Virtual Infrastructure Implementation

Best Practices From A to Z

Edward Aractingi

VMWORLD 2006

About me

- Edward Aractingi
- Marshall University
- Computing Services Systems Administration
- MS, MCSE, RHCT, Security+, A+
- http://edward.aractingi.net

Agenda

- **Marshall University, Data Center & RTI**
- **Virtualization Decision & Planning**
- **Planning Best Practice**
  - Hardware & Network Upgrade
  - Storage and Network Infrastructure
- **Implementation Best Practice**
  - Deployment and Unattended installation
  - Virtual Server High Availability Architecture
  - Upgrade to VI 3
- **Operation Best Practice**
  - SNMP Management and Monitoring
  - Backup Strategies

About Marshall University

- State Funded University
- Huntington, WV
- 24 associate programs,
- 44 baccalaureate
- 46 graduate programs
- Enrolls 16,000 students including 4,000 graduate and medical students

MU Data Center

- Almost 200 Server
- Dell PowerEdge Servers
- EMC CLARiiON SAN arrays
- Microsoft Windows Server 2003
- RedHat Enterprise Linux

Rehall Transportation Institute

- Staff over 50 professors, graduate students and full time employees
- Transportation and Economic development in the Appalachian region
- Database, GIS and Web Apps
- TEDIS Servers environment hosted in MU Data Center

Planning Best Practice for Virtual Infrastructure Implementation

VMWORLD 2006

TEDIS before VMware

- 8 Dell PowerEdge 2650 with 4G RAM
- 2 Dell PowerEdge 1750 with 4G RAM
- 11 TB of Storage on SAN Arrays
- Windows™, Oracle©,SQL Server,
- ESRI© GIS applications
- Centralized Administration

Virtualization Decision

- The growing need for more servers
- End-of-life / Warranty physical servers
- Extremely dynamic OS, DB, Web platforms environment
- Server consolidation and resource optimization.
- Faster server provisioning (3-4 weeks → 5-10 mins)
- Target…….. 100% virtualized environment
- Effective implementation now is 92%

VMware Documentation

- Use what VMWare provides
- Docs /Whitepaper
- User Forums

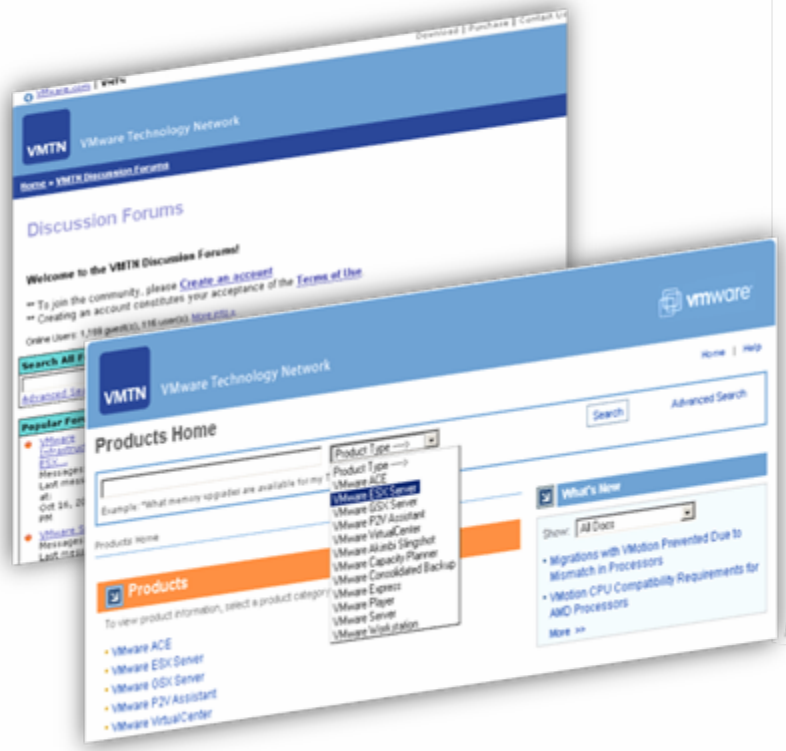

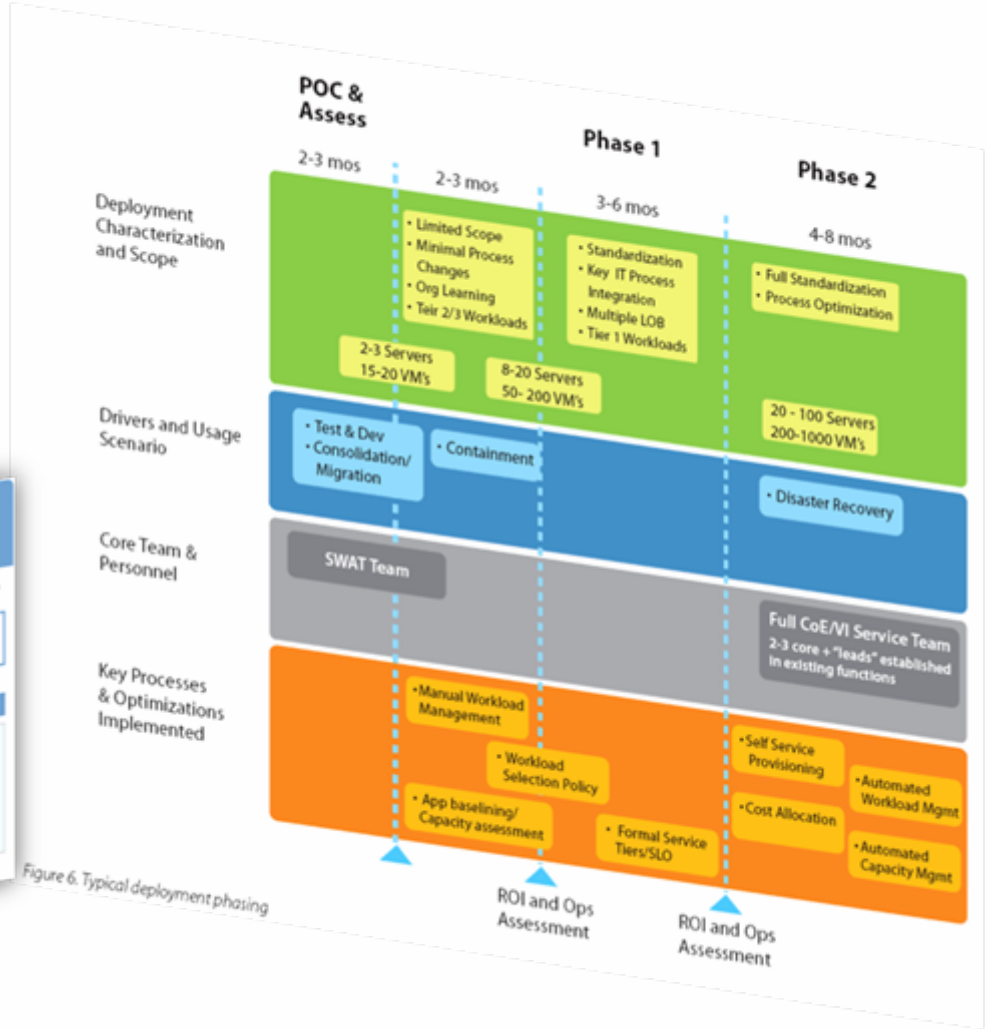

Figure 6. Typical deployment phasing

Resources are available, look for them

Planning / Project Management

ID	Task Name	Duration	Start	Finish
1	Project start	0 days	Mon 1/2/06	Mon 1/2/06
2	Install Memory Modules Upgrades	7 hrs	Mon 1/2/06	Mon 1/2/06
3	Backup Tapes Media & USB Disks arrive	0 days	Fri 1/6/06	Fri 1/6/06
4	Acquire ESX License codes	0 days	Mon 1/9/06	Mon 1/9/06
5	Servers Full Backup	3 days	Mon 1/9/06	Wed 1/11/06
6	Clean and restructure SAN storage	3 days	Wed 1/11/06	Mon 1/16/06
7	Move Current VM guests from App01	2 hrs	Mon 1/16/06	Mon 1/16/06
8	Install ESX server on App01	1 day	Wed 1/11/06	Thu 1/12/06
9	Add Old VM Guests to the new ESX server	1 day	Thu 1/12/06	Fri 1/13/06
10	Install Virtual Center	3 hrs	Thu 1/12/06	Fri 1/13/06
11	**Create Templates**	**1 day**	**Fri 1/13/06**	**Mon 1/16/06**
12	Create Windows Server 2003 Standard Template	1 day	Fri 1/13/06	Mon 1/16/06
13	Create Windows Server 2003 Enterprise Template	1 day	Fri 1/13/06	Mon 1/16/06
14	Create Linux Template	1 day	Fri 1/13/06	Mon 1/16/06
15	Install single instance of Oracle Database	5 hrs	Mon 1/16/06	Tue 1/17/06
16	Install SQL Server on a Windows Template	5 hrs	Mon 1/16/06	Tue 1/17/06
17	Review the first ESX implementation	1 day	Fri 1/13/06	Mon 1/16/06
18	Review SAN Storage structure	2 days	Mon 1/16/06	Wed 1/18/06
19	**P2V Web servers**	**0.47 days**	**Wed 1/18/06**	**Wed 1/18/06**
20	P2V Web01	3.5 hrs	Wed 1/18/06	Wed 1/18/06
21	P2V Web02	3.5 hrs	Wed 1/18/06	Wed 1/18/06
22	Configure CX200 SAN Storage	1 day	Wed 1/11/06	Thu 1/12/06
23	**Install ESX on Web servers**	**1 day**	**Wed 1/18/06**	**Thu 1/19/06**
24	Install ESX on WEB01	1 day	Wed 1/18/06	Thu 1/19/06
25	Install ESX on WEB02	1 day	Wed 1/18/06	Thu 1/19/06
26	**P2V Development Machines**	**1 day**	**Thu 1/19/06**	**Thu 1/26/06**
27	P2V RTIDBX01	1 day	Thu 1/19/06	Thu 1/26/06
28	P2V RTIWEBX01	1 day	Thu 1/19/06	Thu 1/26/06
29	**Install ESX on Development Machines**	**1 day**	**Thu 1/26/06**	**Fri 1/27/06**
30	Install ESX on RTIDBX01	1 day	Thu 1/26/06	Fri 1/27/06
31	Install ESX on RTIWEBX01	1 day	Thu 1/26/06	Fri 1/27/06
32	**P2V RAC**	**1 day**	**Fri 1/27/06**	**Mon 1/30/06**
33	P2V RTIDB01	1 day	Fri 1/27/06	Mon 1/30/06
34	P2V RTIDB02	1 day	Fri 1/27/06	Mon 1/30/06
35	**Install ESX on RAC Database Servers**	**1 day**	**Mon 1/30/06**	**Tue 1/31/06**
36	Install ESX on RTIDB01	1 day	Mon 1/30/06	Tue 1/31/06
37	Install ESX on RTIDB02	1 day	Mon 1/30/06	Tue 1/31/06
38	P2V Commander & Bigbrother	1 day?	Tue 1/31/06	Wed 2/1/06
39	Finish	0 days	Wed 2/1/06	Wed 2/1/06

