AMD on AMD: Production Consolidation using VMware and the AMD Opteron™ Processor

Michael Winslett

Manager of AMD IT Architecture

VMWORLD 2006

AMD: A Leader in Innovation

AMD designs and produces innovative microprocessors and low-power processor solutions for the computer, communications, and consumer electronics industries.

- Founded: 1969
- Headquarters: Sunnyvale, CA
- Employees: 10,750 worldwide
- Sales Mix: 78% international
- 2005 Revenue: $5.8 billion
- 2006 Revenue:
  - Q1 $1.3 billion
  - Q2 $1.2 billion

AMD: A Global Enterprise

… and many sales locations Worldwide

AMD: Market Momentum

- 12 consecutive quarters in which the year-over-year microprocessor sales growth exceeded 20%
- 26% of world wide x86 server processor market in Q206*
- 37% of the world wide and 47% of the USA 4-socket server business**
- AMD-based computers offered by leading hardware providers, including Acer, Dell, Fujitsu, Fujitsu Siemens, HP, IBM, Lenovo, Sun, and NEC

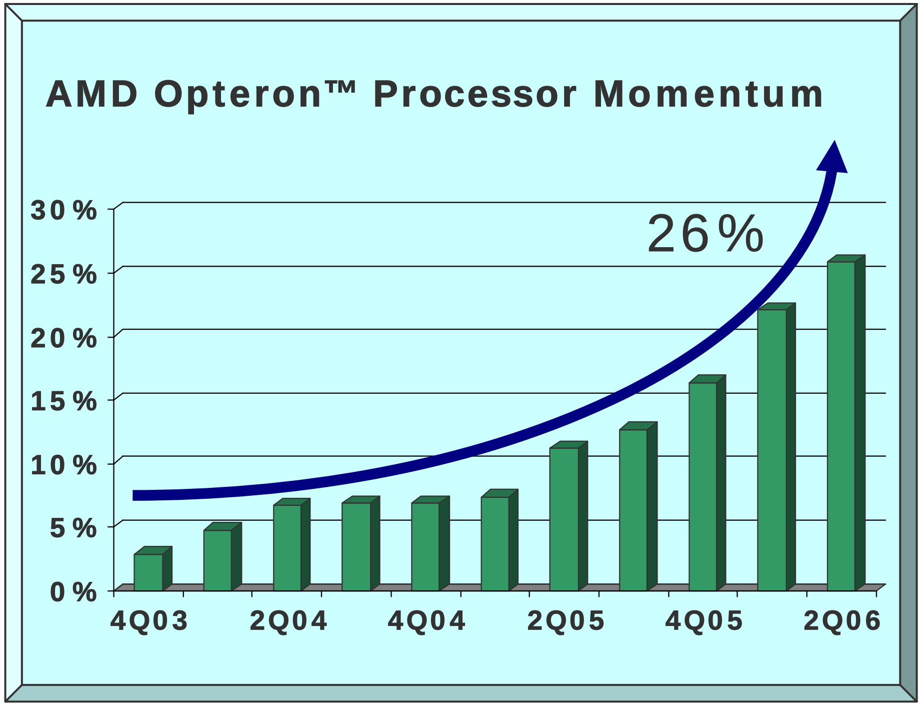

* Mercury Research **Gartner

Why AMD Launched the *AMD on AMD* Project

- Many older servers nearing end-of-life and ready for replacement
  - Even though these servers utilized older technology, they were still in many cases significantly under-utilized
  - Many servers not based on AMD Opteron™ processor technology
  - Desire to modernize this environment while moving to an AMD-based platform
- Power, space, and cooling challenges in the datacenter
  - Our main corporate datacenter was essentially full
  - Squeezing more servers into the datacenter through increased density would require substantial investment in our power and cooling infrastructure
  - Consolidation addressed both of these issues

About AMD IT

- Approximately 150 AMD employees globally
- Most operational services in the US provided through our co-sourcing partner, HCL Technologies, Ltd.
- Centralization of global services
- Regional teams supporting operations and region-specific solutions
- Team locations
  - Austin, TX
  - Sunnyvale, CA
  - Dresden, Germany
  - Singapore
  - Penang, Malaysia
  - Suzhou, China
  - Frimley, UK
  - Markham, Canada

Why AMD Went Virtual

- Although consolidation can be accomplished in the physical server world, this poses significant challenges and risks
  - Re-architecting applications to share servers
  - Lack of flexibility in balancing load
  - Difficulty in recovering from server failures
  - Difficult to manage resource consumption among multiple applications
- Virtualization provided many benefits
  - Allowed consolidation while preserving our current application architecture
  - Transformed to physical servers to commodity computing engines
    - No application-specific configuration exists on the physical boxes, hence they are very easy to replicate and replace
  - Increased redundancy for most systems
  - Quick and easy server provisioning – going from weeks to minutes

AMD on AMD Objectives

- Showcase an enterprise-class virtual server infrastructure based on the AMD Opteron™ processor
- Obtain efficiencies and benefits resulting from virtualization and consolidation
  - Reduced power and space consumption
  - Standard design and tools used globally
    - Initial design developed as part of the implementation in Austin, Texas.
    - This design was replicated to AMD sites in Sunnyvale, Dresden, Singapore, Penang, and Suzhou
  - Reduced per-server cost and labor requirements
  - Improved hardware resource utilization
  - Accelerated server provisioning from weeks to minutes
  - Improved overall availability
  - Improved flexibility for Disaster Recovery
  - Elimination of non-AMD-based servers
- Complete the project worldwide by the end of 2006

Project Methodology

- We identified a partner with a proven track record in large-scale consolidation using VMware: RapidApp
- RapidApp follows a consistent 3-stage methodology for consolidation projects based on VMware – Design, Planning, and Implementation

Project Methodology

- Design
  - Surveyed the server landscape in Austin
    - AMD provided server hardware specifications, CPU utilization, memory utilization, and I/O rates for all systems
  - Factors in selecting candidate systems
    - Type and age of physical server
    - Resource utilization
    - Benefits the server will realize from virtualization (e.g. automatic recovery from hardware failures)
  - RapidApp developed a paper design to support the identified candidate systems and to provide a VMWare ESX infrastructure based on best practices
  - Working together, we identified the key management processes that needed to be developed to operate a large production ESX 3.0 environment successfully

Project Methodology

- Planning
  - With the design as a roadmap, RapidApp worked with AMD to develop a plan for implementation
  - The plan considered available resources, skills, acceptable down times, costs, and overall project objectives
  - We established evening hours each week Monday-Thursday as timeslots for migrations, but we did not assign down times to specific systems at this stage
  - The output was a detailed project plan, identifying resources required, and an estimated cost for implementation
  - Cost included both the hardware and software licensing costs as well as RapidApp consulting resources to assist in the implementation

Project Methodology

- Implementation
  - Implementation included three major phases
    - Process development and testing
    - Migration of non-production systems on ESX 3.0 release candidate code
    - Migration of production systems on ESX 3.0 production code
  - Chose to engage with RapidApp to provide resources and expertise for the migration to augment the AMD team
  - Utilized the period before production ESX 3.0 code was available to refine the management processes for a virtual environment
  - Advertised the available migration timeslots and allowed application owners to sign up for a timeslot on a first-come, first-served basis
  - Migrated approximately 25 non-production systems using ESX 3.0 RC code
  - Upgraded to ESX 3.0 production code and continued migrations of production systems

Why VMware ESX 3.0?

