DataCore™ SOFTWARE

The Storage Hypervisor Company

Automated Storage Tiering:

Optimizing Flash Memory in Today's SAN Environments

STORAGE - A BIG DEAL

- Impacts
  - Users
  - Apps
  - Systems
  - Operations
  - Spending

BUSINESS IMPLICATIONS

Metrics

- **Access**
  - Can I get to it?
  - May I share it?
- **Performance**
  - How quickly?
- **Peace of mind**
  - Is it well protected?

Factors

- Interruptions / outages
- Congestion / slow gear
- Cost to store safely

INTELLIGENT LAYER BETWEEN APPS & STORAGE

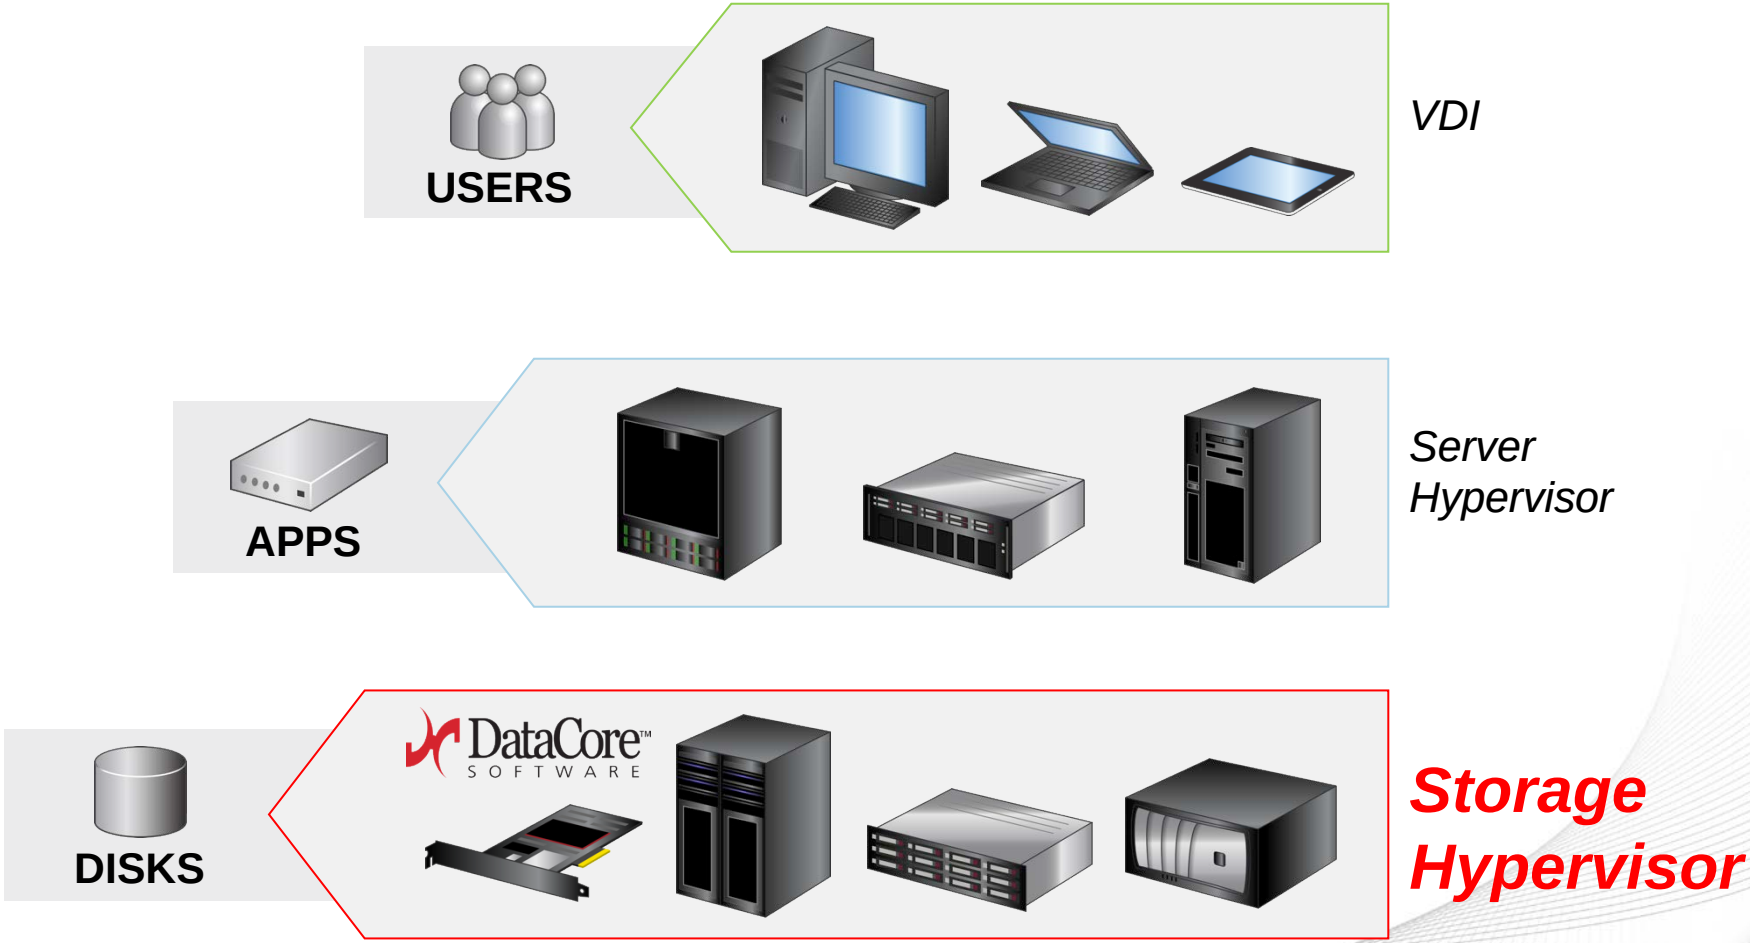

POOLING

- Combines overall capacity

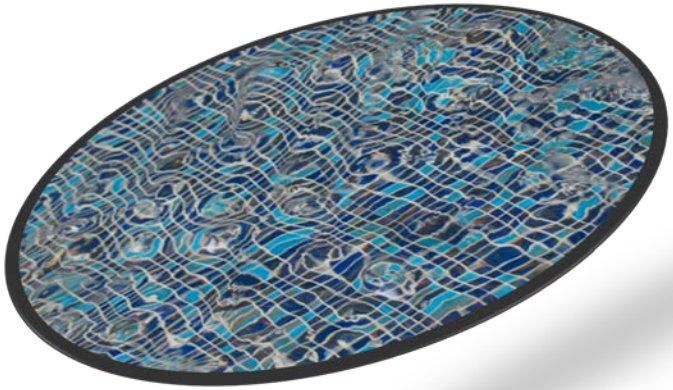

- Makes disks & flash shareable
- Recovers inaccessible or isolated disk space
  - **No longer isolated**
- Hides equipment incompatibilities
  - **Standard server / app connections bridged to device-specific interfaces**

