ITPMG

Achieving Green Storage

Performance Based Consolidations

Alok Jain

January 2009

Agenda

- Why Green Data Centers?
- Why Storage Consolidation?
- Storage Consolidation – Where to begin?
- The Overall Storage Consolidation Strategy
- Defining Storage Consolidation, Performance and Capacity Requirements
- Measuring “Green-ness” of current installed storage base
- Green Storage Measurement Metrics
- Use Case Study Charts
- Green Storage Consolidation Guidelines

Why Green Data Centers?

- Data Centers used 61 billion Kilowatt-hours in 2006
  - equal to power of 5.8 million households
  - Requires 15 power plants
- By 2011 it's expected that 25 power plants will be needed
  - 62 million metric tons of CO2 waste
- On average, 25% of an IT budget is for energy costs
- Recent IDC survey indicated that users expect power consumption to increase by 25% over next 2 years.
- Cost of electricity is increasing steadily
- For each dollar spent on electricity to power the storage and servers, it takes similar amount to get rid of the heat generated
- In many places additional power may not be available at any price
- Power and Cooling expenses have expanded by 8 times over the past decade
- Recent Forrester survey – 54% of respondents ranked IT consolidation as a top-five priority, equal to regulatory compliance
- ***Need to be an Environmentalist, a "Green" Citizen, and Cut Costs!***

Why Storage Consolidation?

- Amount of data being stored is increasing by 56% per year.
- Storage consolidation saves energy and footprints. Leads to reduced electrical cost and real estate
- Highly dense server racks and server virtualization are putting more pressure on storage requirements
- Storage consolidation provides for less assets to manage and administer
- Helps Storage Tiering work better with today's highly densified storage arrays.
  - Newer storage arrays provide for multiple tiers within the array.
  - Newer arrays are more efficient and highly scalable
- Provides for higher storage capacity utilization across the enterprise
- Cost effective with lower TCO
- Operating budgets are not growing as fast as storage requirements
- *One negative is bigger risk domains!*
- *Reduced Electrical, Footprint and Cooling Costs!!!*

Use Case: Example

22 Mid-tier systems consolidated to 2 High-end systems

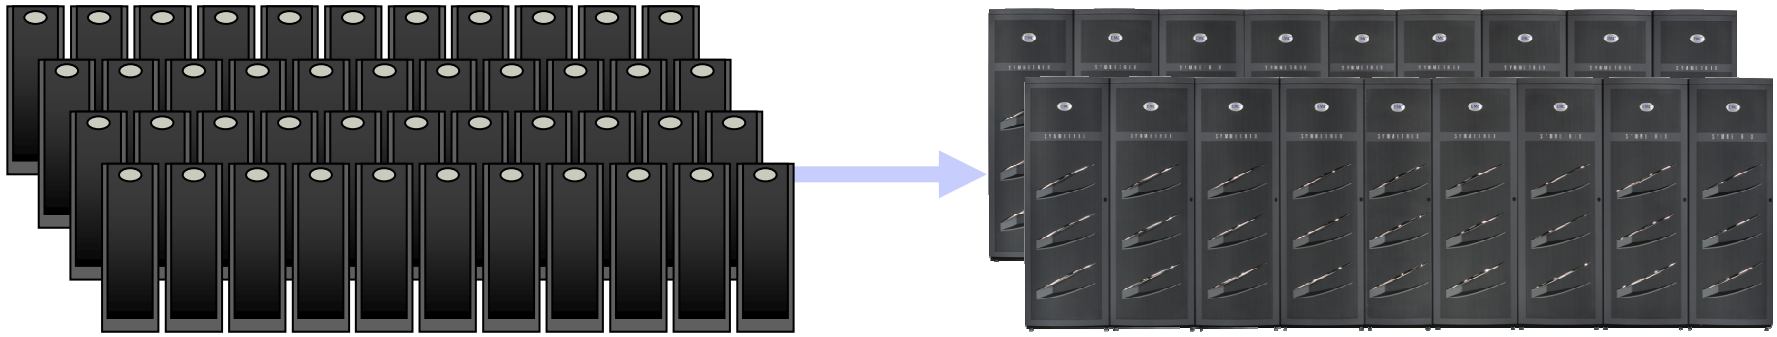

Mid-Tier
885 TB in 146 GB drives

High-end Array
934 TB in 146 GB and 300 GB drives

317 kWh	**Power and Cooling**	**169 kWh**
$416,538	**Annual Cost @.15 kWh**	**$222,066**