- We wanted to be on a current VMware release when the project was done. We did not want to be facing a significant upgrade in the near-term once the project was finished
- Since we did not have an existing virtual environment to migrate to 3.0, the risk of using the newer release was minimal
- We wanted the improved VMFS locking capabilities in 3.0 that support larger ESX farms
- We wanted automatic recovery of VMs from a VMware ESX host failure
- We wanted the larger memory size allowed for VMs

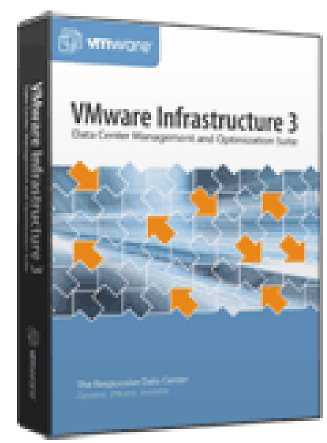

Infrastructure Sizing

- Infrastructure sizing was based on detailed analysis of observed system performance and resource utilization on the original physical systems and the desired consolidation ratios and cost efficiency we wanted to achieve
- Analysis accounted for capability differences between different classes of servers
- Experience during implementation was consistent with what the analysis predicted
  - Very few performance problems once systems were migrated to VMs
  - In the few cases where problems arose, the cause was always a change in resource utilization that occurred after our initial analysis
  - We added a checkpoint in the migration process to avoid this issue

ESX Server Hardware Selection

- The HP DL585 quad-processor, dual-core AMD Opteron™ processor-based server was a "sweet spot" for both consolidation ratio and cost efficiency
  - Large enough to achieve a substantial reduction in servers (potentially greater than 30:1 consolidation ratio based on our average VM size)
  - Not so large that it posed excessive risk from a server failure
  - We had an excellent track record with this server in traditional physical server roles
- The HP DL385 dual-processor, dual-core AMD Opteron processor-based server was chosen for a separate smaller farm in our internet DMZ
  - Smaller number of virtual servers in this environment were well-supported by the smaller system
  - The DL385 incorporates appropriate hardware redundancy for a server in this role
- We chose not to use blade servers because they did not support the number of networking and SAN ports we required

Why AMD64 Technology?

64-bit and Multi-Core

Enhances performance while offering the flexibility to support both 32- and 64-bit applications

Performance-per-watt

Assists data centers in controlling power consumption and heat output

Direct Connect Architecture

Eliminates the 20-year old front-side bus, increasing system efficiency and scalability

AMD Virtualization

Increases utilization by enabling the running of separate, secure operating environments

ESX Server Configurations

- Production/Internal ESX Servers
  - HP DL585 Chassis
  - Quad Processor configuration using AMD Opteron™ 875s
  - 48 GB of Physical Memory
  - 2 - On-board 1000Mb Network Cards
  - 3 - NC7170 PCI-X Dual Port Network Cards
  - 2 - Single Port FC HBAs QLogic QLA2340
  - Expected VM capacity between 24 and 38 VMs per host

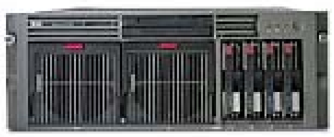

- DMZ ESX Servers
  - HP DL385 Chassis
  - Dual Processor configuration using AMD Opteron 280s
  - 12 GB of Physical Memory
  - 2 - On-board 1000Mb Network Cards
  - 1 - NC7170 PCI-X Dual Port Network Cards
  - Redundant power supply and fan configuration
  - Expected VM capacity between 10 and 16 VMs

Network and SAN Connectivity (Prod)

- 8 Total Physical NICs
  - 2-Production VMs
  - 2-Test/Dev VMs
  - 2-ESX Console
  - 1-VMotion
  - 1-Available
  - Connections split amongst physical cards and PCI Buses
- 2 QLogic HBAs
  - Connections to two fiber switches
  - 4 “paths” to each SAN LUN