SHARE POOLED RESOURCES

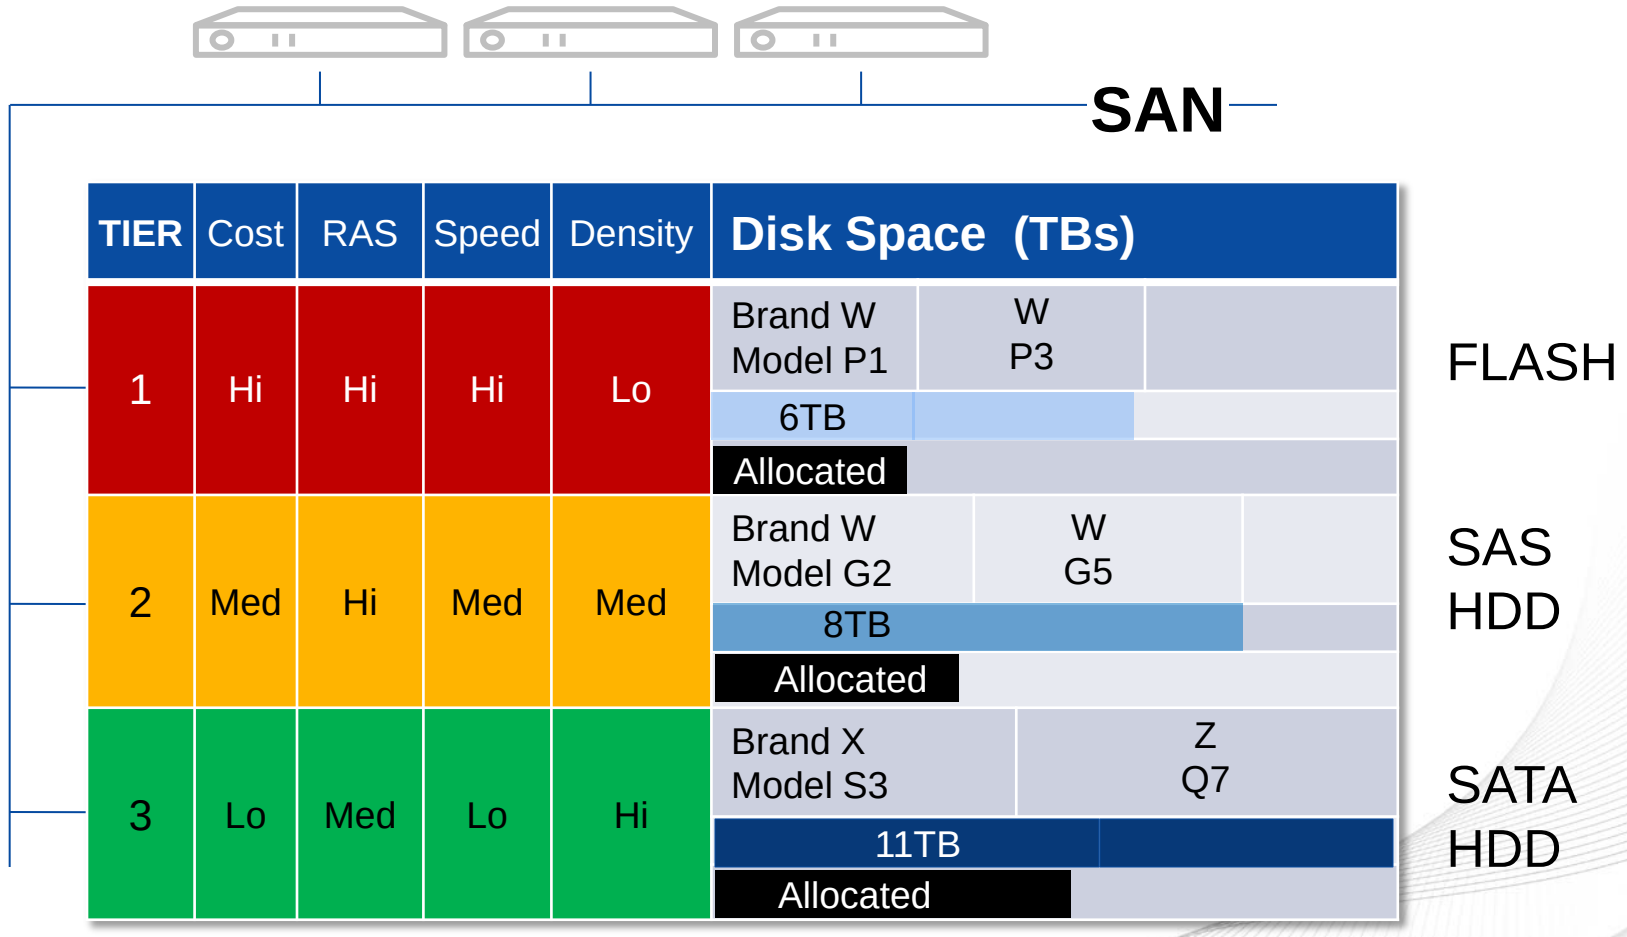

AUTOMATION

- Allocates space just-in-time
  - **Thin provision small groups of disk blocks**
  - **No waste**
- Dynamically directs workloads to most appropriate resource
  - **Auto-tiering**
  - **Determined by access frequency & business rules (high priority override)**
- Device-independent

AUTOMATED STORAGE TIERING

Make best tradeoff between performance & cost

- **Select tiers based on price/performance/capacity**
  - **SSD, SAS, SATA, …**
- **Dynamically migrates blocks among classes of storage**
- **Best location at the time determined by:**
  - **Access frequency**
  - **User preferences**

