

Evolution of virtual infrastructure with Hyper-V

Juraj Sucik, Slavomir Kubacka
Internet Services Group
CERN IT


- 2006
 - ◆ Microsoft Virtual Server 2005
- 2008
 - ◆ Hyper-V
- 2008
 - ◆ SCVMM 2008
- 2009 Sep
 - ◆ Hyper-V 2.0 + SCVMM 2008 R2

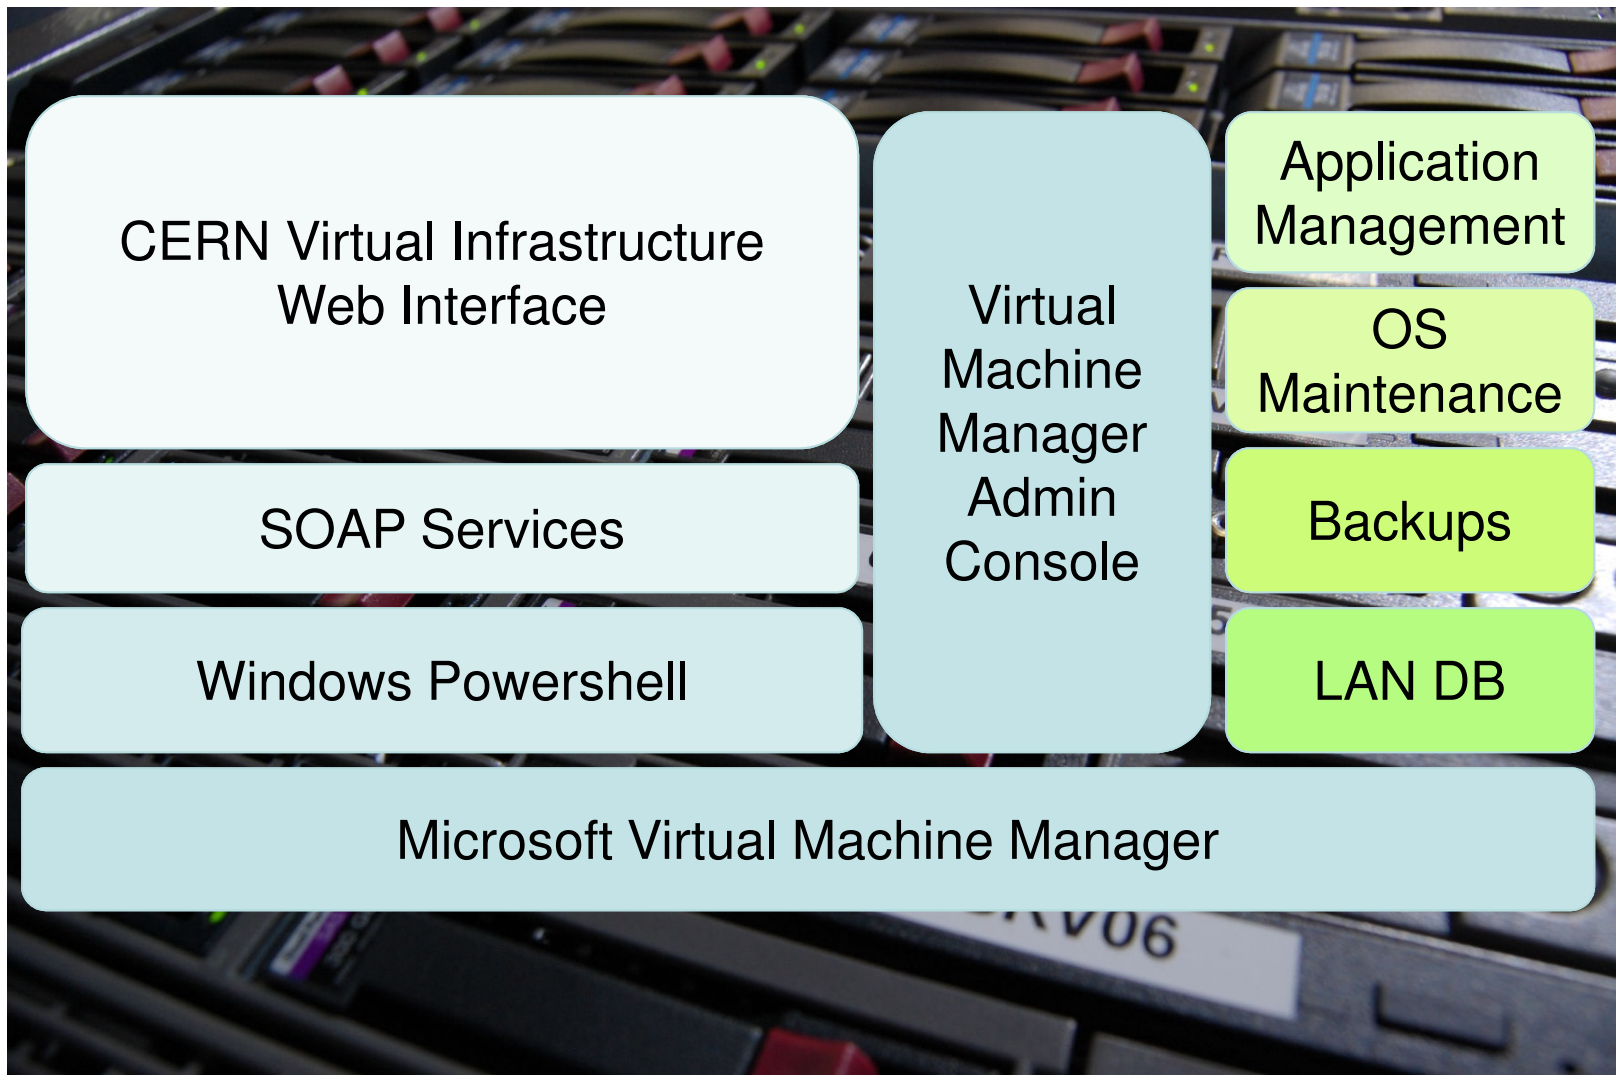
- Hypervisor feature of WS 2008
- 32 and 64-bit virtual machines
- Up to 4 CPUs per VM
- Max 32 GB of memory per VM
- Snapshots
- Failover clustering
- Scriptable interface