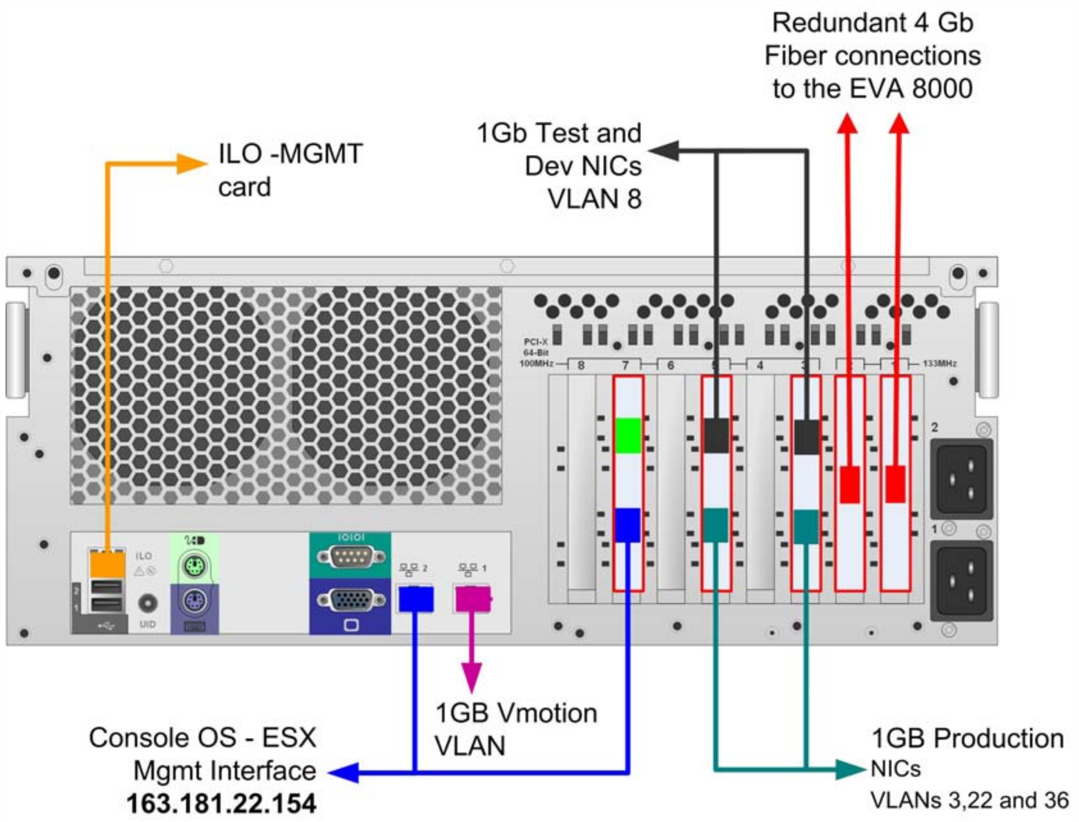

Network and SAN Connectivity (DMZ)

- 4Total Physical NICs
  - 2-DMZ VMs
  - 1-ESX Console
  - 1-VMotion
  - Connections split amongst physical cards and PCI Buses
- 2 QLogic HBAs
  - Connections to two fiber switches
  - 4 “paths” to each SAN LUN

