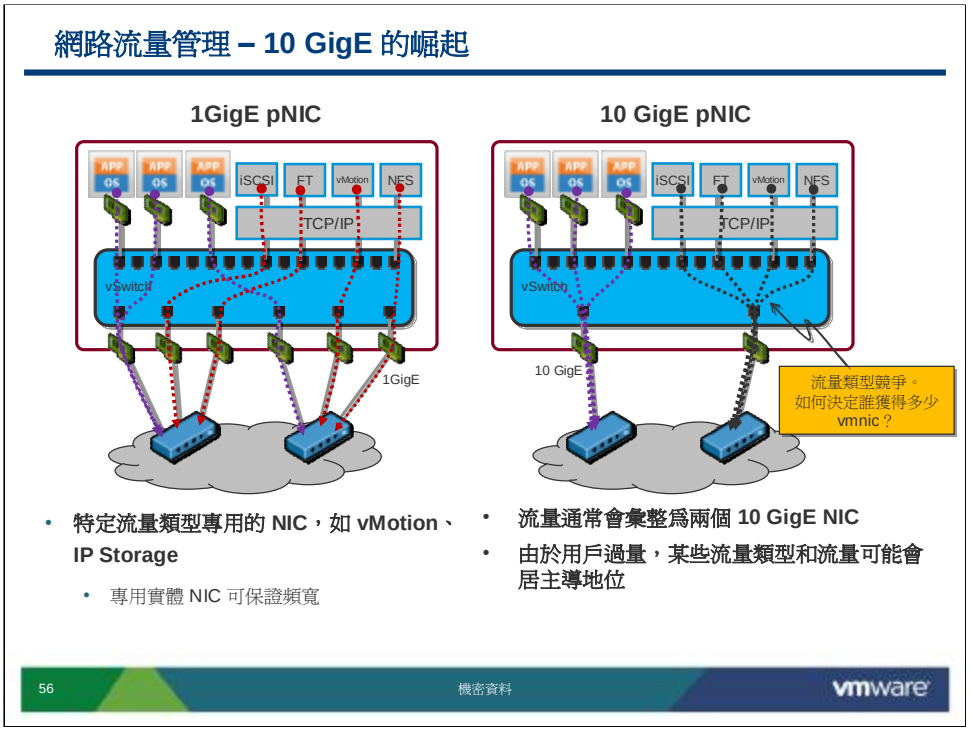


- **Traffic management.**

- **The diagram at left should be familiar to most. When using 1GigE NICs, ESX hosts are typically deployed with NICs dedicated to particular traffic types. For example you may dedicate 4x 1GigE NICs for VM traffic; one NIC to iSCSI, another NIC to vMotion, and another to the service console. Each traffic type gets a dedicated bandwidth by virtue of the physical NIC allocation.**

- **Moving to the diagram at right … ESX hosts deployed with 10GigE NICs are likely to be deployed (for the time being) with only two 10GigE interfaces. Multiple traffic types will be converged over the two interfaces. So long as the load offered to the 10GE interfaces is less than 10GE, everything is ok—the NIC can service the offered load. But what happens when the offered load from the various traffic types exceeds the capacity of the interface? What happens when you offer say 11Gbps to a 10GigE interface? Something has to suffer. This is where Network IO Control steps in. It addresses the issue of oversubscription by allowing you to set the relative importance of predetermined traffic types.**

## 網路 I/O 控制 – 示範

示範個案

- 1GigE pNIC (vmnics)
- 兩種流量類型爭取相同連結 (連結超額分配)
  - iSCSI 設定爲 100 權重 (由 iometer 產生)
  - VM 流量設定爲 25 權重 (由 iperf 產生)
  - 連結可達 890 Mbps
    - iSCSI 獲得 80% à ~710Mbps
    - VM 獲得 20% à ~180Mbps

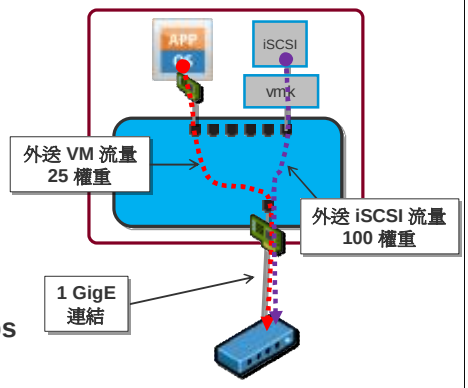




> The slides that follow incorporate short video demonstrations of NetIOC in action. The demonstration involves two traffic types—iSCSI and VM traffic—competing for a 1GigE interface (that in our case can only handle ~890 Mbps).

