Customer Example (Fictitious): Bartok, Inc.

The Environment

- 1000 x86 servers maintained annually worldwide
  - 1 CPU: 300
  - 2 CPU: 500
  - 4 CPU: 200
- Several hundred physical servers deployed annually

The Goals

- Reduce expenses, labor, time, and floor-space required to run hundreds of workloads
- Increase agility of business operations

The Outcome

- 1,000 servers virtualized
- 12:1 consolidation ratio
- $6.5M (NPV) TCO savings over 3 years

Server Consolidation: Server Hardware

- Hardware reduction driven by type of workloads being virtualized
  - Average consolidation ratios typically range from 8:1 to 15:1
  - Average Workload/CPU = 4:1

BEFORE

Before		
Type	**Quantity**	**Price**
1 CPU	300	$4,000
2 CPU	500	$6,500
4 CPU	200	$14,000
8 CPU	0	$30,000

AFTER

After		
Type	**Quantity**	**Price**
1 CPU	0	$4,000
2 CPU	38	$10,000
4 CPU	38	$23,000
8 CPU	4	$45,000

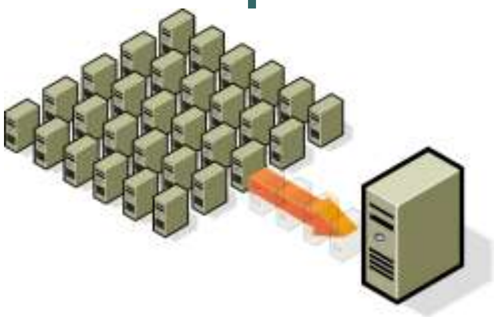

✓ $5,816 saved per workload

Server Hardware Savings				
Year 0	**Year 1**	**Year 2**	**Year 3**	**Total**
$ (1,434,000)	$ 2,416,667	$ 2,416,667	$ 2,416,667	$ 5,816,000

Server Consolidation: SAN Environment

BEFORE

- 1000 servers
- 500 HBAs
- 22 SAN Switches (24 port)
- 5 TB Storage

AFTER

- 80 servers
- 160 HBAs
- 8 SAN Switches (24 port)
- 25 TB Storage

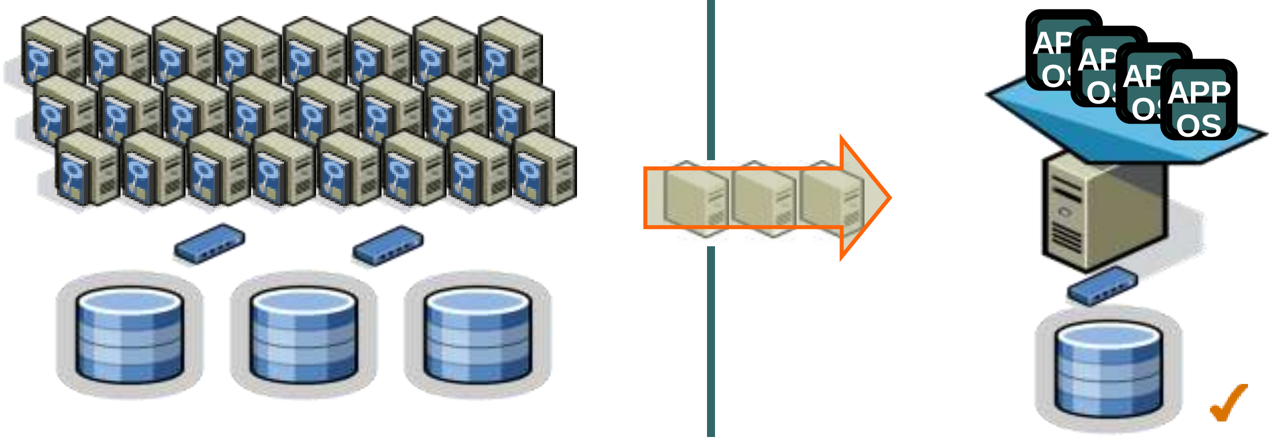

✓ $290 saved per workload

	SAN Environment				
	Year 0	Year 1	Year 2	Year 3	Total
HBA	$ -	$ 113,333	$ 113,333	$ 113,333	$ 340,000
SAN Switches	$ -	$ 23,333	$ 23,333	$ 23,333	$ 70,000
Storage	$ -	$ (40,000)	$ (40,000)	$ (40,000)	$ (120,000)
Total	$ -	$ 96,667	$ 96,667	$ 96,667	$ 290,000

