Optimal Innovations

Hyperformix

LSI

Green Capacity Planning: Theory and Practice

Amy Spellmann, Optimal Innovations
Richard Gimarc, HyPerformix
Charles Gimarc, LSI

St. Louis CMG
May 18, 2010

Agenda

- Green Capacity Planning
  - Methodology
  - Energy Footprint Projection
- Case Study
  - Server Upgrade Scenario
  - Virtualized Scenario
  - Storage Upgrade Options
  - Comparison
- Conclusion

STLCMG

Green Capacity Planning: What is it?

- Extends established methods & best practices
- Introduces an ***energy footprint projection***
- Views energy as a new resource for capacity planners to include in their analysis
- Focuses on cost-effective technologies & practices to manage today's challenges

STLCMG

Green Capacity Planning Combines Tools & Best Practices

Predictive Capacity Planning	***Energy Footprint Projection***	***Cost & Capacity Comparison***

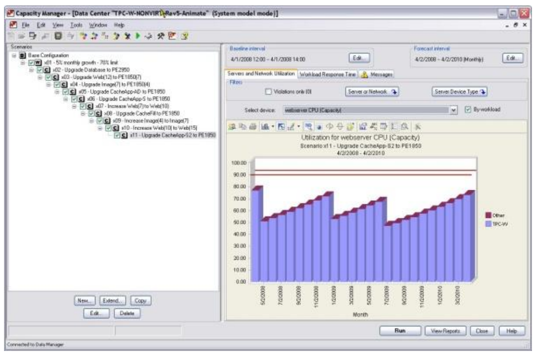

+

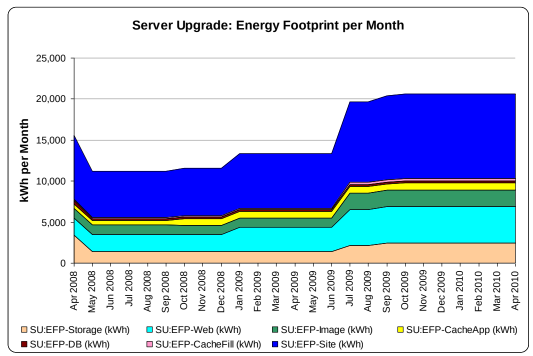

=

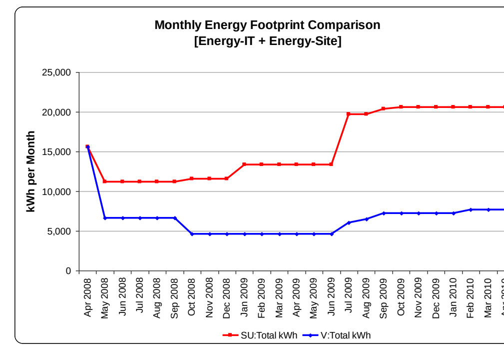

Commercial Capacity Planning Tools	Professional Services	Green Capacity Planning
Established Methods	New Resource to Evaluate	

Green Capacity Planning: Why do we need it?

- By 2010, IDC expects that for every dollar of new server spending, an additional $0.70 will be needed for power and cooling
- 2% of the world's carbon dioxide emissions come from IT equipment--the same amount of pollution as the airline industry creates – Wall Street Journal
- With the advent of high density computer equipment such as blade servers, many data centers have maxed out their power and cooling capacity - Gartner
- New data center construction projects cost $250M and up
- Google spent $600M each for four new data centers in 2007

Green Capacity Planning: Value

- Determines the most efficient way to meet business needs while minimizing computing, power, cooling, space requirements
- Assures SLAs are met
  - Response times
  - Utilization
- Evaluates potential cost & energy savings of virtualization

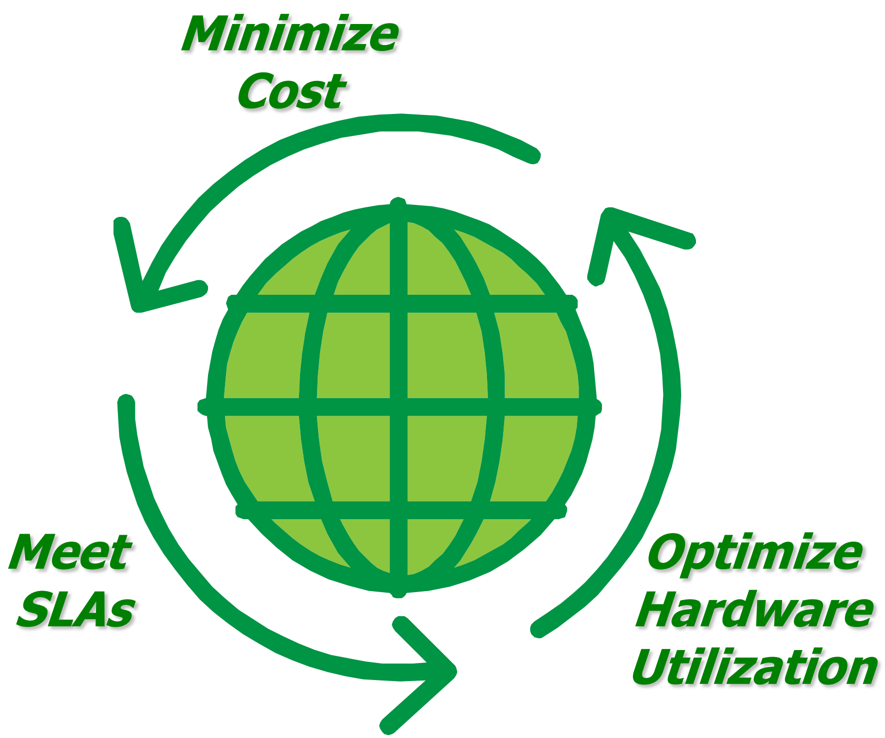

Green Capacity Planning: Factors

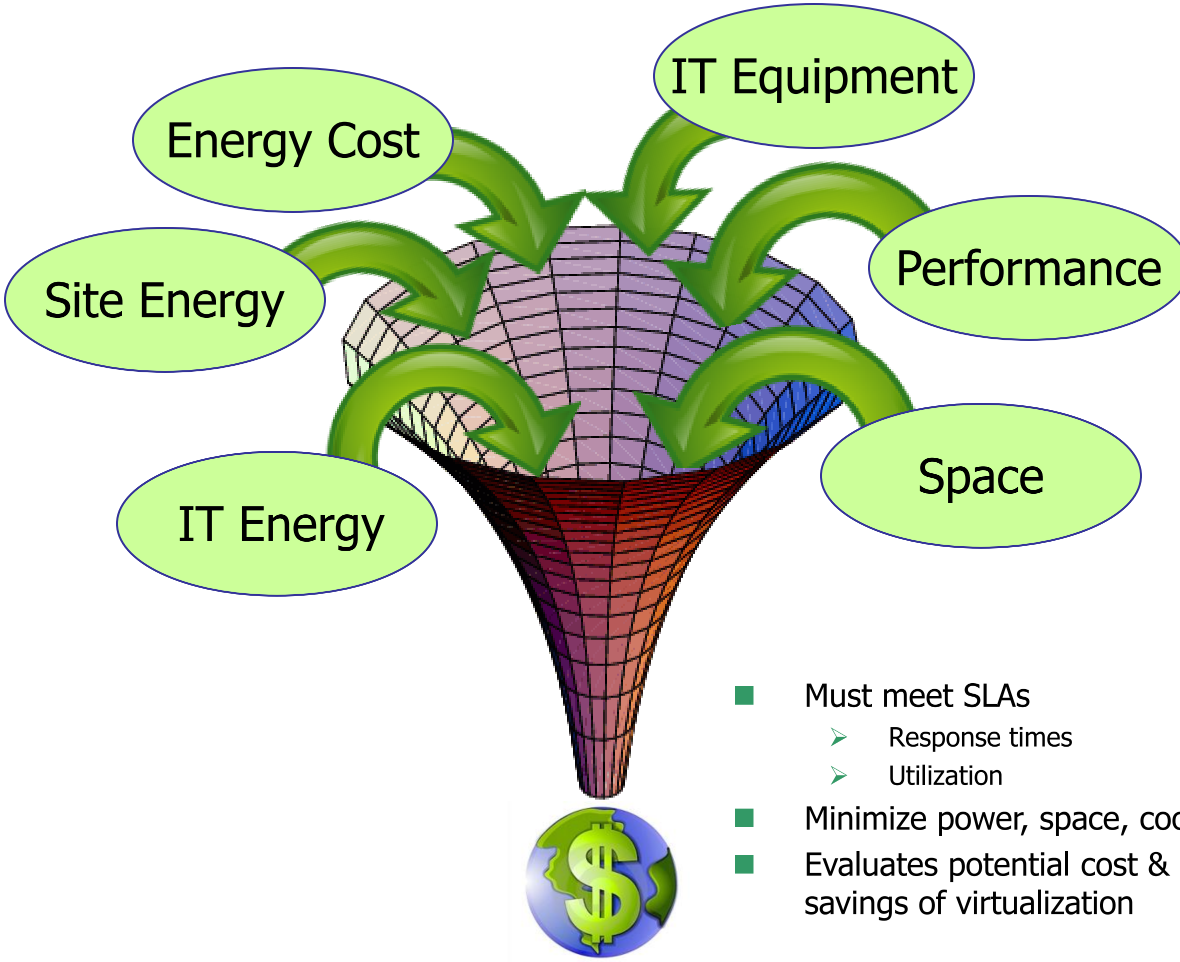

- Must meet SLAs
  - Response times
  - Utilization
- Minimize power, space, cooling
- Evaluates potential cost & energy savings of virtualization

Cost Effective Solution

Methodology: Enhanced Traditional Approach

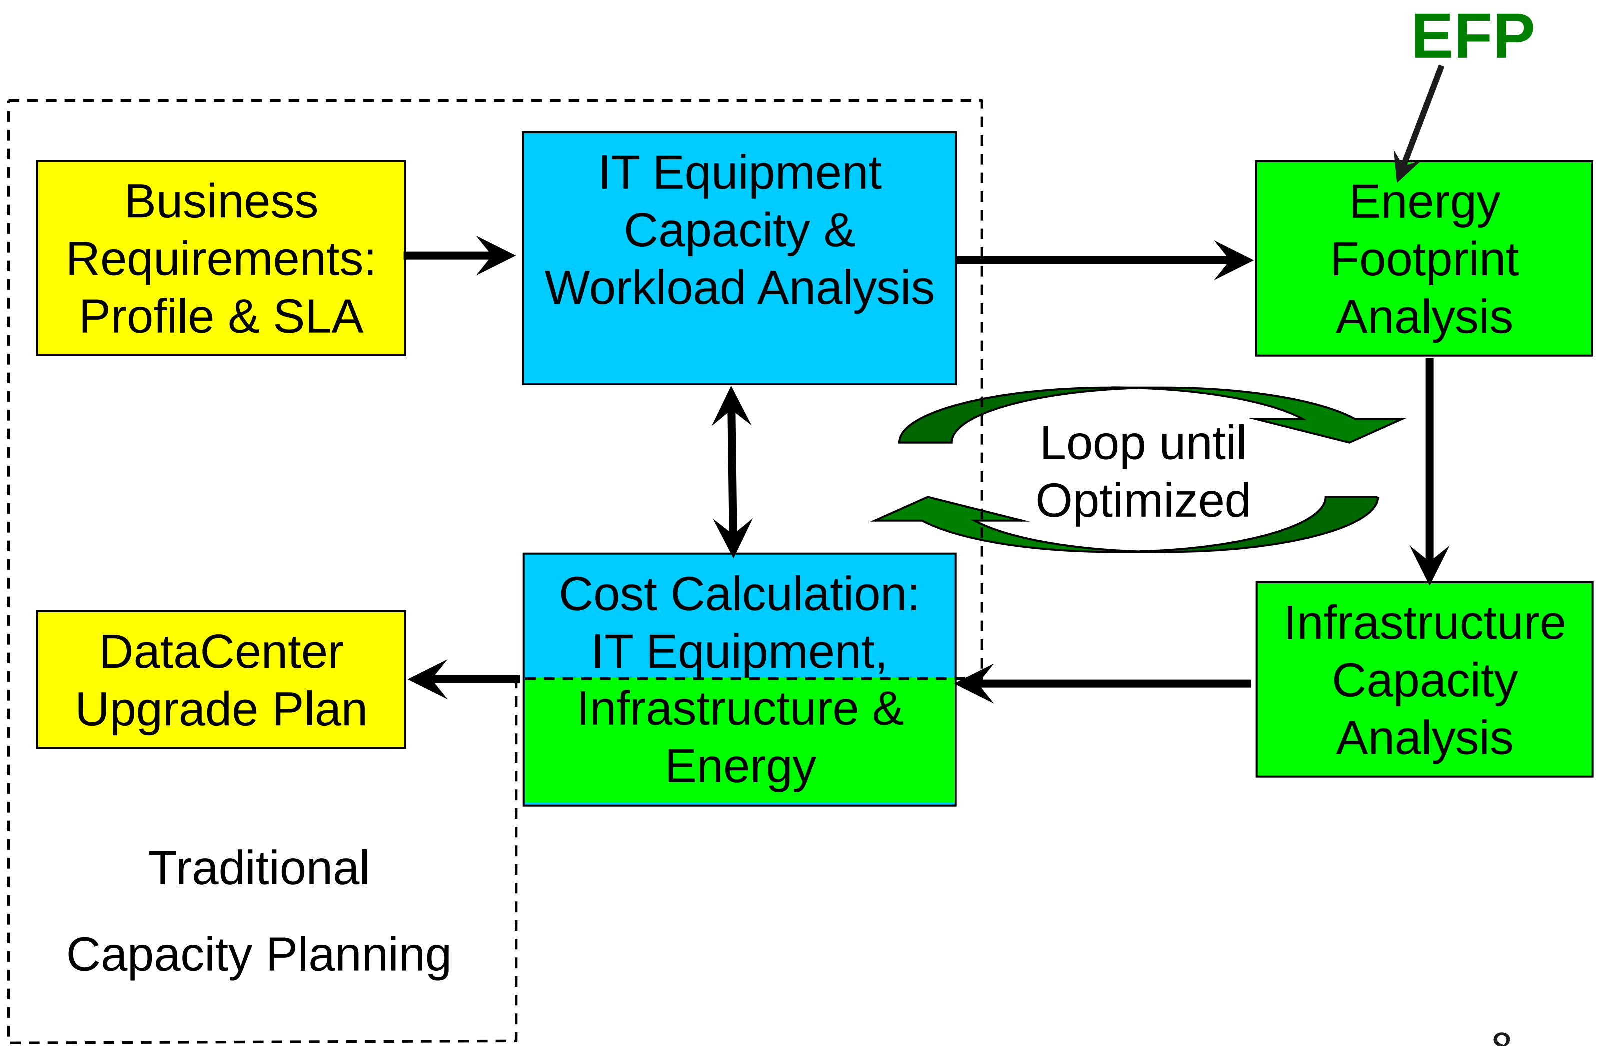

Green Capacity Planning Metrics

- New Inputs
  - Energy Footprint (calculated or measured)
  - PUE
  - Active Idle
  - Max Power
  - Cost of electricity
- Comparison Metrics
  - Energy Footprint Projection
  - Rack Space
  - Cumulative Costs
  - Productivity (Trans per kWh)