BEST RESOURCE FOR THE JOB

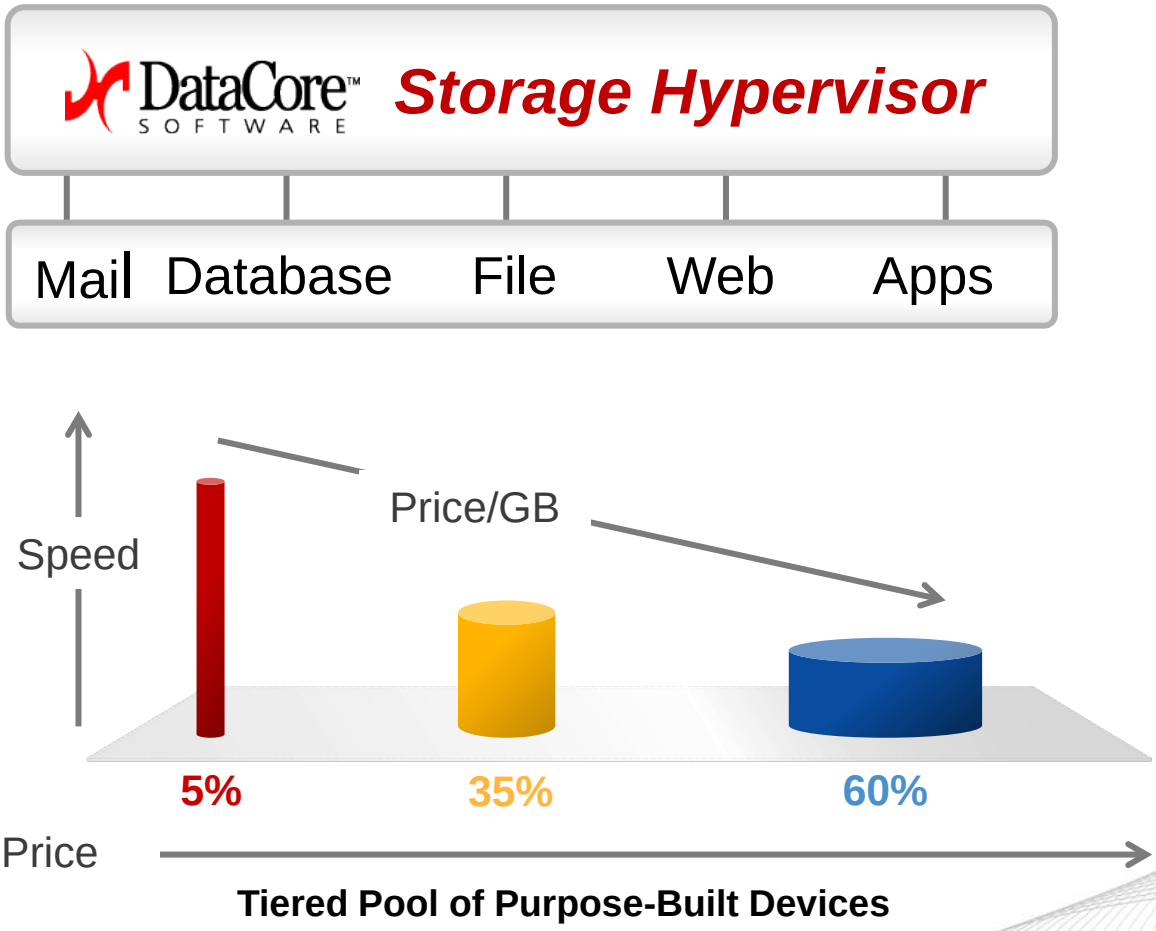

DYNAMICALLY OPTIMIZES RESOURCES

- Based on average access frequency for each tier

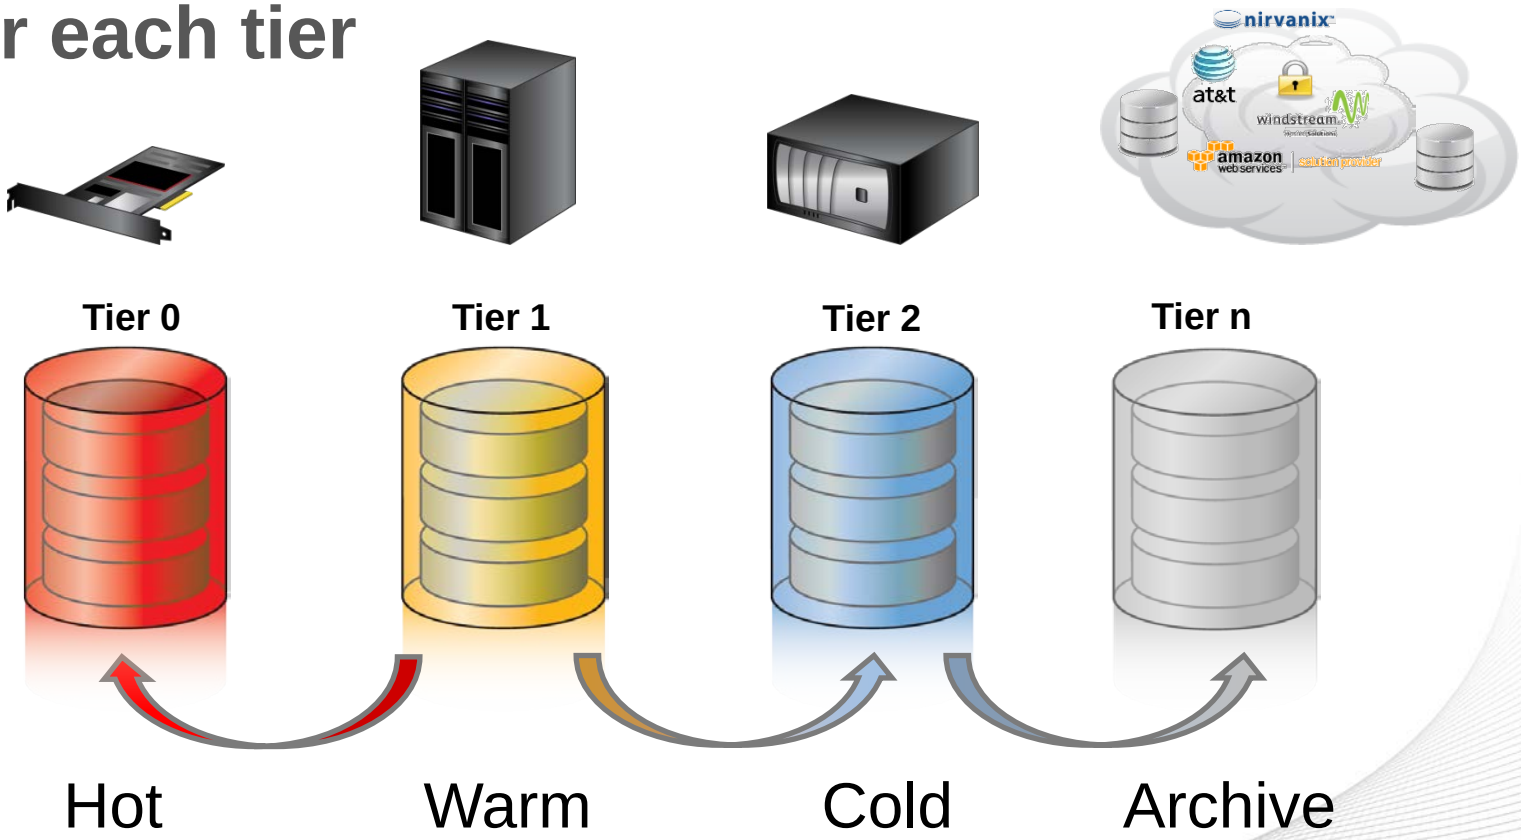

SUB-LUN AUTO-TIERING

- Heaviest hit blocks take advantage of Tier 1 performance
- Other blocks moved to Tier 2 & Tier 3

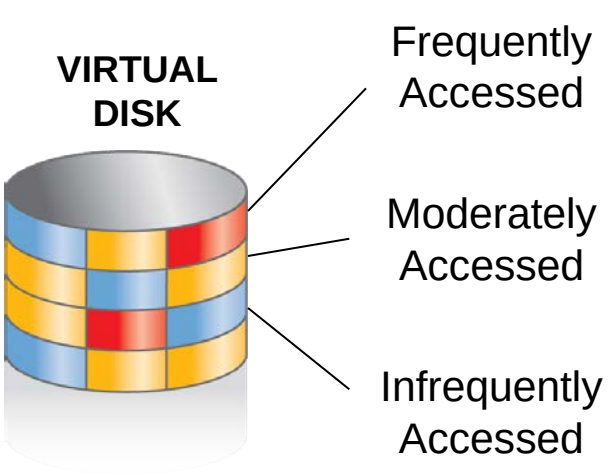

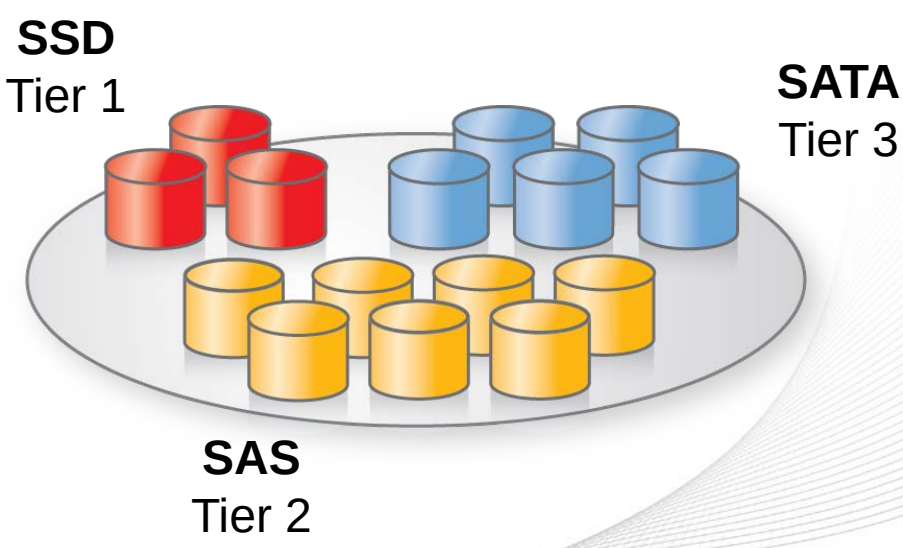

AUTOMATIC FINE TUNING

- Heat maps reveal hot spots of heavy disk activity
  - **SANsymphony-V automatically migrates disk blocks to less active drives**

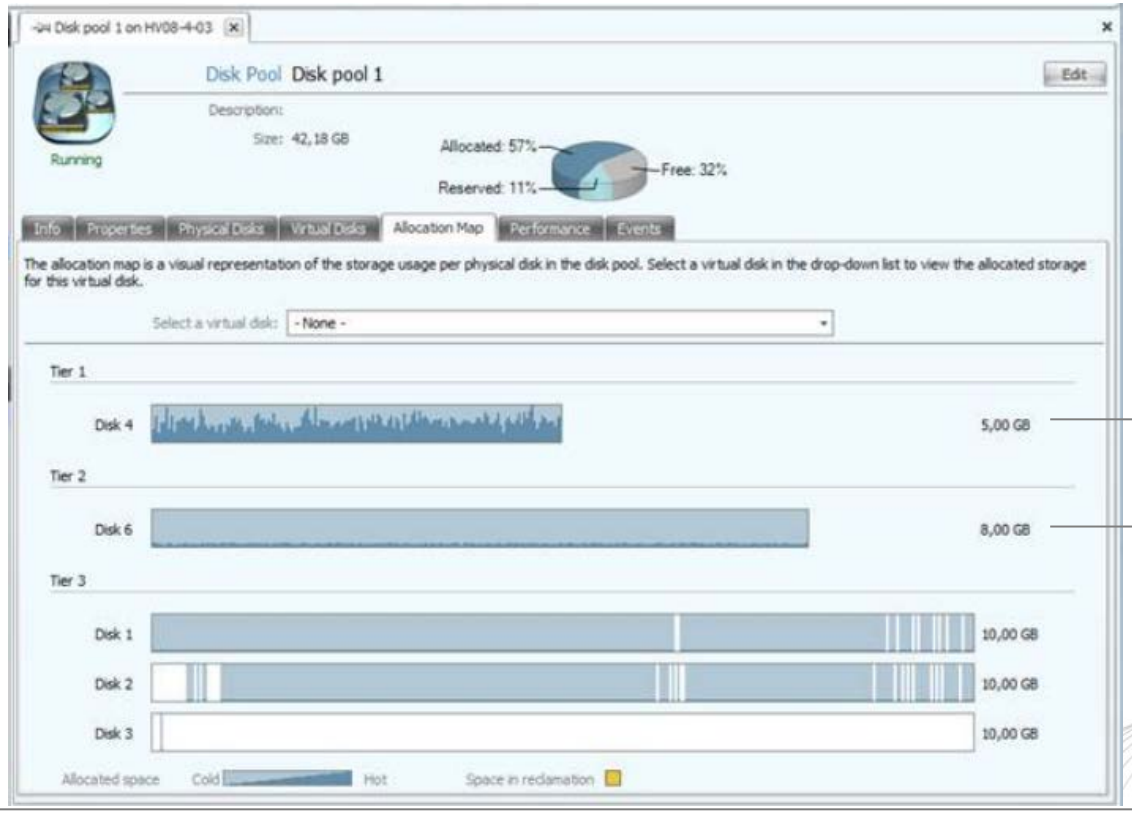

High activity

Low activity

CENTRAL MANAGEMENT

- Common menus & commands
  - Across diverse equipment

- Single console
  - Status dashboard
  - Performance view
  - Consolidated notifications & alerts