Server Consolidation: Network

BEFORE

- 1000 servers
- 2 NICs per server
- 84 Network Switches

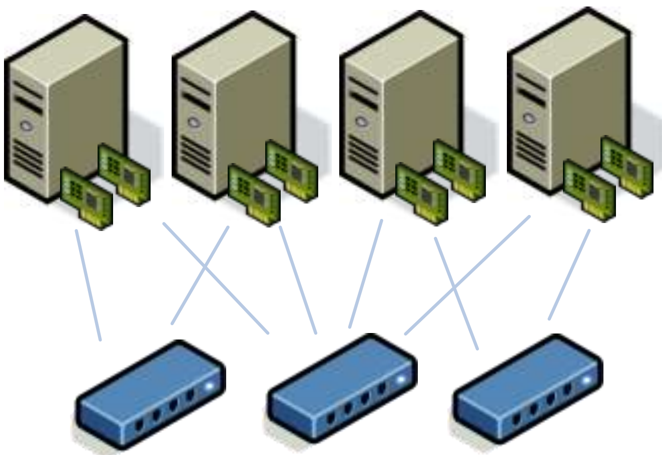

AFTER

- 80 servers
- 3 NICs per server
- 10 Network Switches

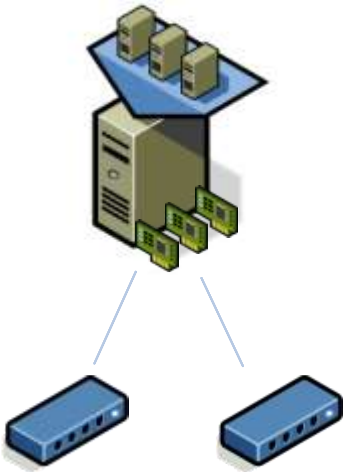

✓ $296 saved per workload

Network				
Year 0	Year 1	Year 2	Year 3	Total
$ -	$ 98,667	$ 98,667	$ 98,667	$ 296,000

vmware

Server Consolidation: Power

- Power consumption calculated from steady-state usage of computing infrastructure, but there are other savings:
  - Other IT Loads: Network switches, SAN components, etc.
  - Non-IT Loads: transformers, uninterruptible power supplies (UPS), power wiring, fans, air conditioners, pumps, humidifiers, lighting
- CPU correlates to U according to the following: 1 CPU = 1U, 2 CPU = 2U, 4 CPU = 4U, 8 CPU = 6U

BEFORE

Before		
Type	**Quantity**	**Power Rating**
1 CPU	300	475W (0.475 kW)
2 CPU	500	550W (0.550 kW)
4 CPU	200	950W (0.950 kW)
8 CPU	0	1600W (1.6 kW)

AFTER

After		
Type	**Quantity**	**Power Rating**
1 CPU	0	550W (0.550 kW)
2 CPU	38	675W (0.675 kW)
4 CPU	38	1150W (1.15 kW)
8 CPU	4	1900W (1.9 kW)

✓ $759 and 355W saved per workload

Power (Computing Infrastructure) Savings				
Year 0	**Year 1**	**Year 2**	**Year 3**	**Total**
$ -	$ 253,160	$ 253,160	$ 253,160	$ 759,481

Density and Thermal Management

- Power consumed by the data center must be evacuated in the form of heat
- Heat dissipation is very difficult, and even more expensive than power consumption
- Data center design is important
  - Airflow inefficiencies are the number one cause for x86 downtime
  - Disk drive failure is #2, and is heavily correlated with airflow inefficiencies
  - Hot-aisle/cold-aisle design (depicted below) more efficient than front-to-back airflow, where heat emission from back of one server is the intake of the front of another server