1,296,480 kWh - Annual Energy Savings - $194,472 (46%)

Source: EMC Corporation

Opportunities for Storage Consolidation

- Primary Opportunities for the consolidation project
  - Tech Refresh
  - Data Center Migration
  - Storage Consolidation – want to be more “Green”
- Primary objective will drive
  - Consolidation decisions
  - Migration Strategies
  - Opportunities for Storage Standardization

Storage Consolidation – Where to Begin?

- Requires understanding and planning
  - What are the workloads?
  - What are the storage array capabilities?
  - What are the performance characteristics of the storage array and workloads?
  - What are the data protection and business continuity requirements?
  - What is the current Storage Area Network Architecture?

Understand the Business Guidelines for Storage Allocations

- What are the business guidelines for the storage provisioning?
  - Dedicated Mode aligned to following – *increases storage consolidation complexity*
    - Development/QA/Production
    - Business Units
    - Applications
    - Tiers
    - Vendor Platforms
  - Shared Mode – *provides better storage consolidation planning and utilizations*
    - Storage is purely allocated based on storage tiers

How to move to Green Storage?

- Identify Data Centers and Source Arrays that are in scope
- Identify provisioned storage data points for each allocation unit
  - Array ID
  - Server names
  - Storage capacity provisioned by server
  - LUN size and LUN counts by server
- Array Level Consolidation – Understand array level IO Profiles
- Host Level Consolidation – Understand application IO Profiles
- Total Storage provisioned by overall array
- *Develop Detailed Migration Plan*
- *Develop Detailed Project Execution Plan*

Storage Consolidation Strategy

- Overall Storage Consolidation Strategy
  - Consolidation Strategy
    - Array level consolidation
      - source array to target array is pre-assigned
    - Host level consolidation
      - Best fit analysis based on performance and capacity requirements
  - Migration Strategy
    - Array based migration
    - Host based migration
    - Tool/Appliance based migration
  - Storage Standardization Strategy
    - Storage Tier Standardization
    - LUN Size Standardization
    - HBA/Driver Layer Standardization

Defining Storage Performance Requirements

- Define storage performance and configuration threshold requirements at critical component levels. Need to incorporate the growth requirements for a defined time period
  - Acceptable IOPS and MBPS
  - Acceptable Fan-Out Ratios
  - Acceptable Masked Device limits
  - Acceptable Risk Domains
- Define critical metrics for each of critical components, For example
  - IO/sec
  - MBPS
  - %Cache hits
  - Platform specific critical metrics (e.g. Write Pending slots at array level and device level for EMC Symmetrix arrays, etc.)
- Create a Component/Metric matrix (CMM) for performance modeling reference

Component-Metric Matrix – Sample (CMM)

Identify components and metrics – Define Max and 95* Percentile thresholds

Component/Metric	IO/sec	MB/sec	%Cache Hits	Metric 1	Metric 2, etc.
Overall Storage Array	20000/15000	1200/800	70% / 80%	…	…
Host Fibre Director	5000/4000	80/65	70% / 80%	…	…
Disk Director	3000/2500	500/400	70% / 80%	…	…
Remote Fibre Adapter Port	2500/2000	1500/1250	NA	…	…
Disks	180/150	20/10	70% / 80%	…	…
Identify Other Components	…	…	…	…	…