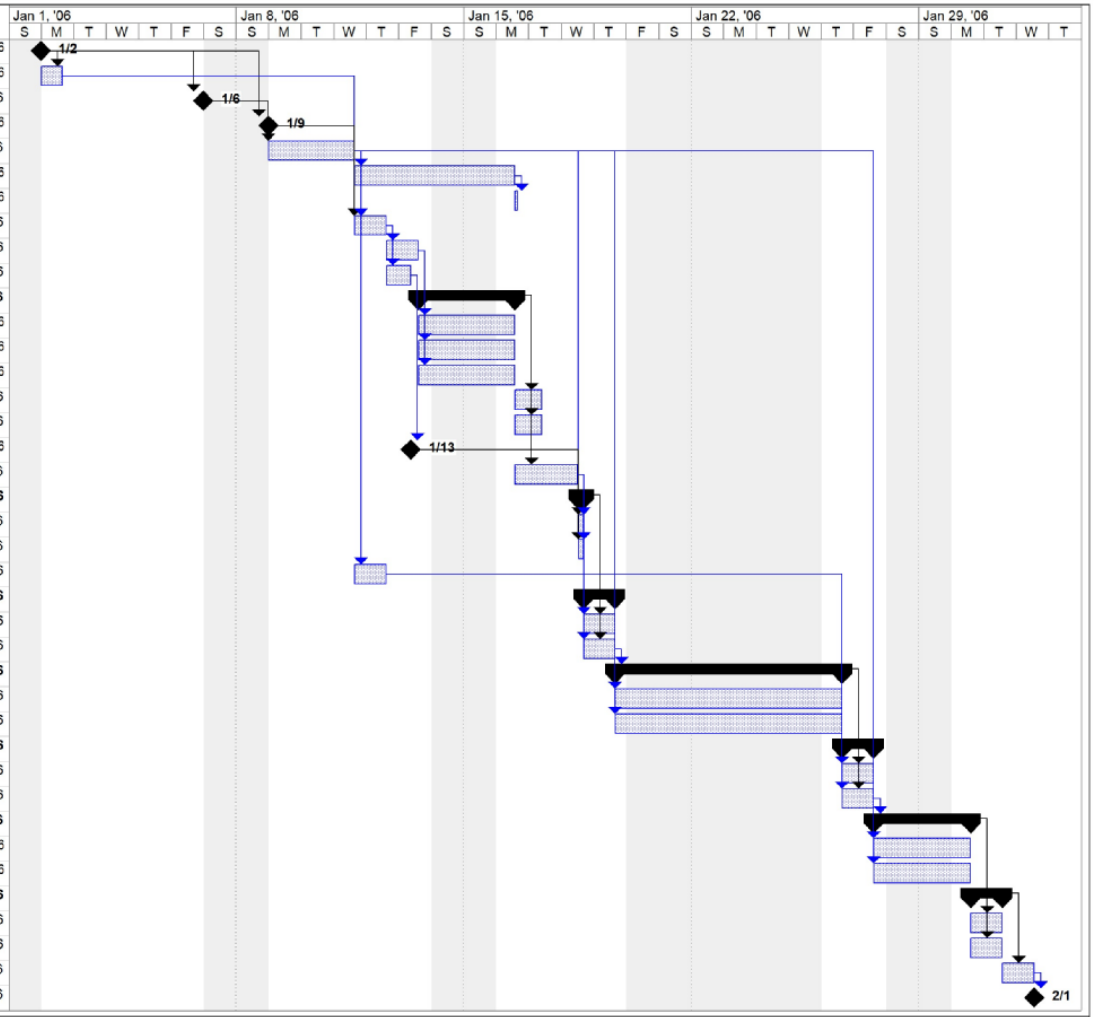

Change Management

- Staging the upgrade
- Targeting powerful servers
- Application and Web servers
- Used TEDISLOG
- Give users access to VMs

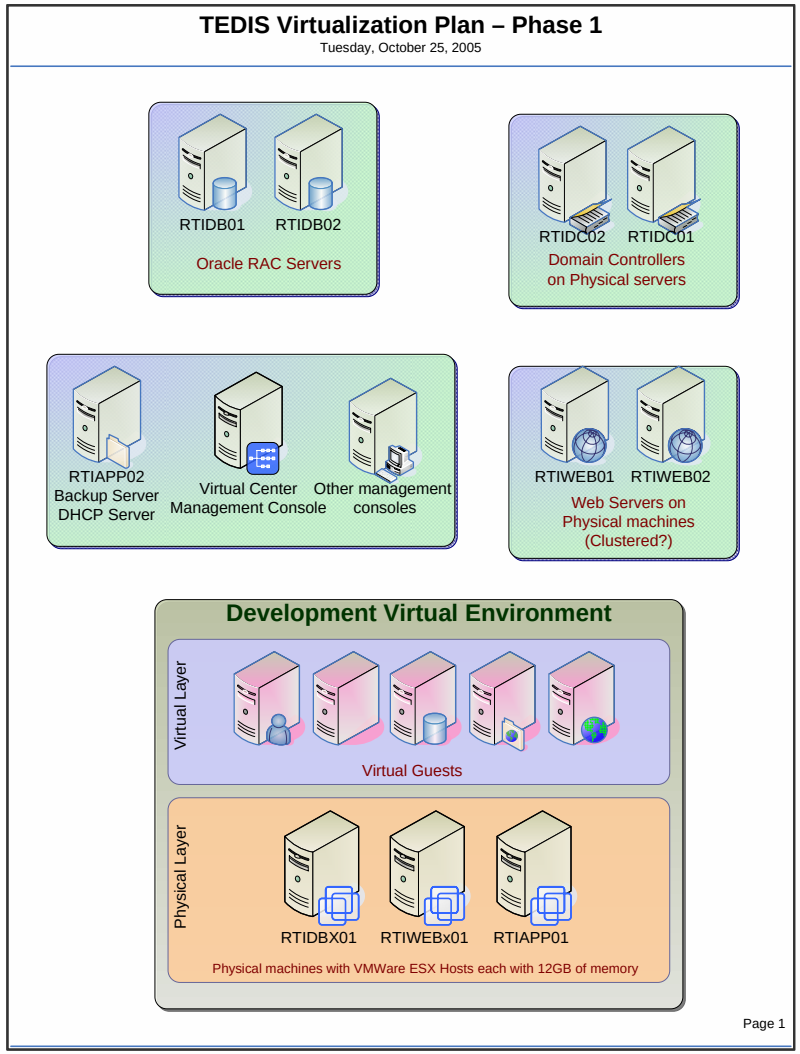

Hardware Expansion

- Part of the plan, we needed more hardware
- Spend on maxing existing server resources rather than buying new servers
- Bumped the memory to 12 GB on every server
- Added two-port NIC card to every server
- Total is 64 GB of memory and 12 CPUs
- Result: Able to go from 8 servers to 36 VMs with minimum cost and maximum flexibility

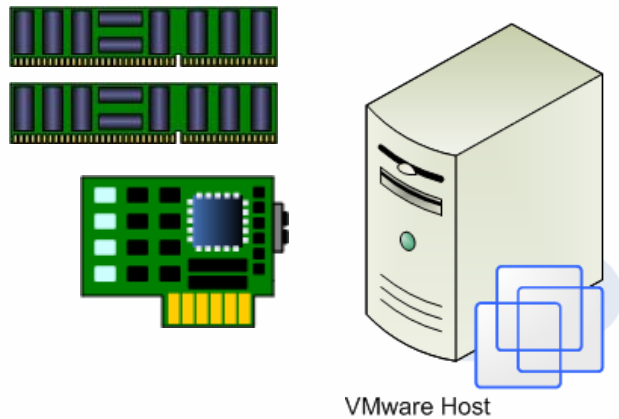

Prepare Configuration Sheets

- Network configurations example

Start IP	End IP		Switch
10.10.10.11	10.10.10.20	ESX Service console	Protected VLan
10.10.10.111	10.10.10.120	DRAC	Protected VLan
192.168.0.11	192.168.0.20	Vmotion	Protected VLan
10.10.20.3	10.10.20.254	Virtual Machines	Public VLan
192.168.10.3	192.168.10.254		N/A