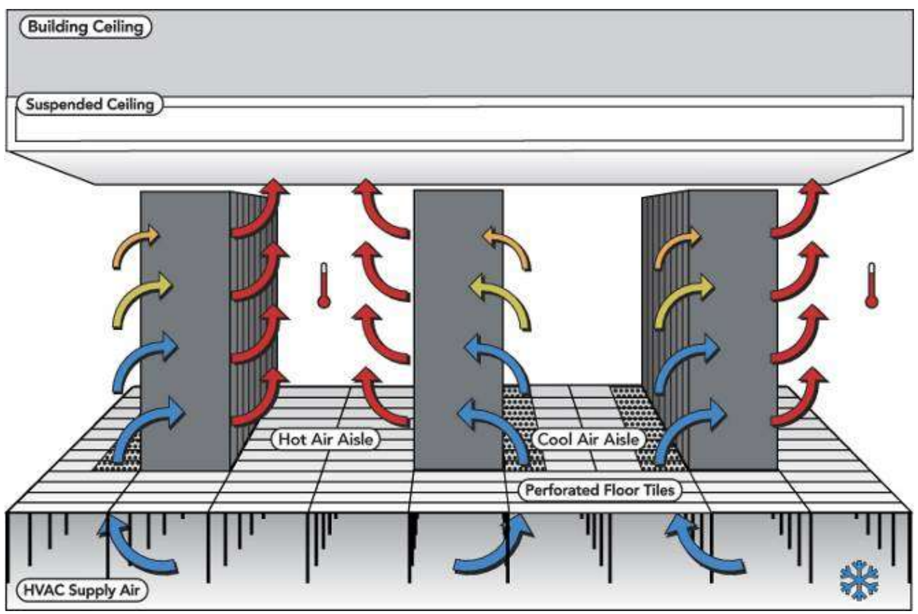

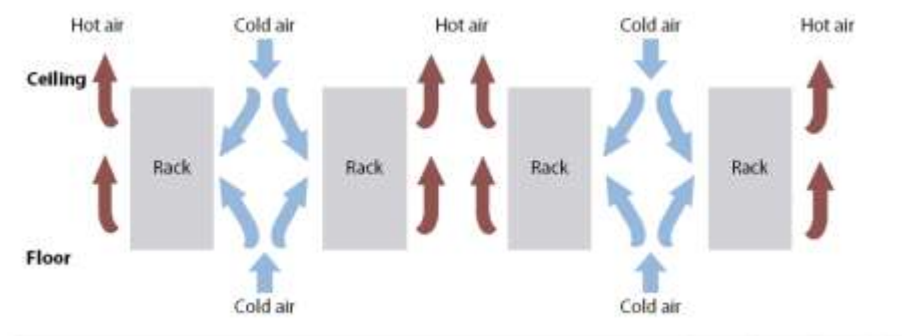

Source: Forrester Research, Inc.

vmware

Density and Thermal Management

- Even with proper design, airflow redundancy is required to achieve proper cooling
  - In the event of CRAC (Computer Room Air Conditioning) unit failure, airflow redundancy can satisfy cooling requirements.
  - Many data centers have areas where heat density is greater than other areas. Focused redundancy satisfies localized data center cooling requirements (see below)
- Additional redundancy is required for humidity
  - Humidification required to reduce potential for static discharge
  - Most air conditioning systems induce humidity loss
  - Supplemental humidification is required, further burdening CRAC units