In addition to traditional capacity planning metrics, e.g., Utilization, Response Time, etc.

STLCMG

Energy Footprint Projection

- Determine energy required by each system component
  - Obtain active-idle, average, maximum power from measurements or vendor specifications
  - Compute energy usage over time (e.g., monthly)
- Add "site" infrastructure energy
  - Overhead energy for running components in the data center
  - Includes: cooling, fans, power distribution units, lighting, etc.

STLCMG

Energy Footprint Projection: Terminology

- Terminology
  - Power
    - Spot measurement, single point in time
    - Measured in watts (W)
  - Energy
    - Power consumption over a period of time
    - Measured in kilowatt hours (kWh)
    - You pay for energy, not power

Power ≠ Energy

- Energy estimation
  - Component-level power requirements (W)
  - Estimate consumption over a period of time (e.g., 1 day)

Power x Time = Energy Usage

Energy Footprint Projection: Data Center Energy Usage

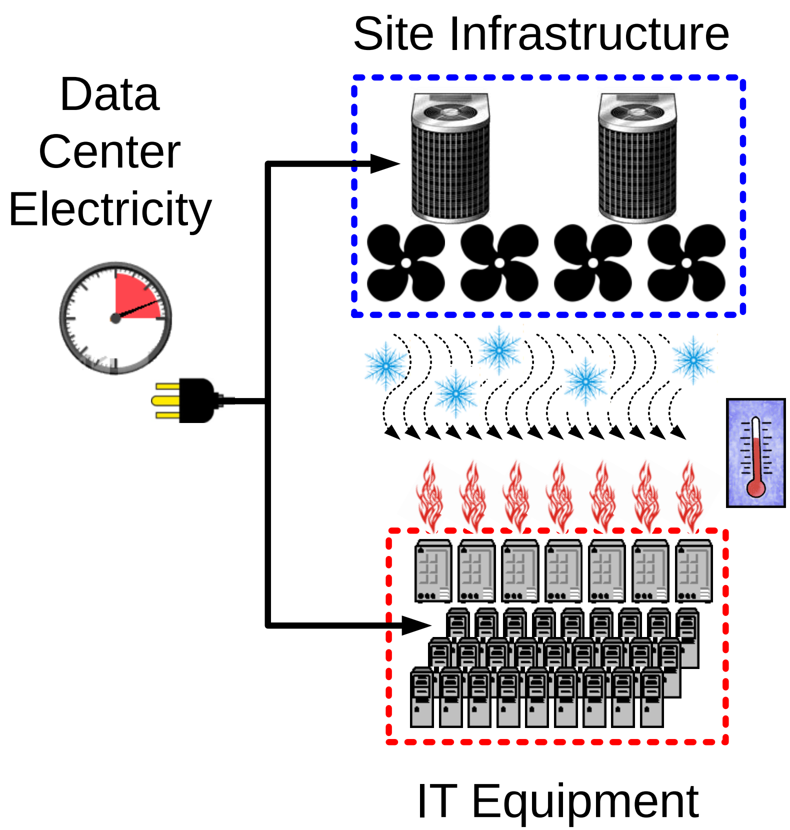

General Flow

- Energy enters in the form of electricity
- IT equipment used to store & manipulate information
- IT equipment generates heat
- Site infrastructure removes generated heat

Questions

- What if your data center's power input is maximized?
- How can you increase your compute power without increasing your power draw?
- Can you improve site efficiency?
- Do you really need to build a new data center?

Energy Footprint Projection: Components

- Two components
  - IT Equipment
  - Site Infrastructure

- Approach
  - *Compute* IT energy
  - *Estimate* Site energy

- Power Usage Effectiveness

$$PUE = \frac{\text{Site Power} + \text{IT Power}}{\text{IT Power}} \geq 1.0$$

  - Today's average of 2.0
  - Use to estimate Site energy
  - Energy = Power X Time

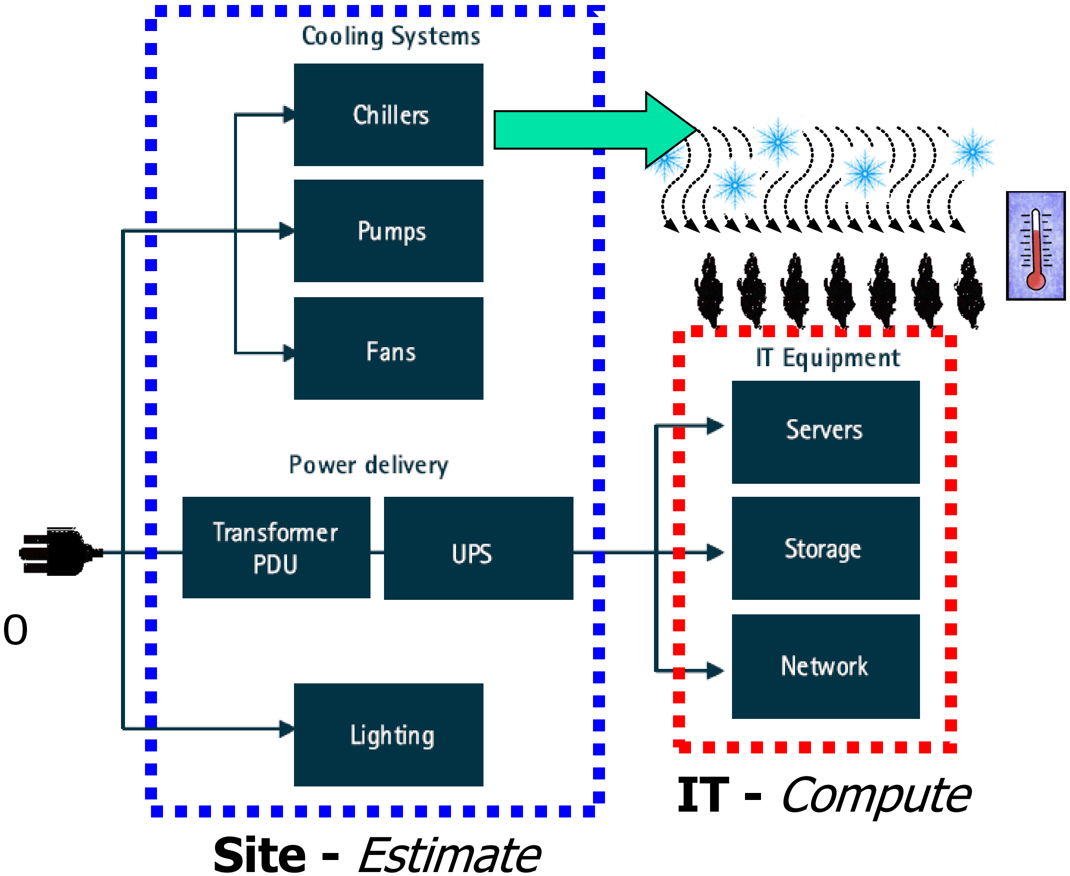

EFP(Day) = $IT\_kWh_{Computed}$ + $Site\_kWh_{Estimated}$

STLCMG

Energy Footprint Projection: Server-Centric Component Contribution

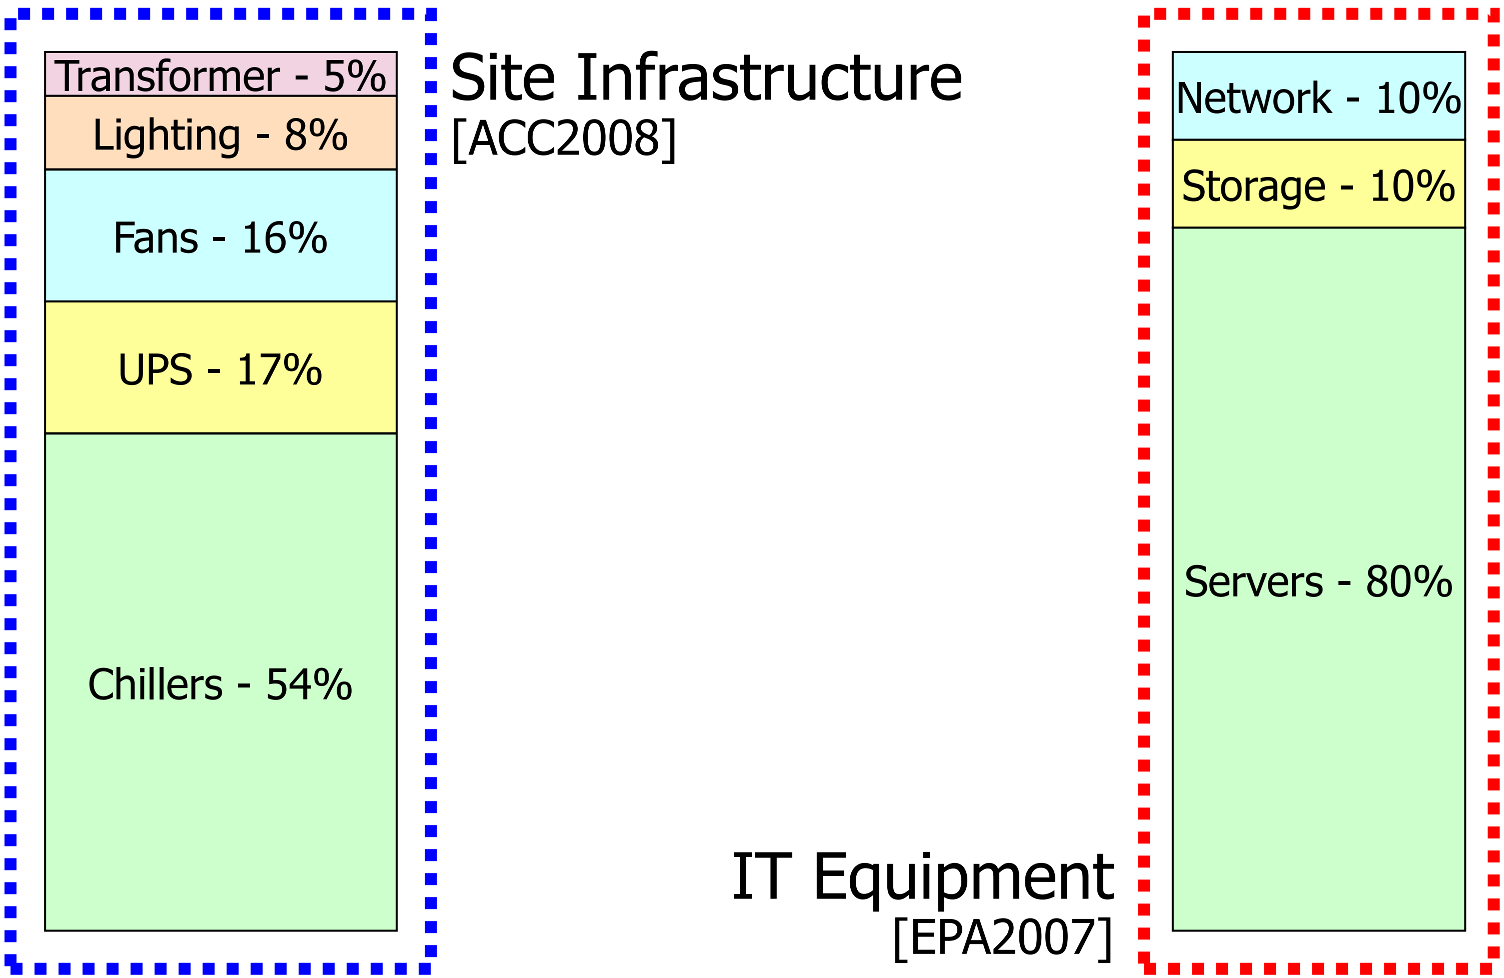

Component Distribution Varies Widely With Business Service

STLCMG

Energy Footprint Projection: Estimating EFP(Day)

$$EFP(Day) = IT\_kWh_{Computed} + Site\_kWh_{Estimated}$$

Given $Power_i$ for each IT device:

$$IT\_kWh_{Computed} = (\sum Power_i) \times (24\ Hours/Day)$$

Power Usage Effectiveness:

$$PUE = \frac{SitePower + ITPower}{ITPower} \geq 1$$

Estimate Site_kWh using assumed PUE:

$$Site\_kWh_{Estimated} = IT\_kWh_{Computed} \times (PUE - 1)$$

$$EFP(Day) = PUE \times IT\_kWh_{Computed}$$

Energy Footprint Projection: Server Power Usage

- How do you estimate the power usage of a server?
- Power specifications for a server (W):
  - Power supply
  - Max Power
  - Active Idle