- External integration point
  - Virtualization

SAME FEATURES SPAN MUTUALLY INCOMPATIBLE DEVICES

NO NEED TO BE A STORAGE EXPERT

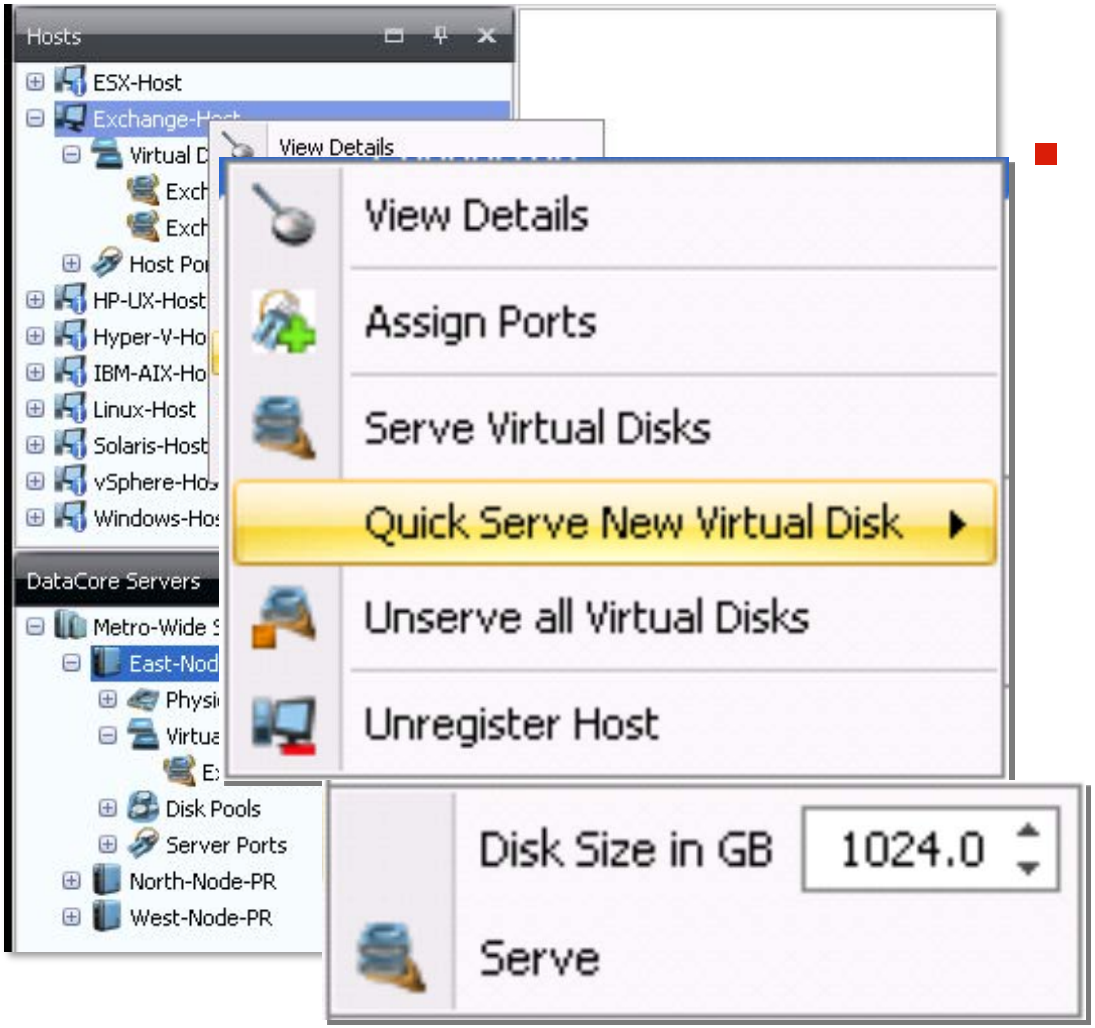

- **Simplify & automate best practices**
  - **Auto-discovers available space**
  - **Thin provisions**
  - **Selects best path**
  - **Sets adaptive cache**
  - **Creates HA mirrors**
  - **Load balances back-end ports**

DEPLOYMENT SCENARIOS

- "Drop-in" on app servers
  a) **Hyper-V parent partition, or**
  b) **Virtual Machine (VM)**

- Install on external "nodes"
  a) **New dedicated servers, or**
  b) **Repurposed app servers**