*The percentile level may need to be adjusted for specific workloads

Measuring the Green-ness – Collect current asset inventory

- Perform a storage inventory of the current installed base.
  - Capture the critical data for each storage array.
  - Best done using a storage capacity analysis and performance base-lining tool, if available (e.g. perfonics™)
- Identify a command set that will export the critical data needed for consolidation analysis and planning. Use the tool/service to analyze the data
  - Capture the following data
    - Storage Array model, microcode, configuration data (cache, disk size, disk speeds, LUN sizes, etc)
    - Storage capacity allocations on each storage array (LUN size, LUN counts, fibre adapters etc.)
    - Collect performance data for each identified critical component

"Green Storage" Measurement Metrics

- In an effort to improve the performance-per-watt, some of the metrics that can be used to measure are:
  - Storage Array level
    - Power Usage
      - Watts per sq. ft.
    - Storage Density
      - TB per sq ft.
      - GB per kW
    - Performance metrics
      - IOPS per kW (or sq. ft.)
      - MBPS per kW (or sq. ft.)
  - Green Grid standardized metrics at DC Level
    - Power Usage Effectiveness (PUE)
      - Total Facility Power/IT Equipment Power (1=100% efficient. 1.6 is very good and achievable)
    - Data Center Infrastructure Efficiency (DCiE, reciprocal of PUE)
      - IT Equipment Power/Total Facility Power x 100%

Symmetrix Performance IOPS & MBPS Profiles

- The following shows the IOPS and MBPS loads (max, 95th percentile, average) on all source array (EMC Symmetrix). These values are used to come up with best fit performance consolidations