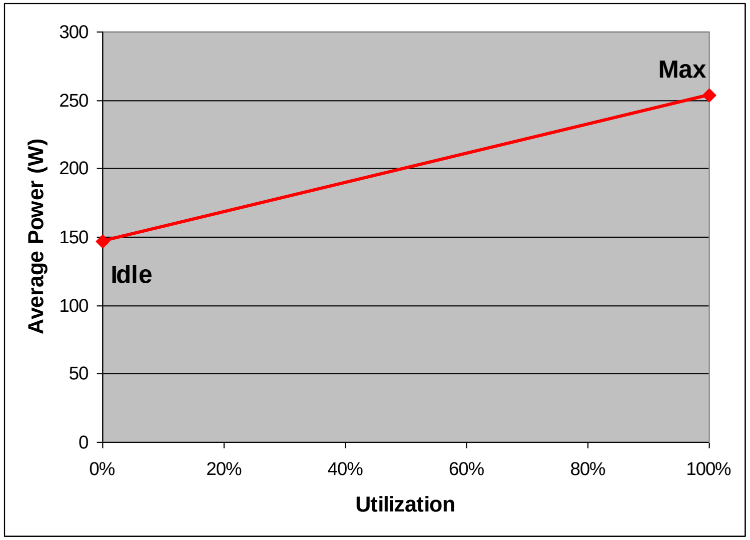

- Approximation
  - Server power usage scales linearly with processor utilization
  - Validated across various servers ± 5% [TGG2008]

$$Power@Util\% = (MaxW - IdleW) \times \left(\frac{Util\%}{100}\right) + IdleW$$

STLCMG

Energy Footprint Projection: Server Power Usage - Example

SPECpower_ssj2008

Dell Inc. PowerEdge 2950 III (Intel Xeon E5440)				SPECpower_ssj2008 = 719 overall ssj_ops/watt	
Test Sponsor:	Dell Inc.	SPEC License #:	55	Hardware Availability:	Mar-2008
Tested By:	Dell Inc.	Test Location:	Round Rock, TX, USA	Software Availability:	Oct-2007
System Source:	Single Supplier	Test Date:	Feb 12, 2008	Publication:	Feb 27, 2008

Benchmark Results Summary

Performance			Power	Performance to Power Ratio
Target Load	Actual Load	ssj_ops	Average Power (W)	
100%	99.2%	298,188	254	1,172
90%	89.7%	269,554	249	1,084
80%	80.4%	241,701	242	997
70%	70.3%	211,354	234	901
60%	60.0%	180,508	225	801
50%	49.5%	148,959	214	695
40%	40.4%	121,477	203	599
30%	30.2%	90,702	190	477
20%	20.1%	60,493	177	342
10%	10.0%	29,990	161	187
Active Idle		0	147	0
			∑ssj_ops / ∑power =	719

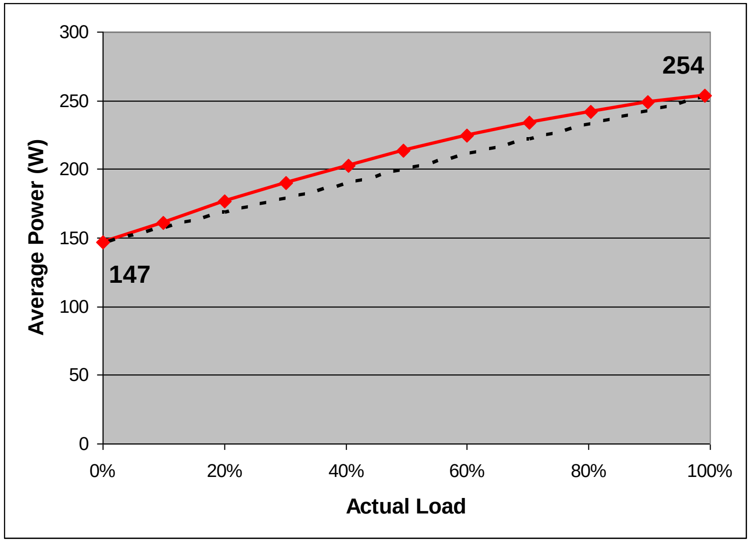

$$Power@Util\% = (MaxW - IdleW) \times \left(\frac{Util\%}{100}\right) + IdleW$$

$$Power@70\% = (254 - 147) \times \left(\frac{70}{100}\right) + 147$$

$$Power@70\% = 221.9$$

Assume server power usage scales linearly with CPU utilization

STLCMG

Demonstrate Green Capacity Planning

- Extend data center lifetime while supporting future business growth
  - Two year planning based on business projections
  - Limited data center power
- Develop a capacity plan for growth/change
  - Upgrade/add servers
  - Leverage virtualization
- Align with business requirements
  - Predict required infrastructure to meet SLAs
  - Reduce today’s energy footprint
  - Select cost-effective solution

STLCMG

Case Study: eCommerce System

- Growing eCommerce system
- Business Requirements
  - 5% monthly growth
  - 2-year upgrade plan to support growth & meet SLAs
  - Reduce IT & site energy costs
  - Optimize hardware utilization
  - Evaluate virtualization as an energy saving solution

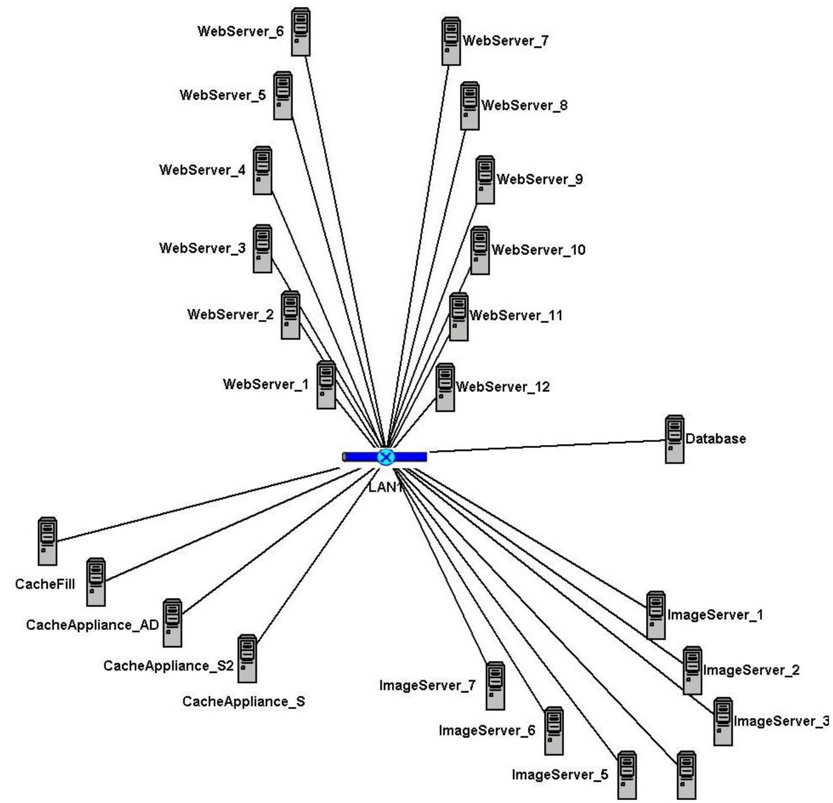

Baseline System
- 24 servers
- 42U Rack space

STLCMG

Case Study: Baseline System

- TPC-W Benchmark
  - Published TPC-W Benchmark Result
    - May 2002
    - Full Disclosure Report - details
  - Simulates Web-based browsing & shopping
  - Provides the baseline specification
    - Application architecture
    - IT equipment
    - Workload
    - Utilization

Benchmark Workload
- 100,000 Items
- 76,000 Users
- 9,709 WIPSb@100,000

Case Study: Baseline Architecture & IT

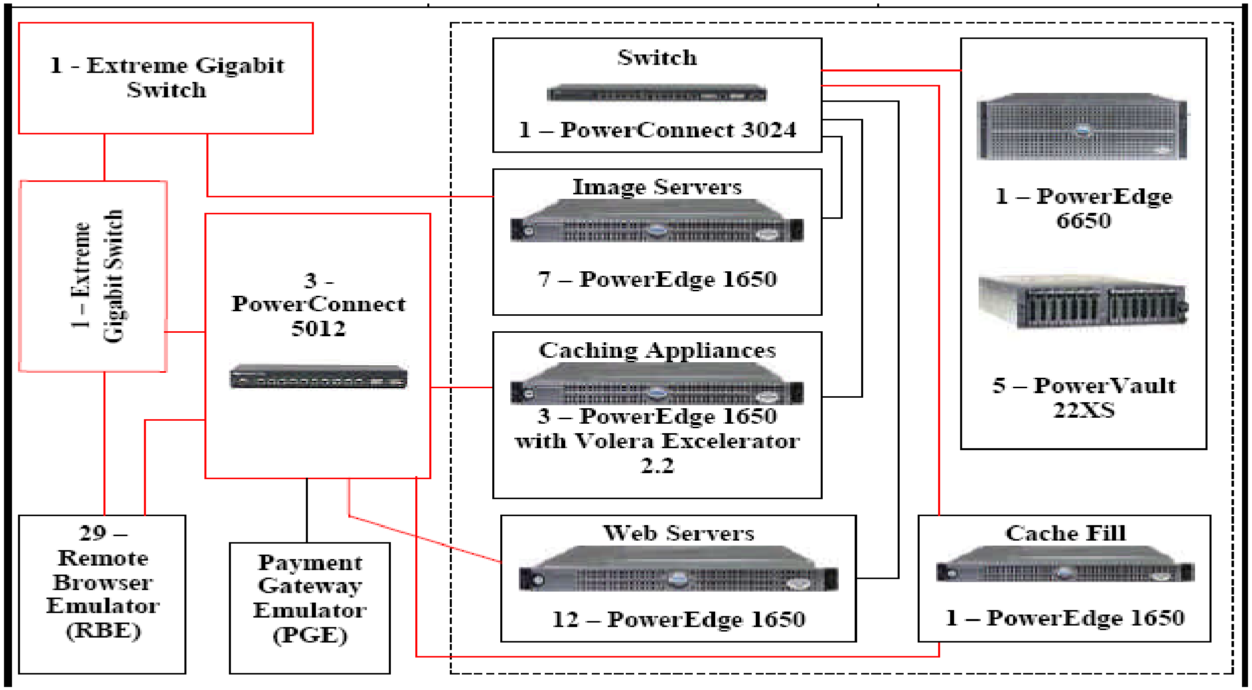

Function	# of Systems	System	OS	Processors	Memory
Web/Image/ Cache Fill Server	18	PowerEdge 1650	Microsoft Windows Powered	2 - Pentium III 1.4GHz/512kB L2	2GB
Image/DNS Server	2	PowerEdge 1650	Windows 2000 Advanced Server	2 - Pentium III 1.4GHz/512kB L2	2GB
Web Cache	3	PowerEdge 1650	Volera Excelerator 2.2	1 – Pentium III 1.4GHz/512kB	2GB
Database Server	1	PowerEdge 6650	Windows 2000 Advanced Server	4 –Intel Xeon MP 1.6GHz/1MB L3	4GB

STLCMG

Case Study: Baseline Database Storage

- 12 HDD per enclosure
- Ultra320 SCSI
- 5 enclosures
- 3U rack space per enclosure
- 41 total disks

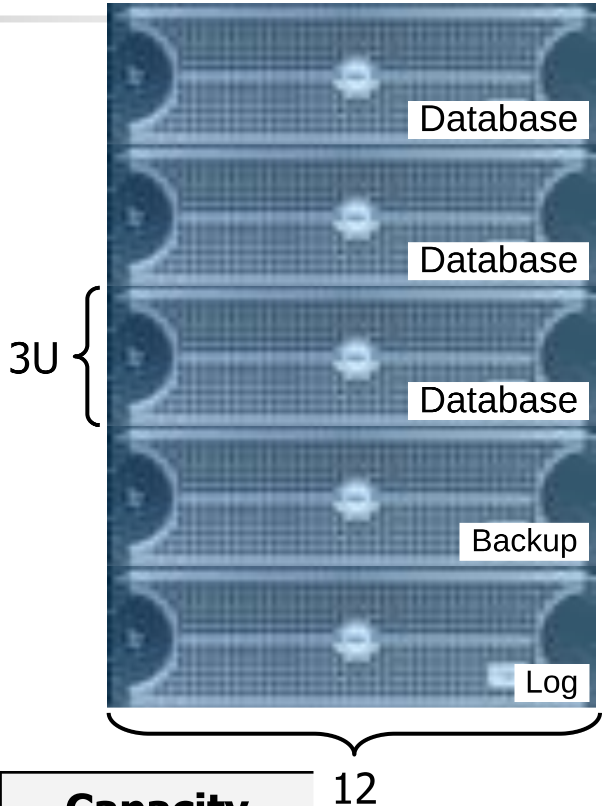

Purpose	Enclosure	# of Disk	RAID Level	Capacity (GB)
Log	1	10	1 / 0	135
Backup	2	7	5	393
Database	3, 4, 5	24	0	812