> We have set iSCSI to 100 shares and VM traffic to 25 shares. iSCSI has 4x importance relative to VM traffic. Each traffic source can fill the link on their own, but in this example they are sharing the link. The shares value translates to 80% bandwidth dedicated to iSCSI and 20% to the VM traffic.

# vCenter 管理功能

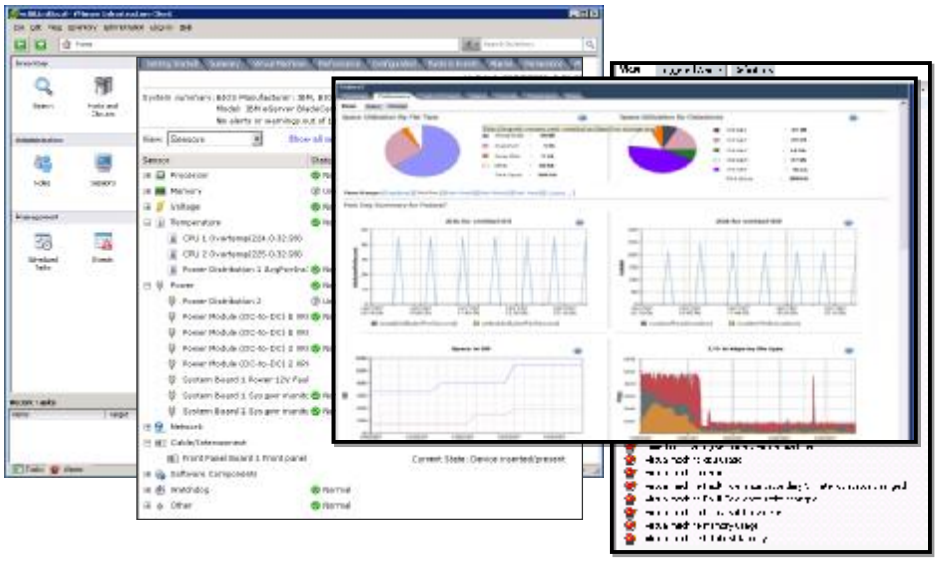

Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

## vCenter 管理功能

Inventory Maps

- Hosts
- VMs
- Networks
- Datastores
- Fault Tolerance relationships

VMotion Resource Maps

- Per VM

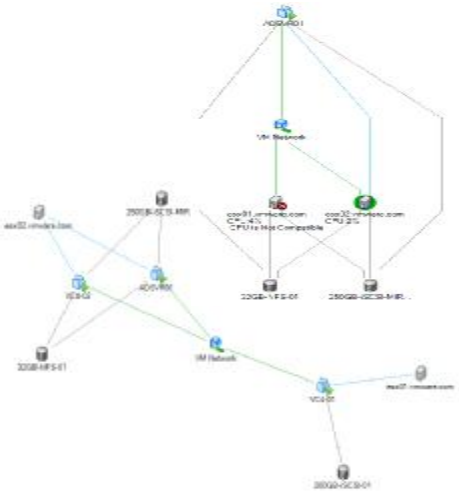

vmware

Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

# Update Manager修正檔管理

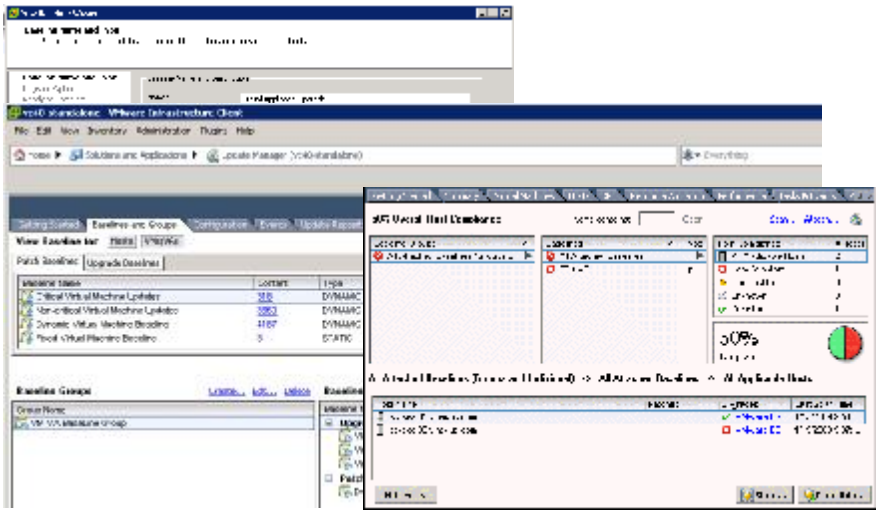

Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

## 主機虛擬化的好處

**1. 改善過低的資源使用率和過高的資訊中心固定成本**(ESX)
- CPU / RAM / Disk / NW
- 電費
-機電空調成本

**2.改善新種服務業務導入衍生的硬體規格採購風險** (ESX/AppSpeed)
- Right-Sizing (上線人數,交易量,資料成長量無法正確預估)
-測試機,開發機,和備援機的採購

**3.降低資料移轉的風險和成本** (P2V)
- 微軟已停止服務的作業系統
-沒有程式碼或廠商停止維護的業務系統

**4.降低維護採購成本** (ESX)
- 主機成本
- 軟體成本

**5.提高業務部署的靈活度** (ESX schedule /vAPP)
- 新商品的檔期促銷方案
- 全球系統的佈署

vmware

Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