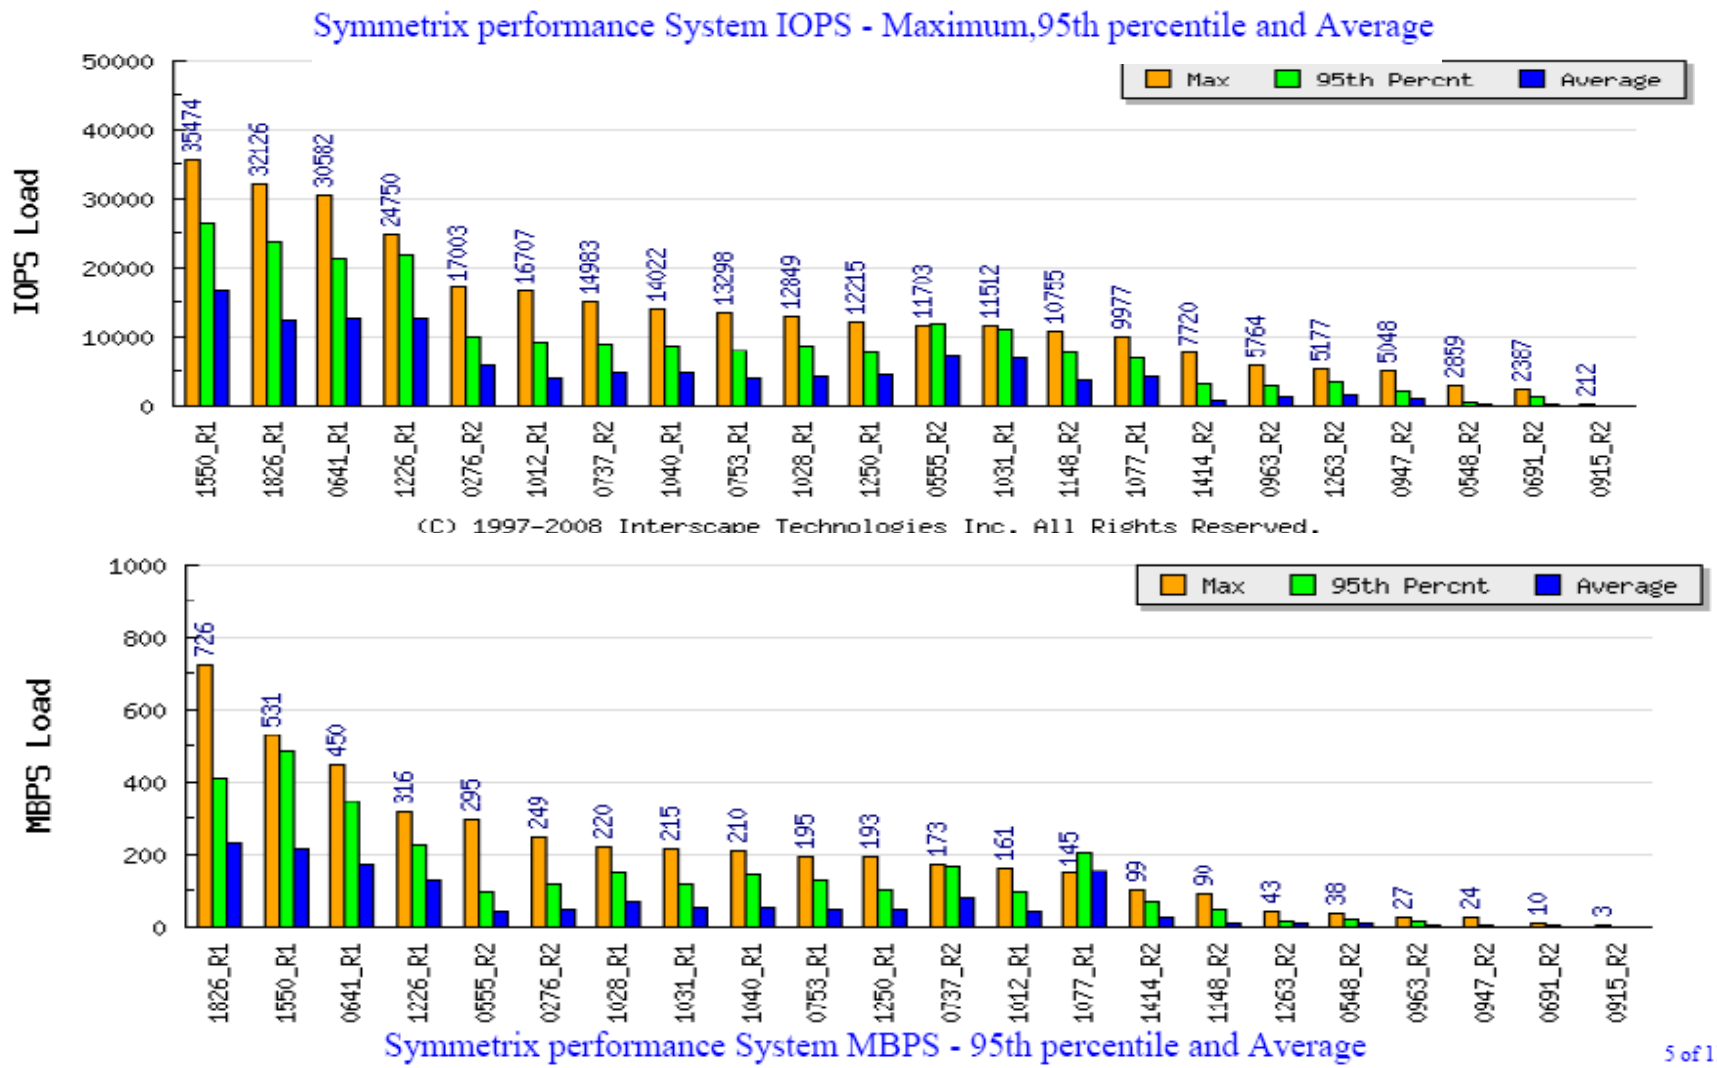

EMC Symmetrix Sample – FA IOPS, MBPS, Masked Devices & FAN-OUT Analysis

- Based on the data collected, the following chart shows the critical data points that needs to be considered for consolidation analysis. Shown are the **max IOPS**, **max Masked Devices**, **max Fan-Outs** on a source array