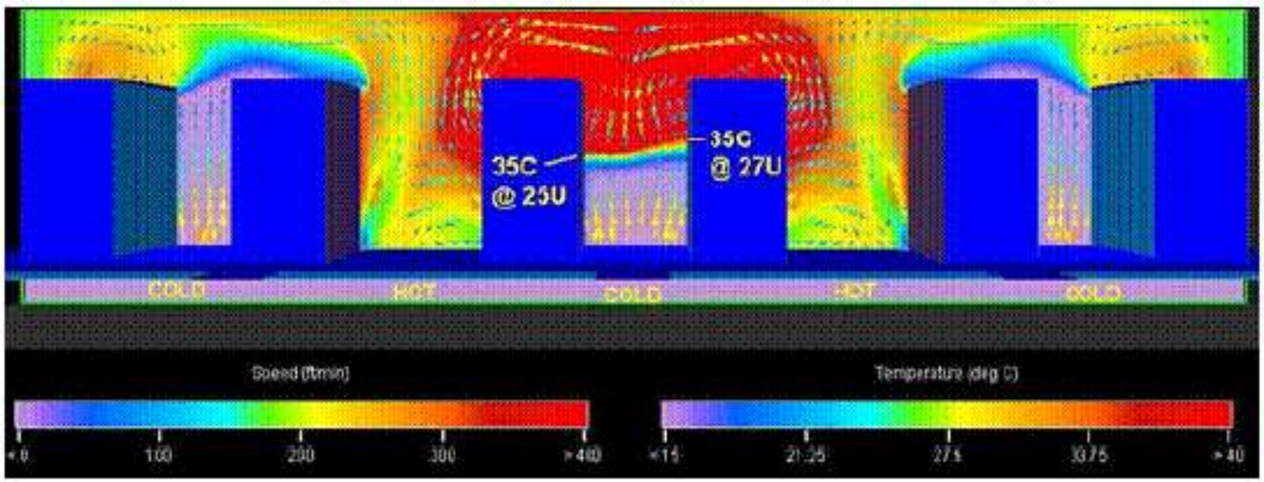

Server Consolidation: Cooling

- Cooling requirements driven by airflow characteristics
  - AC power consumed in the data center is completely converted to heat. Therefore, 1 kW of power consumption = 1kW of heat produced. From power calculations:
    - Total power consumed before VMware = 508.8 kW
    - Total power consumed after VMware = 64.4 kW
  - Other important assumptions
    - 0.8 W of power consumed by cooling equipment for 1W of heat dissipation (empirically determined in HP Labs)
    - 25% redundant airflow required (from previous slides)
    - 25% increase in cooling required to handle inefficiencies such as humidity (from previous slides)

Cooling Savings				
Year 0	**Year 1**	**Year 2**	**Year 3**	**Total**
$ -	$ 316,451	$ 316,451	$ 316,451	$ 949,352

✓ $949 and 444W saved per workload

Server Consolidation: Data Center Real Estate

- Data Center Real Estate depends on several parameters
  - Assume 60kW per sq ft data center (average for most data centers)
  - Cost per square foot to build data center: $600 (conservative estimate)
  - Percent of space used by rack: 30% (Data centers not fully loaded largely because of cooling issues)
  - Units per rack: 24U
    - where U is a rack unit, which is the amount of space taken up by a piece of data center equipment in a rack. 1U is typically 1.75 inches high)
  - Data Center Life: 10 years (assume 6% annual percentage rate)

BEFORE

Before	
Required Space	2,053 sq ft
Total Cost of Data Center	$ 1,232,000
Monthly Payment	$ 13,678

AFTER

After	
Required Space	257 sq ft
Total Cost of Data Center	$ 154,000
Monthly Payment	$ 1,710

✓ $430 saved per workload

Data Center Real Estate Savings				
Year 0	Year 1	Year 2	Year 3	Total
$ -	$ 143,616	$ 143,616	$ 143,616	$ 430,848

vmware

Server Consolidation

VMware Infrastructure not only reduces servers in the data center…

…it reduces racks, too.

Server Consolidation

- Consolidation Results
  - Servers: 1000 → 80
  - HBAs: 500 → 160
  - SAN Switches: 22 → 8
  - Network Switches: 84 → 10
  - Computing Power (kW): 407 → 52
  - Cooling Power (kW): 509 → 64
  - Real Estate (sq ft): 2053 → 257

- Per Workload Savings
  - Servers Hardware: $5,816
  - SAN Components: $290
  - Network: $296
  - Computing Power: $759
  - Cooling Power: $949
  - Real Estate: $431

  - TOTAL: $8541 per workload

vmware

Infrastructure Provisioning

BEFORE

- From 20 hours to build a server and re-load application…
  - Build and configure hardware
  - Load operating system
  - Load configuration tools (Backup, Resource Kit, Monitoring, etc…)
  - Assign 2 IP addresses
  - Build 3 network connections, copper or fiber
  - Turn over to applications team to re-load and re-configure software
  - Test applications
  - Coordinate outage / data migration

