



Automated Storage Tiering

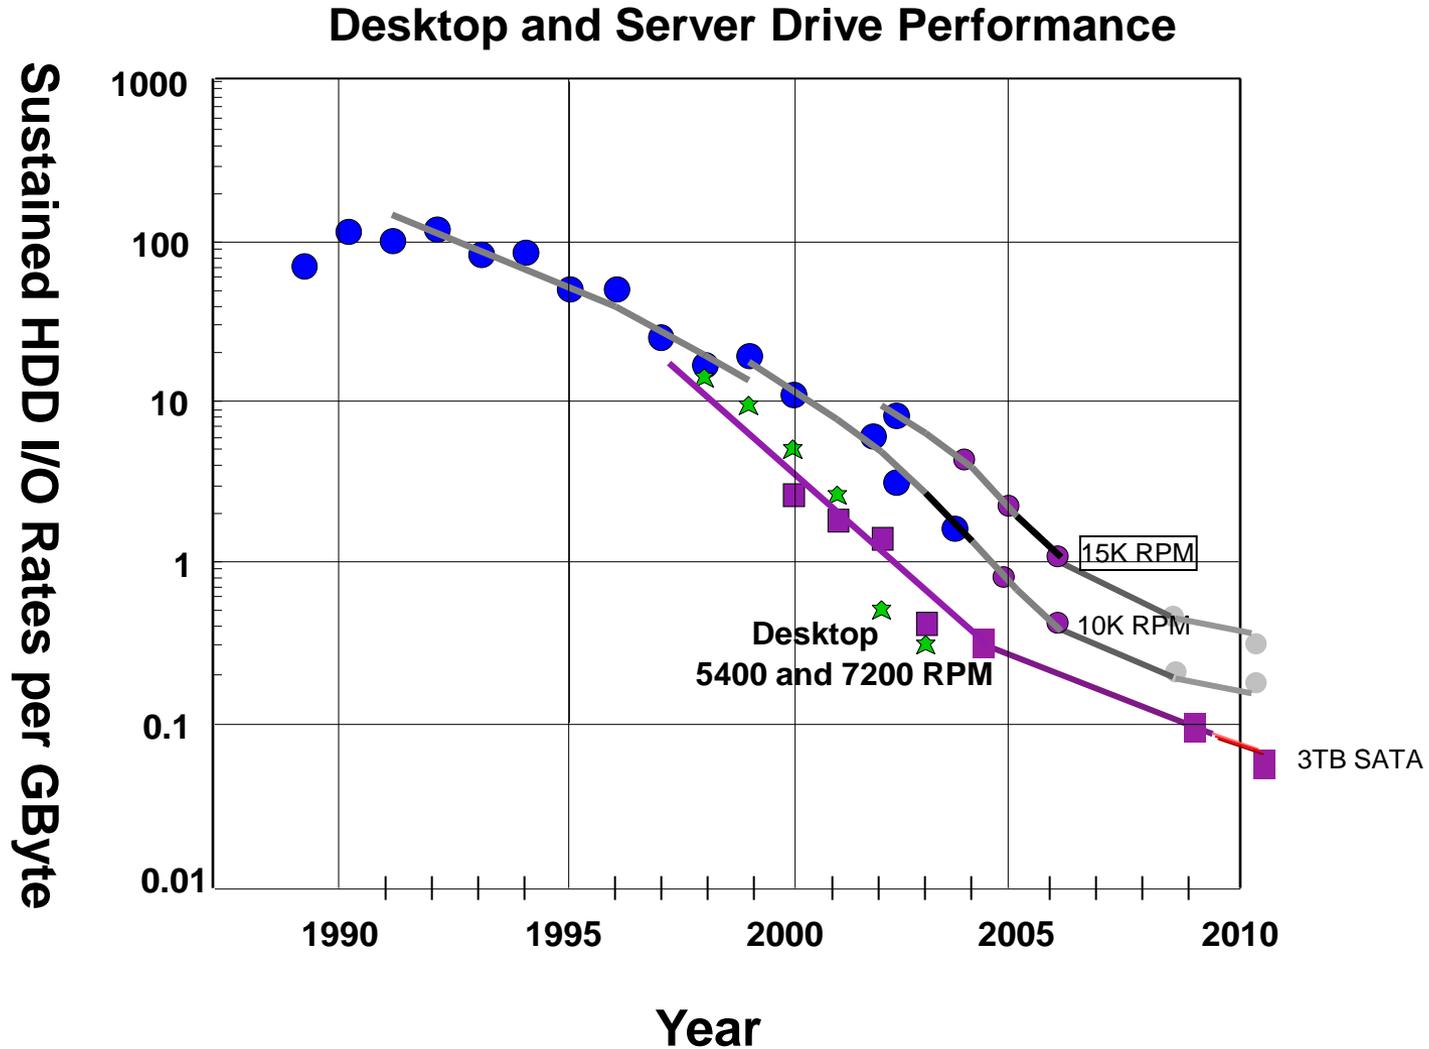
August 2012

Clodoaldo Barrera
Chief Technical Strategist – IBM Storage Systems
barrerac@us.ibm.com

Storage Tiering – Why do we need it?

- **The demands of high I/O Rate workloads and high capacity requirements stress the economics of storage**
- **We have technologies good for IOPs, and technologies good for \$/Gbyte (but not both).**
 - We need to address a mix of workloads
 - And the workloads change over time
- **Static assignment of data to a storage type, and human managed placement will not address the dynamic need, or the scale and complexity.**
- **We need multiple tiers of storage and automated data placement and movement**

IO Performance / Capacity losing ground



Overview of Easy Tier

- **Easy Tier provides automatic volume management capabilities**