- Then scale up and scale out as needed

ADAPTIVE STORAGE IaaS LAYER

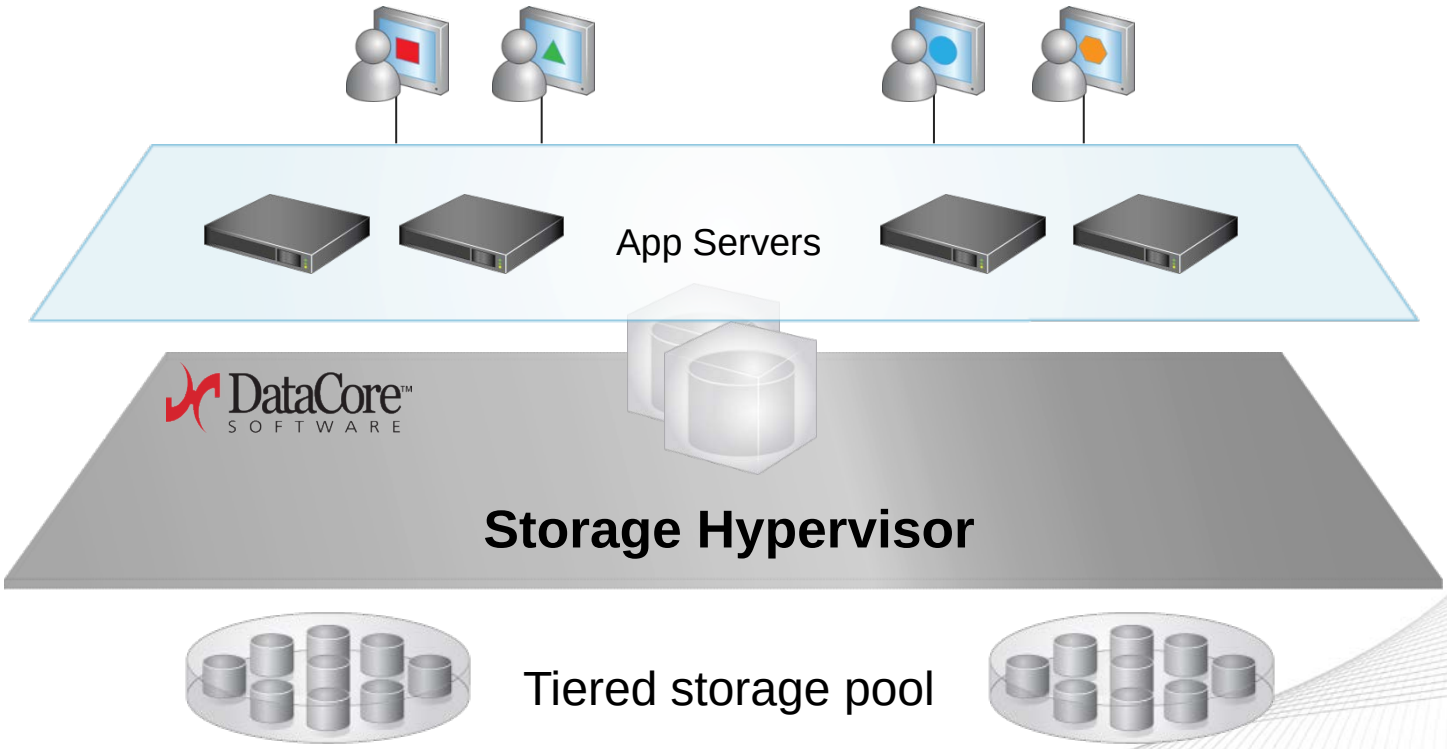

CONFIGURED EXTERNALLY

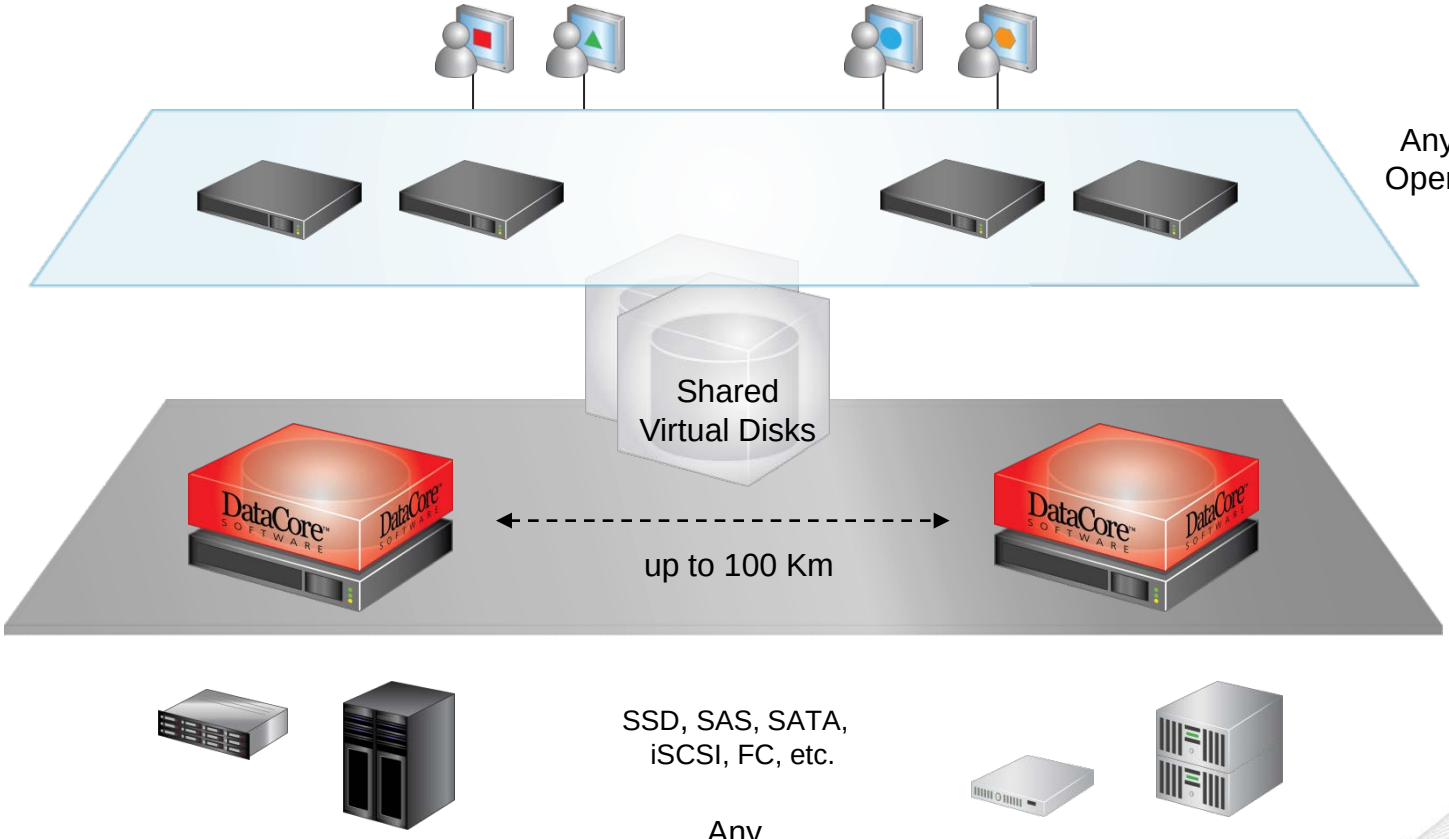

RESIDENT IN APP SERVERS

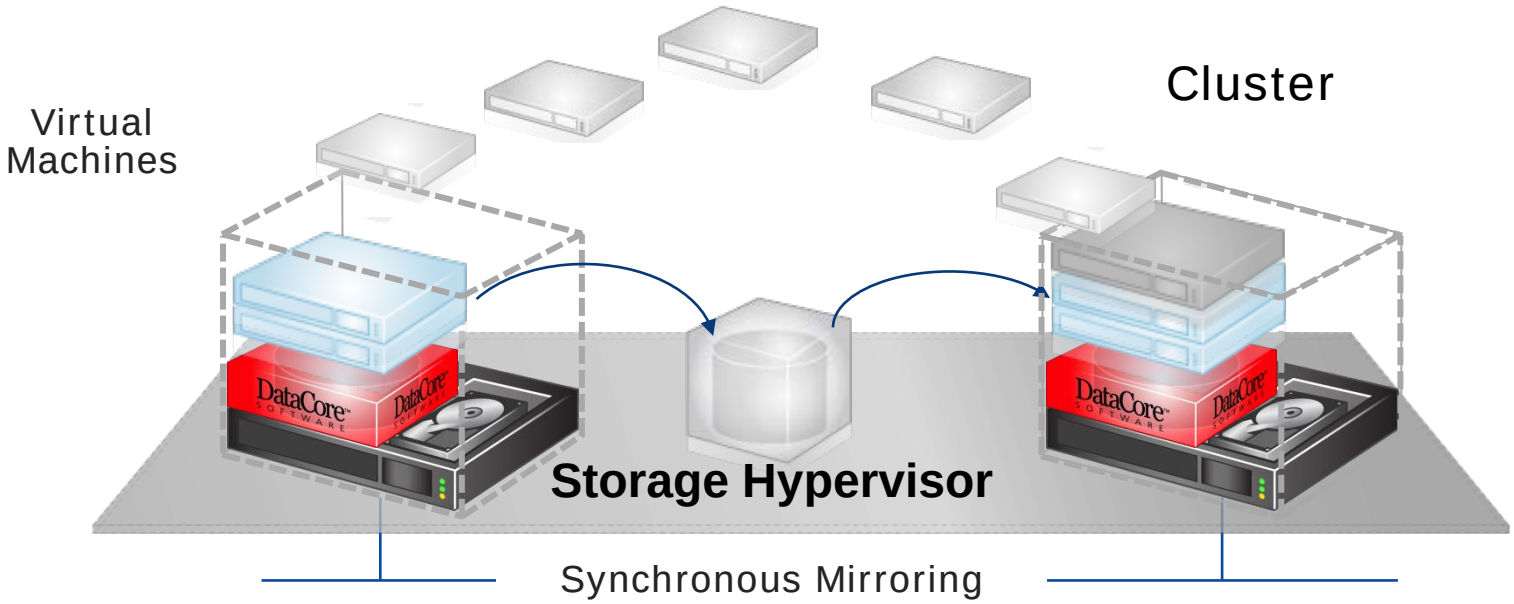

- Leverage server internal drives & SSD / Flash
- No SAN storage required to get started
- Expand outboard later if outgrown