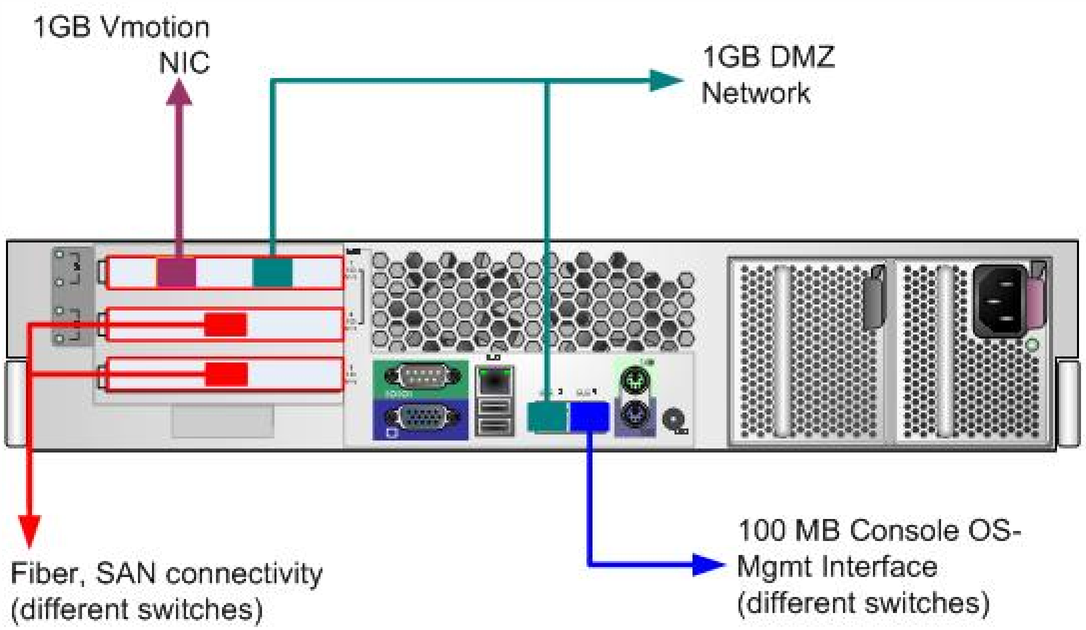

Automated VM Network Failover

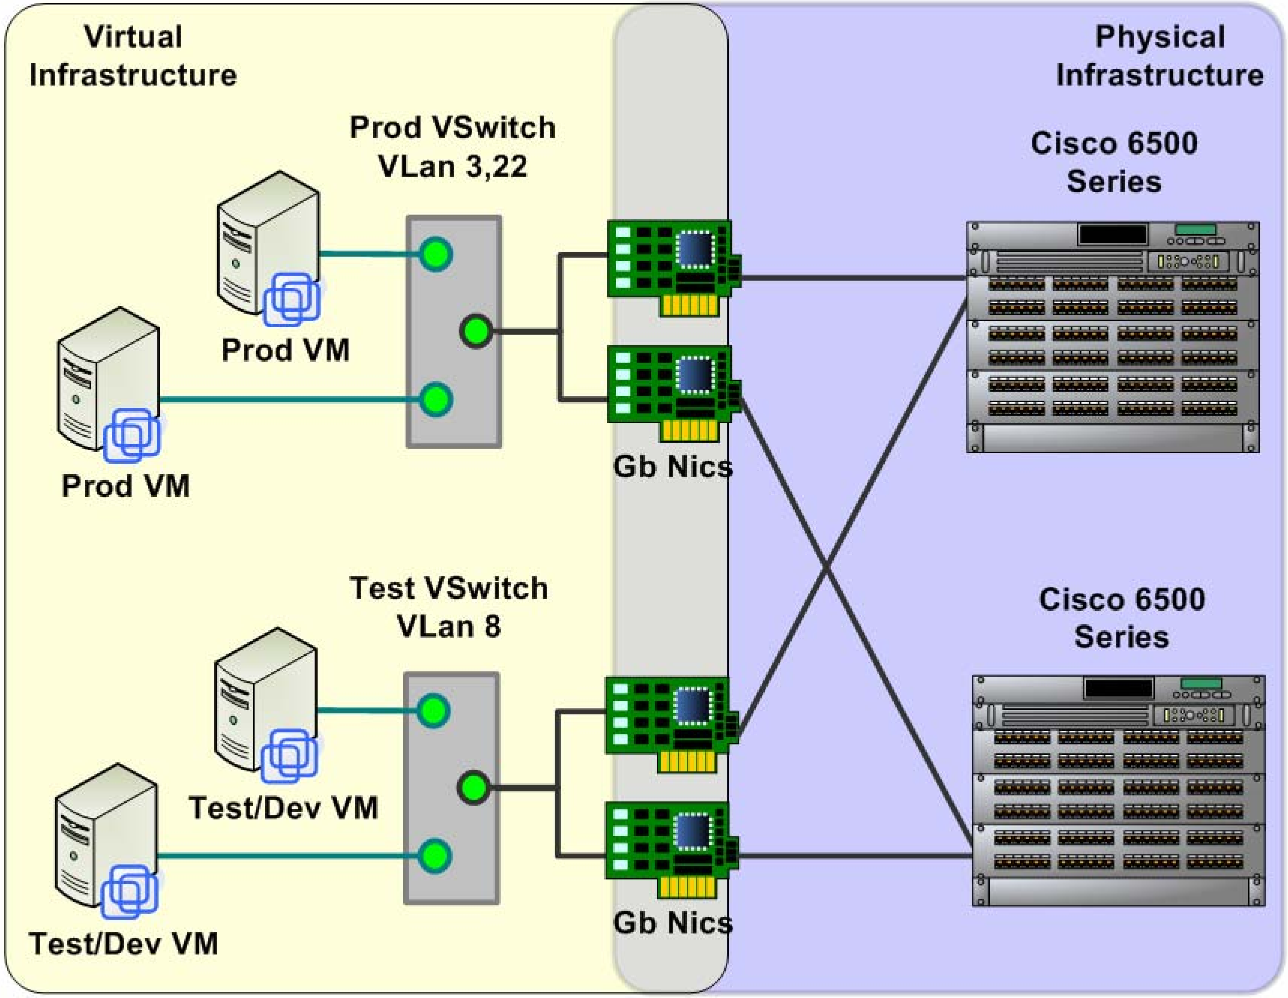

Automated VM Network Failover

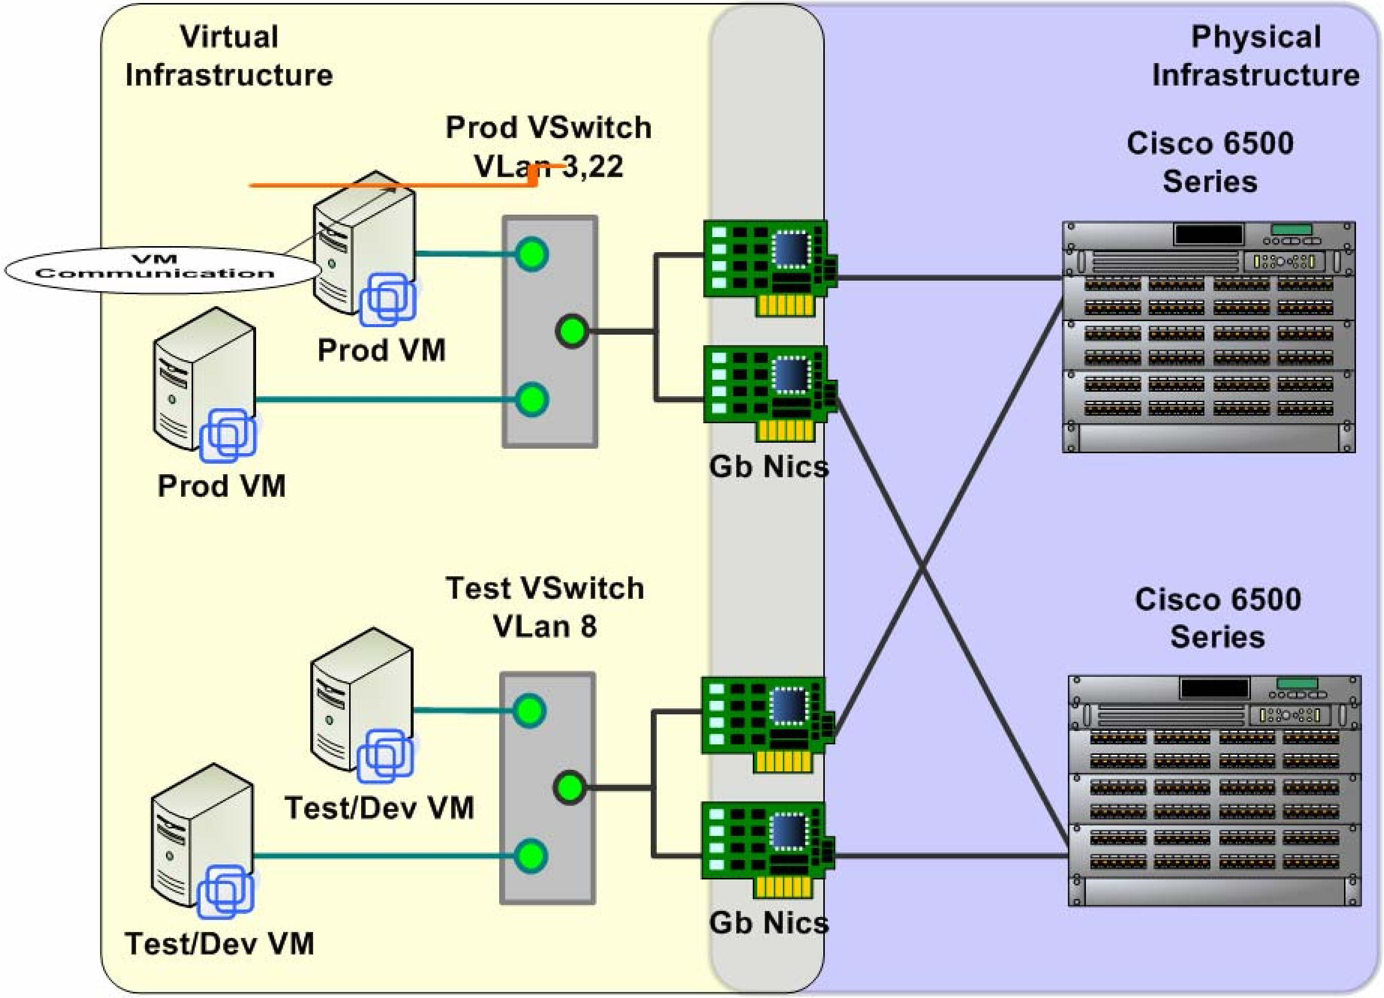

Automated VM Network Failover

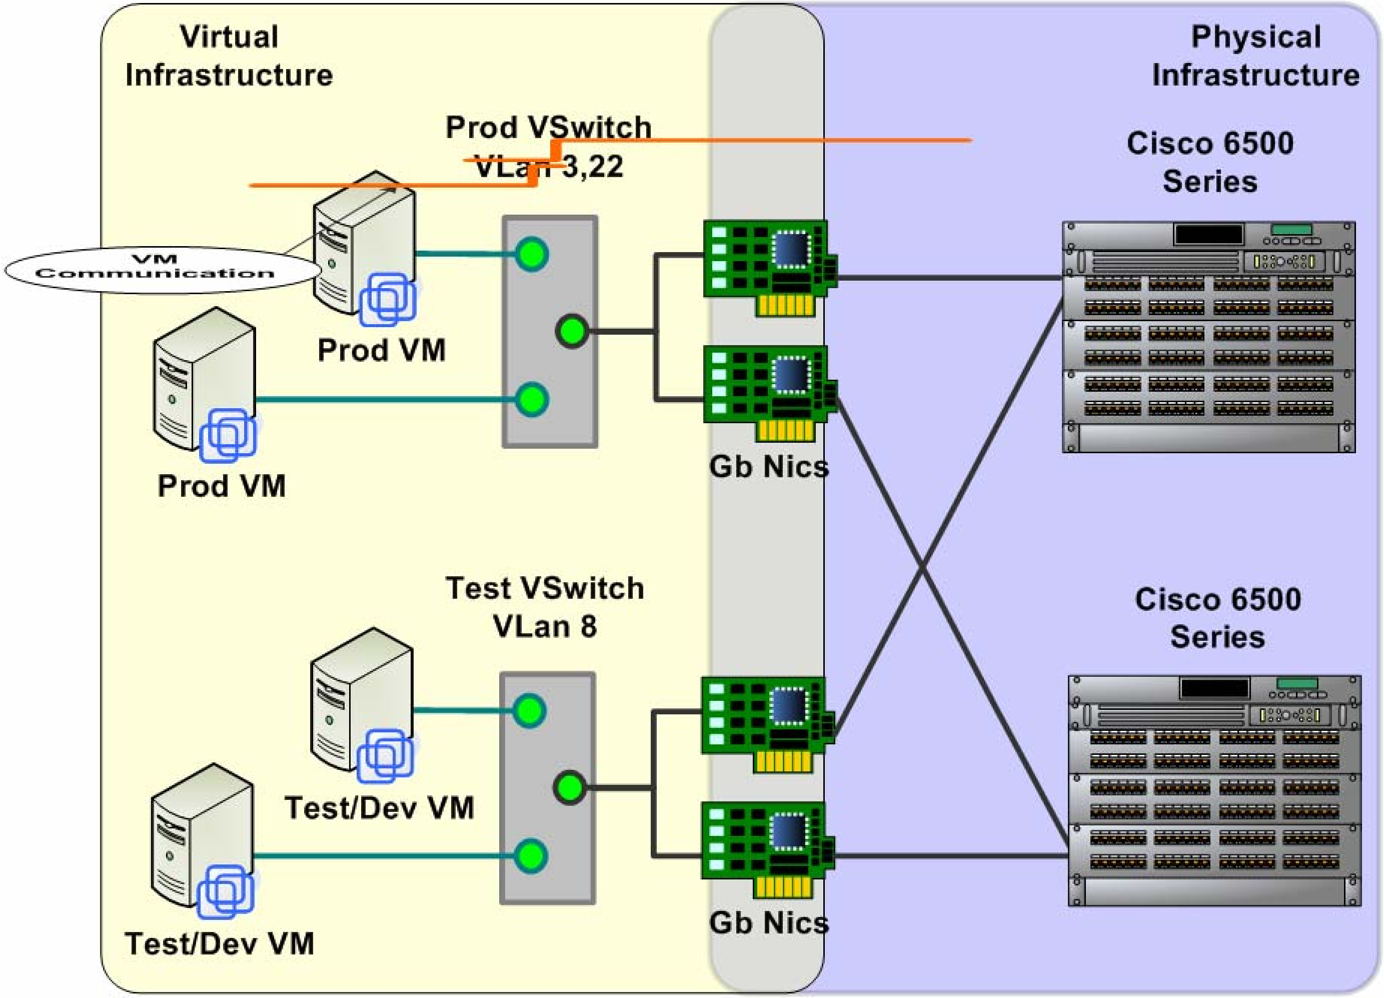

Automated VM Network Failover

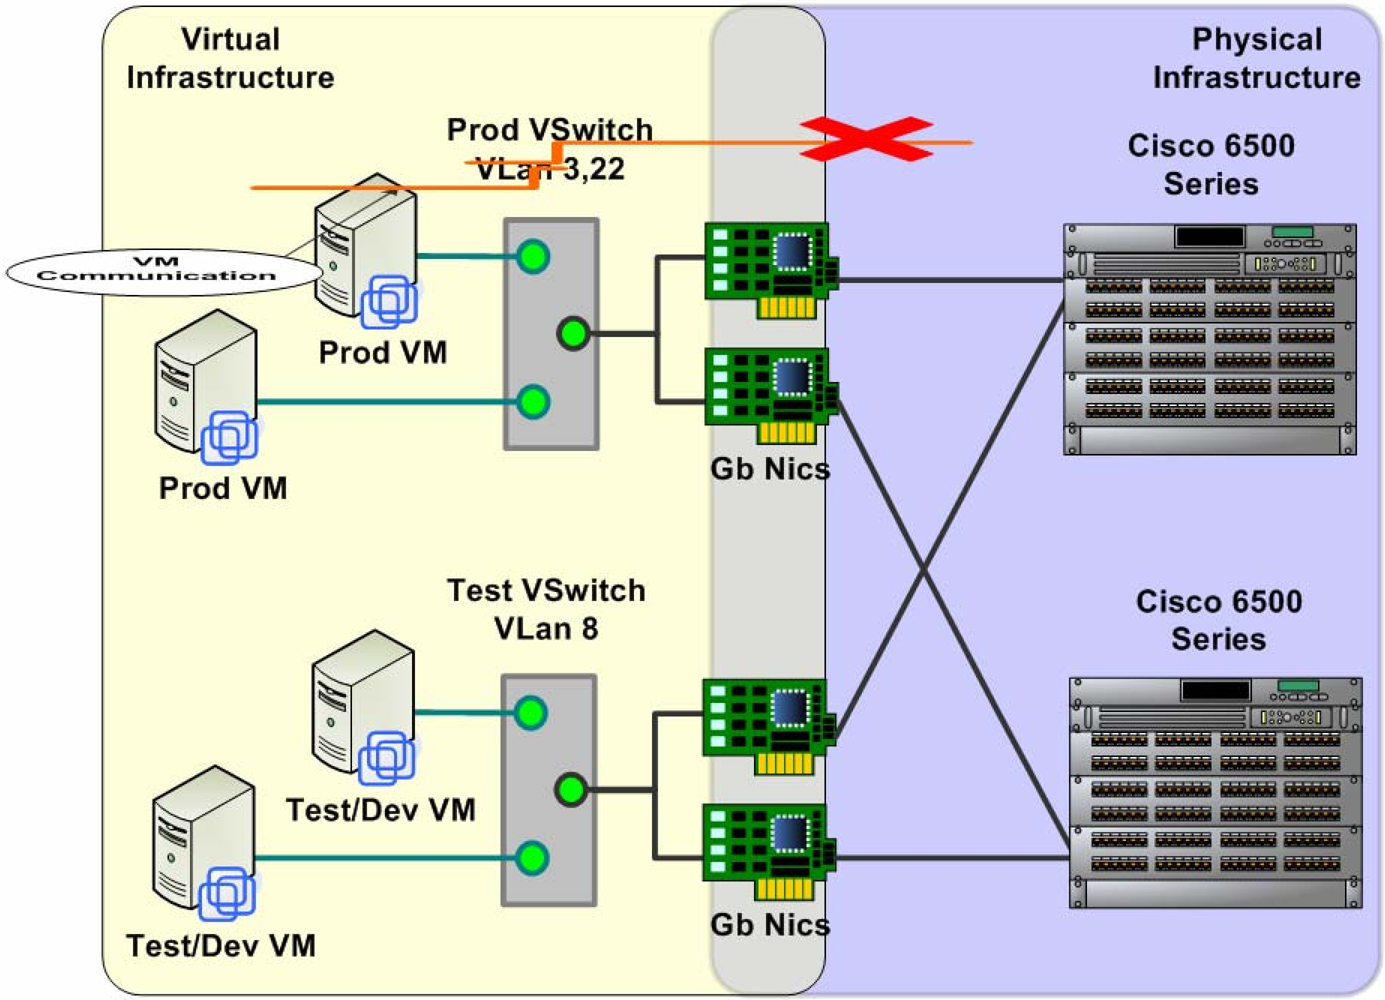

Automated VM Network Failover

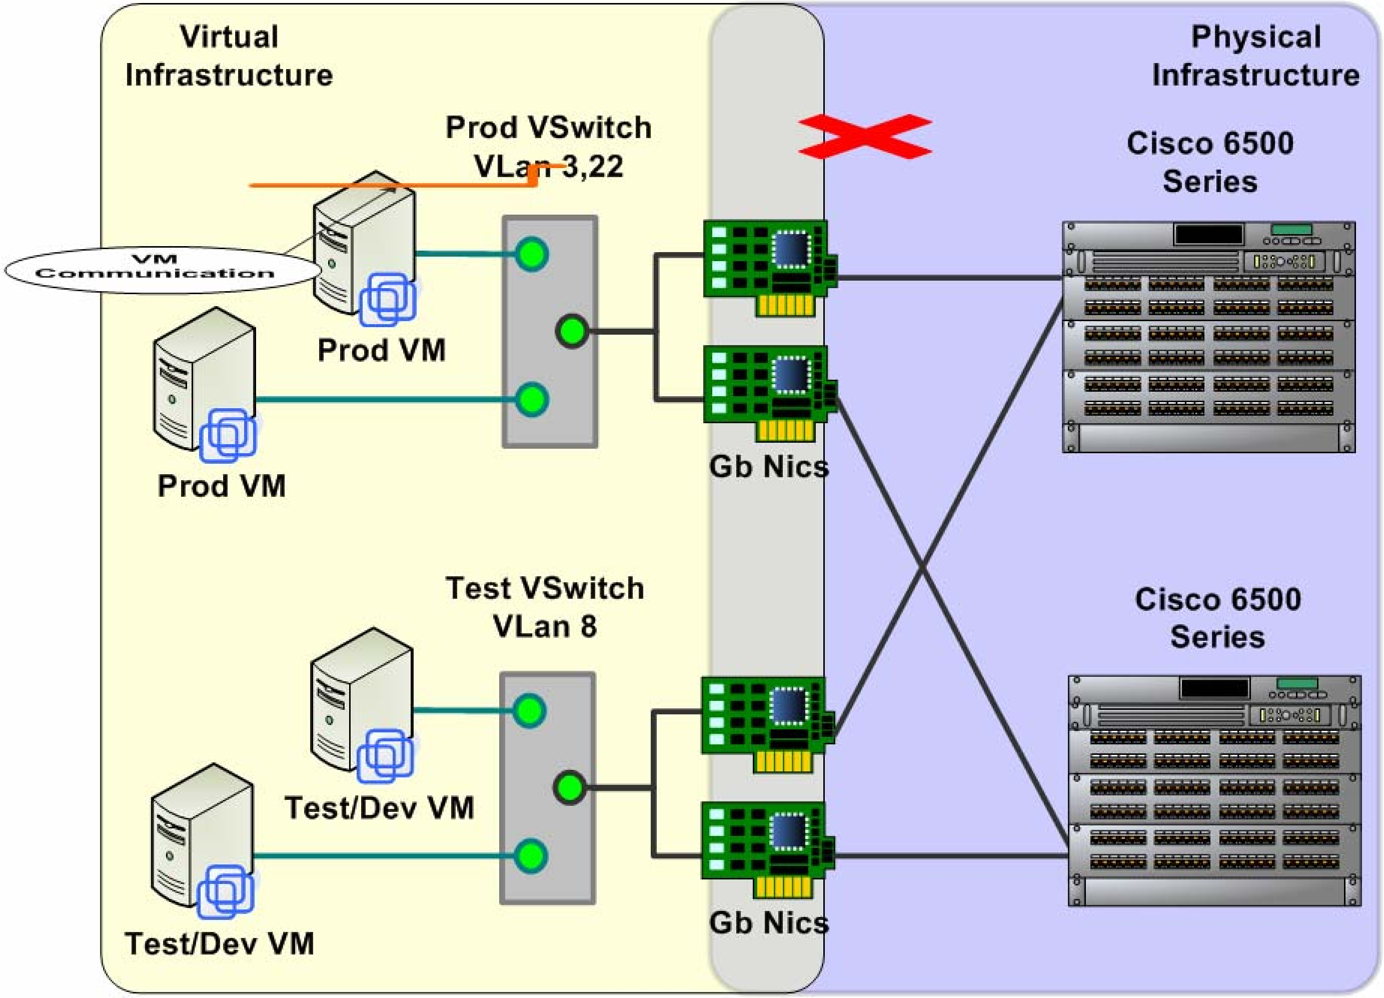

Automated VM Network Failover

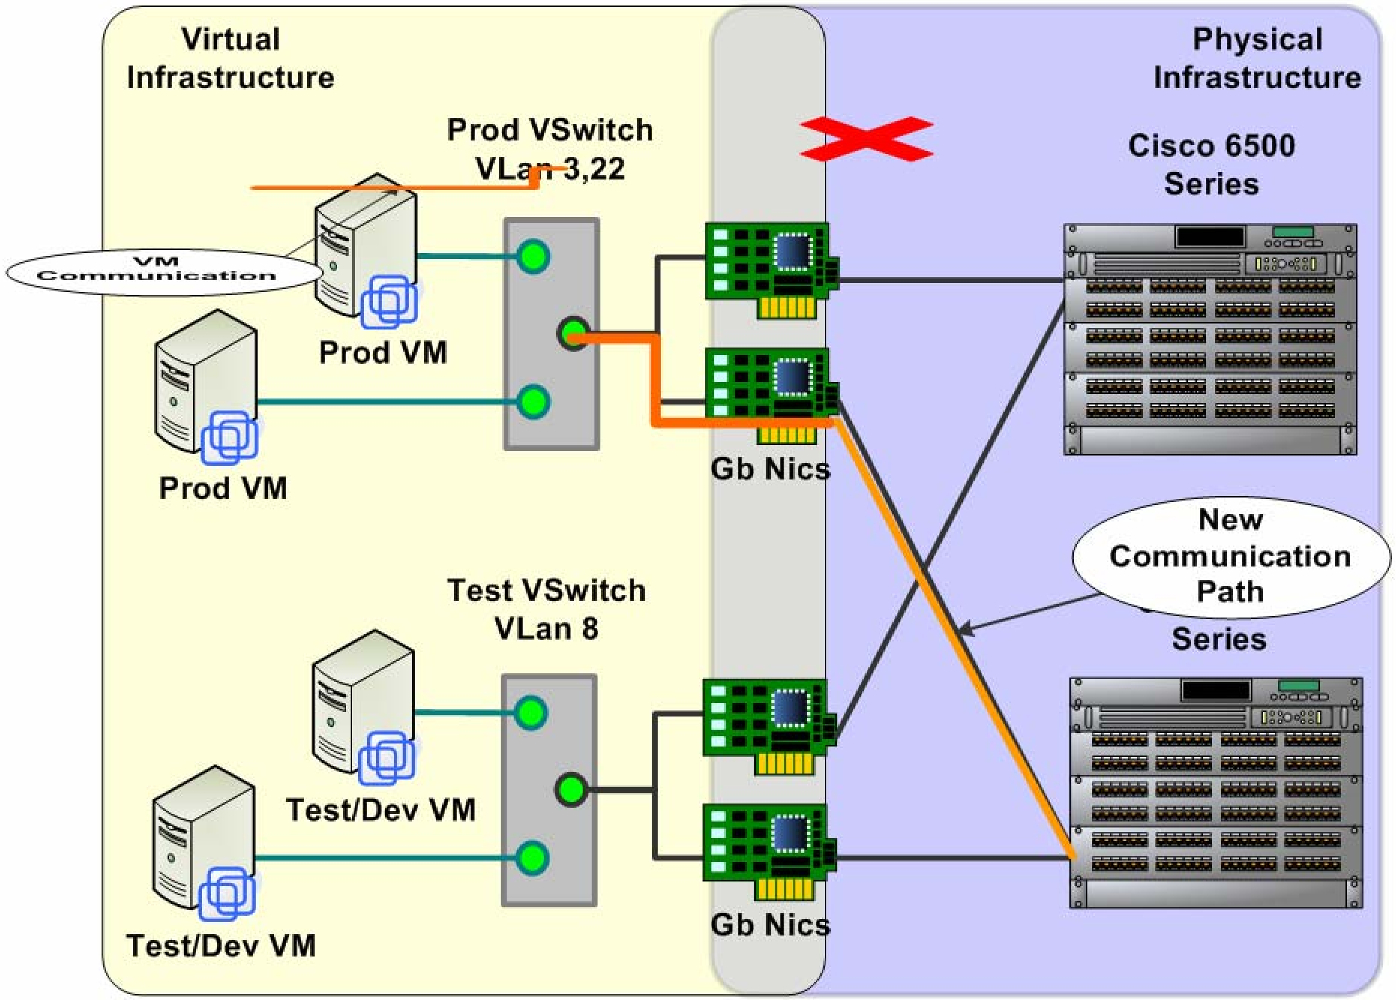

Processes and Procedures