 - Sub-volume drive tiering - Data movement and placement based on historical data access patterns
 - Automated drive utilization balancing to remove hot spots or populate new, empty ranks
 - No manual input required to optimize performance as workloads change
- **Easy Tier planning and management**
 - Storage Tier Advisor Tool (STAT) for I/O analysis and projected benefits
 - A command line interface and a GUI for setup and management
 - Manual command based volume relocation, extent pool merge, rank depopulation
- **Easy Tier is a licensed feature for the DS8700/DS8800**
 - No charge – but you still have to order the feature and activate the license
 - Supported by all server platforms with no additional software or host configuration

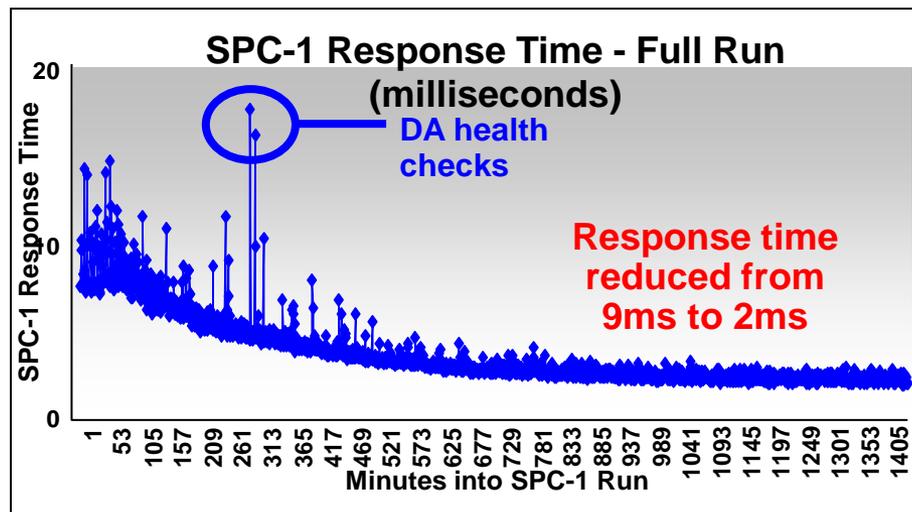
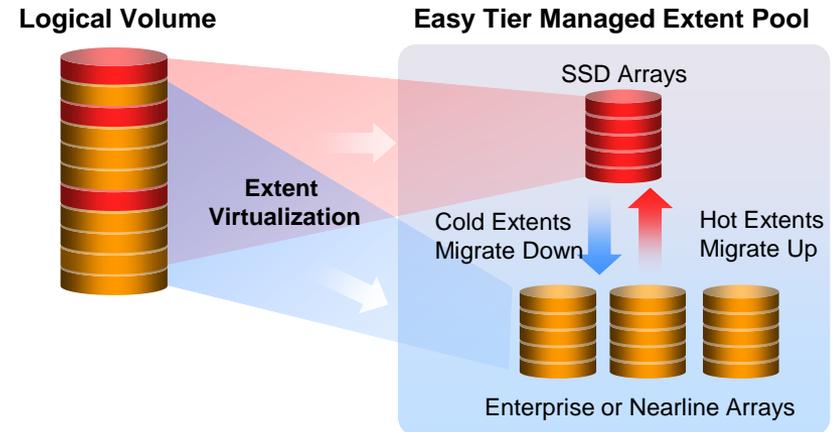
DS8000 Easy Tier at a Glance