- DHCP Scopes
- SAN Storage Groups

Logical Network Architecture

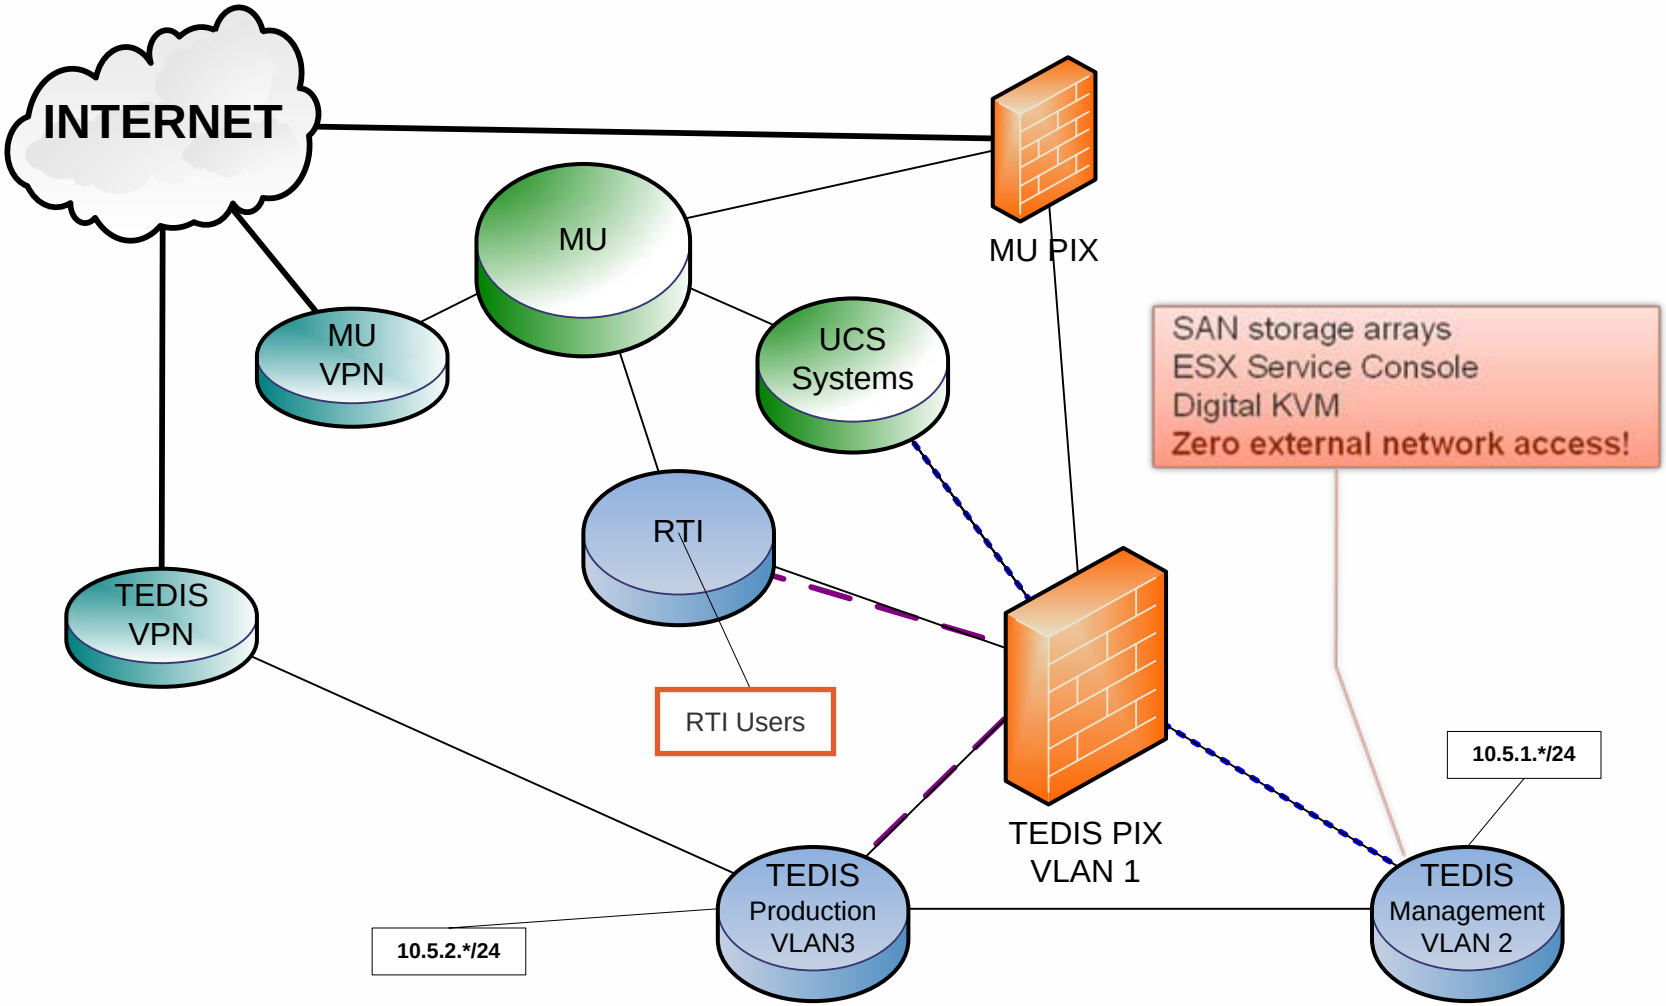

Consider securing service console

Implementation Best Practice for Virtual Infrastructure Implementation

VMWORLD 2006

Physical Network Infrastructure

- The best security practice is at the network layer
- Configure network equipments to block external access to service console's network COMPLETELY

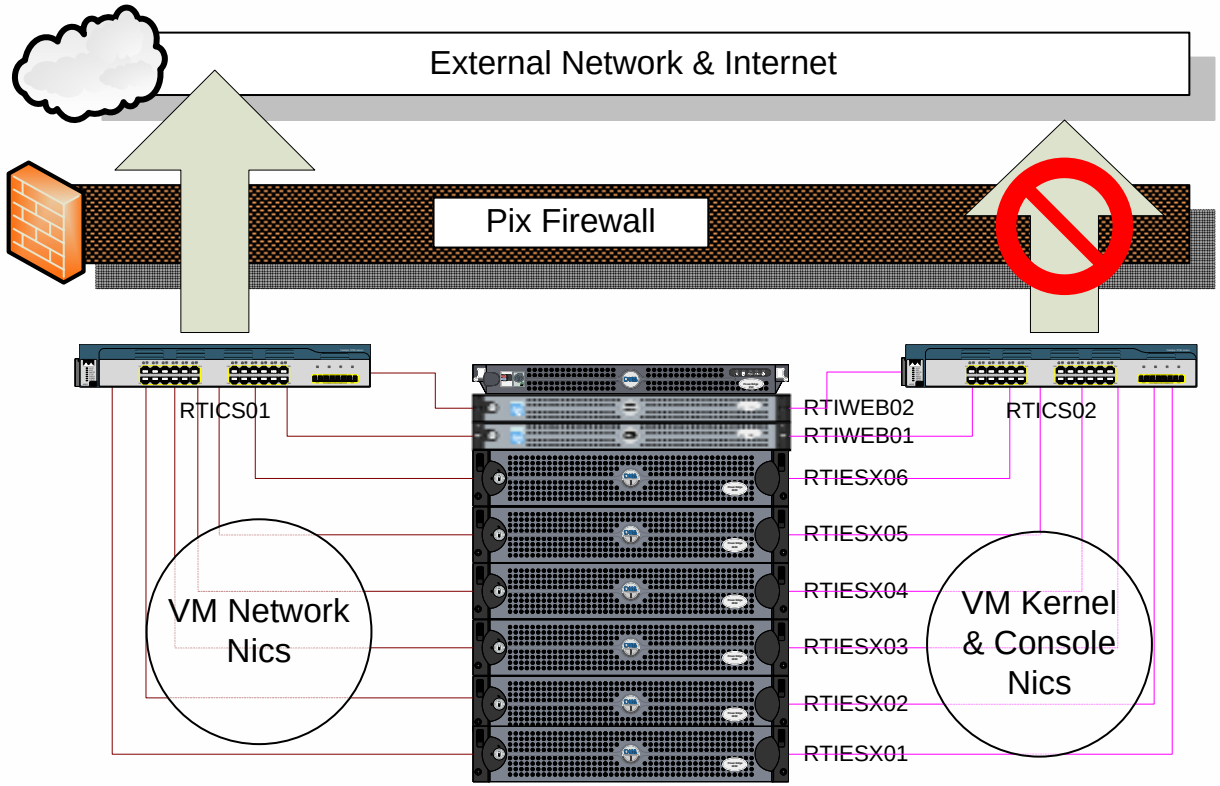

Virtual Network Infrastructure

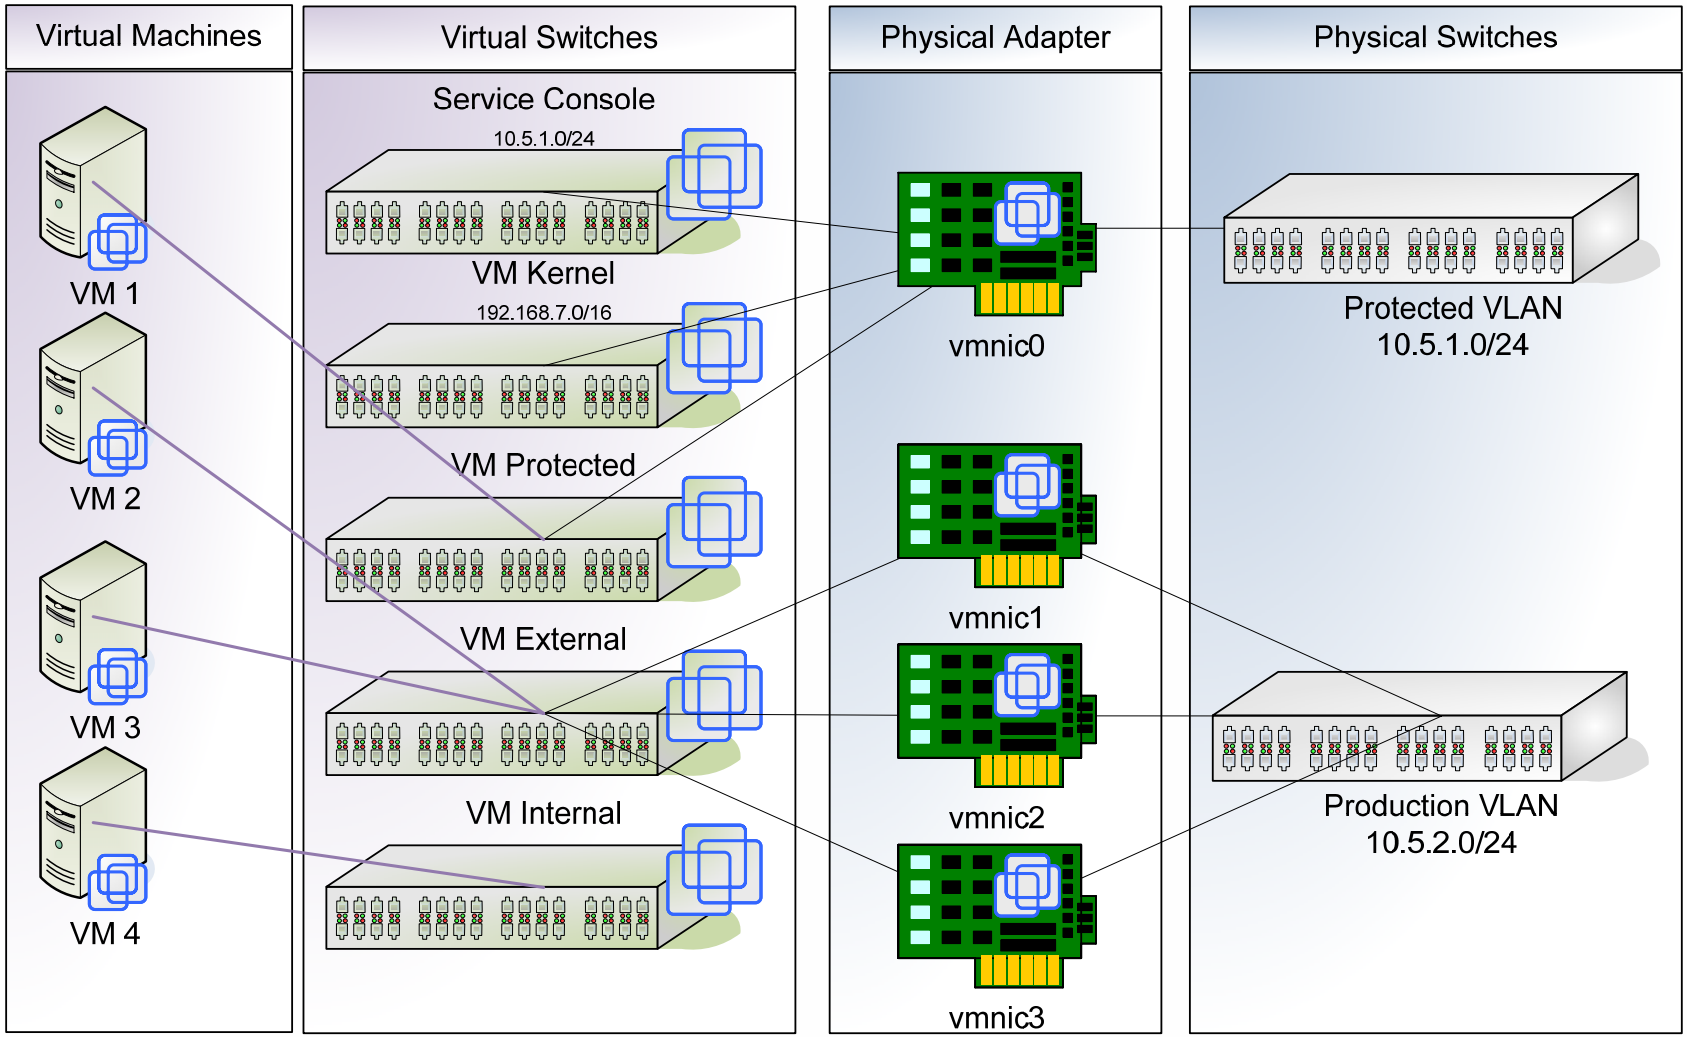

VM will benefit from multiple NICs

Storage Area Network

- EMC CLARiiON CX700 (10 Terabytes)
- EMC CLARiiON CX200 (1 Terabyte)
- SAN Management with EMC Navisphere
- 8 servers connected to CX700 with production data
- A specific storage group for all ESX Servers LUNs
- EMC documentation includes extensive support for ESX