- Enterprise management solution
- Windows Powershell API
- V2V and P2V capabilities
- Web portal
- Intelligent placement
- Library and templates
- Delegated management roles
- Job history
- Support for highly available VM
- VM Migration







CERN Virtual Infrastructure - Mozilla Firefox

File Edit View History Bookmarks Tools Help

cern.ch https://winservices.web.cern.ch/winservices/services/cvi/Request.aspx

CERN Virtual Infrastructure

CERN home > IT Department admsuck - admsuck@cern.ch [Logout] [Details]

NICE Services

Home Help Services Other

Menu

- [Request a Virtual Machine]
- [Manage my Virtual Machines]
- [Costs]

Owner: CERN\admsuck

Delegated Administrators: web-services

*You can select a mailing list that **you own** or **you are administrator of**. Members of this mailing list will be able to manage this computer. Beware that you should be a member of this list to be able to manage this machine.*

Computer Name: webafs20

Description: AFS web server

Physical Host: -- Central Service --

Expiration: 30/10/2009 (in a format of: dd/mm/yyyy)

Operating System: SLC5 x64

Memory: Windows Server 2008 R2 x64

CPUs: Windows Server 2008 x86

Payment

Cost:

Budget Code:

Request

14 days, but... **first 31 days are free of charge** (0.96 CHF / day) 0 days

(0.685 CHF / day / each 1 GB above 1 GB) 0.00 CHF

(0.685 CHF / day / each CPU above 1) 0.00 CHF

Last update: October, 2009 - Contact: helpdesk@cern.ch

Images: 50/50 Loaded: 82 KB Speed: 66.29 KB/s Time: 1.238 Done

CERN Virtual Infrastructure - Mozilla Firefox

File Edit View History Bookmarks Tools Help

cern.ch https://winservices.web.cern.ch/winservices/services/cvi/ManageVM.aspx

CERN Virtual Infrastructure

CERN home > IT Department suck - Juraj.Suck@cern.ch [Logout] [Details]