- More than a dozen documented processes/procedures created to manage the new environment including:
  - Automated ESX Server builds using Altiris
  - VM request and deployment process
  - Host and VM recovery procedure
  - Patching and patch testing processes for hosts and VirtualCenter
  - Weekly and Monthly preventative maintenance tasks for ESX
  - Snapshot process for VMs including standard SLAs for snapshot delivery
  - QA Checklists for both Host and VM builds
  - Process / procedure to create and update VM templates in the environment
  - Procedures for granting and denying access to the MGMT tools
  - Monitoring configuration standards for VMs and Hosts
  - VMFS/LUN provisioning process

Tools Used in Managing the Environment

- Extensive use of VMotion
- Distributed Resource Scheduler (in manual mode)
- Distributed Availability Services
- ESX hosts are monitored via VirtualCenter
- Virtual Machines are monitored via HP OpenView agent
- Veritas NetBackup used for backing up VMs at the VM level
  - Considering more advanced backup processes as tools are released for ESX 3.0
- Tool for Physical to Virtual migration (P2V): UltimateP2V

Consolidation Achieved

- In Austin, 117 servers consolidated to 7 active ESX 3.0 servers plus 2 swing servers
- In Sunnyvale, approximately 33 servers consolidated to 2 active ESX 3.0 servers plus 1 swing server
- New VMs have been added in Austin to bring the total number to 180
- Overall Physical/Virtual ratio between the two sites
  - 23:1 (not including swing servers)
  - 17:1 (including swing server
- Consolidation ratios will increase as new VMs are added to the environment