Unattended Installation

- Scripted Installation
- Using one ESX to host the installation files
- NFS Share
- DHCP reservation
- Recommended for:
  - 5 or more servers
  - Planning to replace hardware
  - Staff got hit by a bus

Dell OpenManage

- Dell Remote Access Card (DRAC)
- TFTP Server
- Remote floppy disk image
- No server room access

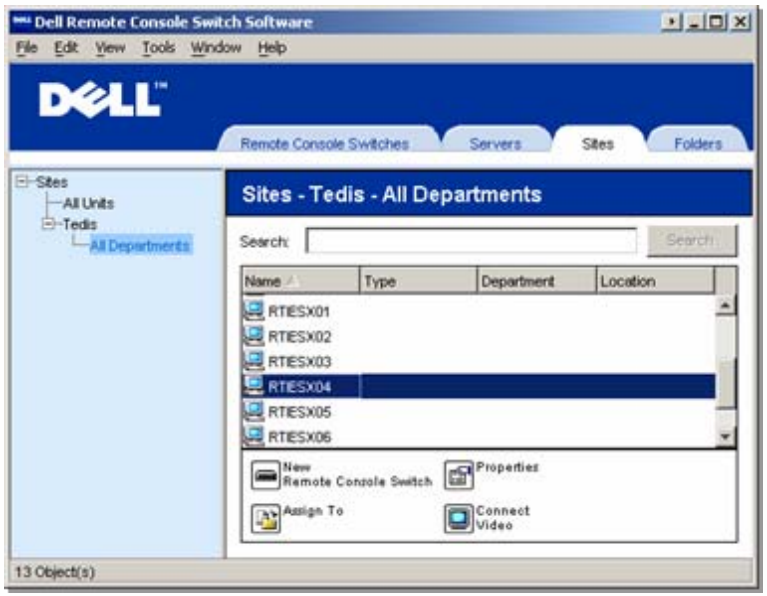

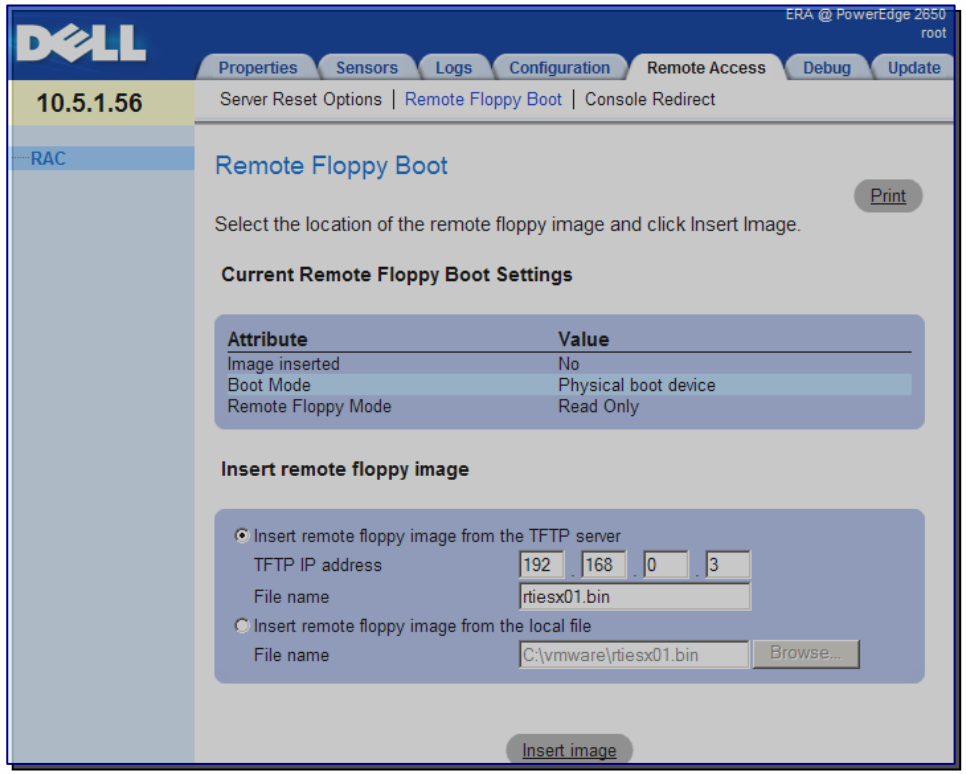

KickStart CD

- Looking for faster methods of ESX Deployment
- Facilitate replacement of failed or retired hardware
- Once CD to recover and install ESX unattended
- Integrated with PXE *

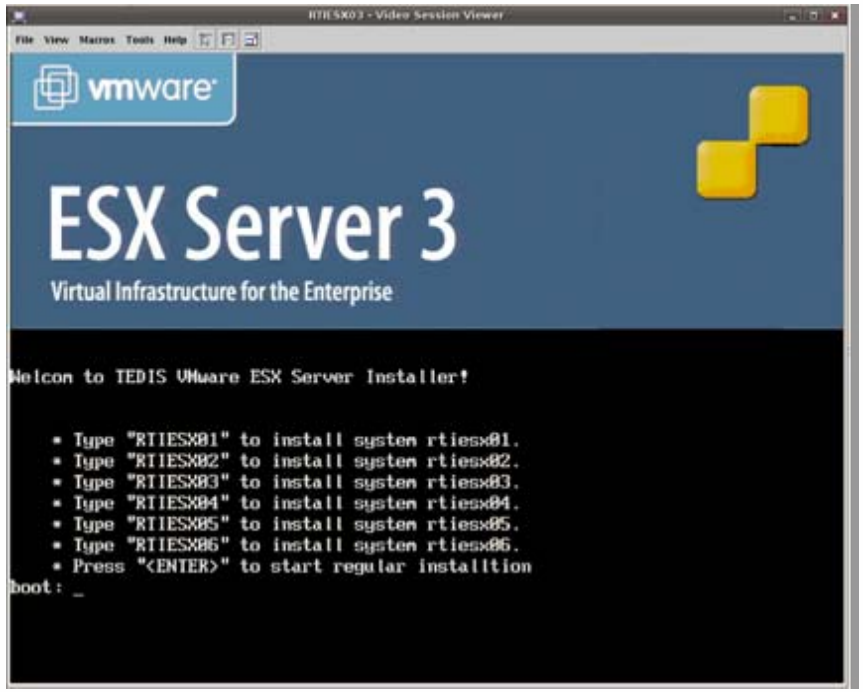

Kickstart file example

```
# This file is used for VMware ESX Server Scripted Install Deployment
# Installation Method
cdrom
# root Password
rootpw --iscrypted  $$9el$eUieKkFf$mW75WnAq9nkmJELvBUi9a/

# Authconfig
auth --enableshadow --enablemd5

# BootLoader ( The user has to use grub by default )
bootloader --location=mbr

# Timezone
timezone America/Los_Angeles

# X windowing System
skipx

# Install or Upgrade
install

# Text Mode
text

# Network install type
network --bootproto static --ip 10.5.1.73 --netmask 255.255.255.0 --gateway 10.5.1.2 --nameserver
10.5.2.12 --hostname rtiesx03.tedis.local --addvmportgroup=1 --vlanid=0

# Language
lang en_US

# Langauge Support
langsupport --default en_US
```

P2V Utilization

- Physical Machine connected directly to LUNs with over a terabyte of data.
- Use Raw Device Mapping (RDM)
- Deployed in few hours
- Saved many days

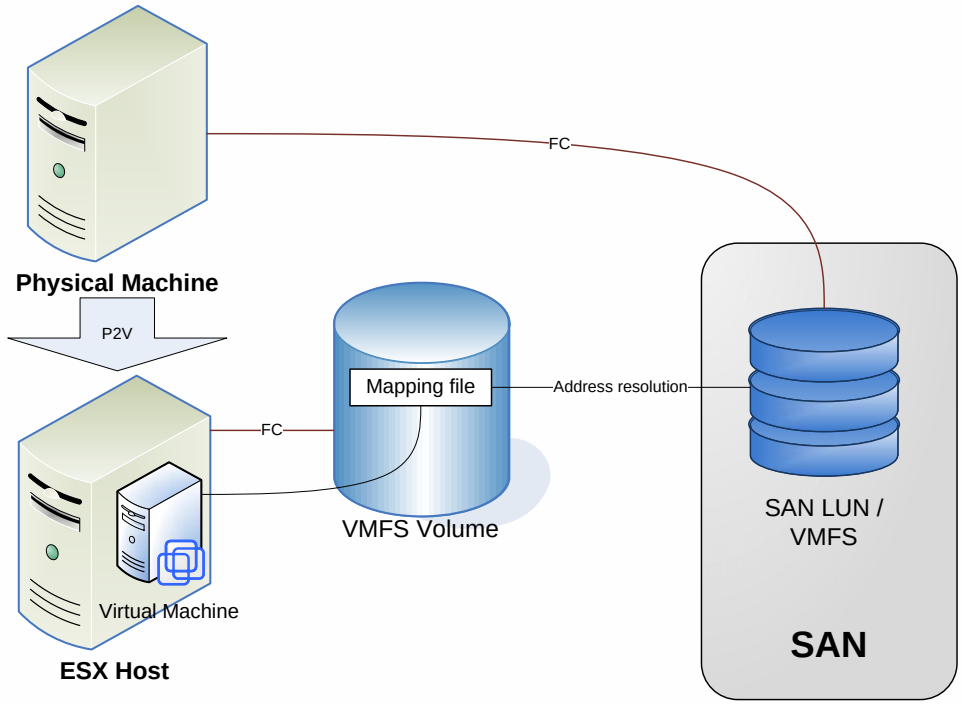

Upgrade to VI3

- Created a new LUN for VMFS v3
- Migrated all VMs from ESX01 before upgrade
- Used two environments for a while
- VMotion for minimum VM downtime
- Migrating other ESX servers
- Destroy old VMFS 2