- Monitors real time performance of each 1GiB extent (sub-volume) to determine the data 'temperature'
- Creates extent migration plan for optimal data placement every 24 hours based on performance statistics
- Migrates extents one tier at a time across heterogeneous tiers and within a homogeneous tier according to migration plan(s) over 24 hour period
- A limited number of extents are chosen for migration every 5 minutes to avoid performance impact
- **DS8000 value = better performance optimization and ease of use**



Easy Tier Benefits

- Relocating just 5% of the data from HDDs to SSDs, reduced average I/O response time from 9ms to 2ms!
- Optimal use of expensive SSDs

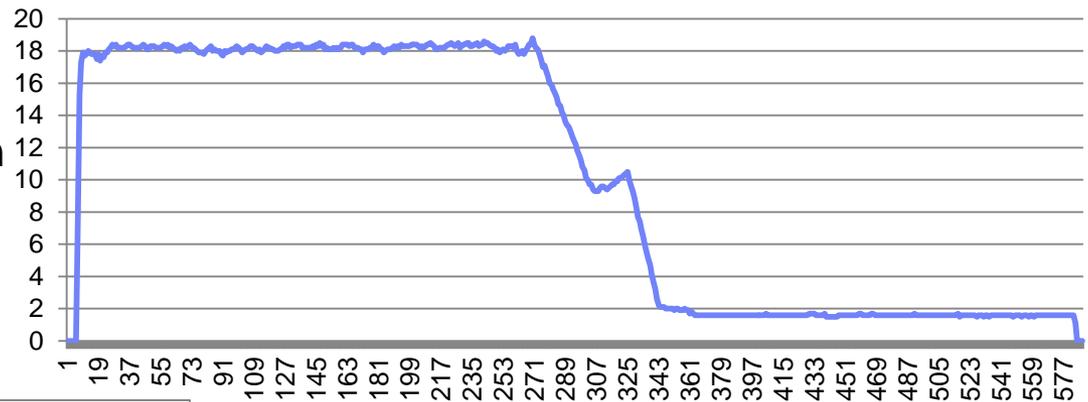


Source: Storage Performance Council, April 2010: http://www.storageperformance.org/results/benchmark_results_spc1#a00092

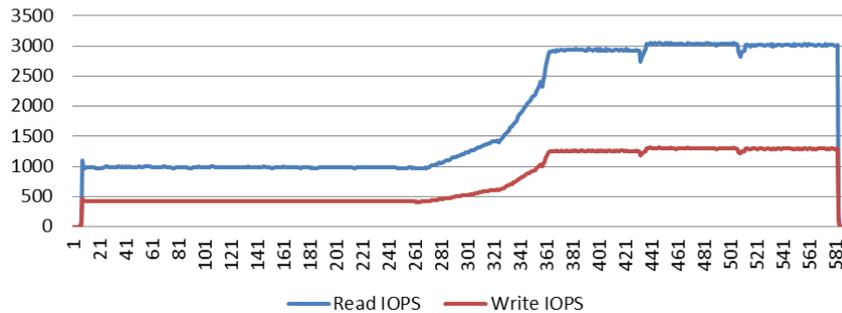
DS8000 Easy Tier Impact: Large Banking Customer

- **Effects of Easy Tier (7.8 % SSD Capacity added, 1.5 TB SSD to 21 TB SAS Pool)**
- **Read Response Time reduced by 91 % (18 ms → 1.6 ms)**
- **Read IOPS +597 %**
- **Write IOPS +30 %**
- **Traffic Profile**
R:W 70:30, 8K blocks, random

Test AIX002
(589 * 5 Minutes = 49h)