## 主機資訊現況和未來問題分析

**6. 提高備援機制彈性** (Converter/ESX/SRM)
- RTO & RPO過長包含主機和終端機
-1:1的備援主機架構費用成本過高

**7. 提高計劃性系統維護的可靠度**
-硬體停機升級 (VMotion)
-作業系統修正檔更新 (Update Manager)
-業務系統過版 (Snapshot Manager/Clone)

**8. 提高資訊安全**
-機房主機權限管理 (vCenter Permission)

**9. 提高資訊中心搬遷的靈活度**
-企業整併需作資料中心轉移 (Converter/ESX/)

**10. 節能省碳** (ESX)
-綠色公民及綠色供應鏈
-碳交易

vmware

Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

# 效益分析

| Cost & Benefit | AS-IS | TO-BE | Gap Realization |
|---|---|---|---|
| **可量化目標(財務目標)** | | | |
| **資料中心維運管理成本** | | | $NT2,786,080 |
| 主機用電費 | $NT1366560 (74700W)<br>( 3元 * 365 * 24 * 52度) | $NT420,480 (22890W)<br>(3元 * 365* 24 * 16度) | $NT946,080 |
| 主機維護費用 | 52 * 10,000 =520,000 | 18 * 10,000 =180,000 | 降低65%的維護費用340,000 |
| 空調散熱耗能 | 182229 BTU | 110920BTU | 7噸冷氣空調<br>約NT350,000 |
| 新購主機和備援主機 | | 可增加34新主機最爲其他運用如備援機或新服務主機 | $NT1,700,000 |
| 機房空間費用 | 約需佔用81 U | 約個37U | 降低54%機架空間 |
| 主機資產利用率 | CPU整體平均利用率低於2% | 提高至50%以上 | |
| **非量化目標** | | | |
| 資訊營運管理 | 1:1 分散式架構資源不可共享 | 資訊架構可隨業務擴展彈性擴充 | |
| | 缺乏系統管理工具 | 可透過內建vCenter管理工具監控主機效能和服務健康狀況 | 如果採購至少須花費NT600,000以上 |
| | 備援機制未完整 | 新系統架構提高系統整體可用度 | 可提高SLA |
| | 缺乏開發和測試環境 | 未來Java核心系統轉換可配合虛擬化技術提高軟體開發的品質 | |

Copyright © 2009 VMware Inc. All rights reserved. Confidential and proprietary.