AFTER

- To 1 hour to copy a virtual machine and restart
  - Redirect virtual disk to new VMware virtual machine instance
    [Tools already loaded]
    [Application already loaded, configured]
  - Done

vmware

Infrastructure Provisioning

Recall: 1000 servers before, 80 after, 3 year refresh rate

- Over a 3 year period, 333 servers replaced per year before virtualization, 27 servers replaced per year after virtualization (ignoring new workloads)
- 2 hours required to perform ESX installation across 80 servers = 160 hours
- $60/hour administrator rate

BEFORE

- 333 servers replaced per year
- 20 hrs per workload: 20*333 = 6660 hours

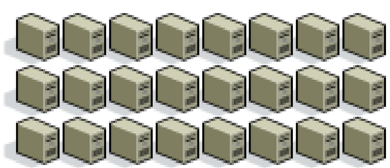

AFTER

- 27 servers replaced per year
- 1.5 hrs per workload (conservative): 1.5*333 = 500 hours

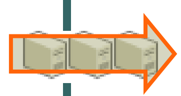

✓ 333 servers replaced per year = ~ 6,000 labor hours saved

Infrastructure Provisioning Savings				
Year 0	Year 1	Year 2	Year 3	Total
$ -	$ 370,326	$ 370,326	$ 370,326	$ 1,110,979