STLCMG

Case Study: Database Technology Upgrade

- Move to current SAS technology that will allow scaling for 2 years
- Replace original Ultra 320 SCSI storage subsystem with newer SAS (Serial Attached SCSI) technology
- Reduce power by changing from 3½” to 2½” SFF disks
- Triple storage density by moving from 12 disks in a 3U space to 24 disks in a 2U space
- Achieve higher performance to meet the demands of the system after 2 years of workload growth
- Decrease power per enclosure by 50%

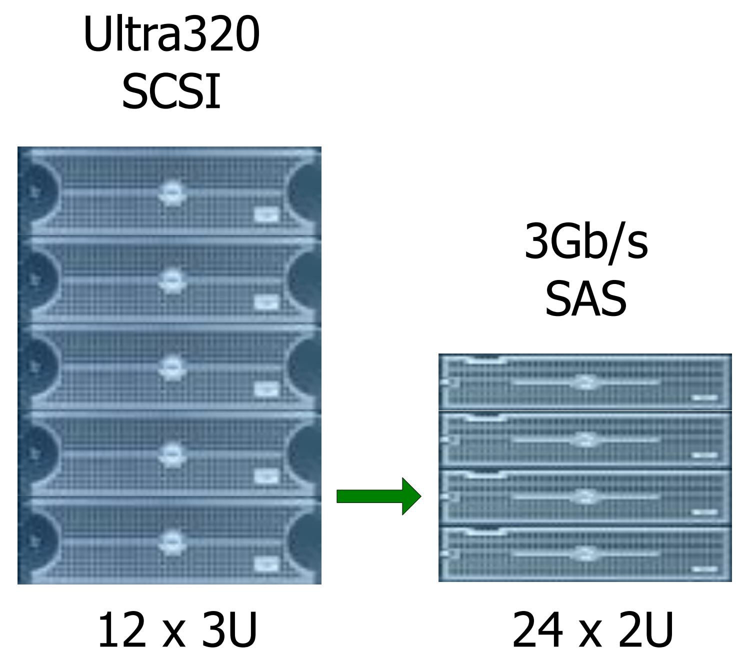

Case Study: Modeling Approach

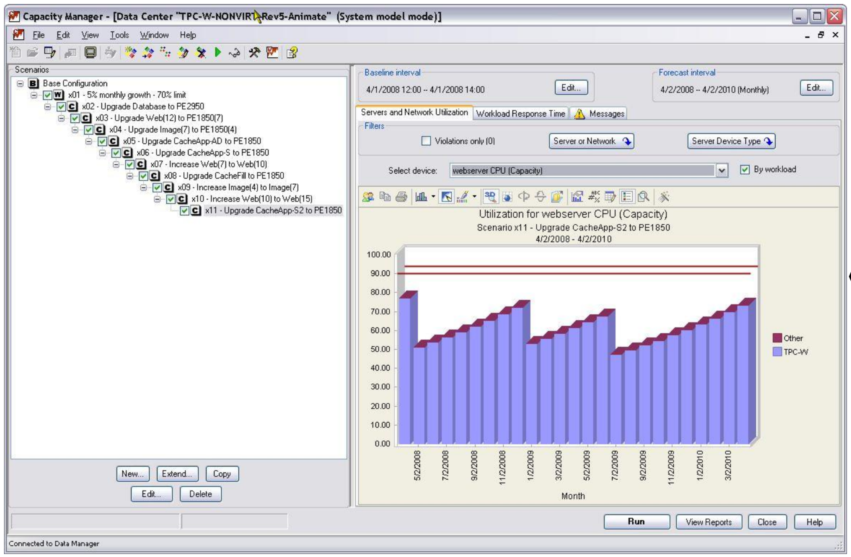

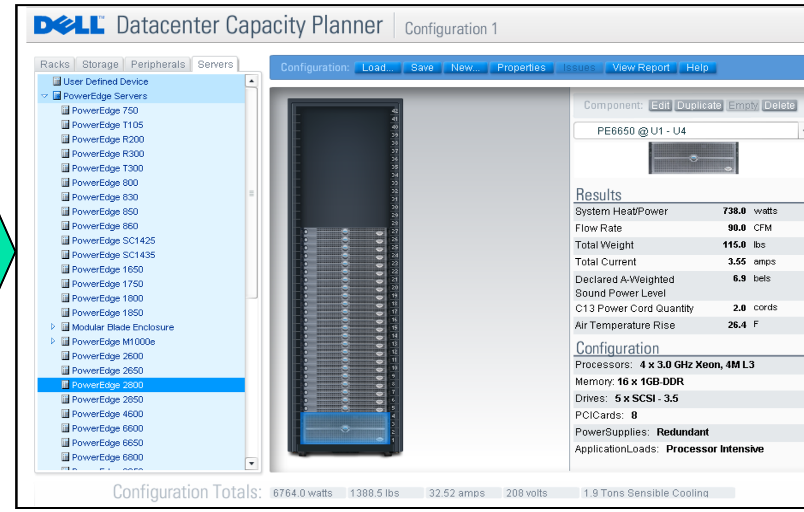

- Utilizing HyPerformix's Capacity Manager for Capacity Analysis
- Dell's Datacenter Capacity Planner for Power Metrics
- Integrating results for Server/Resource Capacity, Power Requirements & Business Metrics

STLCMG

Case Study: Modeling Scenario Strategy

Server Upgrade

- 5% monthly growth for 2 years
- 70% capacity threshold
- Optimize the number of upgrades
- Dell PE 2950 for the database
- *Dell PE 1850 for all other servers*

Virtualization

- 5% monthly growth for 2 years
- 70% capacity threshold
- Optimize the number of upgrades
- Dell PE 2950 for the database
- *Virtualize all other servers on a pool of Dell PE 1950s*
- *Optimize use of virtual hosts*

For each scenario

1. Determine infrastructure required to support workload growth (Capacity Manager)
2. Estimate energy footprint
3. Calculate cost of IT equipment & energy

Scenario - Server Upgrade: Database Storage Provisioning

- Growth Assumption:
  - Database storage grows at same rate as workload (5% per month)
- Baseline Assumptions:
  - Use same capacity disks
    - 36GB for Log & Database
    - 73GB for Backup
  - Same disk fill factor
  - One RAID adapter per enclosure
- Method:
  - Estimate required disk space based on growth
  - Compute number of disks
  - Compute number of enclosures & adapters

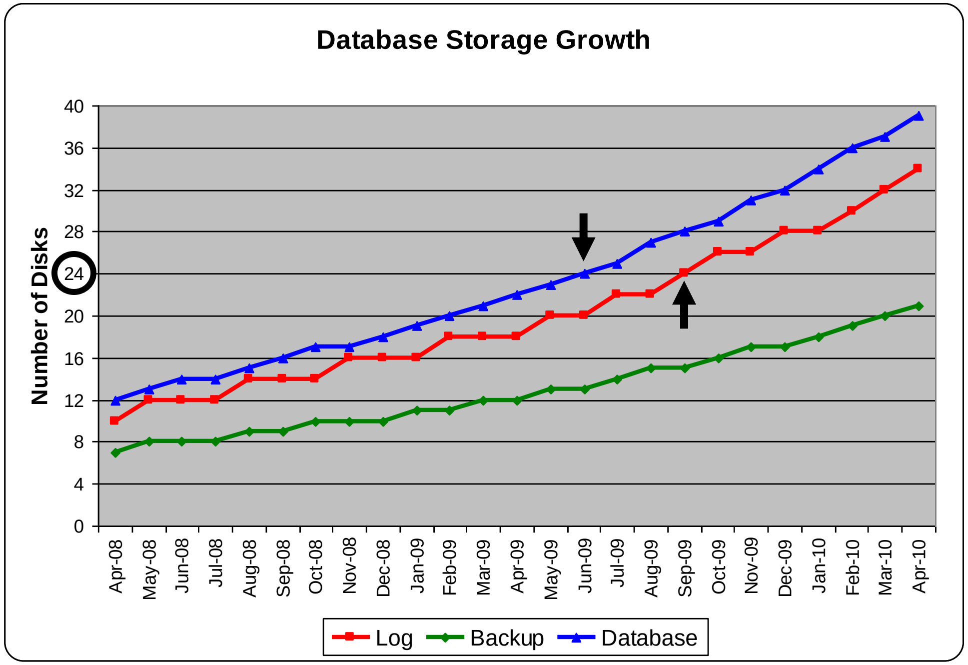

- Results:
  - Start with 4 enclosures
  - Add new enclosure when disk count reaches 24

Scenario - Server Upgrade: Timeline

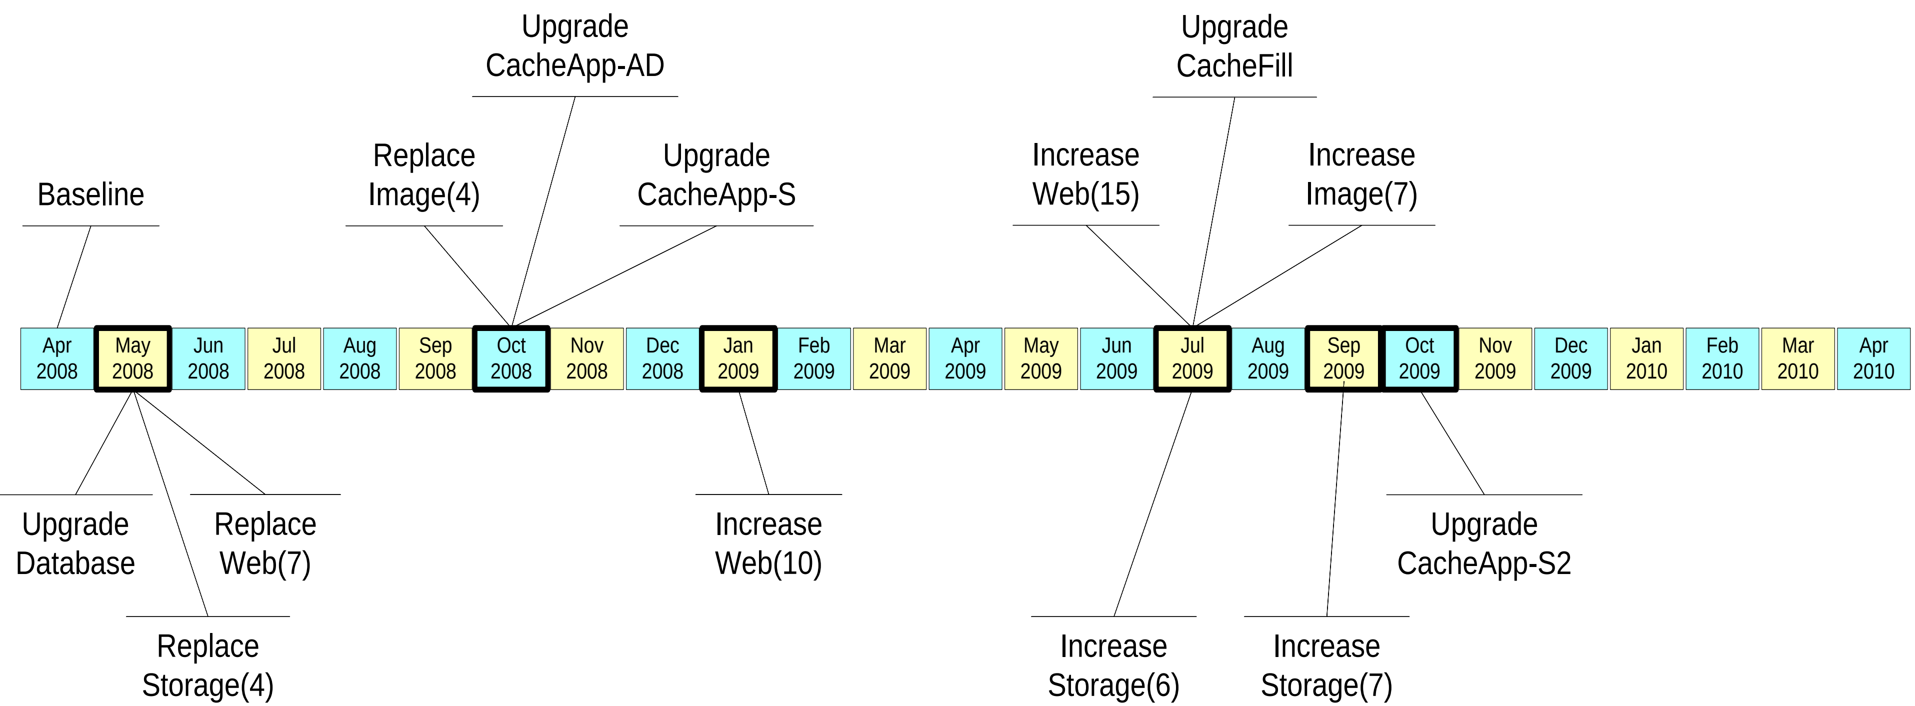

- Minimize number of server upgrade points
- Servers upgraded once per 3-month period
- Retire old servers (replace with faster more efficient servers)