NICE Services

Home | Help | Services | Other

Menu

- [Request a Virtual Machine]
- [Manage my Virtual Machines]
- [Costs]

VM Name	Operating System	Owner	Administrators	CPU	Memory	Expiration Date	State	Operation Progress
ablinuxagent	Running	CERN\basciano	CERN\basciano	1	1.00 GB	17/10/2009	Running	
ablinuxintroscope	HostNotResponding	CERN\	CERN\unknown	1	1.00 GB	31/12/9999	HostNotResponding	
adfs2	Running	CERN\ormancey	CERN\ormancey	2	2.00 GB	31/01/2019	Running	
adkvl01	Running	CERN\baehler	CERN\baehler	2	2.00 GB	31/12/9999	Running	
adkvl02	Running	CERN\baehler	CERN\baehler	1	2.00 GB	31/12/9999	Running	
adkvl03	Running	CERN\baehler	CERN\yt-ce-srvadmin-smt-admin	2	2.00 GB	09/07/2012	Running	
adkvl05	Running	CERN\baehler	CERN\yt-ce-srvadmin-adsk	2	2.00 GB	17/07/2012	Running	
aecert01	Running	CERN\aelwell	CERN\aelwell	1	1.00 GB	28/10/2009	Running	
alicedaqwvs	Running	CERN\ornella	CERN\alice-daq-virtualserver	2	4.00 GB	05/06/2010	Running	
bdswin2003	Running	CERN\admbsilv	CERN\admbsilv	1	1.00 GB	01/12/2009	Running	
bdswin2008	Running	CERN\bsilvade	CERN\bsilvade	1	1.00 GB	01/12/2009	Running	
becodev01	HostNotResponding	CERN\	CERN\unknown	2	2.00 GB	31/12/9999	HostNotResponding	
blenski	Running	CERN\blenski	CERN\blenski	1	1.00 GB	22/10/2009	Running	
bltest	Running	CERN\admblens	CERN\admblens	1	1.00 GB	16/10/2009	Running	

- Connect
- Console
- Start
- Stop
- Save State
- Properties
- Delete


Images: 51/51 | Loaded: 67 KB | Speed: 238.43 KB/s | Time: 0.327 | Done

IS



Internet Services

CERN IT
Department


Enhancements



- Live migration
- Cluster Shared Volume (CSV)
 - ◆ Enables multiple nodes in a cluster to access a single shared LUN
 - ◆ Dynamic I/O redirection
- Network optimizations
 - TCP/IP Traffic in a VM can be offloaded to a physical NIC on the host computer
- Processor compatibility mode
 - Allows live migration across different CPU versions within the same processor family



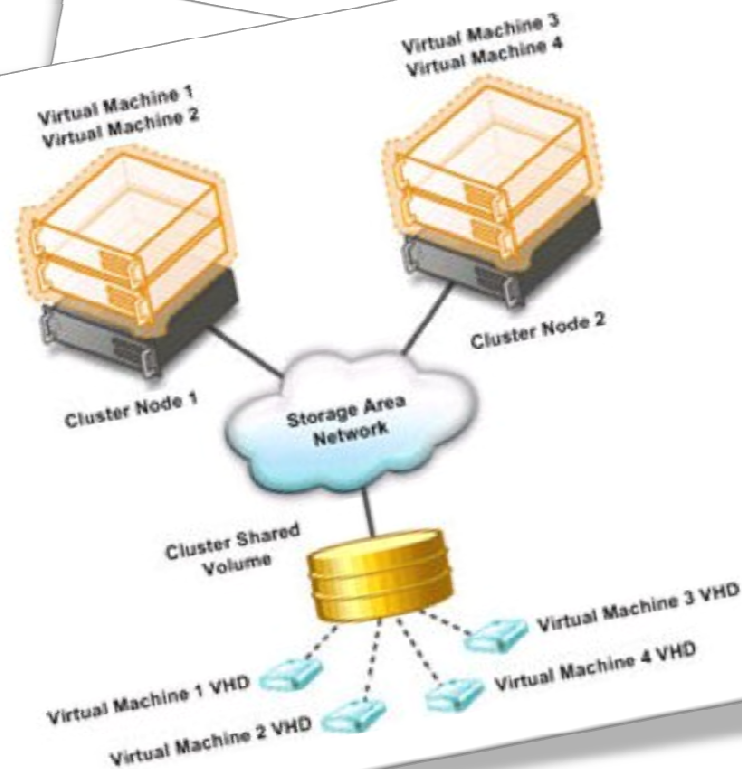
- Logical Processor Support
 - ◆ Support for 32 logical processors on host computer
- Hot Add/Remove Storage
 - ◆ Add and remove VHD disks to a running VM without requiring a reboot
- Second Level Translation (SLAT)
 - Leverage new processor features to improve performance and reduce load on Windows Hypervisor
- Better SMP support for Linux