# 雲端之旅步驟一

*Tackling the Operational Challenges of Managing a Growing Virtualized Environment*

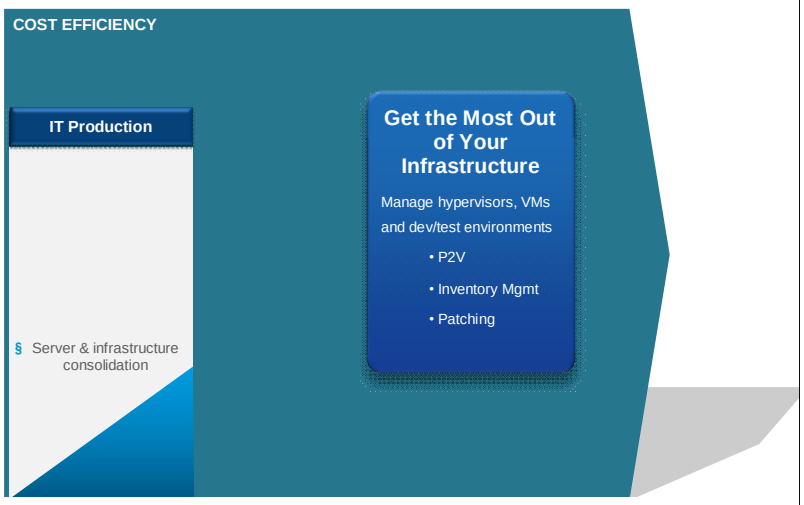


64

vmware

# 雲端之旅步驟二

*Tackling the Operational Challenges of Managing a Growing Virtualized Environment*