FULLY REDUNDANT ARCHITECTURE

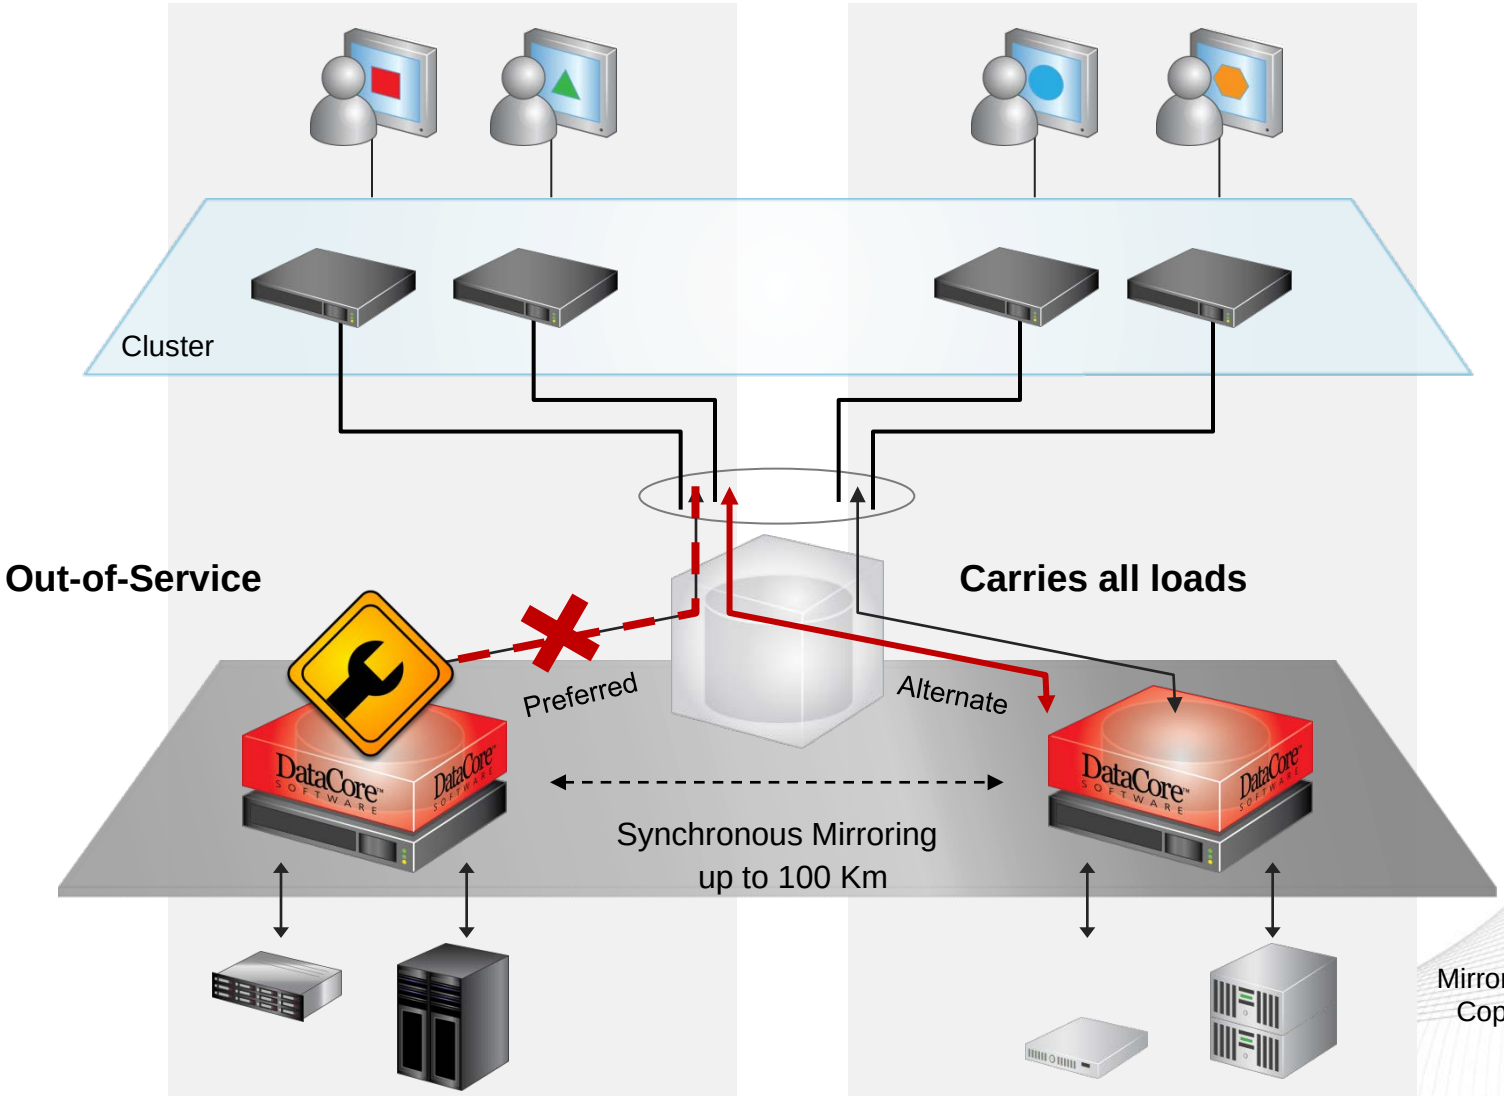

AUTO-RESTORES ORIGINAL CONDITIONS

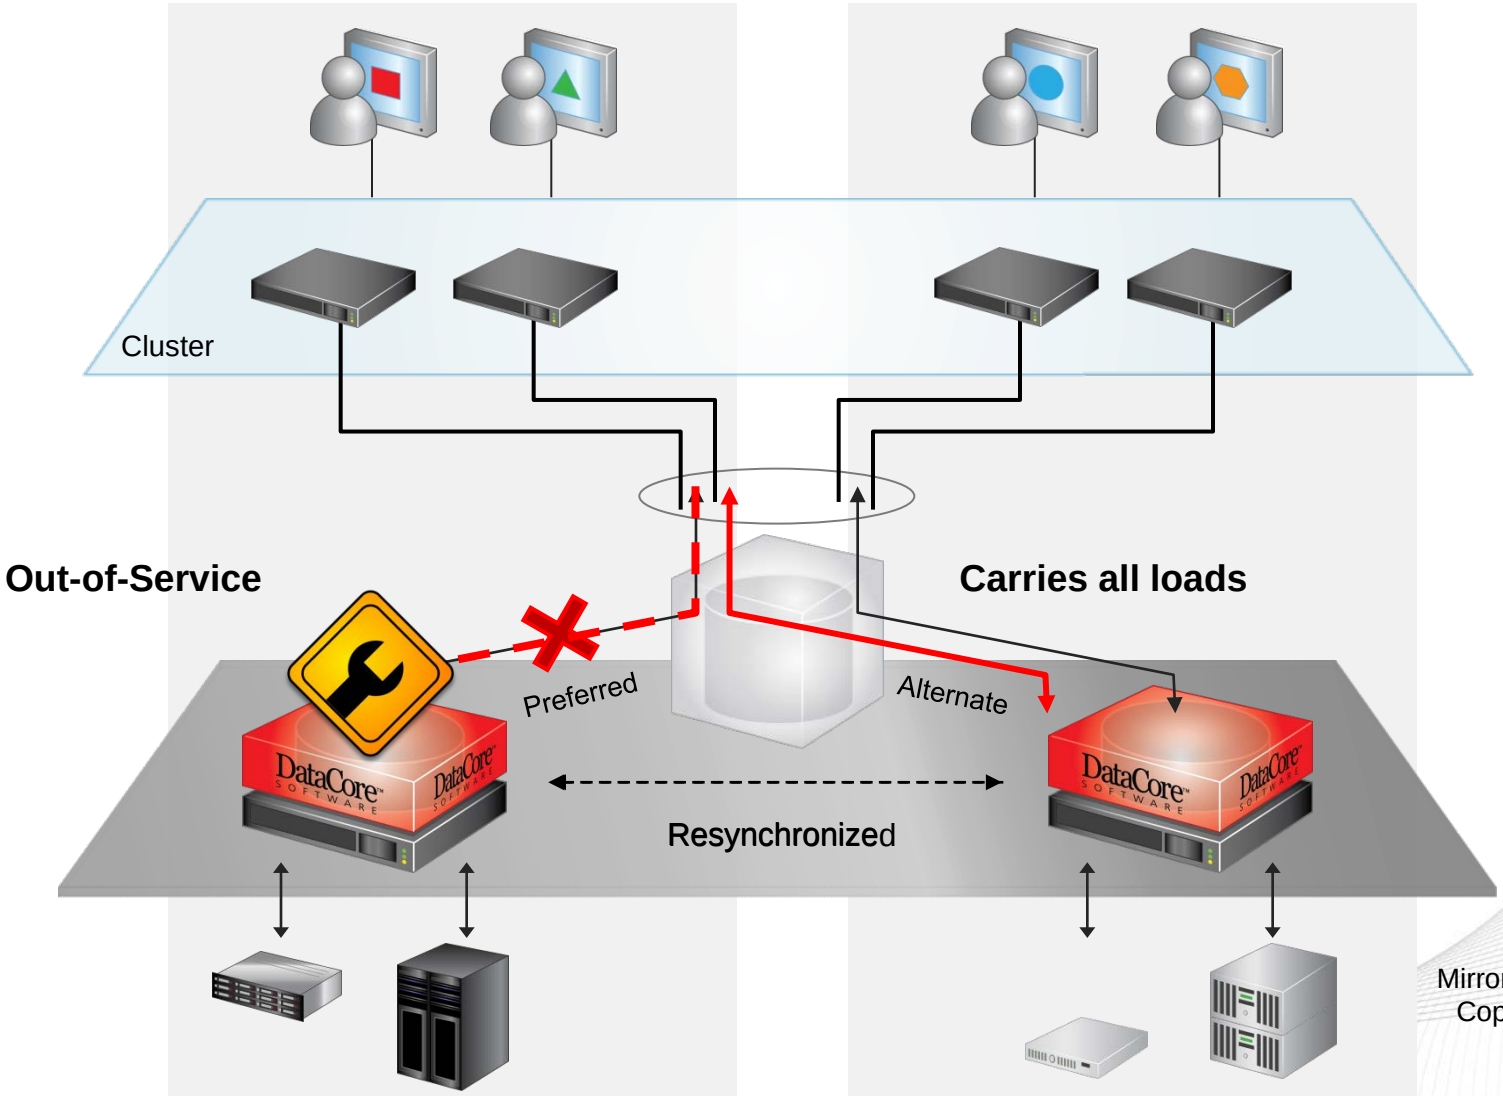

STORAGE HYPERVISOR
FOR DATA CENTERS OF ALL SIZES

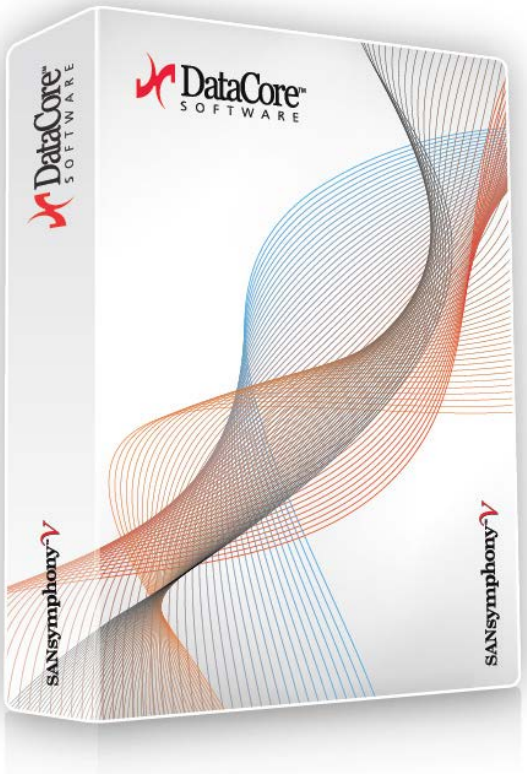

Large Enterprises

Cloud Providers

Midsize Corporations

Smaller Businesses

HEALTHCARE EXAMPLE

NewYork-Presbyterian
The University Hospitals of Columbia and Cornell

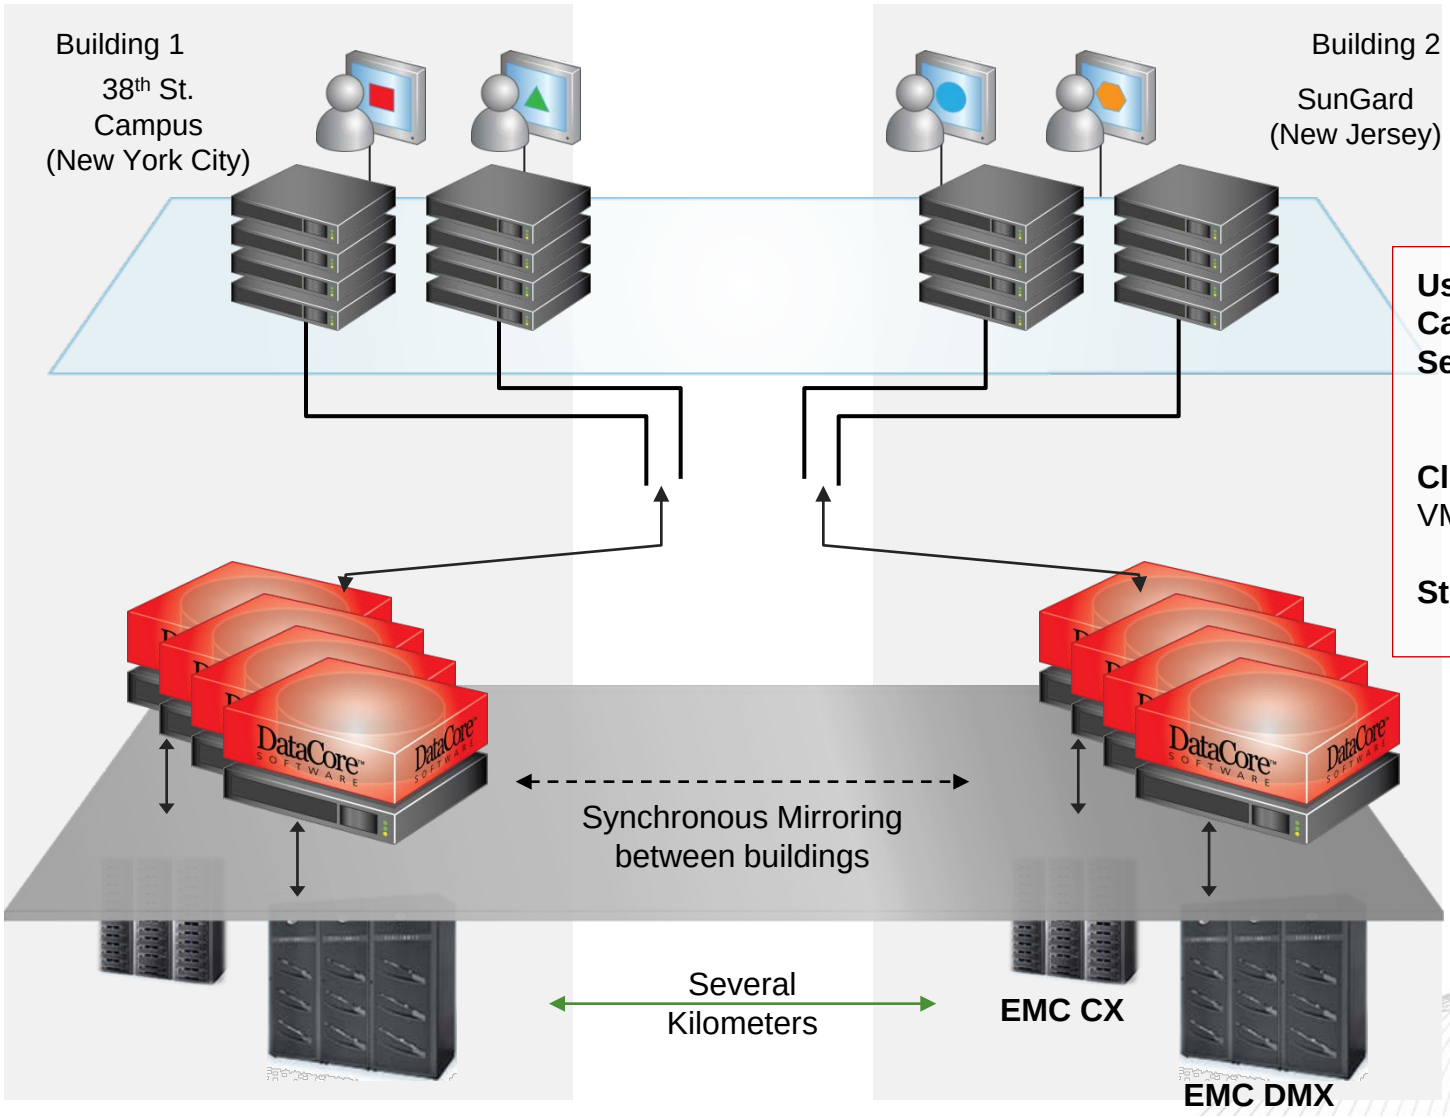

Users: 5,000 doctors
Capacity: >500 TB mirrored
Servers: 300+ Windows, Solaris, AIX servers

Clusters: MSCS, HACMP. VMware

Storage: EMC DMX, CX, VNX

HIGH-SPEED CACHING

Speeds up performance

- **Accelerates disk I/O response from existing storage**
- **Uses x86-64 CPUs and memory from DataCore nodes as powerful, inexpensive “mega caches”**
- **Anticipates next blocks to be read, and groups writes to avoid waiting on disks**

TURBO-CHARGED I/O

Avg. Response (*µS*)	
DataCore L1 Cache	*<20*
Array L2 Cache	*250 – 300*
Disk Drive	*4000 – 6000*