- Manage WS 2008 R2 Hyper-V
- Live Migration
 - Detects if Live migration can be done
- Maintenance mode
 - Placement of new VM not allowed
 - Existing VMs migrated off or saved
- Multiple VM per LUN using CSV
 - Supports CSV feature of HV 2.0
- V2P feature

- **SAN related enhancements**
 - Promote non-HA VM to HA VM by migrating it to a clustered host, and vice versa to “demote” the VM
- **Network optimizations**
 - If enabled, VMM will configure the VM to use VMQ or Chimney, if available on the host
- **Rapid provisioning**
 - Avoids copying VHD from library
- **VDI integration**

- Maintenance reasons
- Load balancing
- Green IT
- Fast migration
- SOAP interface




- No dropped network connections
- No perceived loss of service
- Clustered Shared Volumes facilitates LM
- Leverages Failover Clustering



Quick Migration

(Windows Server 2008 Hyper-V)


- 1. Save state**
 - Create VM on the target
 - Write VM memory to shared storage
- 2. Move virtual machine**
 - Move storage connectivity from source host to target host via Ethernet
- 3. Restore state & Run**
 - Take VM memory from shared storage and restore on Target
 - Run



Live Migration

(WS08R2 Hyper-V)

- 1. VM State/Memory Transfer**
 - Create VM on the target
 - Move memory pages from the source to the target via Ethernet
- 2. Final state transfer and virtual machine restore**
 - Pause virtual machine
 - Move storage connectivity from source host to target host via Ethernet
- 3. Un-pause & Run**




Aspect	vSphere 4	Hyper-V R2
# CPU core	64	64
Memory	1TB	2TB
# nodes in cluster	32	16
# virtual CPU	8	4
# guest per host	256	192
Virtual memory	256GB	64GB
Hot-add disk	Yes	SCSI only
VM move	Live	Live
# of snapshots	32	50
HA via clustering	Yes	Yes
Market share	44%	23%

Source:

[Login, USENIX Magazine, Oct 2009](#)

- RHEL supported as guest OS
- Open source drivers (GPL) in 2.6.32
- CPU Benchmark

HEP - SPEC Benchmark



- Time synchronization
 - ◆ Kernel parameters, e.g. notsc divider=10
- Virtual serial console
- Admin privileges
- 5 Linux templates

```
\\caevmrsv02\pipe\jura\slc5 - PuTTY
Found volume roup "VolGroup0" using metadat type lvm2
Activating logical volumes
  2 logical volume(s) in volume group "VolGroup00" now active
Trying to resume from /dev/VolGroup00/LogVol01
No suspend signature on swap, not resuming.
Creating root device.
Mounting root filesystem.
kjournald starting. Commit interval 5 seconds
EXT3-fs: mounted filesystem with ordered data mode.
Setting up other filesystems.
Setting up new root fs
no fstab.sys, mounting internal defaults
Switching to new root and running init.
unmounting old /dev
unmounting old /proc
unmounting old /sys
type=1404 audit(1256231324.597:2): enforcing=1 old enforcing=0 auid=4294967295 ses=4294967295
type=1403 audit(1256231325.218:3): policy loaded auid=4294967295 ses=4294967295
INIT: version 2.86 booting
       Welcome to Scientific Linux CERN SLC
       Press 'I' to enter interactive startup.
Setting clock  (utc): Thu Oct 22 19:08:52 CEST 2009 [ OK ]
Starting udev: [ OK ]
Loading default keymap (us): [ OK ]
Setting hostname localhost.localdomain: [ OK ]
Setting up Logical Volume Management:  2 logical volume(s) in volume group "VolGroup00" now active
[ OK ]
Checking filesystems
Checking all file systems.
[/sbin/fsck.ext3 (1) -- /] fsck.ext3 -a /dev/VolGroup00/LogVol100
/dev/VolGroup00/LogVol100: clean, 154789/9427968 files, 1103883/9420800 blocks
[/sbin/fsck.ext3 (1) -- /boot] fsck.ext3 -a /dev/hdal
/boot: clean, 44/26104 files, 27162/104388 blocks
[ OK ]
Rem
```