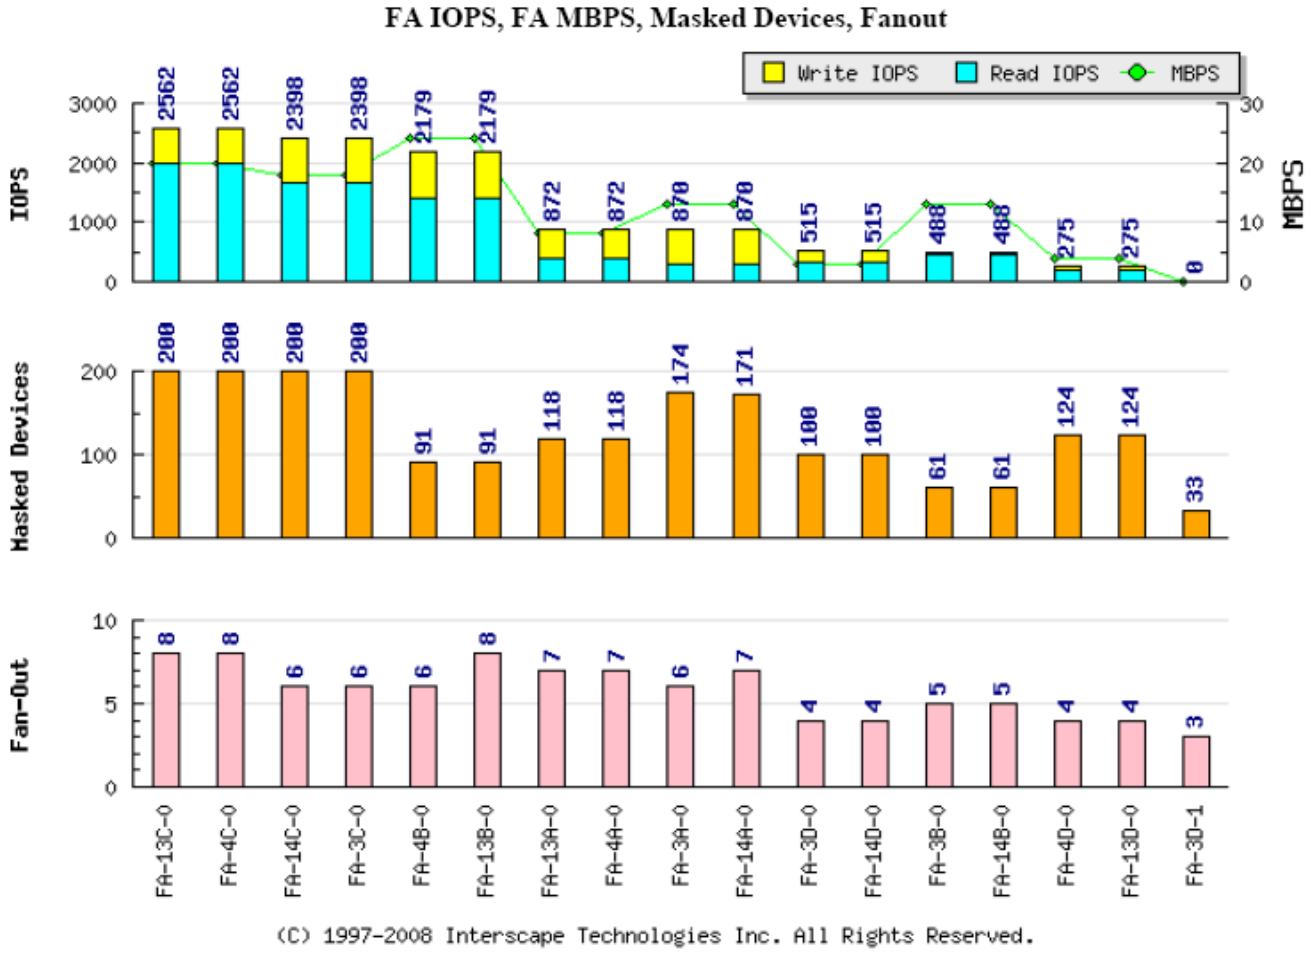

Target Storage Array – Performance Heat Chart (Model)

- Sample shows modeled performance heat chart based on source IO profiles for combined load. Shown are the various critical array components modeled max or 95th percentile load on the Target Array (EMC Symmetrix)