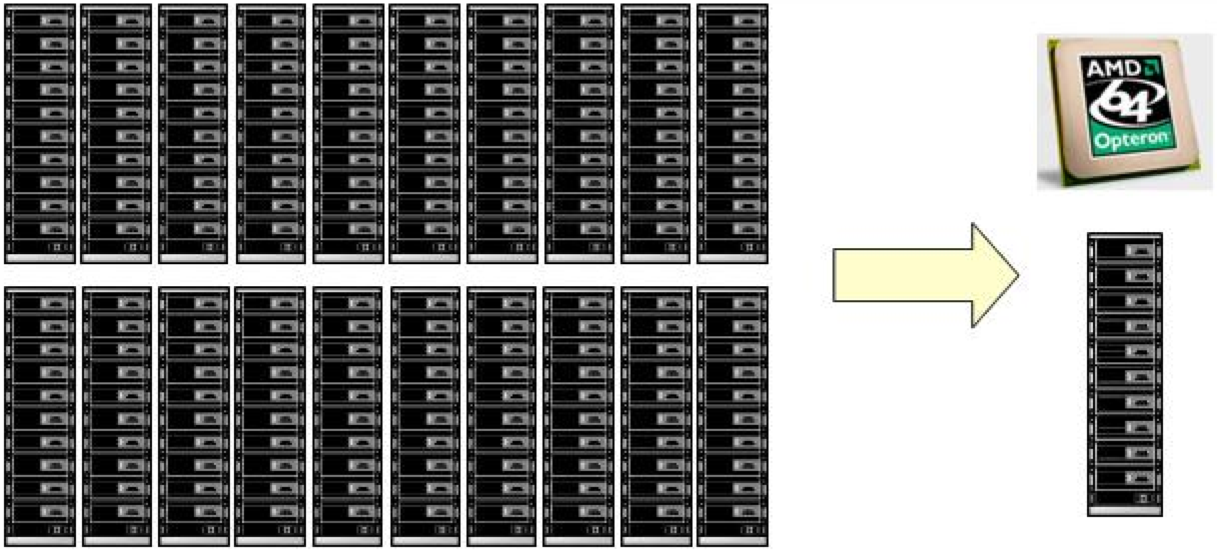

Power Savings (Including Cooling)

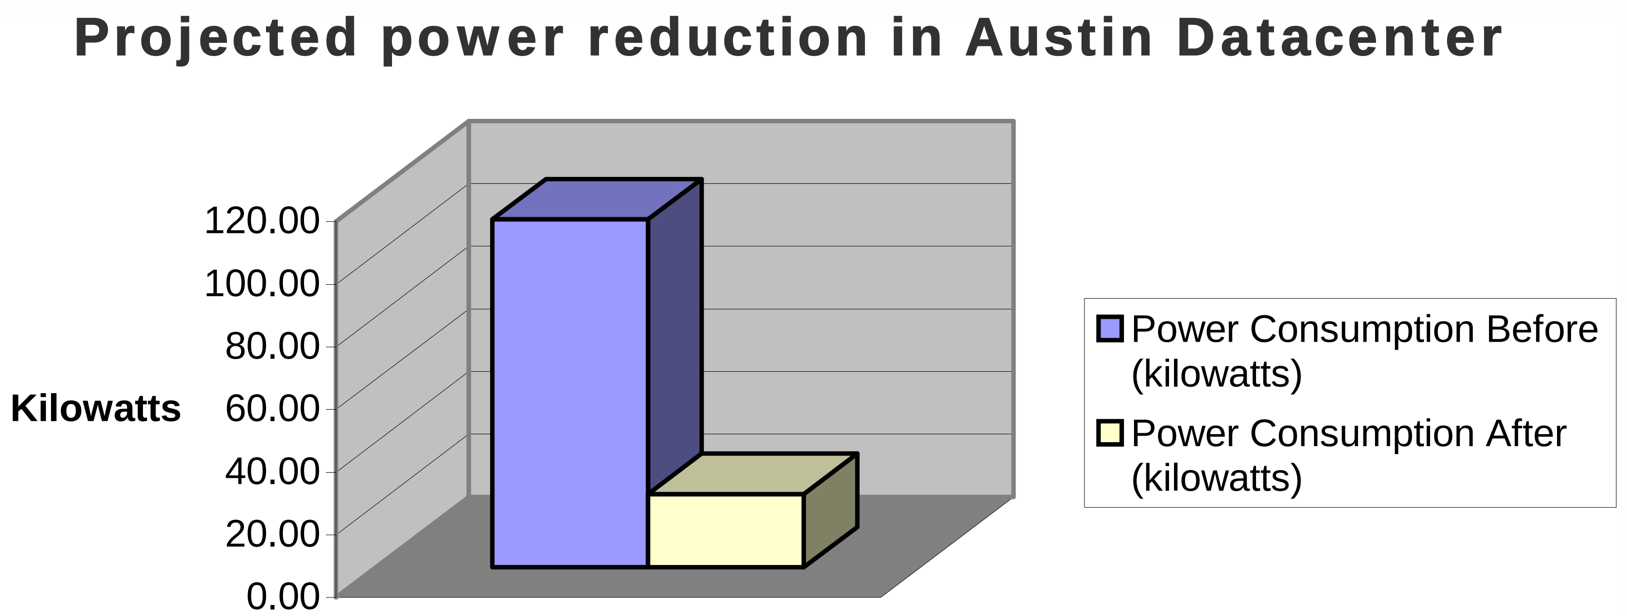

79% reduction in power consumption

$69K/year in power savings in Austin

Estimated $100K/year in total power savings globally.

AMD Smarter Choice

AMD on AMD – Financial Assumptions

- Analysis covers B3 Data Center (Phase 1 of Project)
- Consolidation (Virtual machines to Physical Server): 22 to 1
- Capacity to allow for 200 Virtual Machines (135 in initial scope)
- Hardware Refresh cycle every 3 years, one third per year
- Moving to virtual environment will free up B3 power capacity (cost avoidance)
- Cost opportunity to reduce B3 Data Center Kilowatt consumption
- Cost Opportunity for reduced support costs (TBD)
- Internal discount rate: 13%

AMD on AMD – Financial Summary

- 3 Year Net Present Value: **$1.7M**
- Payback Period: **<1 Year**

Timeframe	Savings Area	Financial Implication
Year 1 - 2006	Purchase of less servers to meet AMD on AMD goal	Reduced capital request in 2006
Year 2 - 2007	Purchase of less physical servers for increased capacity requirements	Reduction in 2007 capital request for server capacity
Year 2 - 2007	Cost Avoidance – B3 DC Capacity	No increase in B3 DC power costs
Year 4 - 2009	Reduced cost for server refresh	Reduced capital request in 2009
On Going	Reduced B3 DC Support Costs	Reduction in DC support expense

Other Benefits

- Improved redundancy for many systems
- Automatic recovery of VMs from ESX host hardware failures
- Greater flexibility in Disaster Readiness options
  - Assuming replicated data, ESX servers at a DR site could be quickly repurposed to run mission-critical systems
- Improved standardization of management processes

Next Steps

- We expect to complete rollout worldwide by the end of 2006
- We will utilize this design in ATI datacenters as we work to integrate AMD and ATI
- We will continue to track developments in the virtualization space and incorporate those into our standard design as appropriate

Conclusion

- *AMD on AMD* has been very successful for us
  - Significant consolidation achieved
  - Relieved power and cooling stress in Austin
    - Without this program, we would have had to do significant datacenter infrastructure upgrades
  - Achieved the expected benefits of virtualization
    - Improved redundancy and DR flexibility
    - Greatly reduced time to provision servers
    - Improved efficiency of Operations staff
  - Financial ROI

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