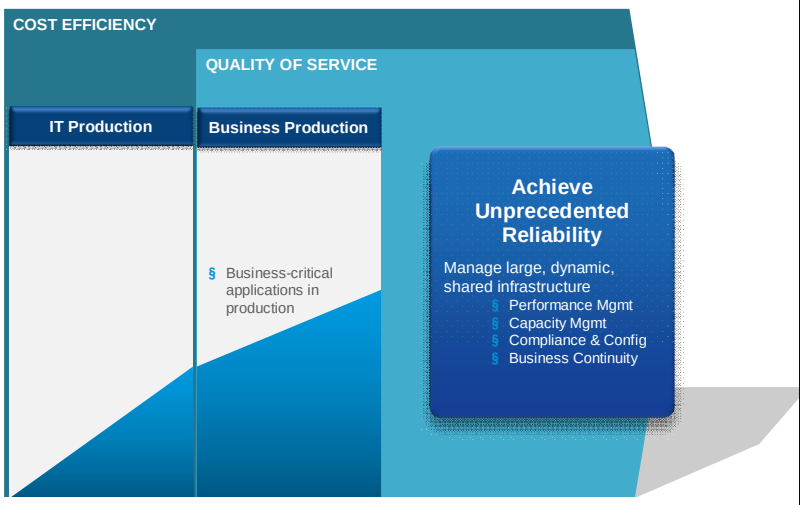

## 雲端作業平台將逐漸成為資訊中心基礎架構...

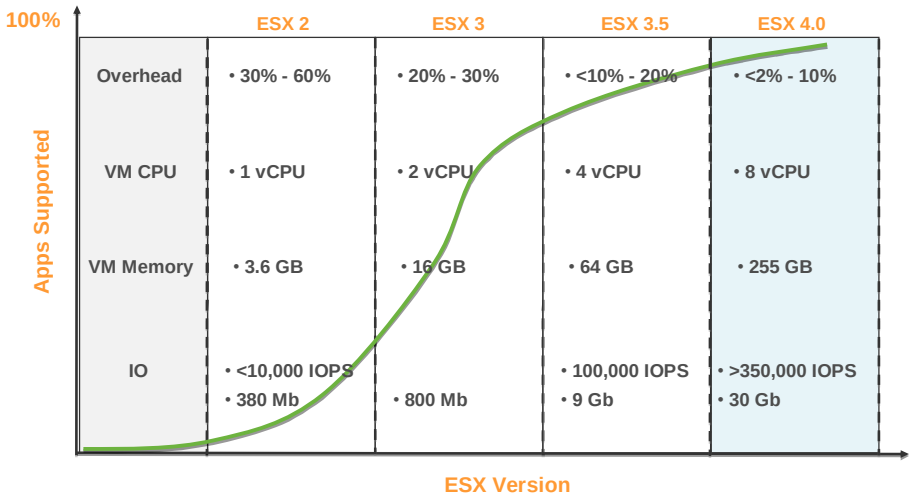


Source: VMware Capacity Planner analysis of > 700,000 servers in customer production environments

## 雲端平台將逐漸成為資訊中心基礎架構...

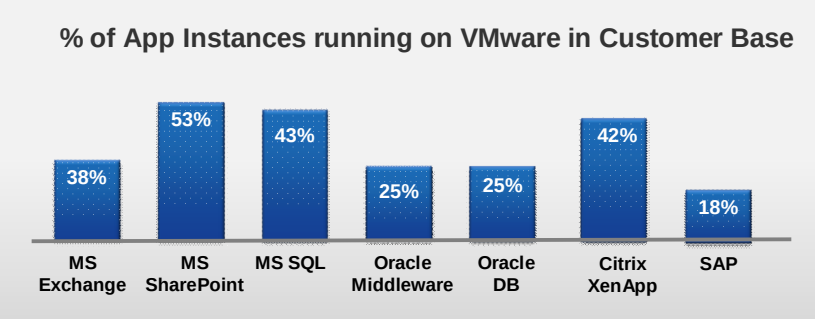


Source: VMware customer survey, January 2010, sample size 1038

Data: Total number of instances of that workload deployed in your organization and the percentage of those instances that are virtualized

***In a recent Gartner poll, 93% respondents used x86 virtualization for applications in production***

Source: Gartner Data Center Poll 2009

"2009 Data Center Poll Results for Virtualization Initiatives" , 9 March 2009

# 雲端之旅步驟三

*Tackling the Operational Challenges of Managing a Growing Virtualized Environment*

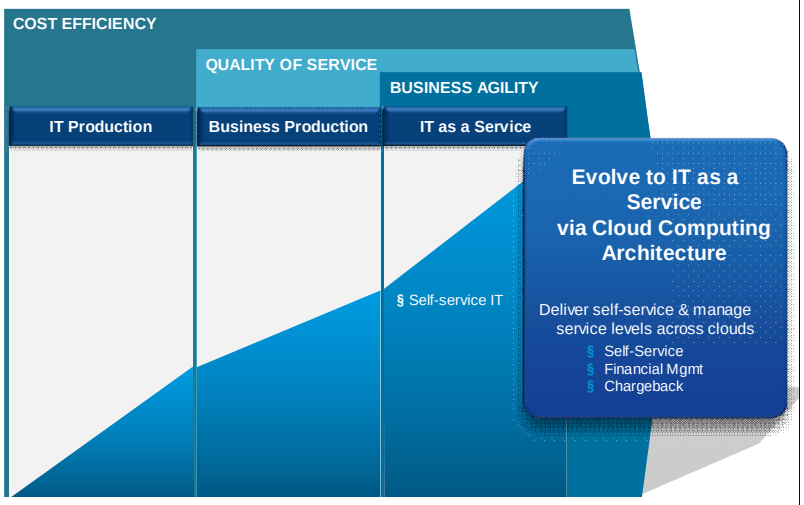

## 問題與討論


- **Notes for presenter:**
  - This presentation is meant for net new customers who are being exposed to desktop and application virtualization for the first time.
  - For customers who have a good understanding of VMware View, we recommend using the “What’s new with VMware View 4” customer facing presentation.
  - Please use the VMware View in-depth technical presentation for a deep dive into each of the components and features high lighted in this sales presentation.
  - Recommend using only 2 case studies as part of the overall presentation. Additional VMware View customer use cases can be found at http://www.vmware.com/customers/