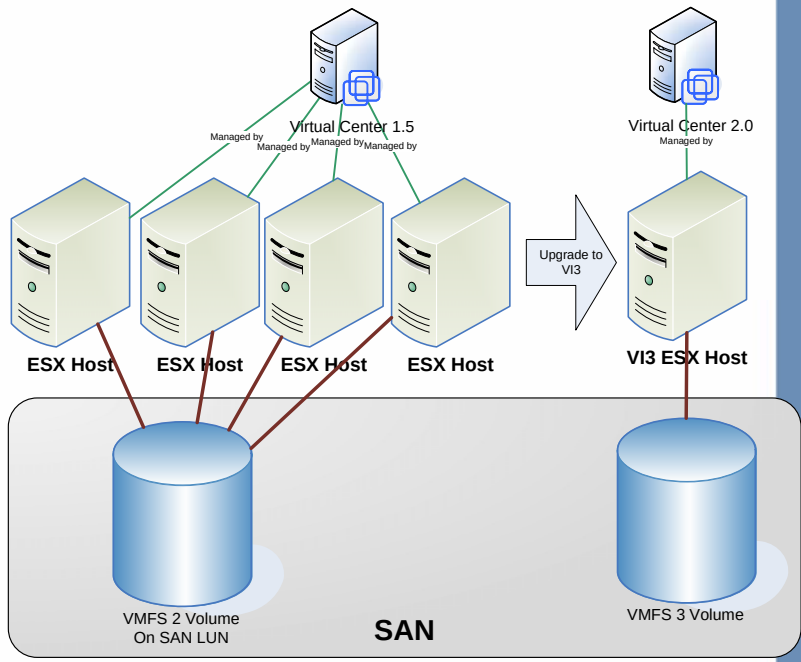

VirtualCenter Servers

- Distributed, clustered redundant servers
- Two Cross-host clustered SQL DB servers
- Two Cross-host clustered VC service servers
- 2 SQL Servers
- 2 VCenter servers

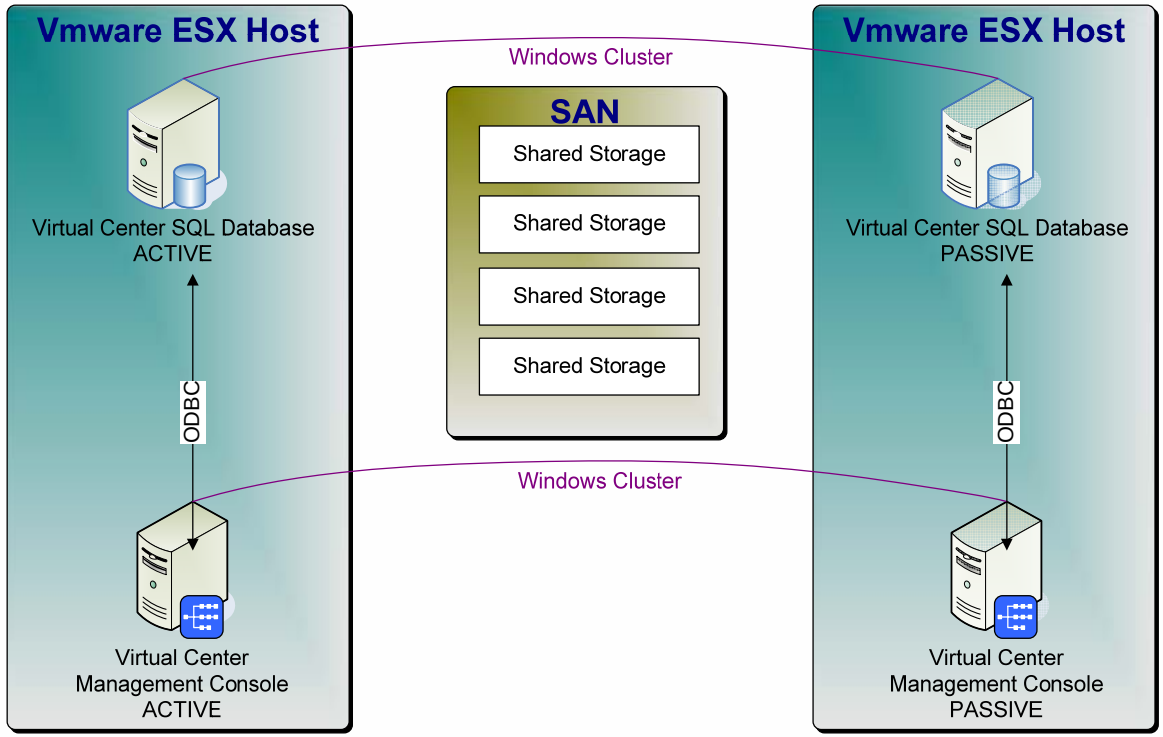

Virtual Center Security

- Role-based
- Active Directory
- Across domain
- Extremely flexible

- No host SSH access
- No remote root
- Subnet restrictions
- SSL for http access
- C:\Documents and Settings\All Users\Application Data\VMware\VMware VirtualCenter\SSL\ /etc/vmware/ssl

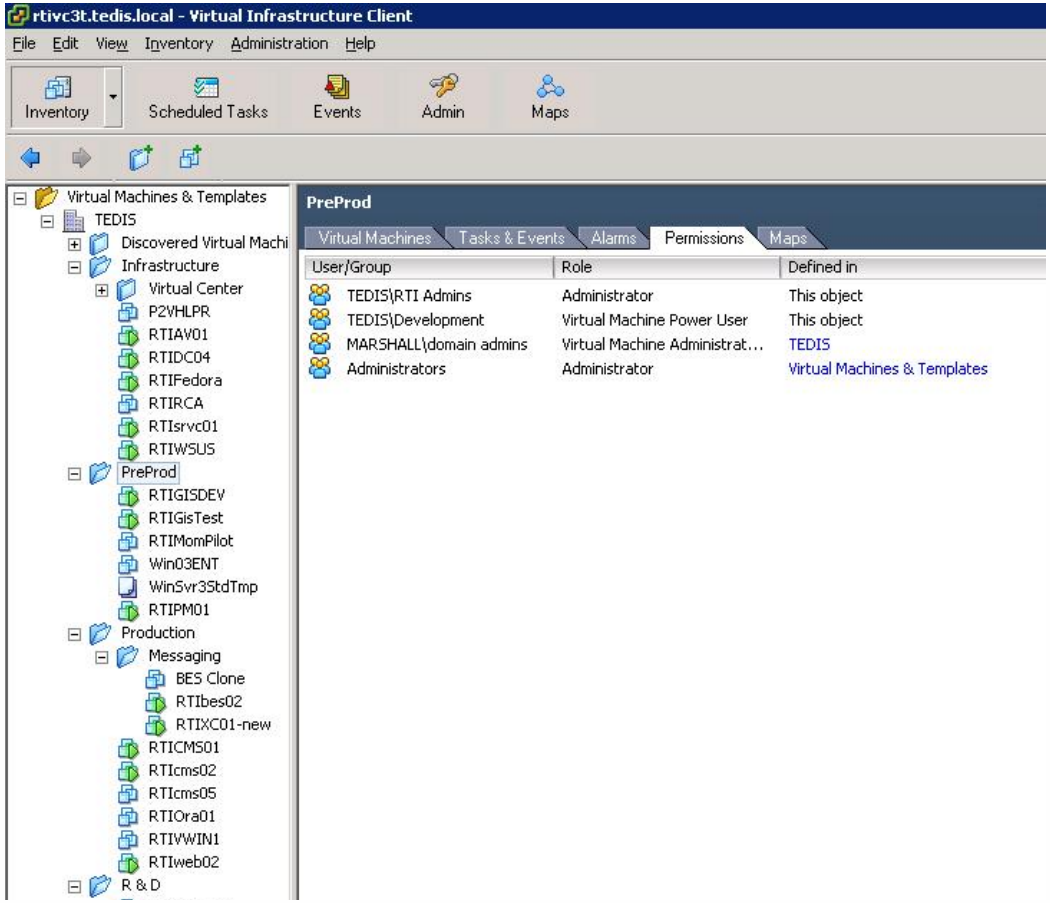

Operation Best Practice for Virtual Infrastructure

VMWORLD 2006

Third Party Tools

- Dell OpenManage Server Administrator
- EMC Navisphere Agent for Linux
- Commvault iDataAgent for Linux
- Changed fstab, sshd and other linux services
- IPSwitch's Whatsup Pro for monitoring SNMP
- Centralized log management: Splunk, Syslog-ng *

Install Dell OpenManage

- http://www.dell.com/downloads/global/solutions/Installing_Dell_OpenManage_50_on_ESX_3.pdf
- `./srvadmin-openipmi.sh install-force dkms install --force -m openipmi -v 35.13.RHEL3`
- `./srvadmin-install.sh`
- `esxcfg-firewall -o 1311,tcp,in,OpenManageRequest`
- `srvadmin-services.sh start`
- Centralized management with ITA

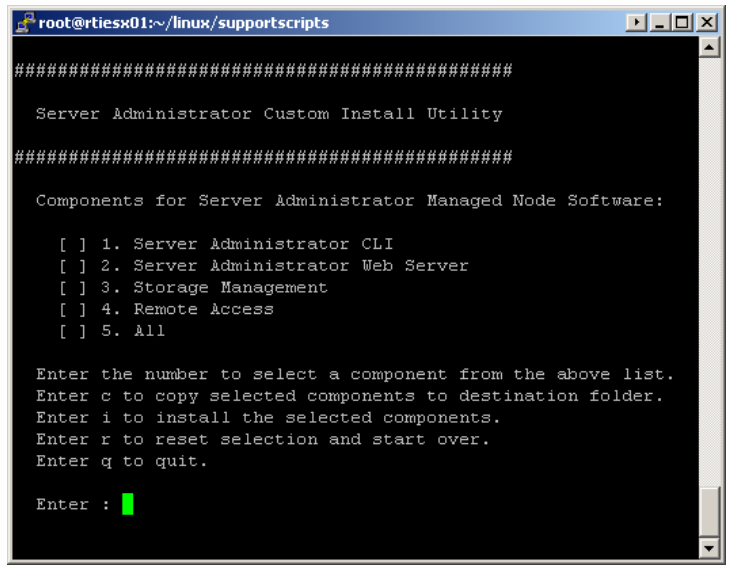

Dell OM Server Administrator

- Installed on the client
- Provide Hardware info
- Update firmware drivers
- Send alerts