Business Continuity: The Cost of Downtime

Millions of dollars per hour in lost revenue by industry

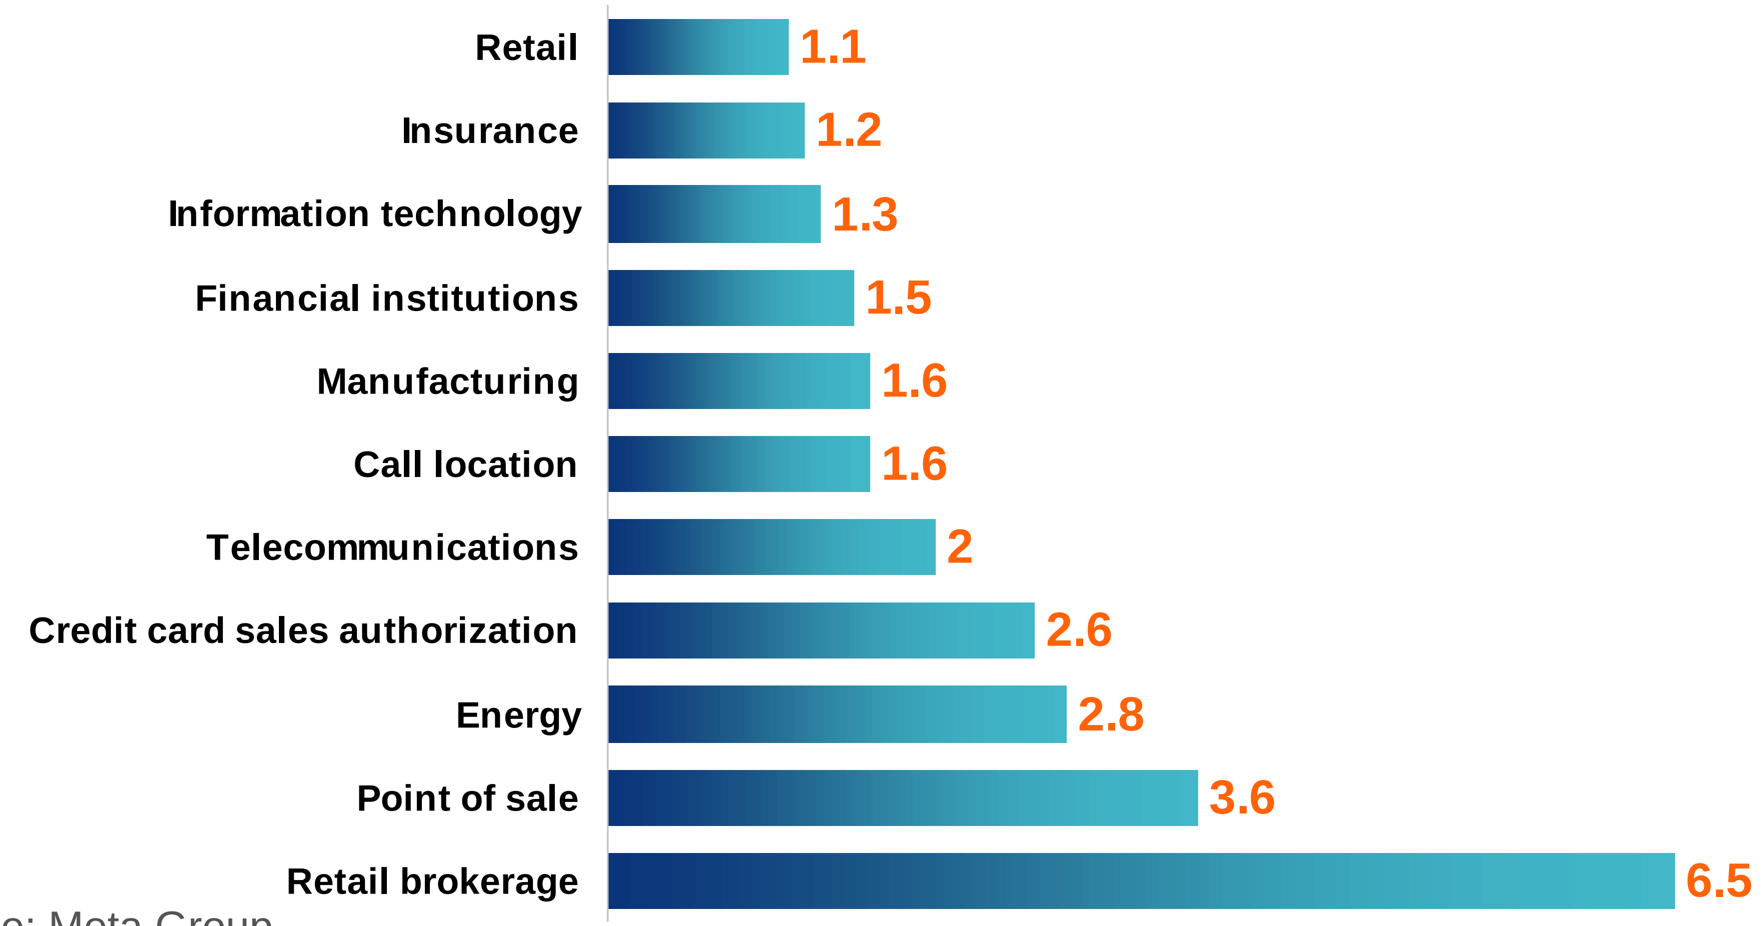

Source: Meta Group

Business Continuity: Unplanned Downtime

- Unplanned downtime can be very costly, but VMware Infrastructure mitigates these costs:
  - VMware can conservatively reduce downtime by 75% (Sample of customers)
    - Virtual Machines can be proactively moved from overloaded hosts
    - Virtual Machines can be quickly restored on other servers if a host experiences hardware failure
  - Yankee Group study reports approximately 15 hours of x86 downtime per year (not entirely dedicated to hardware, but still can be used as first-order approximation)
  - Assume $20,000 cost of downtime (Emerson Study[1])

BEFORE

$$Costs_{UnplannedDowntime(Before)} = (\$20,000)(15) = \$300,000$$

AFTER

$$Costs_{UnplannedDowntime(After)} = (\$20,000)(3.75) = \$75,000$$

Unplanned Downtime Savings				
Year 0	Year 1	Year 2	Year 3	Total
$ -	$ 225,000	$ 225,000	$ 225,000	$ 675,000

[1] http://www.continuitycentral.com/news02575.htm

Business Continuity: Recovery

- In event of failure, VMware can offer not only savings but survival based on quick and easy recovery
  - Complex server configuration unnecessary with VMware Infrastructure due to virtual machine hardware independence
  - Multi-step restorations are simplified to single-step file recovery
  - 94% of companies that experience a disaster do not experience long-term survival (Cummings, Haag & McCubbrey 2005)