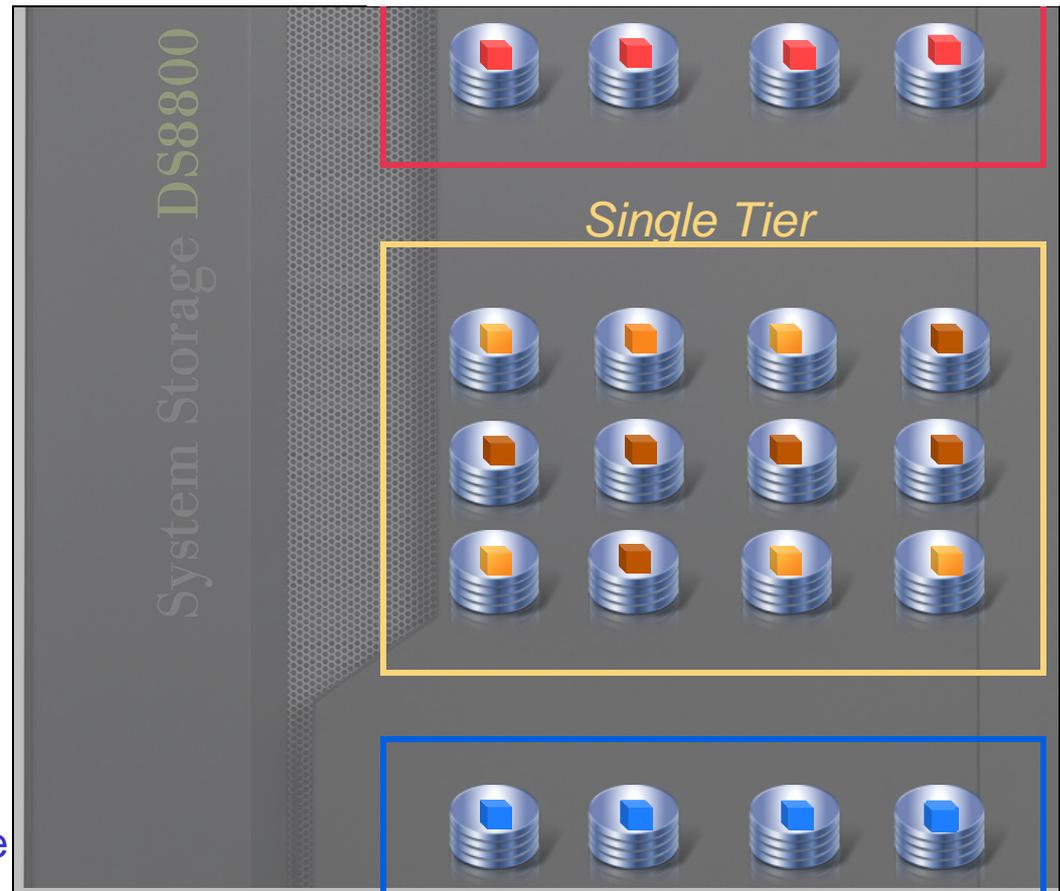
Test AIX002
(589 * 5 Minutes = 49h)



— Read Response Time (ms)

Easy Tier intra-tier Rebalancing

- Continuous rebalancing *within* each tier maintains peak performance across all drives
- Automatic rebalance of data across all ranks when additional capacity added to the pool
- Easy Tier complements other performance optimization features
 - Advanced caching algorithms
 - Storage Pool Striping (wide striping)
 - I/O Priority Manager for QoS
- DS8000 value = automatic performance balance and ease of use