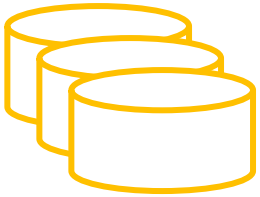

Virtual Disks

Up to 1 TB RAM cache per node

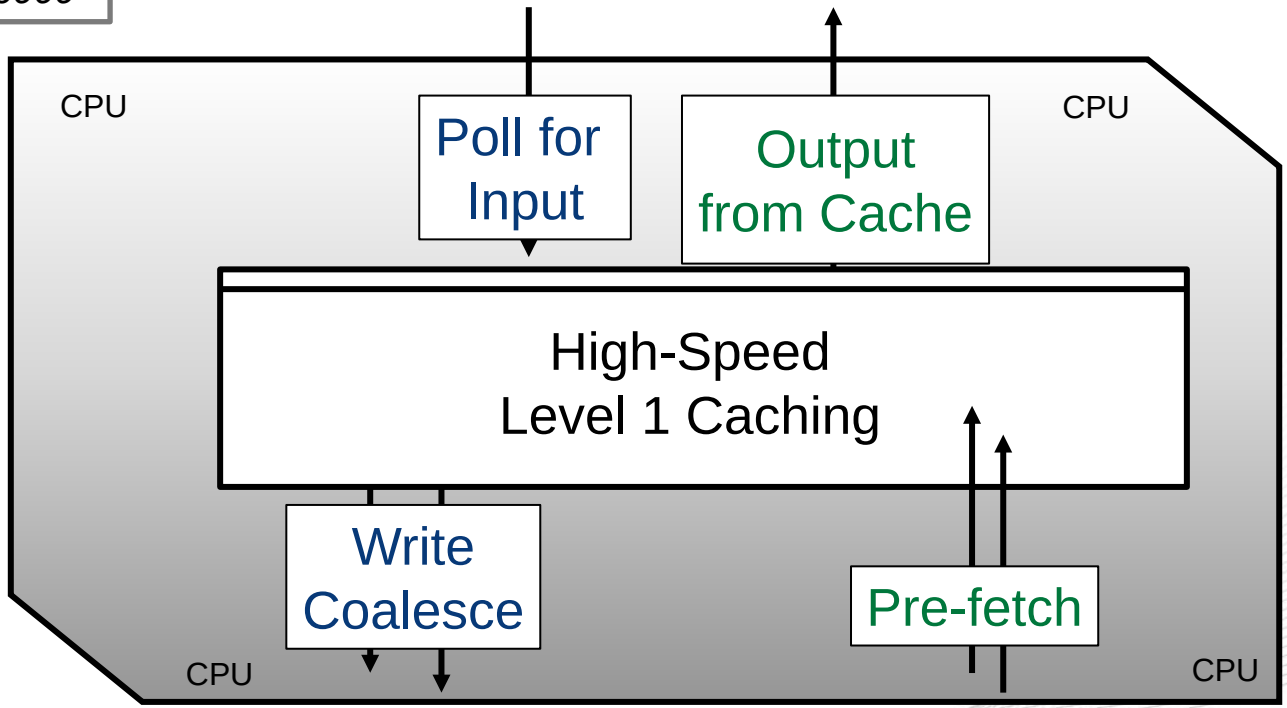

disk drives

VIRTUAL DISK MIGRATION

Move contents from one disk to another without impacting applications

- **Eliminates downtime for hardware disk upgrades**
- **Clears & reclaims space occupied by original**
- **Provides pass-through access to drives previously connected directly to hosts**

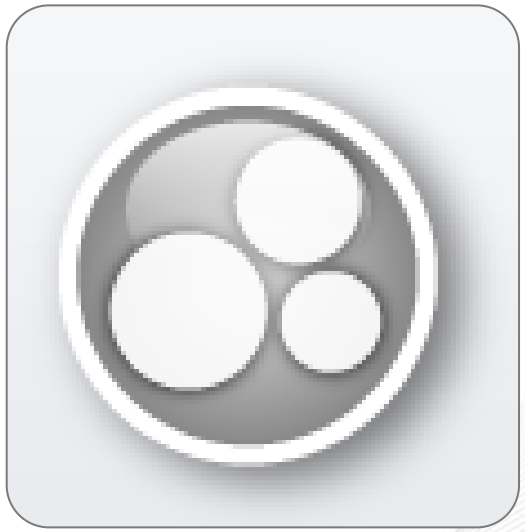

NON-DISRUPTIVE DISK UPGRADES

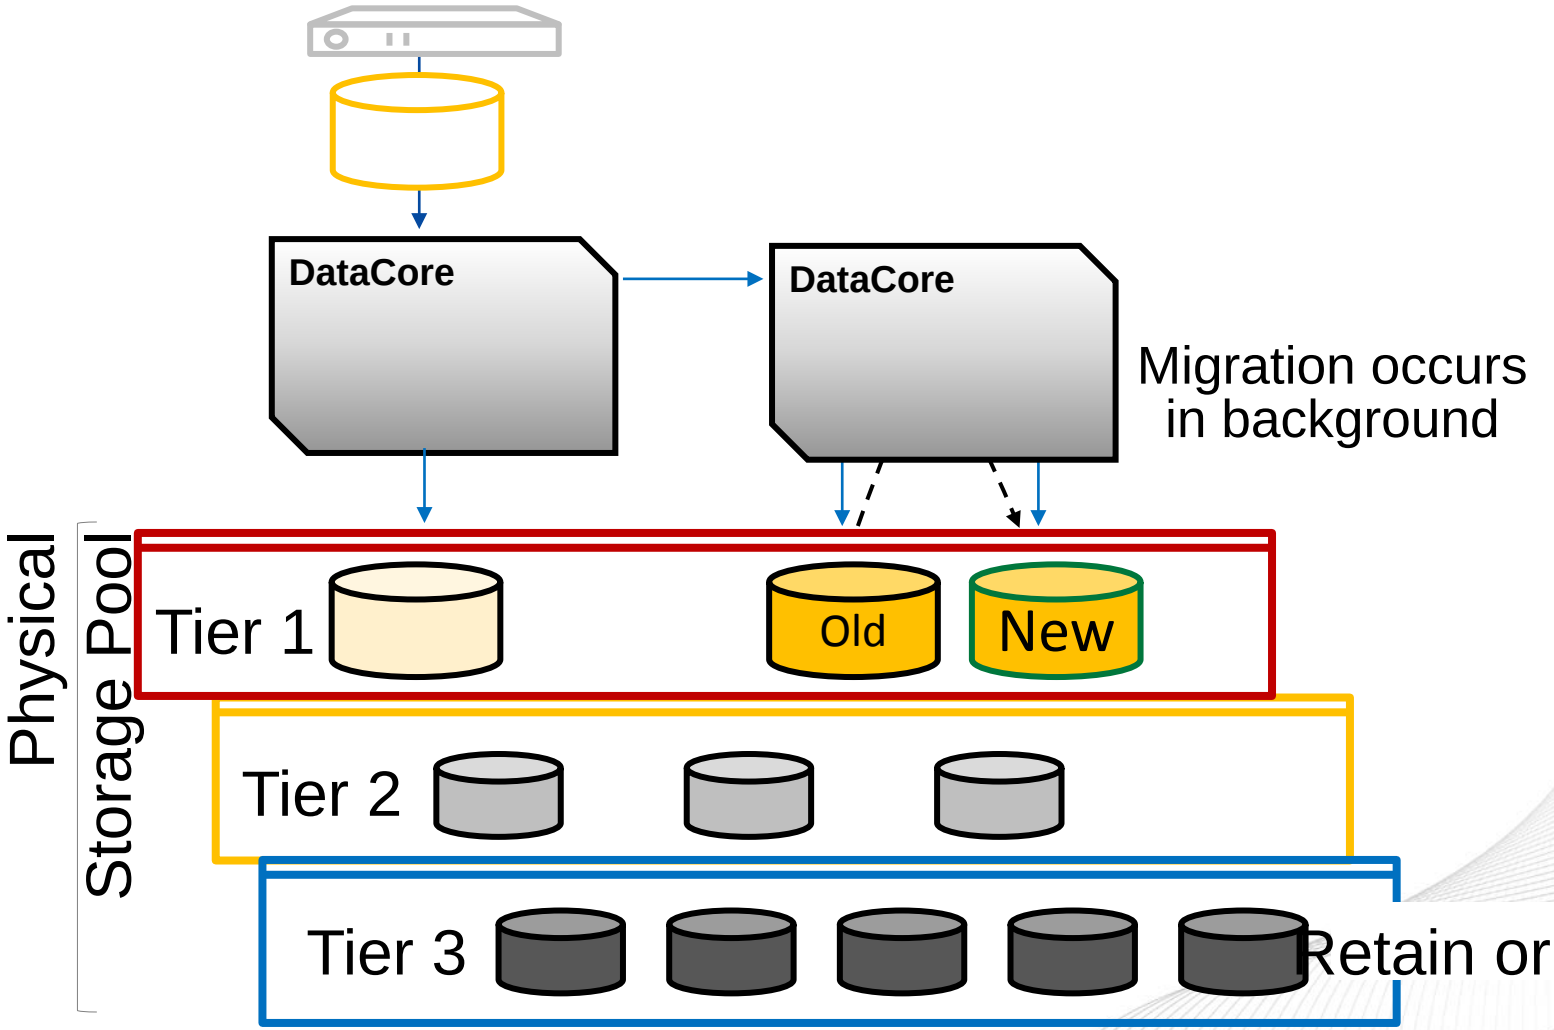

BENEFITS

- Maximize combined value of storage assets
- Take cost out of storage infrastructure
- Simplify management & migration
- Overcome device-specific differences
- Provide higher levels of service

Customers report up to:

- 60% COST SAVINGS
- 90% UTILIZATION
- 100% UPTIME
- 200% FASTER

DataCore™ SOFTWARE

The Storage Hypervisor Company

www.datacore.com