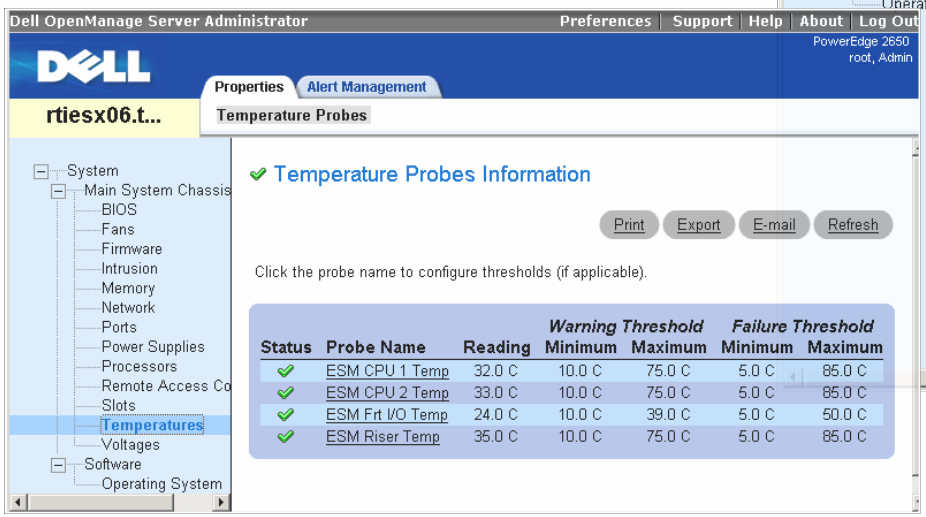

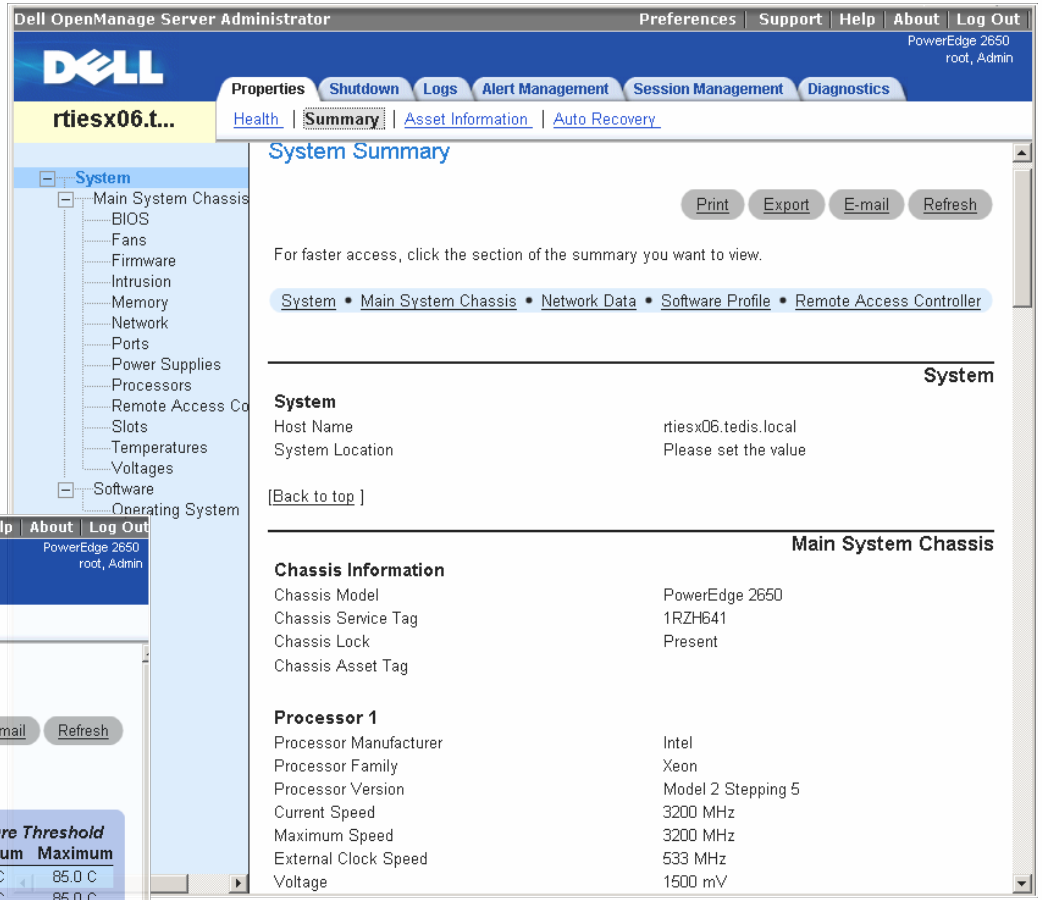

Use SNMP for Remote Management

- Edit snmpd.conf
  `# vi /etc/snmp/snmpd.conf`

  Configure it to point to the management server IP address
  use a community name (here it's public)
  `trapsink *.*.*.* trapcommunity public`

  Then start snmpd service
  `# service snmpd start`

  Configure it to autostart
  `# chkconfig snmpd on`

  Test it on local machine
  `#snmpwalk -v 1 -c public localhost system`

  Test it on another system (x.x.x.x is the esx server's IP)
  `#snmpwalk -v 1 -c public x.x.x.x system`
  `Compile ESX MIBs from ESX CD media`
  Configure your management server to receive SNMP and act upon.

Register MIBs

- /usr/lib/vmware/snmp/mibs/ on ESX
  - VMWARE-ESX-MIB.mib
  - VMWARE-RESOURCES-MIB.mib
  - VMWARE-ROOT-MIB.mib
  - VMWARE-SYSTEM-MIB.mib
  - VMWARE-TRAPS-MIB.mib
  - VMWARE-VMINFO-MIB.mib

IPSwitch Whatsup Pro

- Monitoring
- Send alerts
- Forward SNMP
- Log history
- Reporting
- Shortcuts
- Monitors:
- Ping
- Http
- 902
- VC service

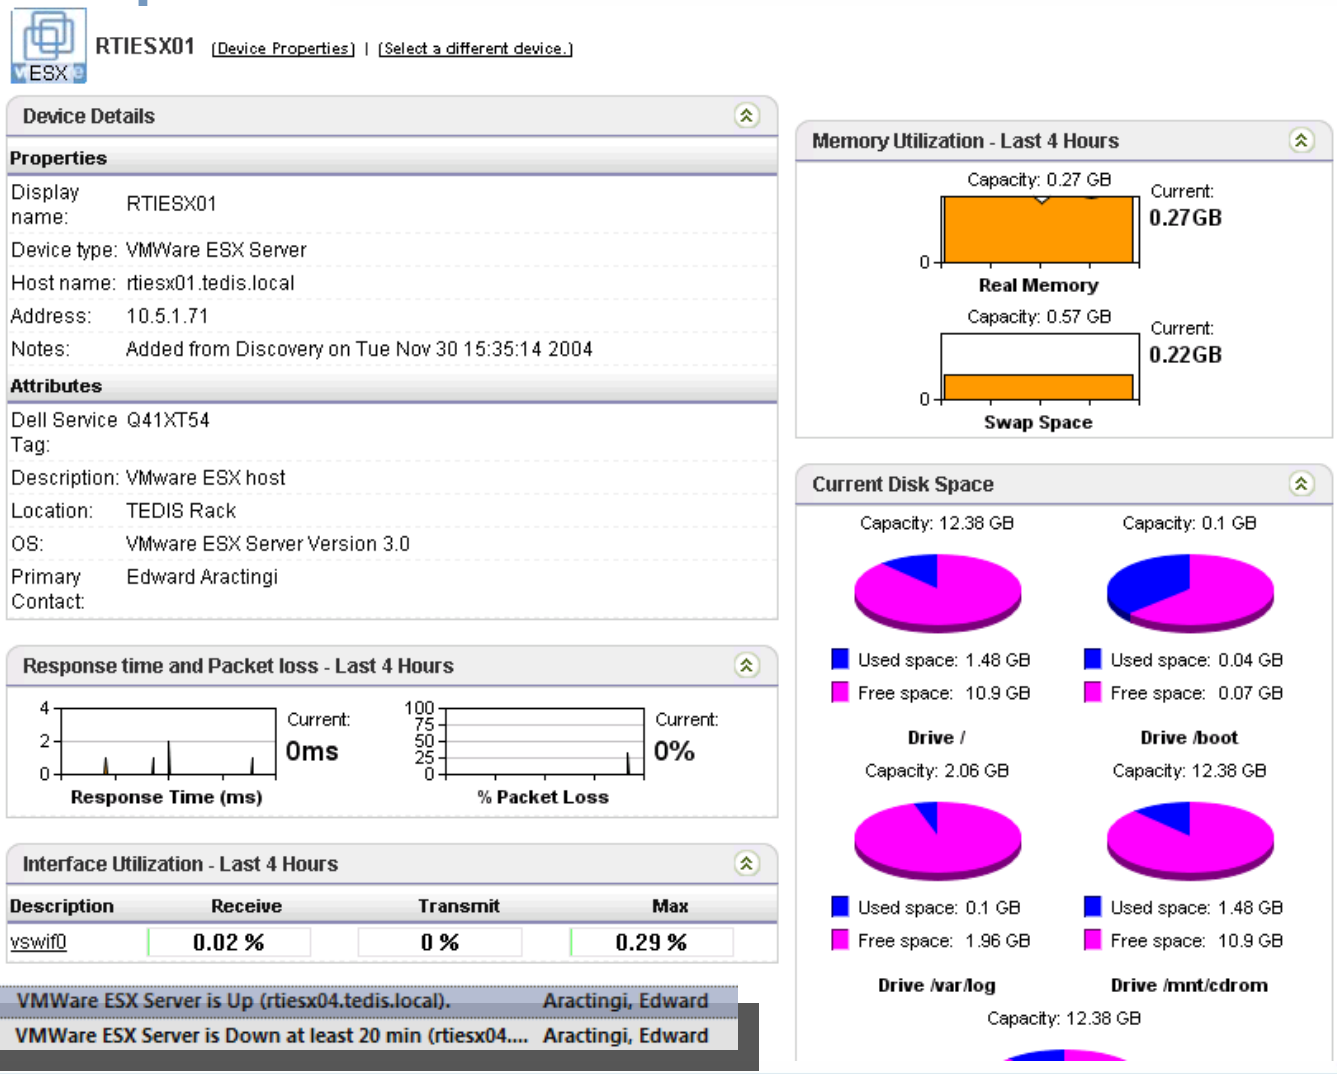

EMC Navisphere Agent

- # esxcfg-firewall -o 6389,tcp,in,EMCNaviAgent
- # esxcfg-firewall -o 6389,tcp,out,EMCNaviAgent