STLCMG

Scenario - Server Upgrade: Infrastructure - Final Plan

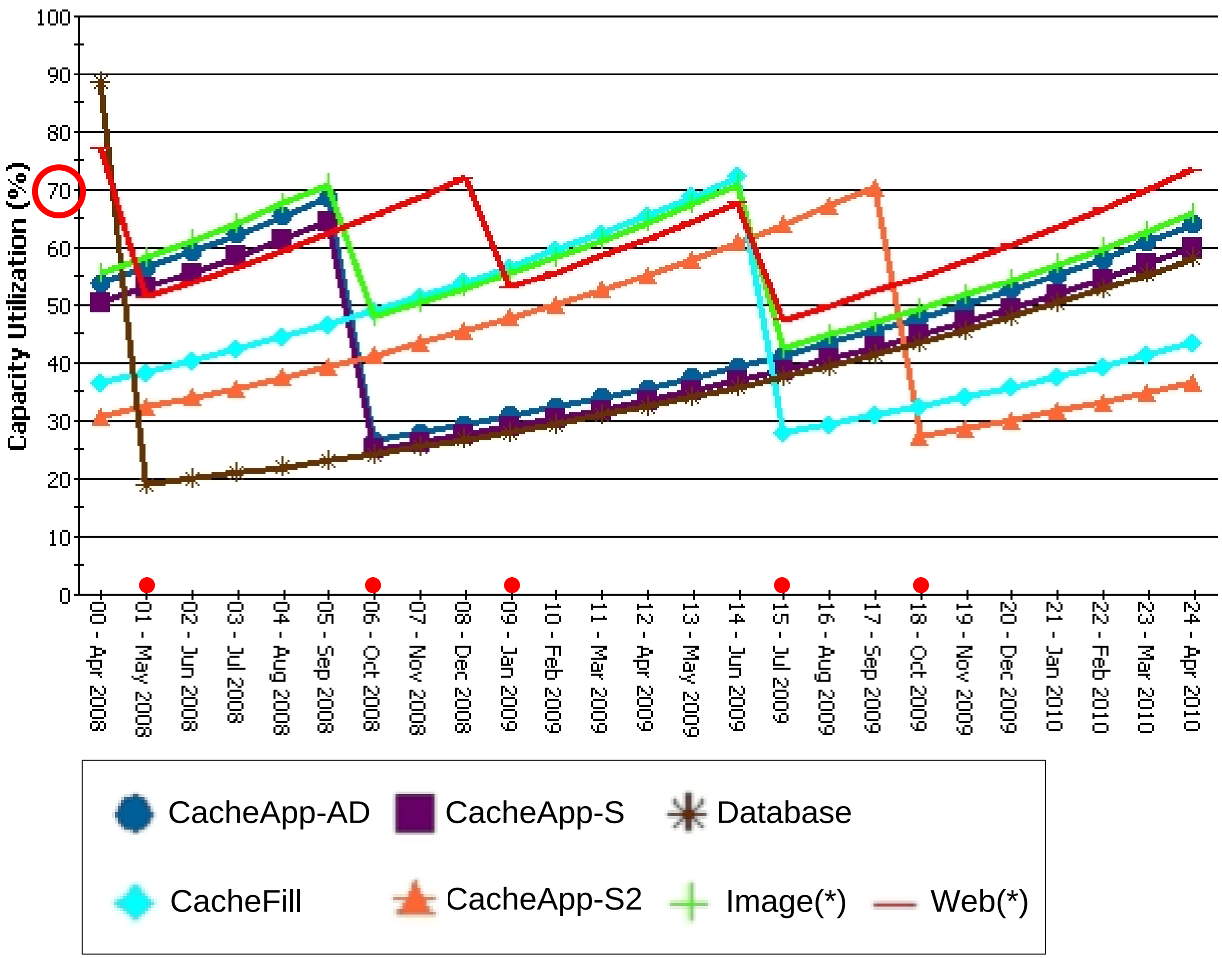

- Final 2-year capacity plan
- Infrastructure supports 5% growth over 24 months
- Server upgrades grouped to reduce downtime (one upgrade per quarter)

STLCMG

Scenario - Server Upgrade: Infrastructure Components

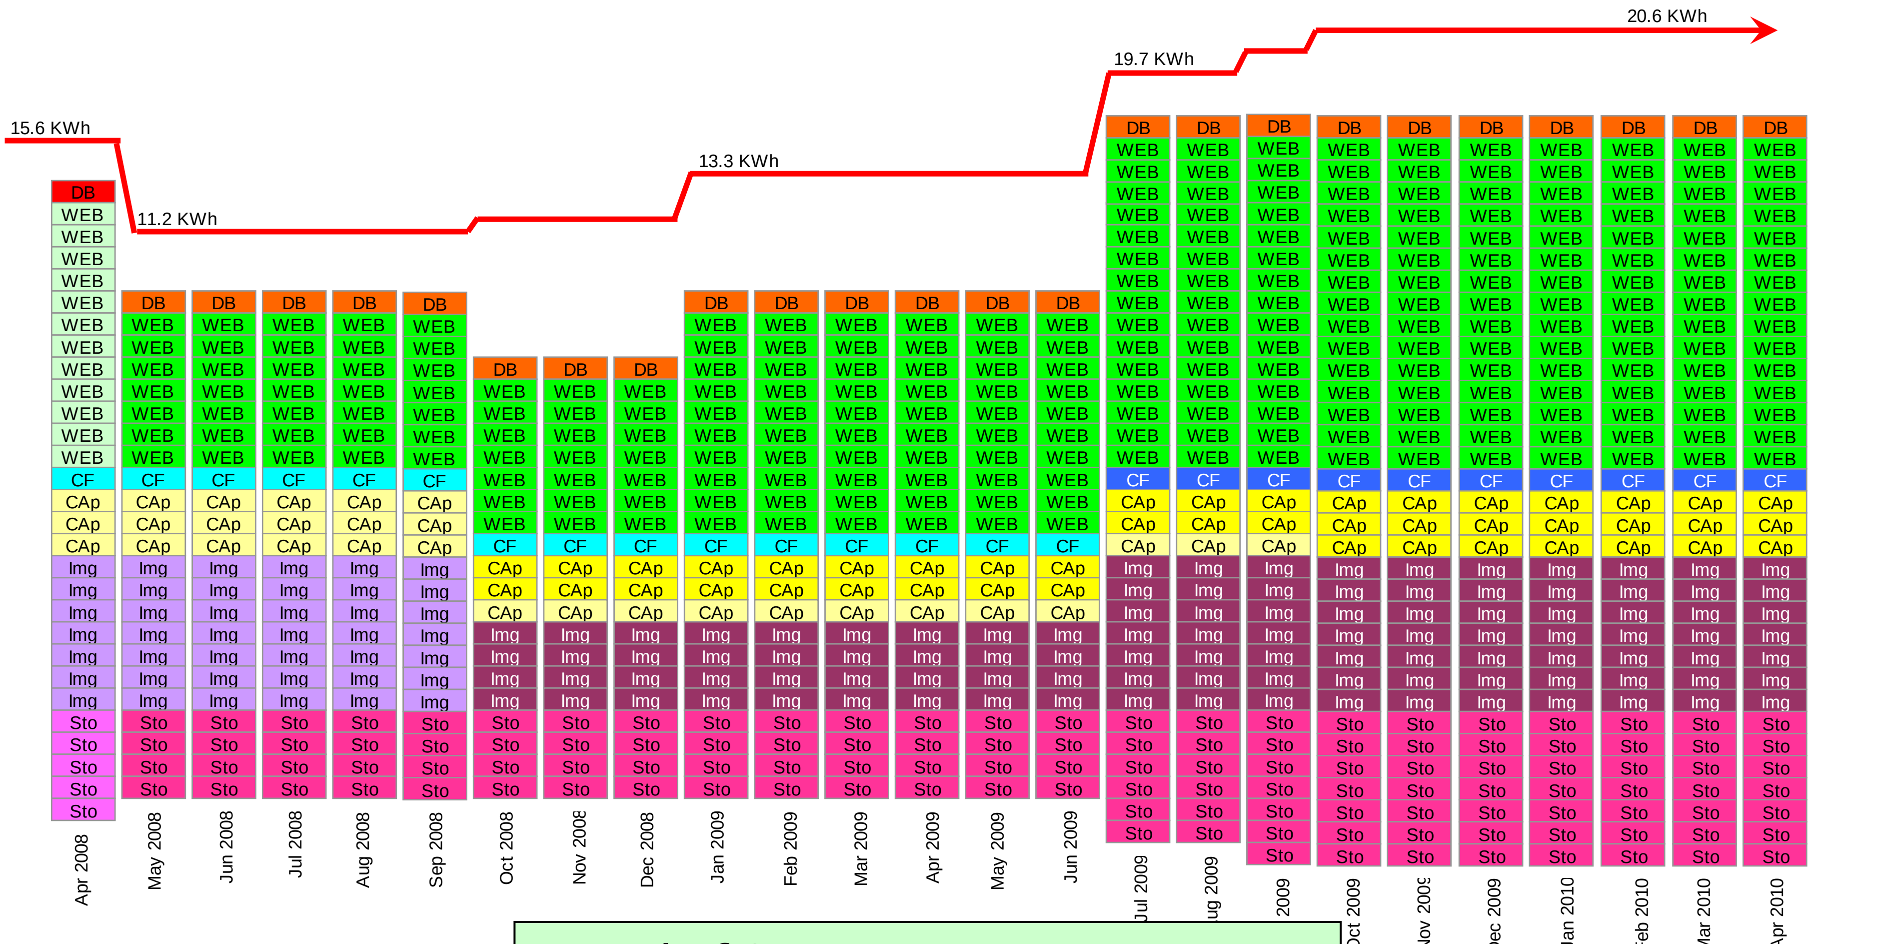

At end of 2-years

- 27 Servers 21 kWh / mo
- $150K 42U rack

Scenario - Server Upgrade: Monthly Energy Footprint

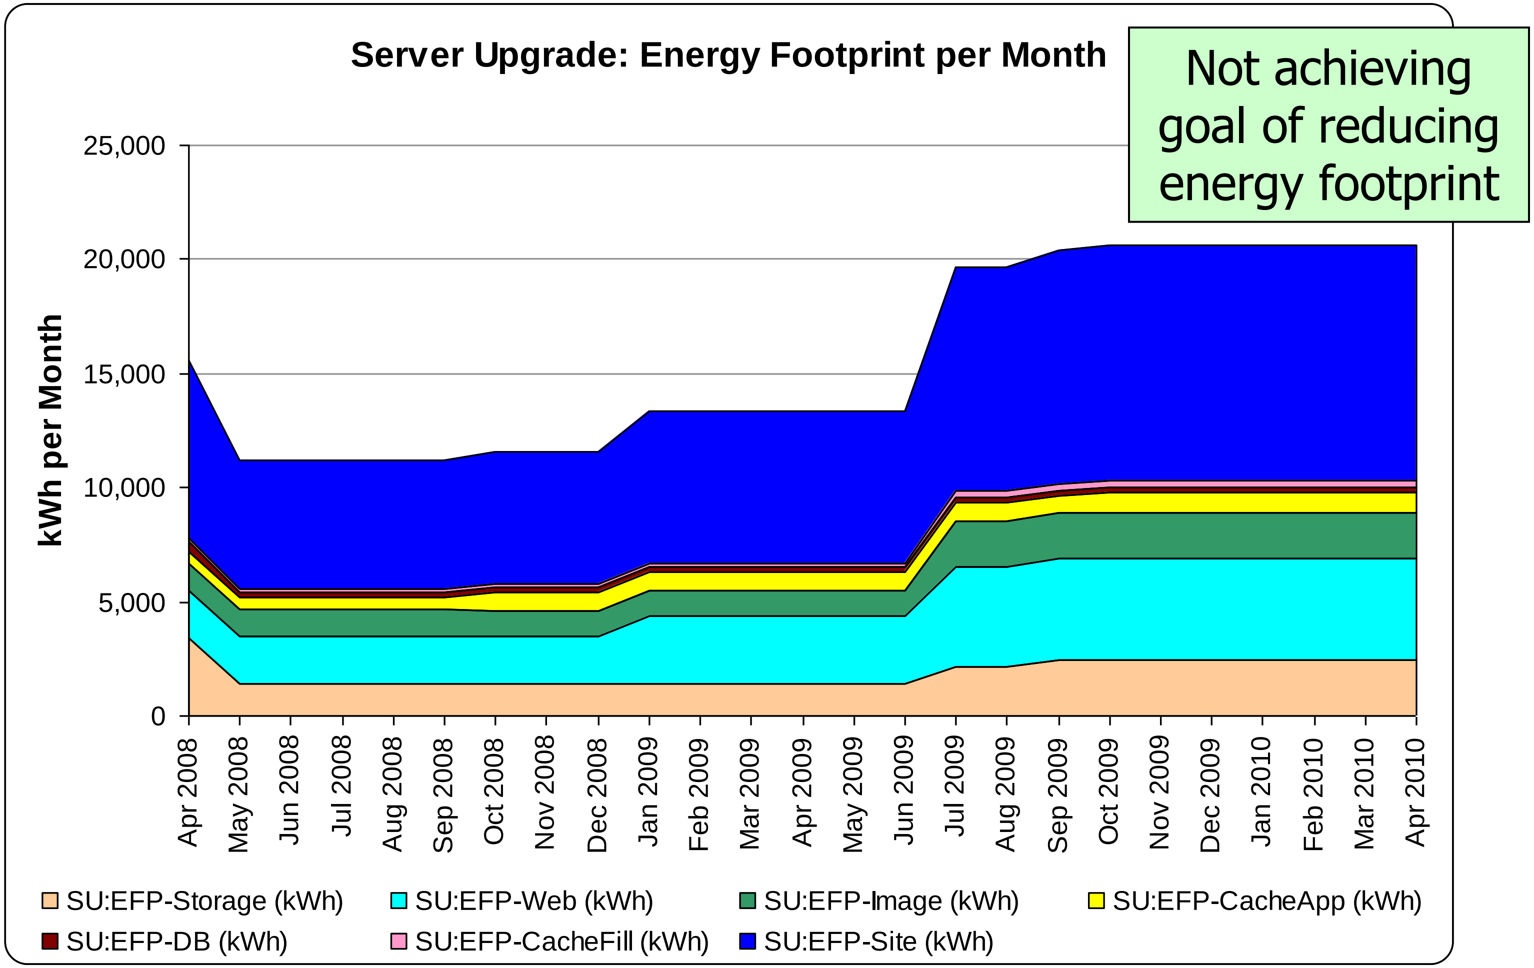

Case Study
Modeling Scenario Strategy

Server Upgrade

- 5% monthly growth for 2 years
- 70% capacity threshold
- Optimize the number of upgrades
- Dell PE 2950 for the database
- *Dell PE 1850 for all other servers*