BEFORE

- 40 hours for full recovery

$$Cost_{Recovery(Before)} = (\$20,000)(40) = \$800,000$$

AFTER

- 10 hours for full recovery

$$Cost_{Recovery(Before)} = (\$20,000)(10) = \$200,000$$

✓ $600,000 saved if disaster occurs

Summary

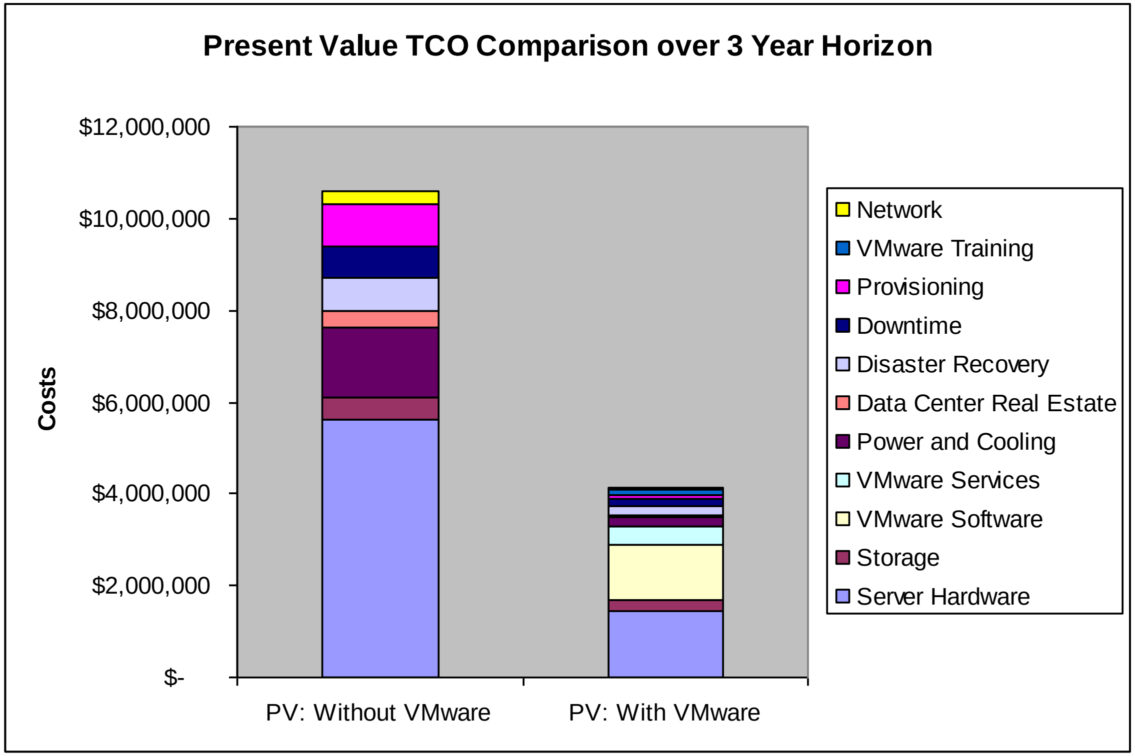

Excluding Soft Costs

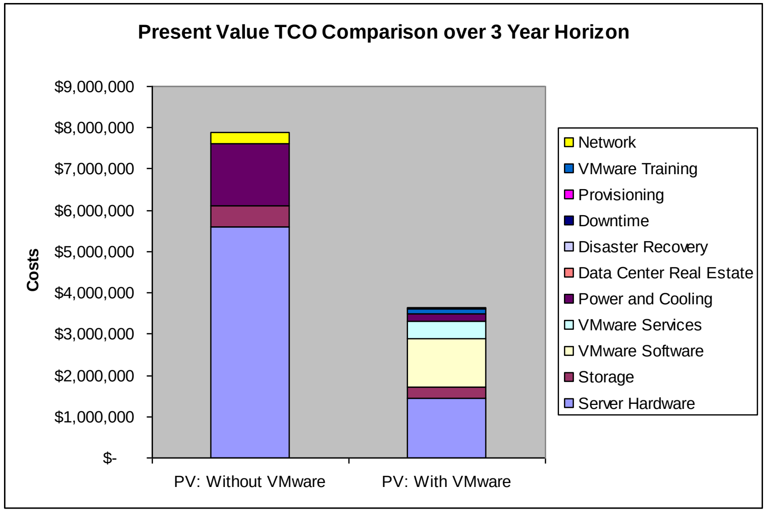

✓ 3 Year TCO Reduction (NPV): $6,461,244