Aspect	Service consolidation	Batch virtualization
Scale (machines)	~ 100	~ 1000
CPU usage	Little	High
Hardware	Reliable	Cheap
Services	Critical	Non-critical
Migration	Live	Not required
VM life time	Long	Limited

Perl SOAP client
HMS
AIMS
Lemon

SOAP Services

Windows Powershell

Microsoft Virtual Machine Manager

Virtual
Machine
Manager
Admin
Console


Application
Management

OS
Maintenance

Backups

LAN DB

- VOBox service – dedicated servers for experiments: 222 and growing rapidly!



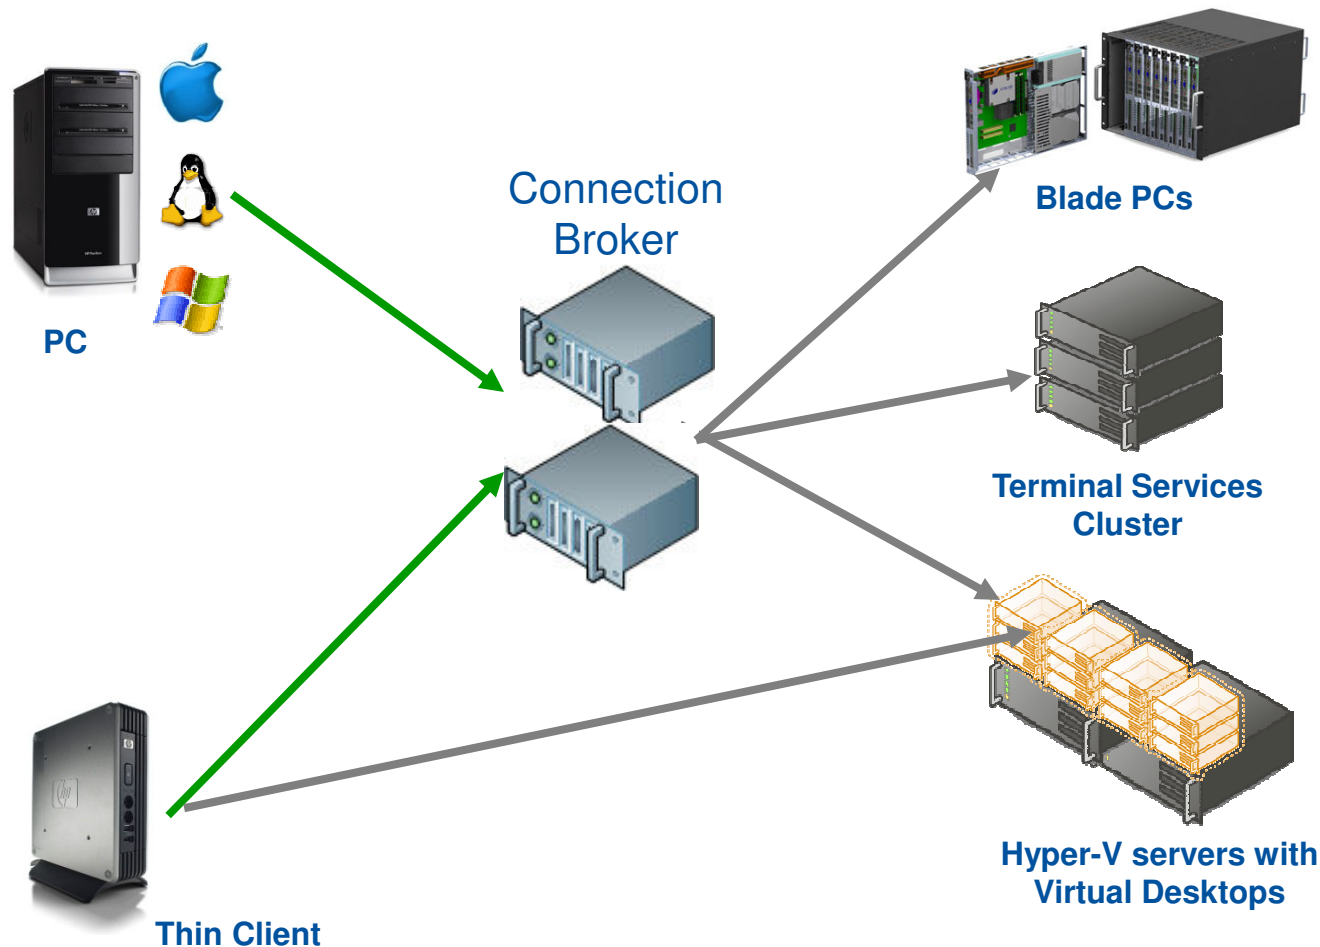
- Consolidation of servers on critical power supply as the power is very limited
- Development resources for IT-FIO