Virtualization

- 5% monthly growth for 2 years
- 70% capacity threshold
- Optimize the number of upgrades
- Dell PE 2950 for the database
- *Virtualize all other servers on a pool of Dell PE 1950s*
- *Optimize use of virtual hosts*

For each scenario

1. Determine infrastructure required to support workload growth (Capacity Manager)
2. Estimate energy footprint
3. Calculate cost of IT equipment & energy

Scenario - Virtualization: Server Utilization Results

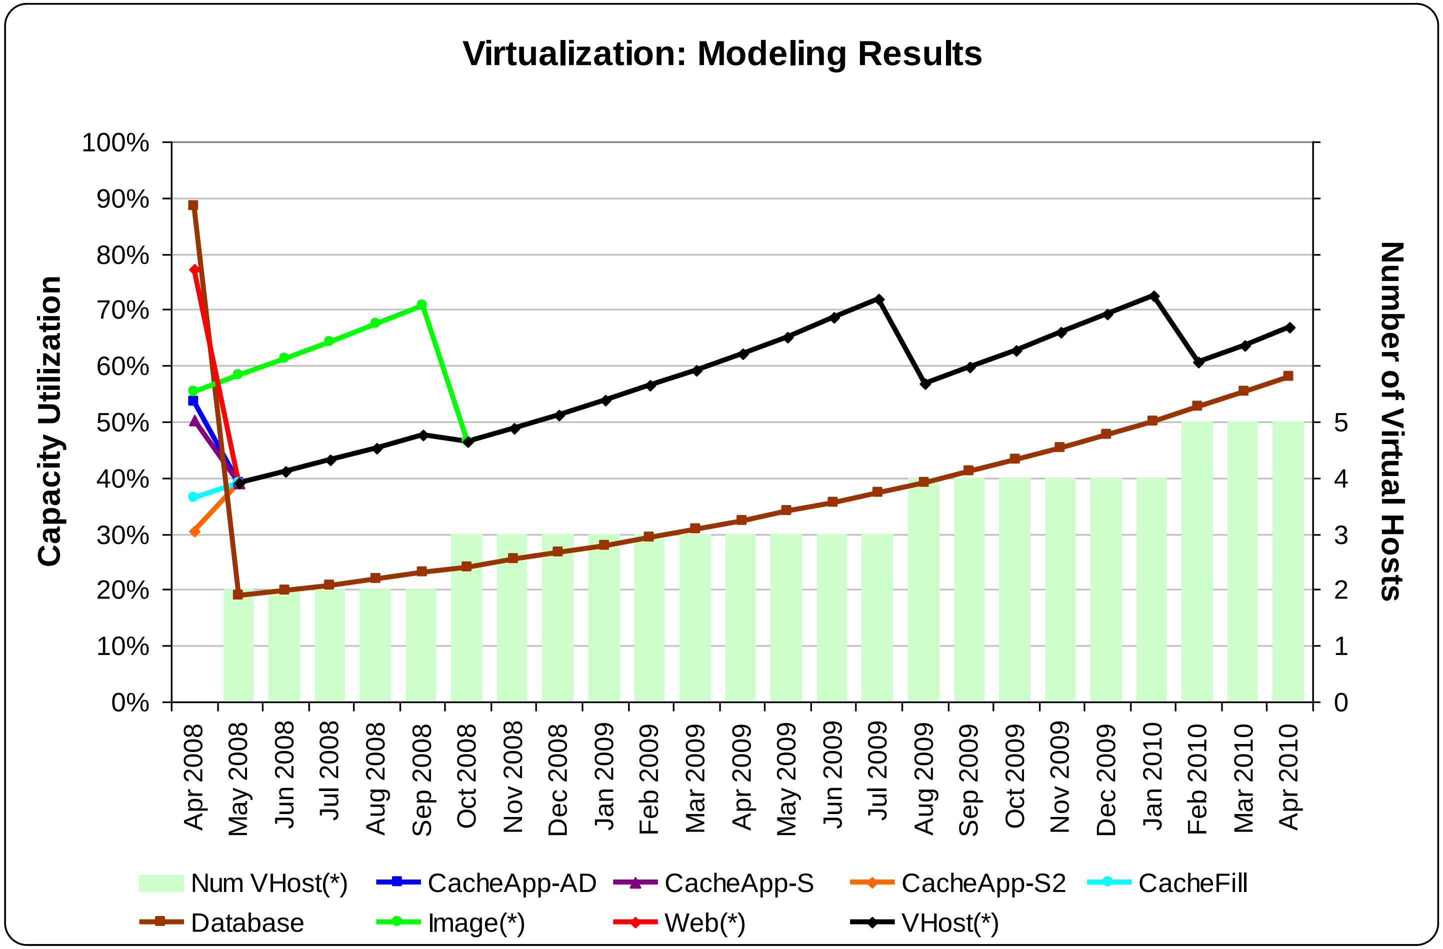

STLCMG

Scenario - Virtualization: Infrastructure Components

At end of 2-years

- 6 Servers
- 7.7 kWh / mo
- $120K
- 21U rack

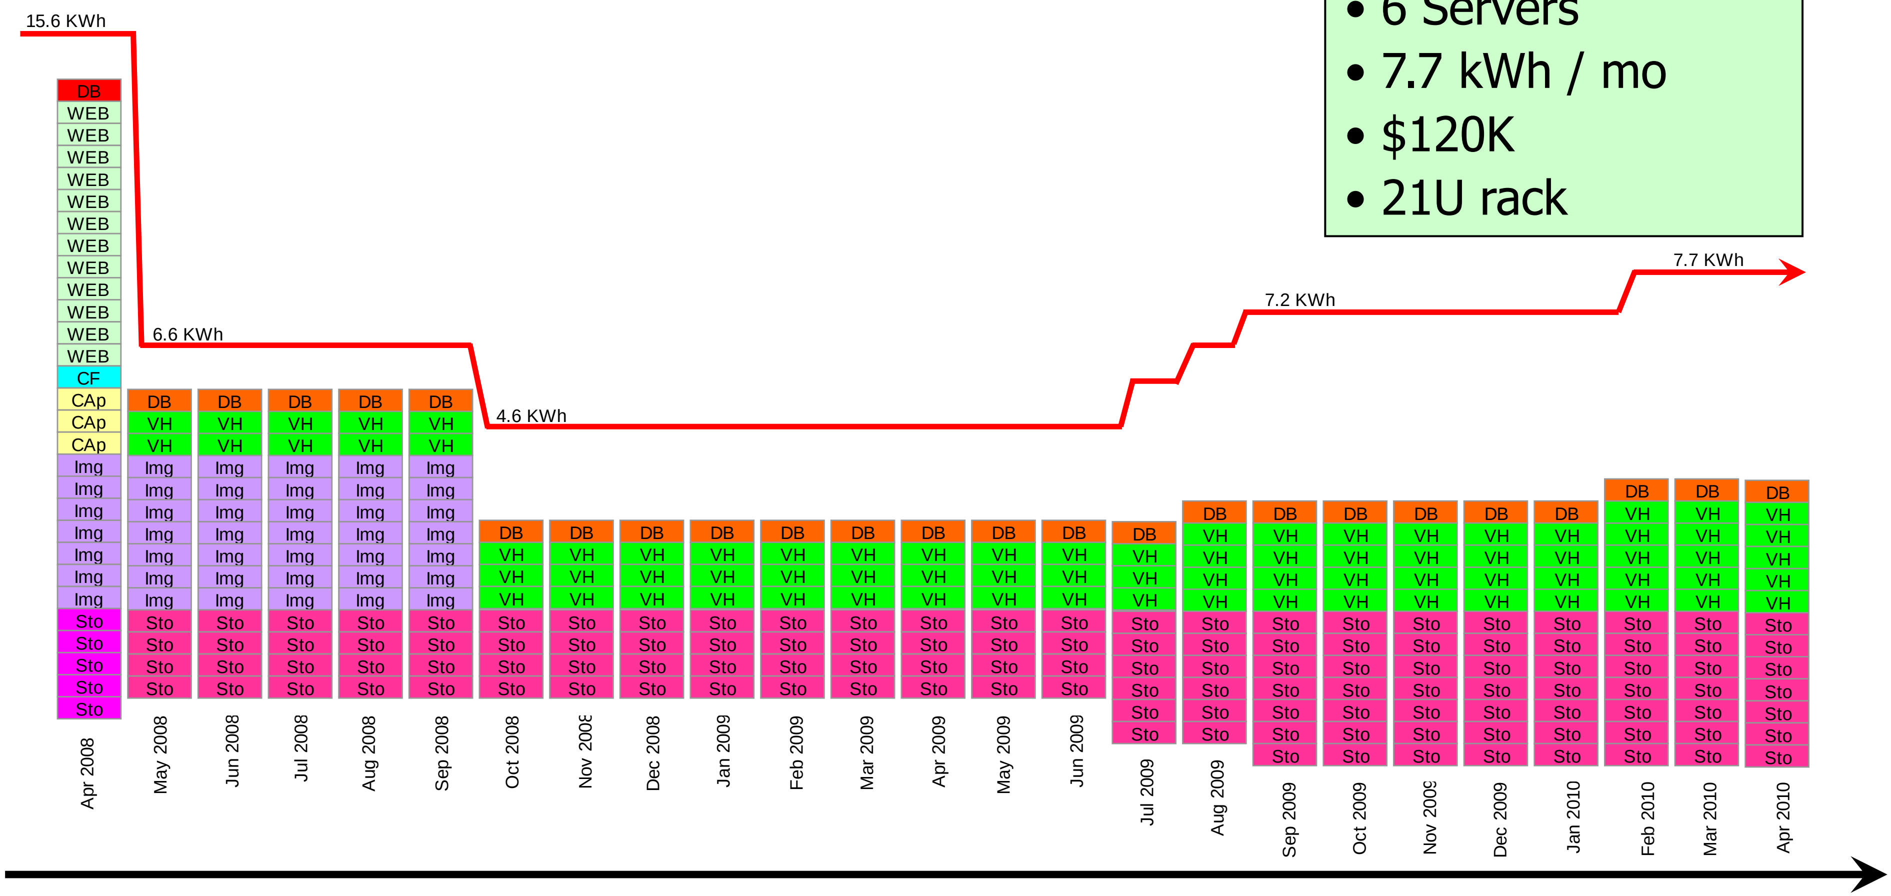

STLCMG

Example - Storage Upgrade: Replace HDD with SSD

- Replace Database rotating disks with SSDs
  - Random workload
  - Other stores (Log, Backup) are mostly sequential
- Keep same data throughput and latency levels
  - Maintain SLA
- Reduce
  - DataCenter footprint
  - Power
  - Cost (CAPEX, Operating)
- Compare Storage at month 24 only

Example - Storage Upgrade: Replace HDD with SSD

- Replace 4 HDD with one SSD
  - Similar capacity, will be sufficient to maintain performance
  - Reduced space, power & cost (capex)

Feature	HDD (3G SAS)	SSD (3G SATA)
Capacity	36 GB	128 GB
Size	2.5" SFF	2.5" SFF
Active Power, W	7.9	2.8
Idle Power, W	5.8	0.1
Cost, $/disk (est)	300	400
Sequential IOPs	30K – 100K	50K – 100K
Random RD IOPs	300	4.2K – 30K *
Random Wr IOPs	150	2.6K – 25K *
Sequential RD MB/s	120	240