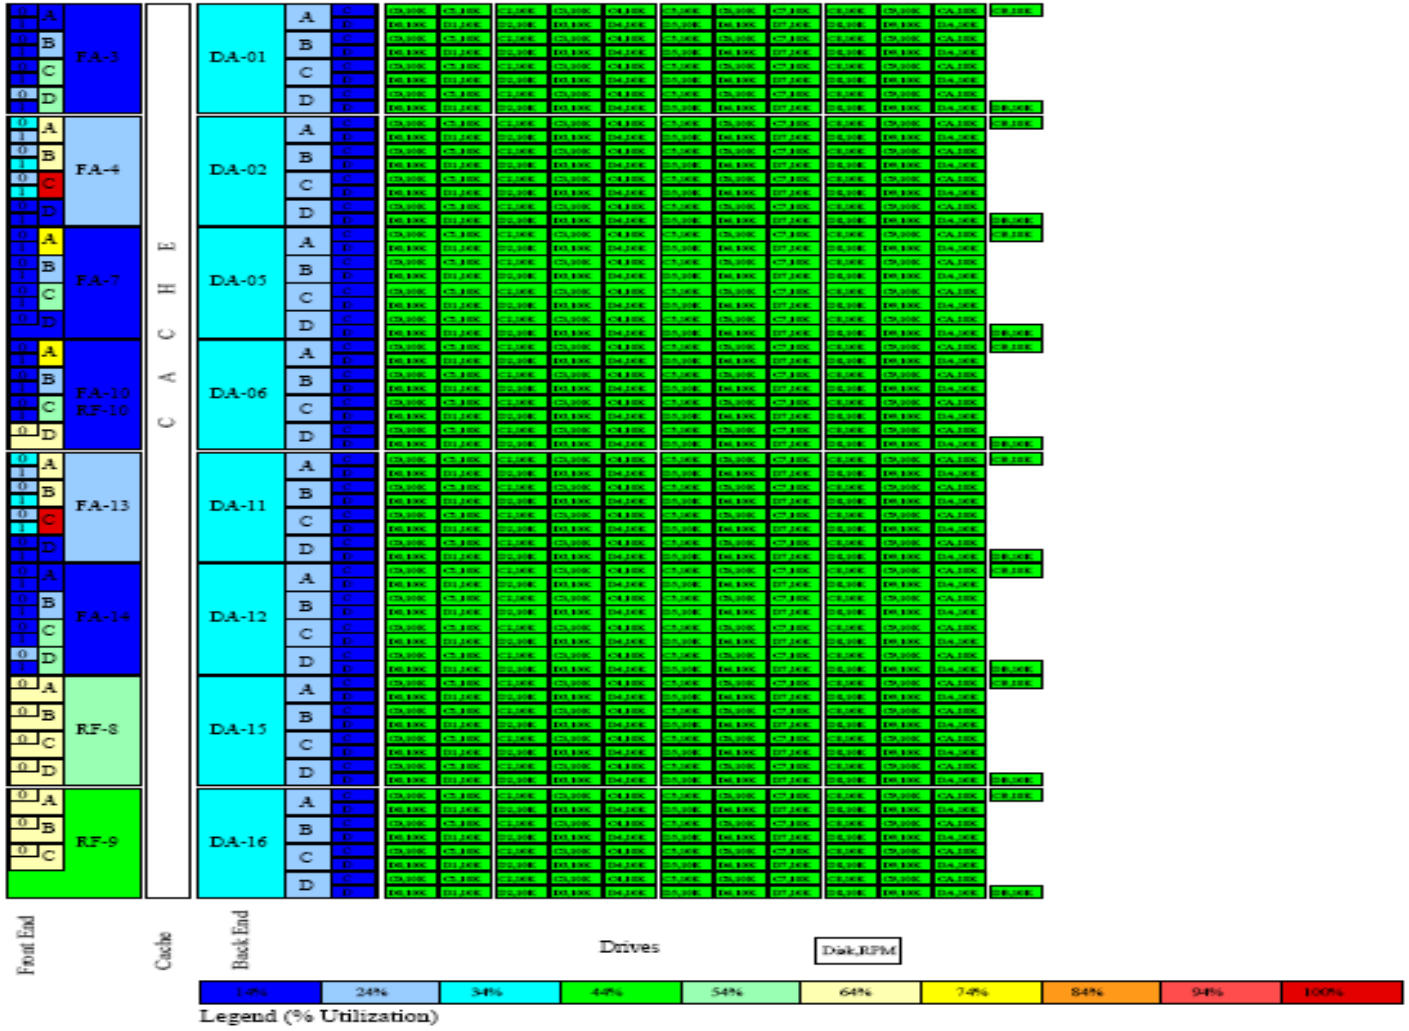

Energy Cost vs. Disk Density

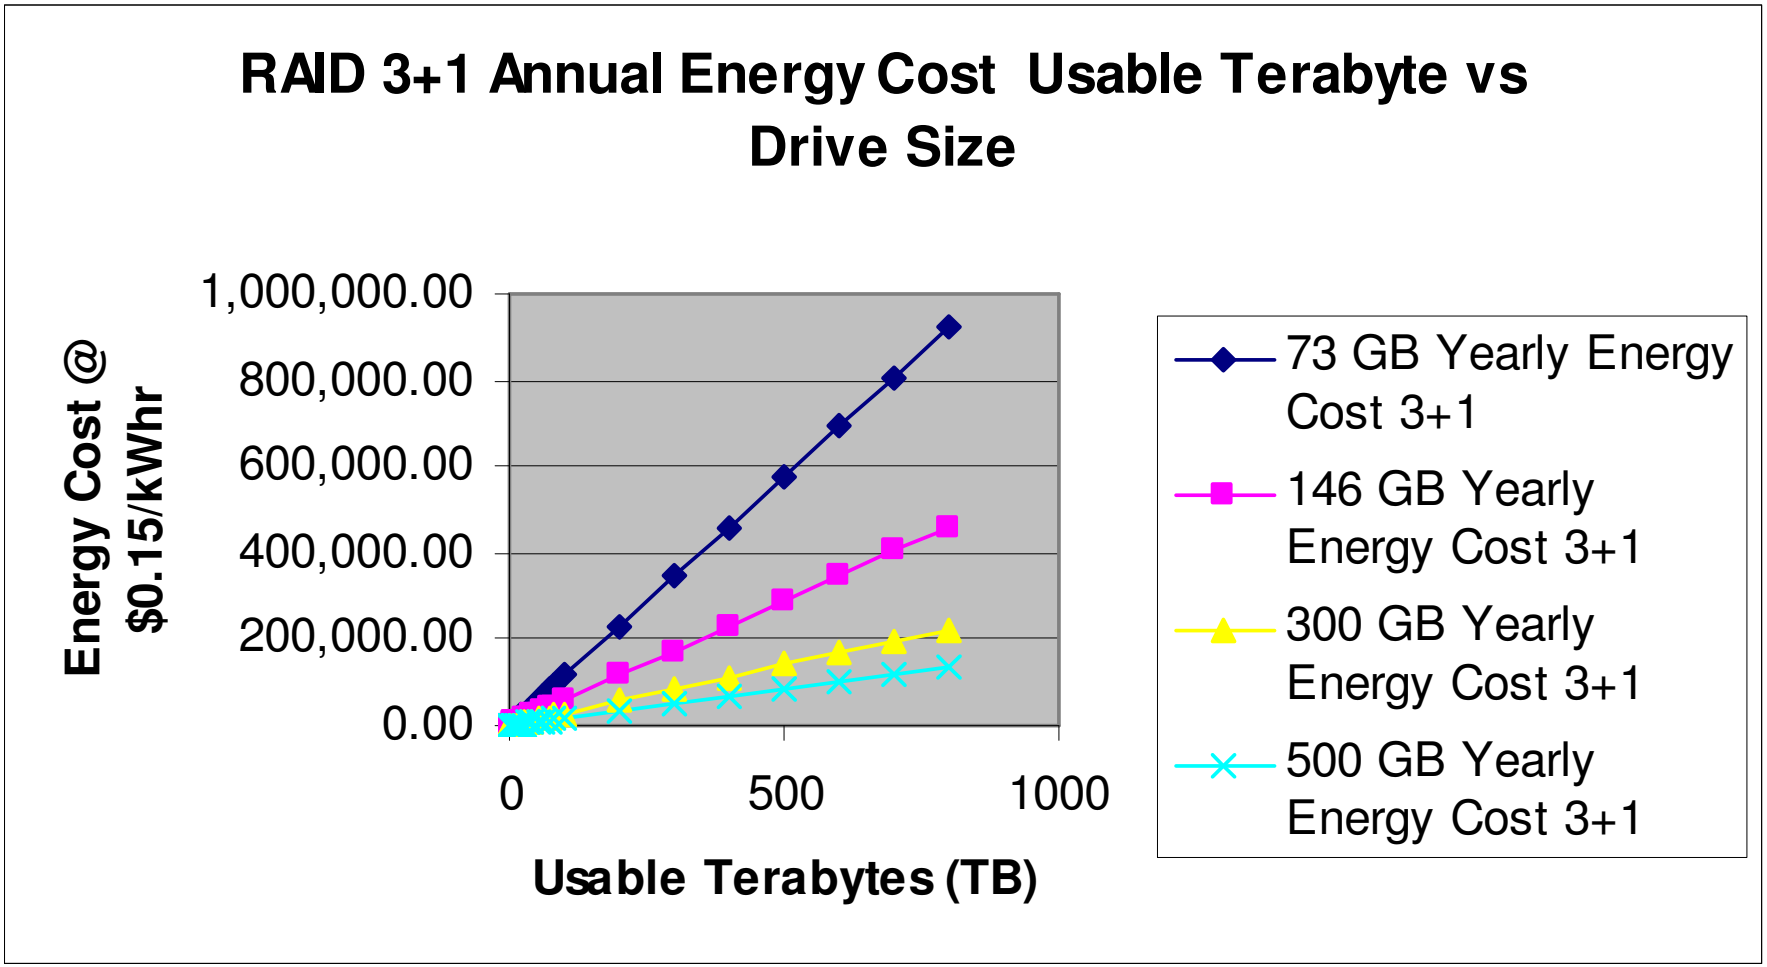

Higher density disks brings the energy cost down

Source: EMC Corporation

Strategies for Maximizing Storage Efficiency – To Achieve Green Storage

- Consolidate storage arrays which will lead to lower TCO overall
- Install **storage arrays that will scale well** for next 3 years to avoid forced migrations. Arrays should provide linear performance scalability
- Implement **Thin Provisioning** as part of storage tiers
- Implement **Application Tiering** within arrays
- Make right business decisions on **Application RAID Levels**
- Perform **pro-active application profiling** for on-going IO profile analysis.
- Use **higher density disks** for better footprint density
  - 1TB of data on 7200 RPM, 1TB SATA drive is 94% more efficient then storing it on 15000 RPM, 73GB Fibre Channel drive
- Buy just enough physical disks that will satisfy the performance requirements. Every **spinning disk adds to the power usage** and footprint requirement

Strategies for maximizing storage efficiency – contd.

- Implement **multiple tiers** with-in large storage arrays
- **Solid state disk drives** use up to 98% less power compared to lower density disk drives.
- **Information Life Cycle Management** can be used. Consider using Massive Array of Idle Disks technology (MAID). Archival data can be moved from nearline storage to MAID. This can introduce higher complexity for data management
- **Energy usage** should be considered part of Operational Expense for buying decisions besides looking at Capital Expenditure
- Use **array based data replication** technologies to help streamline the data migration process for high speed and data integrity and minimize on administrative migration overhead
- Incorporate a **Storage Capacity Planning & Trending** tool that includes Performance Trending

To receive additional materials on Green Storage, contact ITPMG at:

by email:
information@itpmg.com

or phone:
+1 843.377.8228