- CERN fabric management integration
 - ◆ LEAF
 - ◆ Lemon
 - ◆ Quattor
 - ◆ SLS
- Integrate Hyper-V drivers with SLC
- Rapid provisioning

Office

Computer Centre



- Propose Virtual Desktop self service
 - ◆ for experiment developers
 - ◆ as an alternative to dual-boot
 - ◆ as an alternative to Terminal Services
- Evaluate a thin client technology, which could be solution for
 - ◆ public computers
 - ◆ basic office users



Jack PC

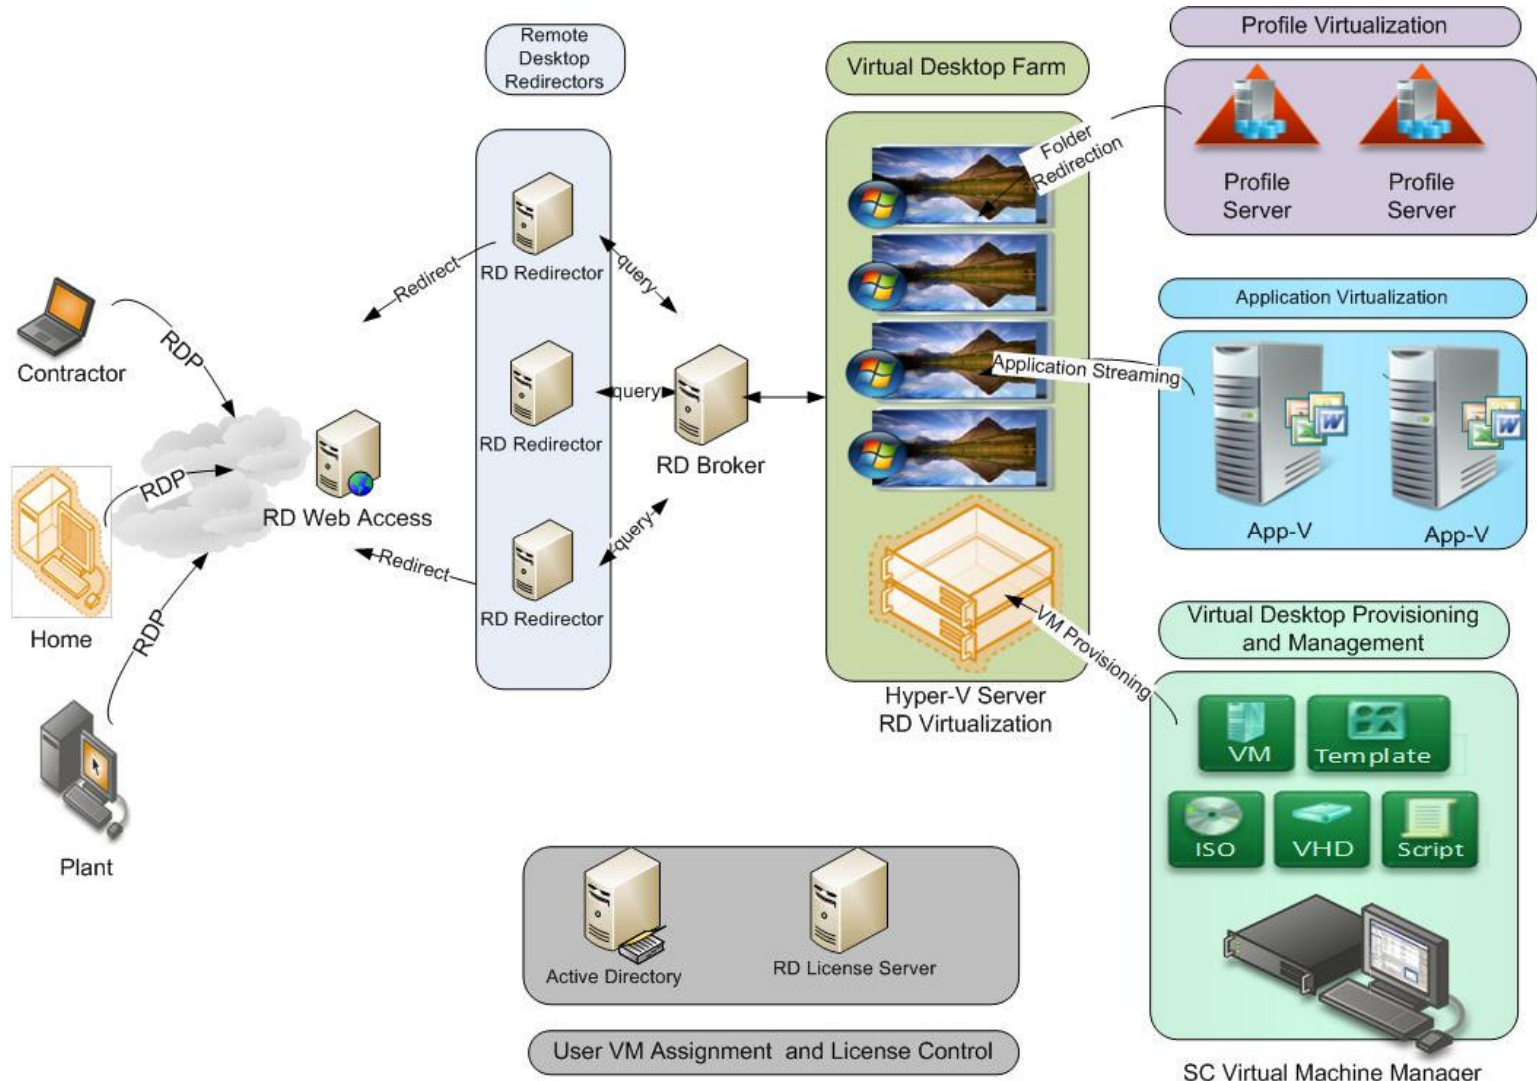


Thin Client

- Latest editions of Hyper-V + SCVMM in production
- Better Linux support
- Live migration
- Integration with CERN IT services
 - ◆ Fabric management tools

Visit our website

CERN Virtual Infrastructure: <http://cern.ch/cvi>



Quick Migration

(Windows Server 2008 Hyper-V)

1. Save State / Memory Transfer


- Create VM on the target
- Move VM memory pages shared storage via Ethernet

2. Move state and virtual machine restore

- Move storage machine to target host via Ethernet
- Move storage Ethernet activity from source host to target host via Ethernet

3. Restore state & Run

- Take VM memory from shared storage and restore on Target
- Run



- 1) *Request VM*
- 2) *Delete VM*
- 3) *Start VM*
- 4) *Stop VM*
- 5) *Save State of VM*
- 6) *Migrate*