* - Workload Dependent

Example - Storage Upgrade: Replace HDD with SSD

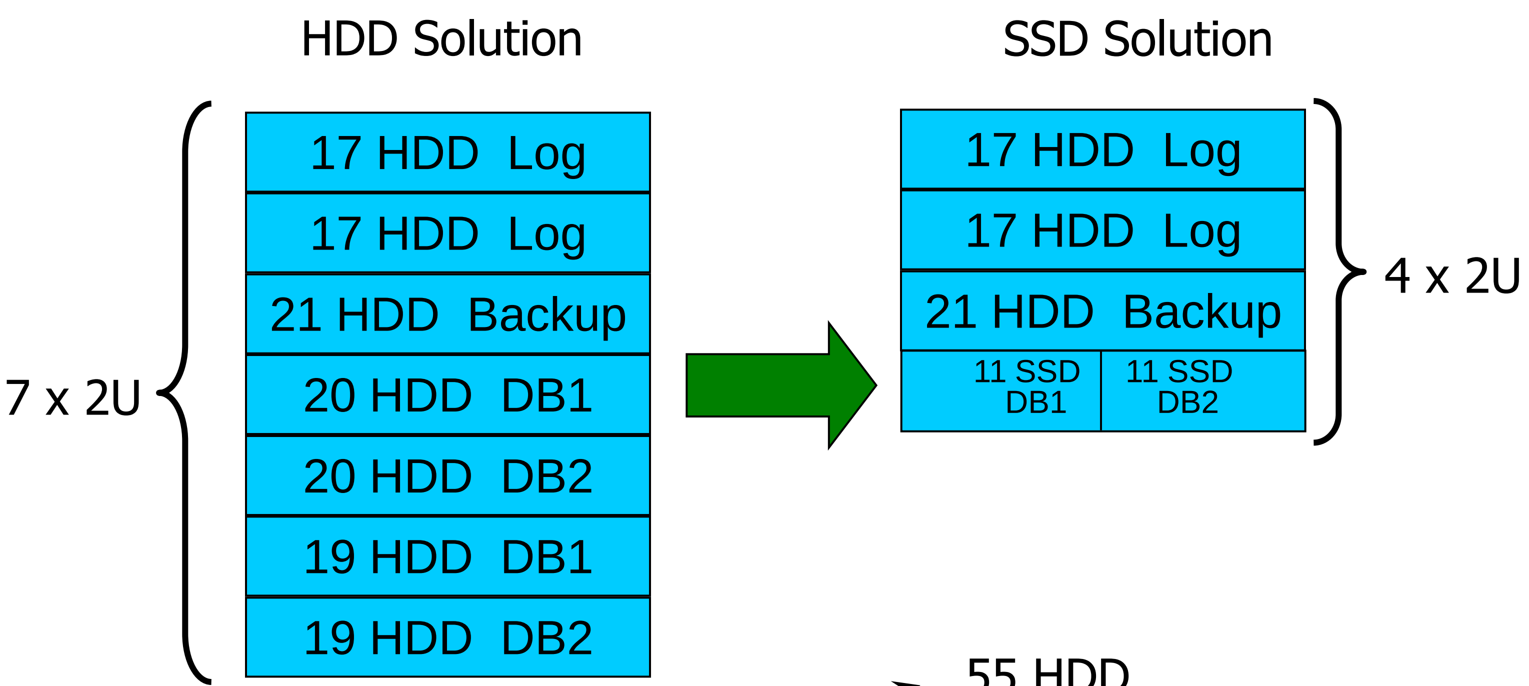

HDD Solution	SSD Solution	
133 HDD	55 HDD	
	22 SSD	42%
3395 W Active	1579 W Active	53%
$68,040 capex	$58,161 capex	15%
14U rack	8U rack	43%

STLCMG

Scenario Comparison: Server Upgrade vs. Virtualization

	Server Upgrade	Virtualization
Server Count	27	6
Energy Footprint	**21,000 kWh**	**7,700 kWh**
Cumulative Cost	$150,000	$120,000
Rack Space (U)	42	21

STLCMG

Scenario Comparison: Infrastructure Sprawl

Baseline

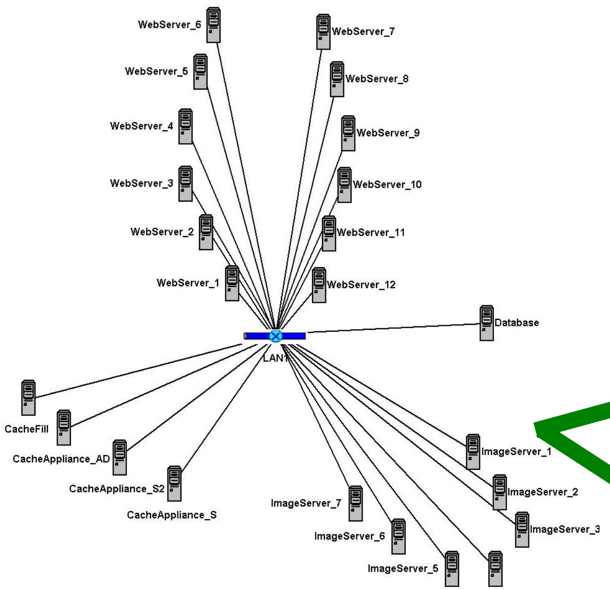

24 svr, 42U
15.6 kWh/mo
9,700 WIPSb

Server Upgrade

27 svr, 42U
21 kWh/mo
31,300 WIPSb

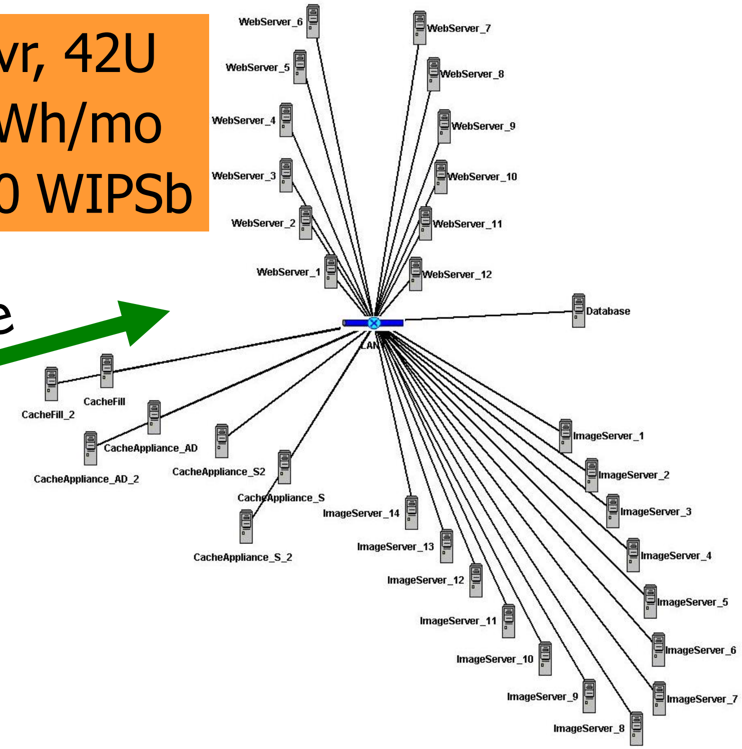

Virtualization

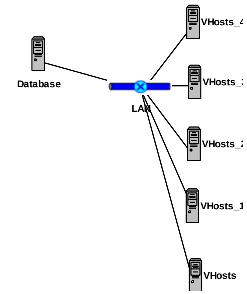

6 svr, 21U
7.7 kWh/mo
31,300 WIPSb

Scenario - Virtualization: Monthly Energy Footprint

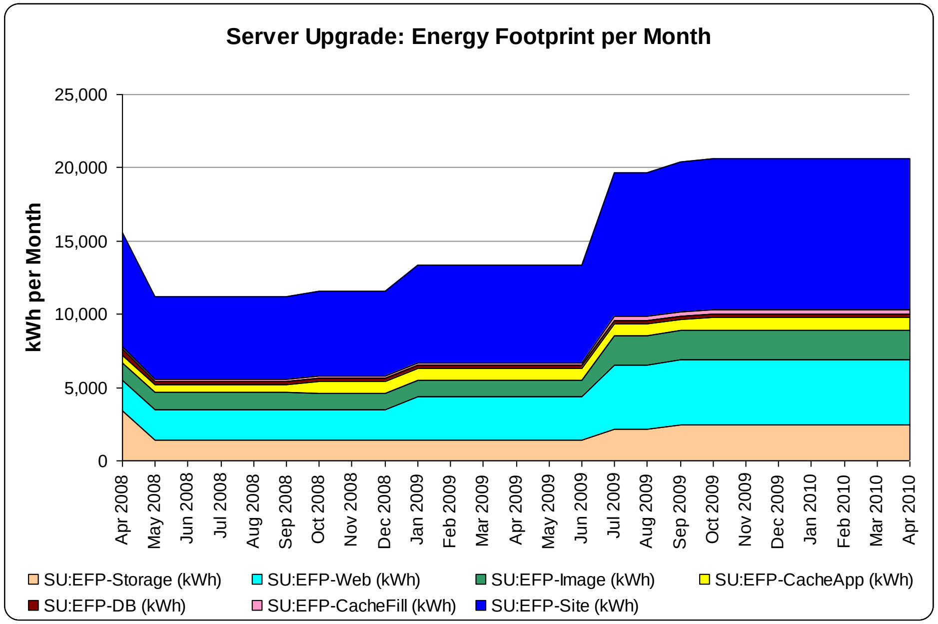

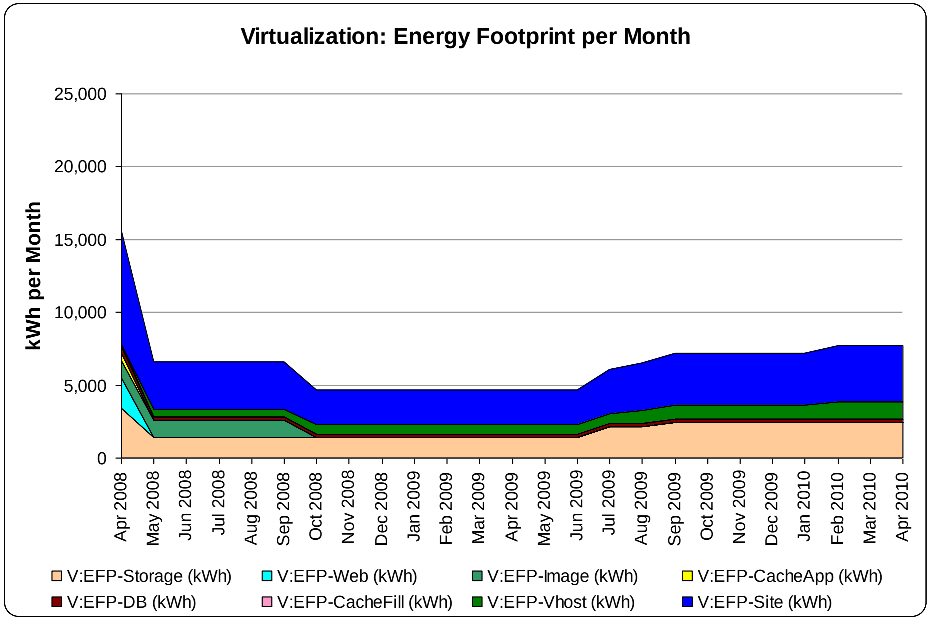

Significant reduction:

Energy Footprint (50%)