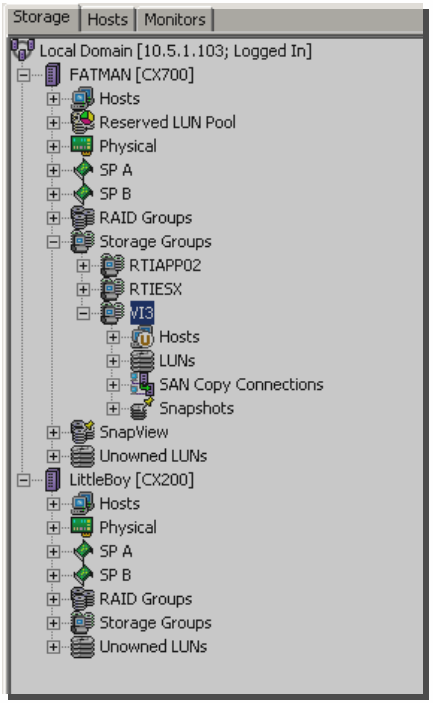

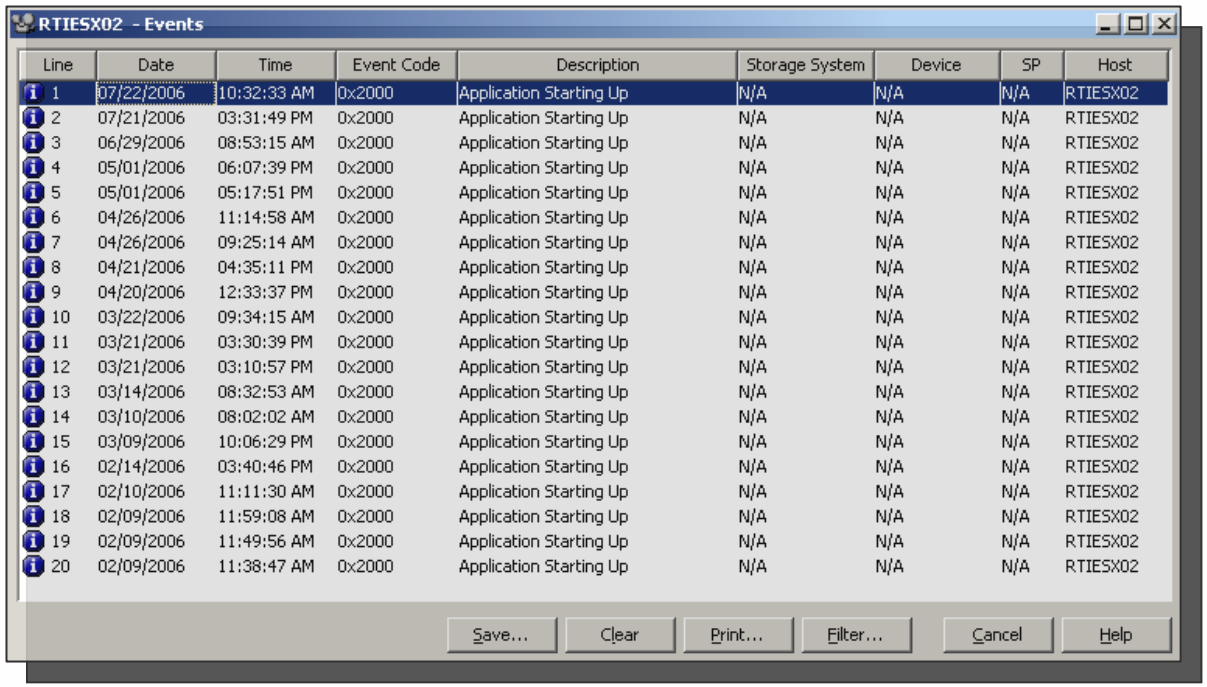

Patch Management

- Hardware Drivers and Firmware
  - Dell OpenManage IT Assistant
- ESX Servers updates
  - esxupdate (Setup local depot) *
- VMs OS updates
  - WSUS server and RedhatUp2date

Operation Best Practice
Backup Strategies

VMWORLD 2006

Backup Strategies

- Commvault host agent to backup the entire VMFS volume
- SnapView clones for off-network backup
- Individual backup agents on VMs
- Consolidated backup on a physical machine

Install Commvault Agent on VI 3

- Install Commvault Agent for linux
- Open Commvault Static and dynamic ports
- Configure subclient for pre/post-backup scripts

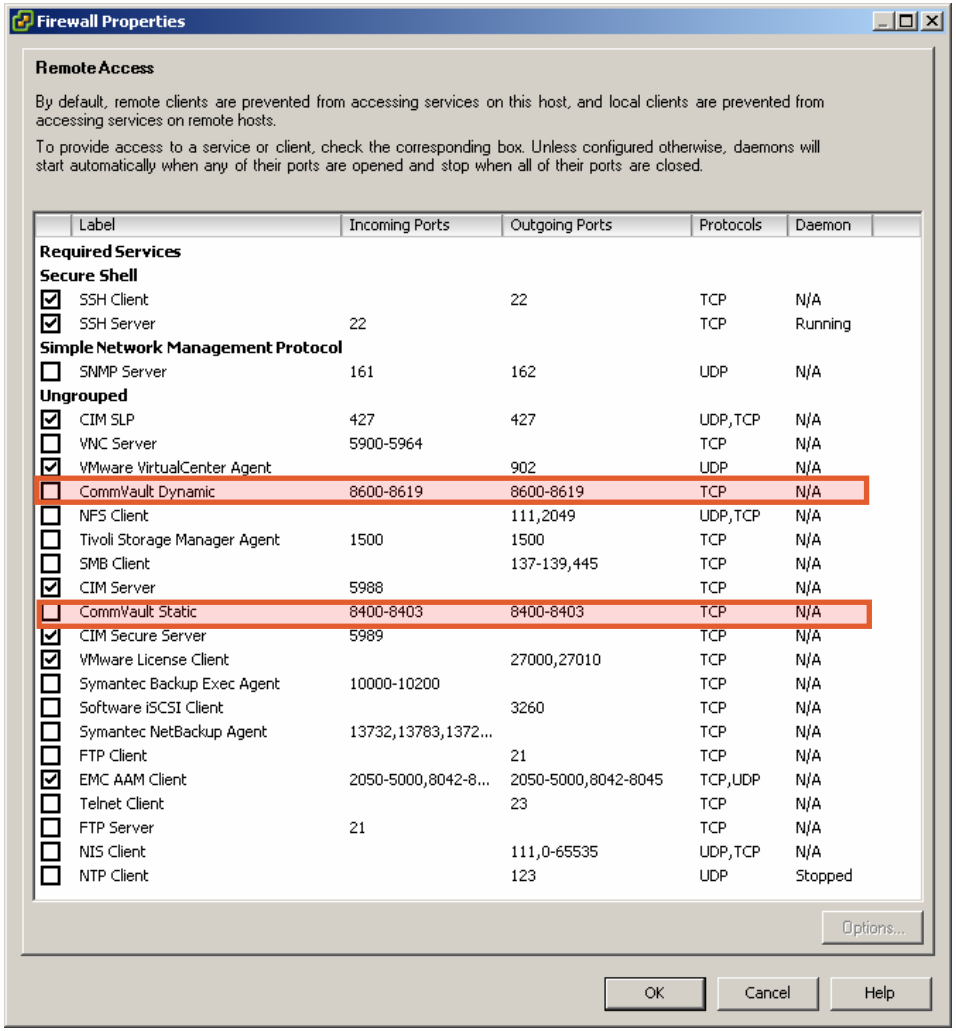

SAN Clone

- Commvault iDataAgent with Pre-backup and Post-backup scripts
- Use SnapView to clone the entire VMFS volume
- Minimize recovery time and scheduled downtime
- Reduce network traffic for backup
- Creating frequent point-in-time images
- Use clone on a different machine to backup the LUN

Consolidated Backup with VI3

- Used our Commvault Media Agent
- Has HBA and connected to the SAN
- Hosts the magnetic store for our D2D2T
- Minimize network traffic
- Optimize storage utilization

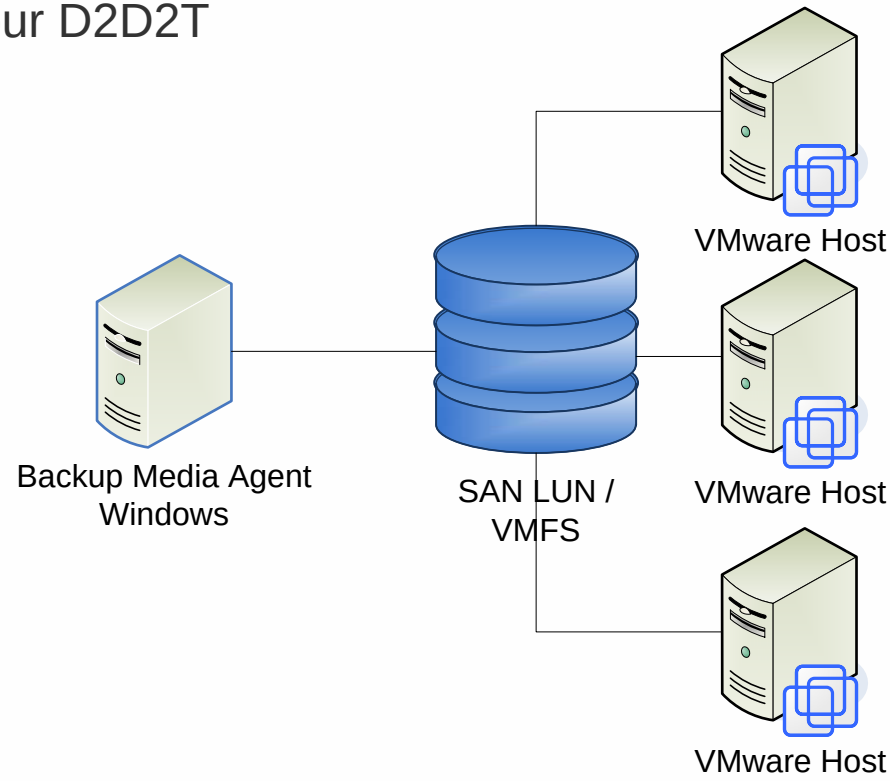

VCB SAN and Backup Configuration

- Backup VMs at file level.
- Save Backup Agents
- Used Commvault Media Agent as VCB Proxy
- Subclient for every VM