STLCMG

Scenario Comparison: Rack Space

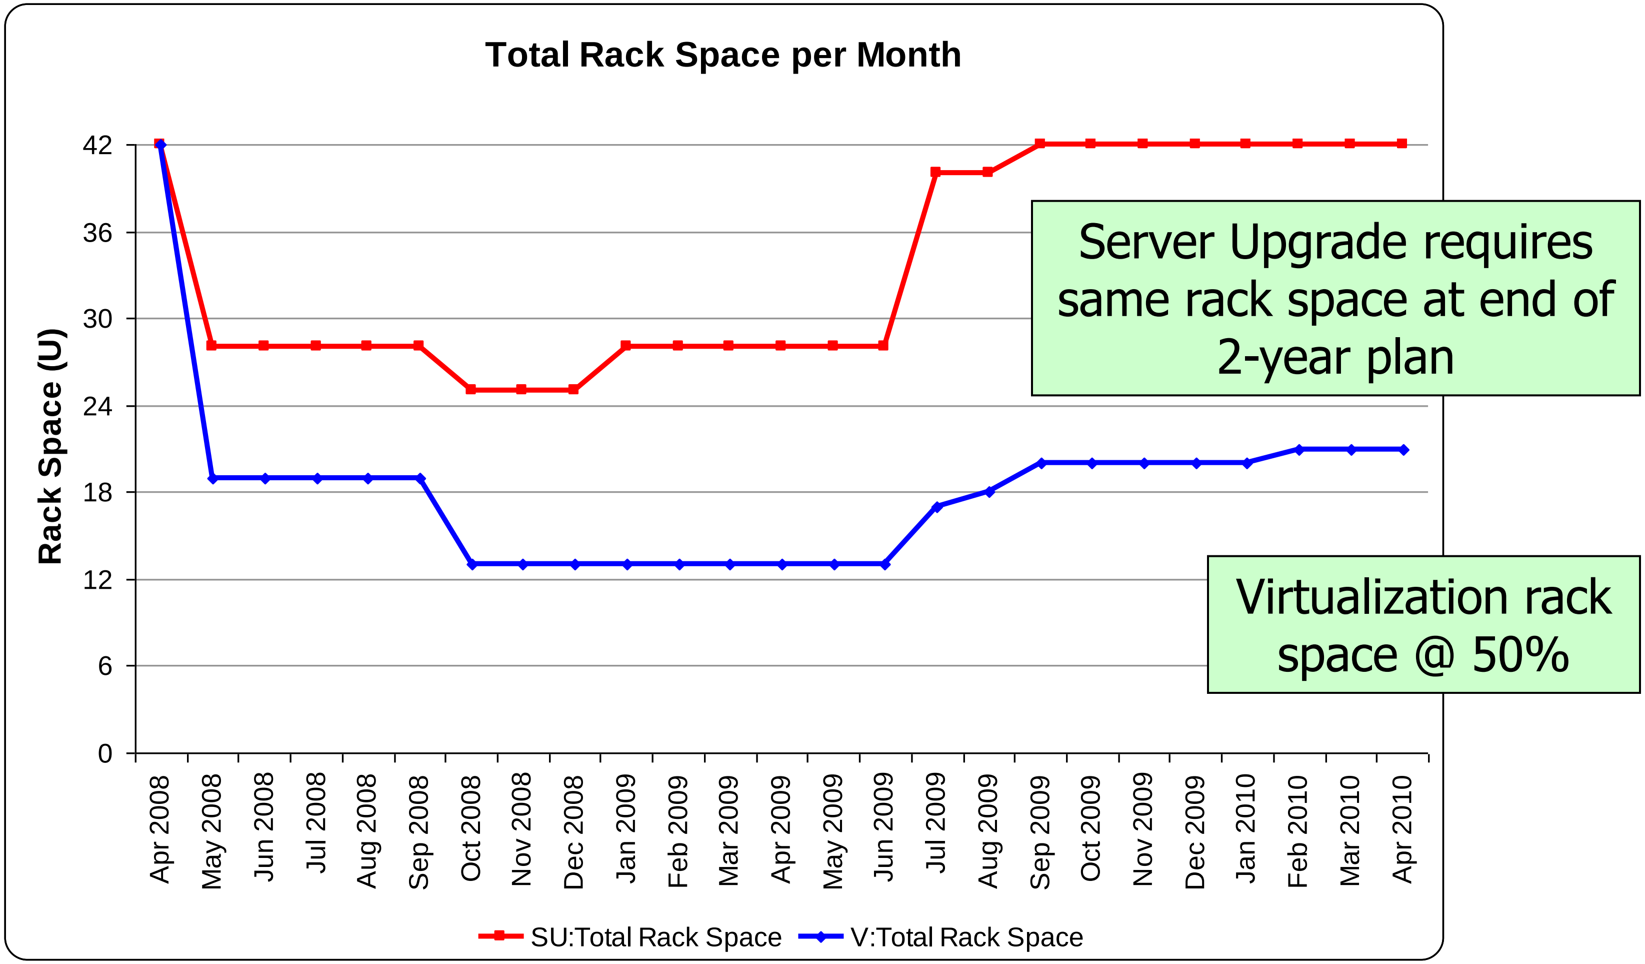

STLCMG

Scenario Comparison: Detailed Cumulative Cost

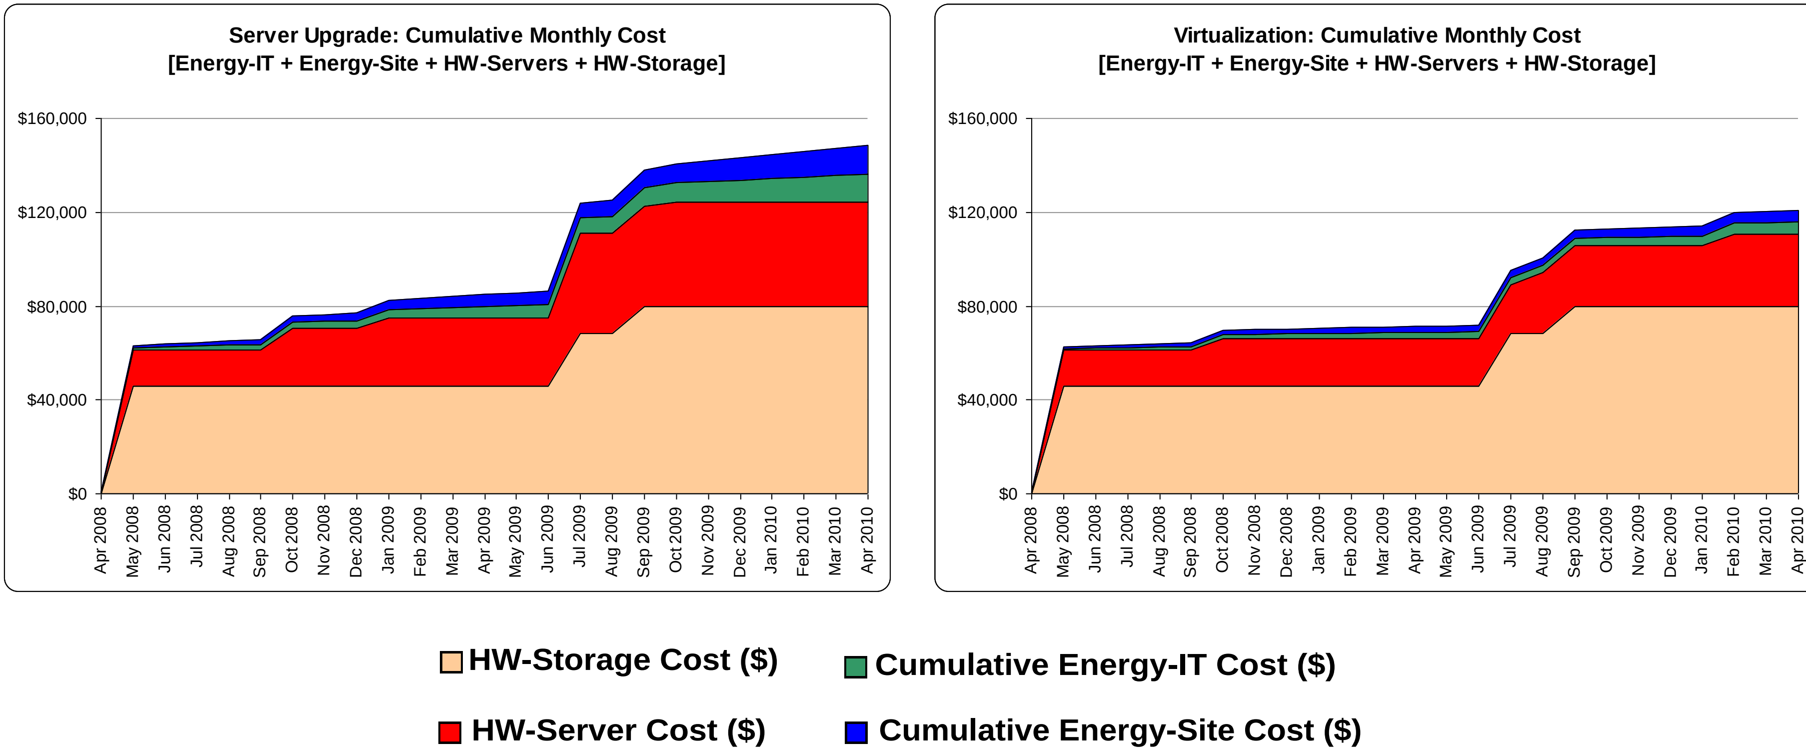

- Identical storage costs
- Virtualization server cost 30% lower
- Energy cost 60% less

Looking Ahead

- Emerging technologies are already impacting Green Capacity Planning in the DataCenter
  - SSD
    - Reduce # of disks & power
  - Multicore Processors
    - Higher CPU capacity in same physical & power footprint
  - Dynamic Power Management
    - Manage power of subsystems & components relative to current load
  - Virtualization
    - Combine multiple servers into one server
- Facilities improvements
  - Creative Cooling
    - Fluid cooling, higher temp operation
  - Decreasing PUE
    - Driving PUE very close to 1.0 (lights, fans, chiller)
  - Climate-driven locations

STLCMG

Energy & PUE Reductions

Reducing PUE doesn't necessarily mean reduction in energy utilization

First 3 bars all have same PUE, but different monthly KWH

Applying PUE reductions reduce energy utilization

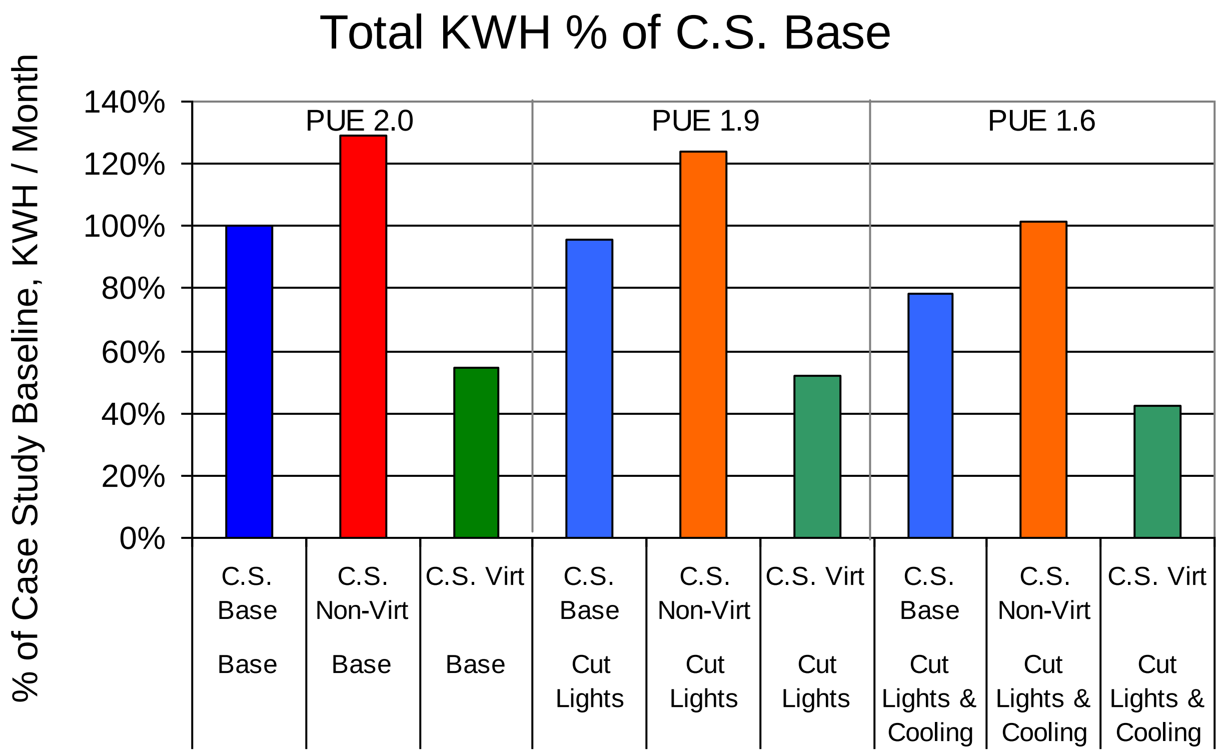

C.S. Base = Case Study Baseline system

C.S. Virt = Case Study system, virtualized, after 24 months of upgrades

STLCMG

Conclusion

- Green Capacity Planning
  - Energy is a resource
  - Introduce Energy Footprint Projection
  - Expands traditional capacity planning
- Case Study
  - Demonstrates the methodology
  - Analyzes alternative infrastructures
  - Reduces energy footprint with Virtualization
- Value
  - Actionable metrics
  - Quantifies IT and Site energy usage & costs
  - Supports responsible decision making
  - Bridges the gap between IT & Facilities

STLCMG

Primary Reference

This presentation is based upon a 2008 study by:

Amy Spellmann, Optimal Innovations

amy@optimalinnovations.com

Richard Gimarc, Hyperformix

rgimarc@hyperformix.com

Charles Gimarc, LSI Corporation

Charles.Gimarc@lsi.com

And a paper presented at the 2008 Computer Measurement Group conference, December 2008. The paper may be downloaded from:

http://www.optimalinnovations.com/pdf/green_capacity_planning.pdf

Reference

- [ACC2008] "Data Center Energy Forecast Report, Final Report, July 2008
- [EPA2007] "Report to Congress on Server and Data Center Energy Efficiency", U.S. Environmental Protection Agency, August 2008
- [KOOM2007] "Estimating Regional Power Consumption By Servers", Jonathan G. Koomey, LBNL Tech Note
- [SPEC2008] SPECpower_ssj2008, Feb 27, 2008, Dell PowerEdge 2950 III (Intel Xeon E5440)
- [TGG2007] "The Green Grid Data Center Power Efficiency Metrics: PUE and DCiE", 2007, The Green Grid, Page 6
- [TGG2008] "Five Ways to Reduce Data Center Server Power Consumption", 2008, Mark Blackburn, The Green Grid, Page 5
- [GCP2008] "Green Capacity Planning: Theory and Practice", Spellmann Gimarc & Gimarc, CMG 2008
- [GCP-b2008] "Green Capacity Planning", 2008, C. Gimarc, http://wikibon.org/?c=wiki&m=v&title=Green_Capacity_Planning