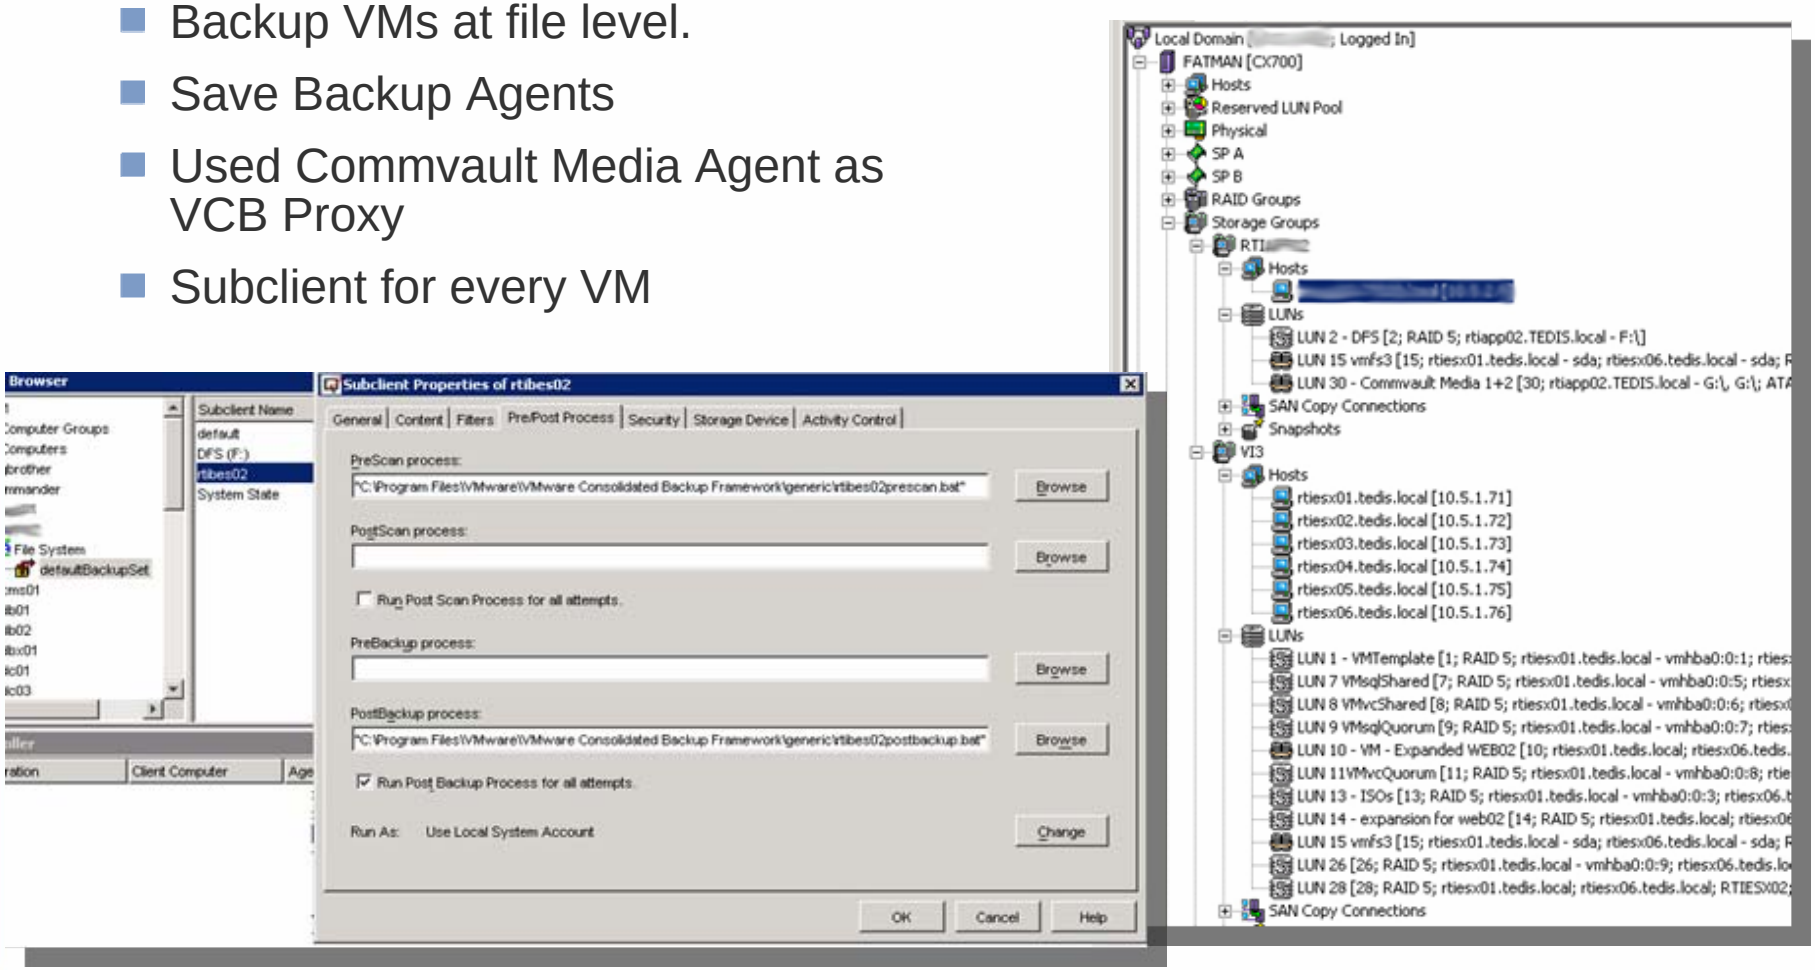

HA / DRS

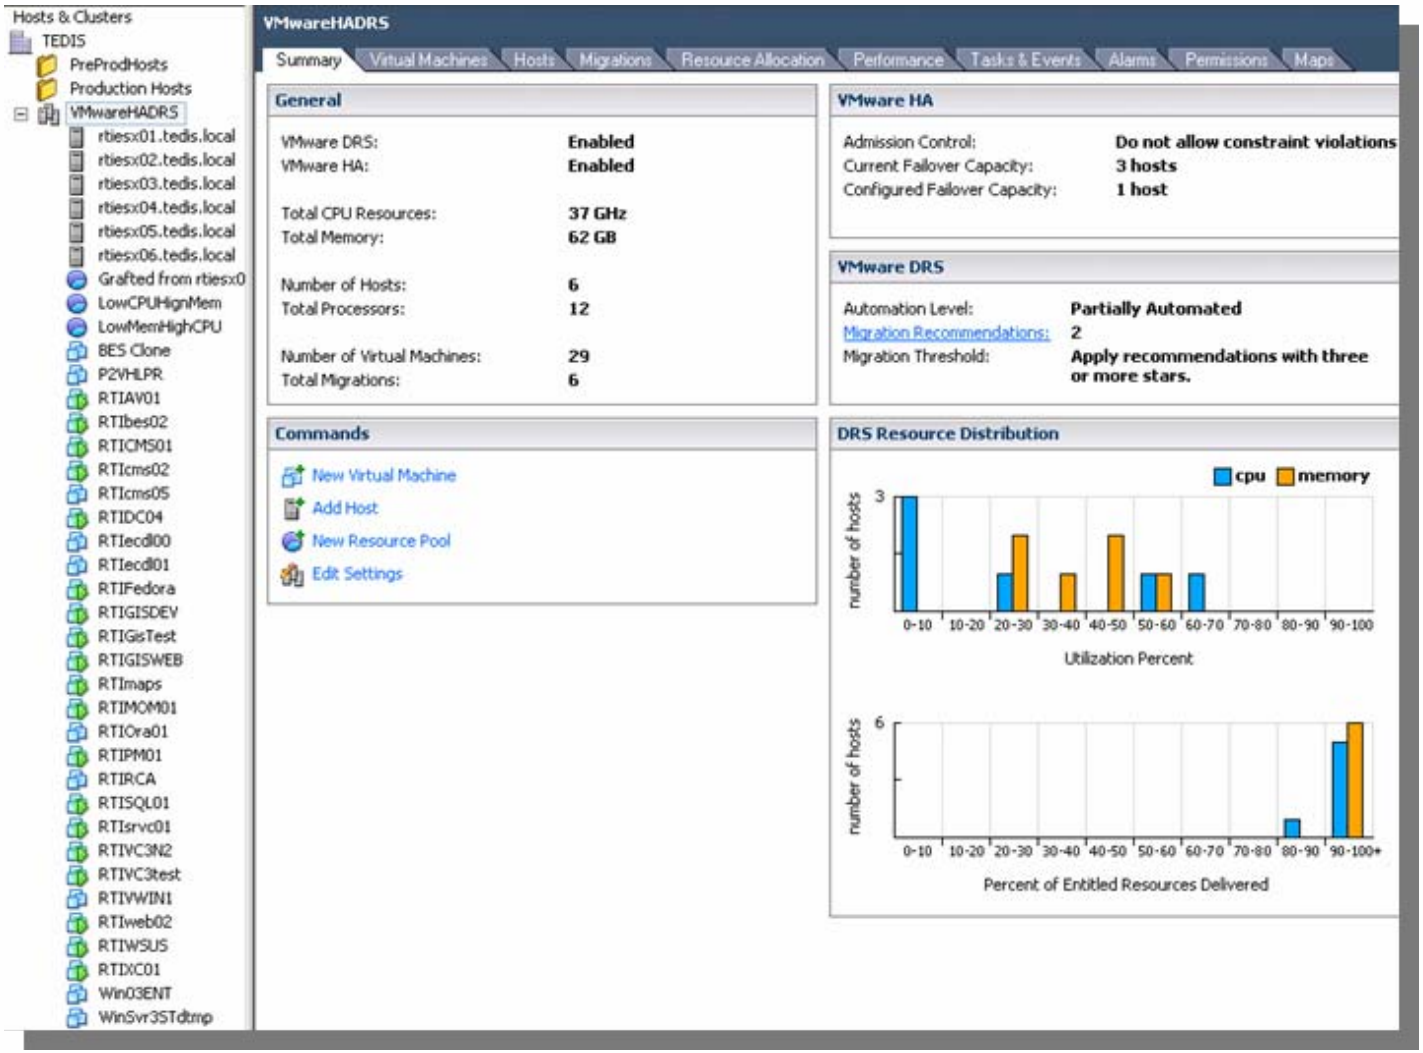
Hosts & Clusters
TEDIS
PreProdHosts
Production Hosts
VMwareHADRS
rtiesx01.tedis.local
rtiesx02.tedis.local
rtiesx03.tedis.local
rtiesx04.tedis.local
rtiesx05.tedis.local
rtiesx06.tedis.local
Grafted from rtiesx0
LowCPUHighMem
LowMemHighCPU
BES Clone
P2VHLPR
RTIAV01
RTIbes02
RTICMS01
RTIcms02
RTIcms05
RTIDC04
RTIecdl00
RTIecdl01
RTIFedora
RTIGISDEV
RTIGisTest
RTIGISWEB
RTImaps
RTIMOM01
RTIOra01
RTIPM01
RTIRCA
RTISQL01
RTIsrvc01
RTIVC3N2
RTIVC3test
RTIVWIN1
RTIweb02
RTIWSUS
RTIXC01
Win03ENT
WinSvr3STdtmp
VMwareHADRS
Summary
Virtual Machines
Hosts
Migrations
Resource Allocation
Performance
Tasks & Events
Alarms
Permissions
Maps
General
VMware DRS: Enabled
VMware HA: Enabled
Total CPU Resources: 37 GHz
Total Memory: 62 GB
Number of Hosts: 6
Total Processors: 12
Number of Virtual Machines: 29
Total Migrations: 6
VMware HA
Admission Control: Do not allow constraint violations
Current Failover Capacity: 3 hosts
Configured Failover Capacity: 1 host
VMware DRS
Automation Level: Partially Automated
Migration Recommendations: 2
Migration Threshold: Apply recommendations with three or more stars.
Commands
New Virtual Machine
Add Host
New Resource Pool
Edit Settings
DRS Resource Distribution
cpu
memory
number of hosts
0-10 10-20 20-30 30-40 40-50 50-60 60-70 70-80 80-90 90-100
Utilization Percent
number of hosts
0-10 10-20 20-30 30-40 40-50 50-60 60-70 70-80 80-90 90-100+
Percent of Entitled Resources Delivered

Challenges Summary

- User involvement at early stages might be a barrier
- In some cases, there is a need to rearchitect the data center
- Network performance is as important as other resources
- Securing the hosts is very critical
- The risk for proliferating number of VMs
- Misuse features (clone VM as backup instead of Snapshot)
- Licensing issues (OS, Backup Agent, Apps)
- SAN Disk usage (vmdk files are partially used)

Conclusion Continued

- 80% Planning /Project Management with 20% implementation
- Leverage existing management tools
- Automate any possible activity
- Spare resources are VERY useful
- Remember ESX implements Linux commands and concepts
- VMware VI provides many features, use them when possible
- Virtualization is an exciting technology, enjoy it!

Presentation Download

Please remember to complete your
session evaluation form
and return it to the room monitors
as you exit the session

The presentation for this session can be downloaded at
http://www.vmware.com/vmtn/vmworld/sessions/

Enter the following to download (case-sensitive):

Username: cbv_rep
Password: cbvfor9v9r

VMWORLD 2006