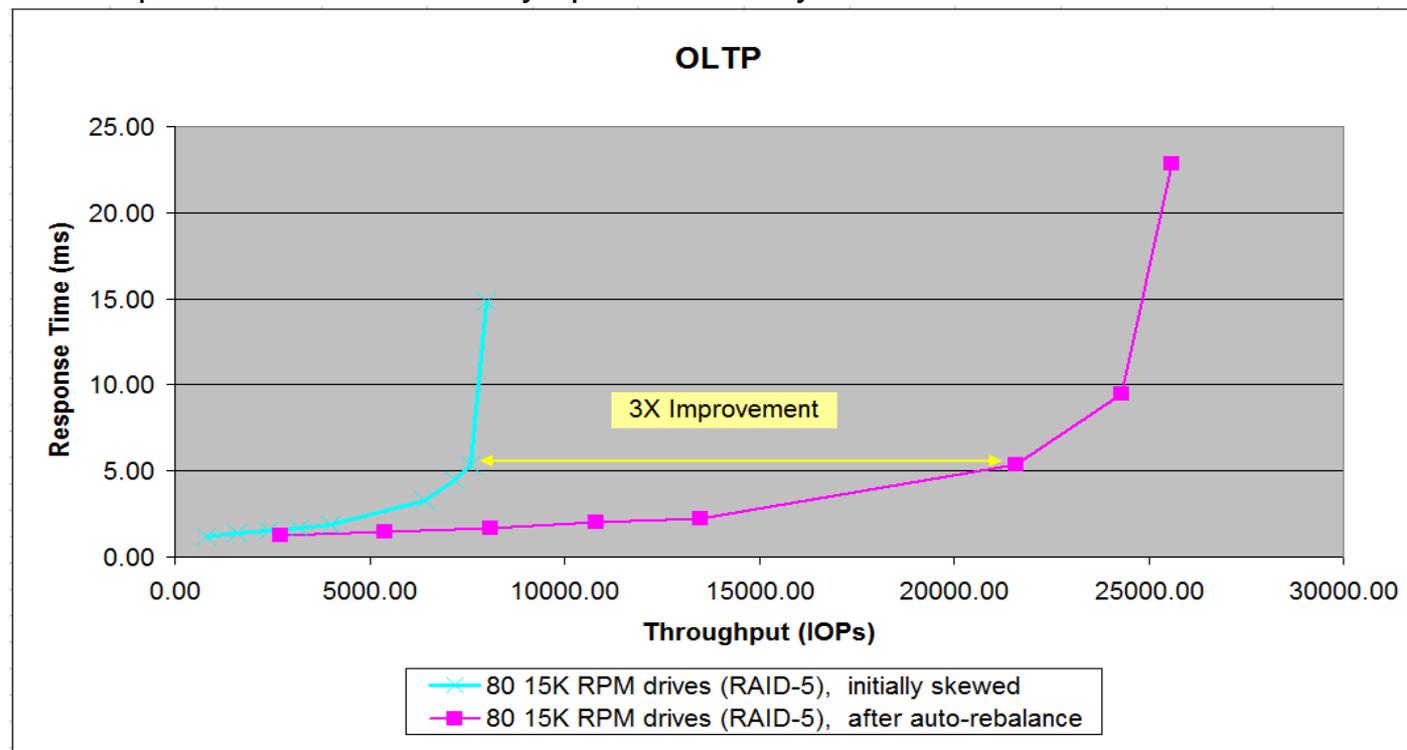


Easy Tier intra-tier Rebalancing – Improved Throughput

Supports single-tier and multi-tier systems

Performance gain on single-tier pool

- The re-balanced system showed a **3x** throughput improvement and equal or better response times without any operator activity!



All performance data contained here were obtained in the specific operating environment and under the conditions and/or configuration described above and is presented as an illustration. Details are available in the "IBM System Storage DS8800 and DS8700 Performance with Easy Tier 3rd Generation" whitepaper. Performance obtained in other operating environments may vary and customers should conduct their own testing

Easy Tier Supported Versions

- **Easy Tier 1 (DS8700 R5.1)**
 - Automated cross-tier performance management for *SSD/HDD hybrid pools*
 - Manual mode management support for *dynamic extent pool merge* and *dynamic volume relocation*
- **Easy Tier 2 (DS8700/DS8800 R6.1)**
 - Automated cross-tier performance or storage economics management for hybrid pools with *any 2 tiers* (SSD/ENT, SSD/NL or ENT/NL)
 - Automated *intra-tier performance management* (auto-rebalance) in hybrid pools
 - Manual mode management support for *rank depopulation* and optimized *volume restriping* within non-managed pools (manual volume rebalance)
- **Easy Tier 3 (DS8700/DS8800 R6.2)**
 - Automated cross-tier performance and storage economics management for hybrid pools with *3 tiers* (SSD/ENT/NL)
 - Automated *intra-tier performance management* in both hybrid (multi-tier) as well as *homogenous (single tier) pools (auto-rebalance)*
 - *Thin Provisioning support* for Extent Space Efficient (ESE) Volumes
- **Easy Tier 4 (DS8800 R6.3)**
 - Support for *encryption capable environments*
 - DS8800 has FDE capable SSD and nearline

Storage Tier Advisor Tool (STAT)

- System Summary Report
 - Support pool configurations with up *3 tiers* (SSD – ENT – NL)
 - Tier status: *Existing Tier SSD/ENT/NL is IOPS/BW overloaded/skewed*

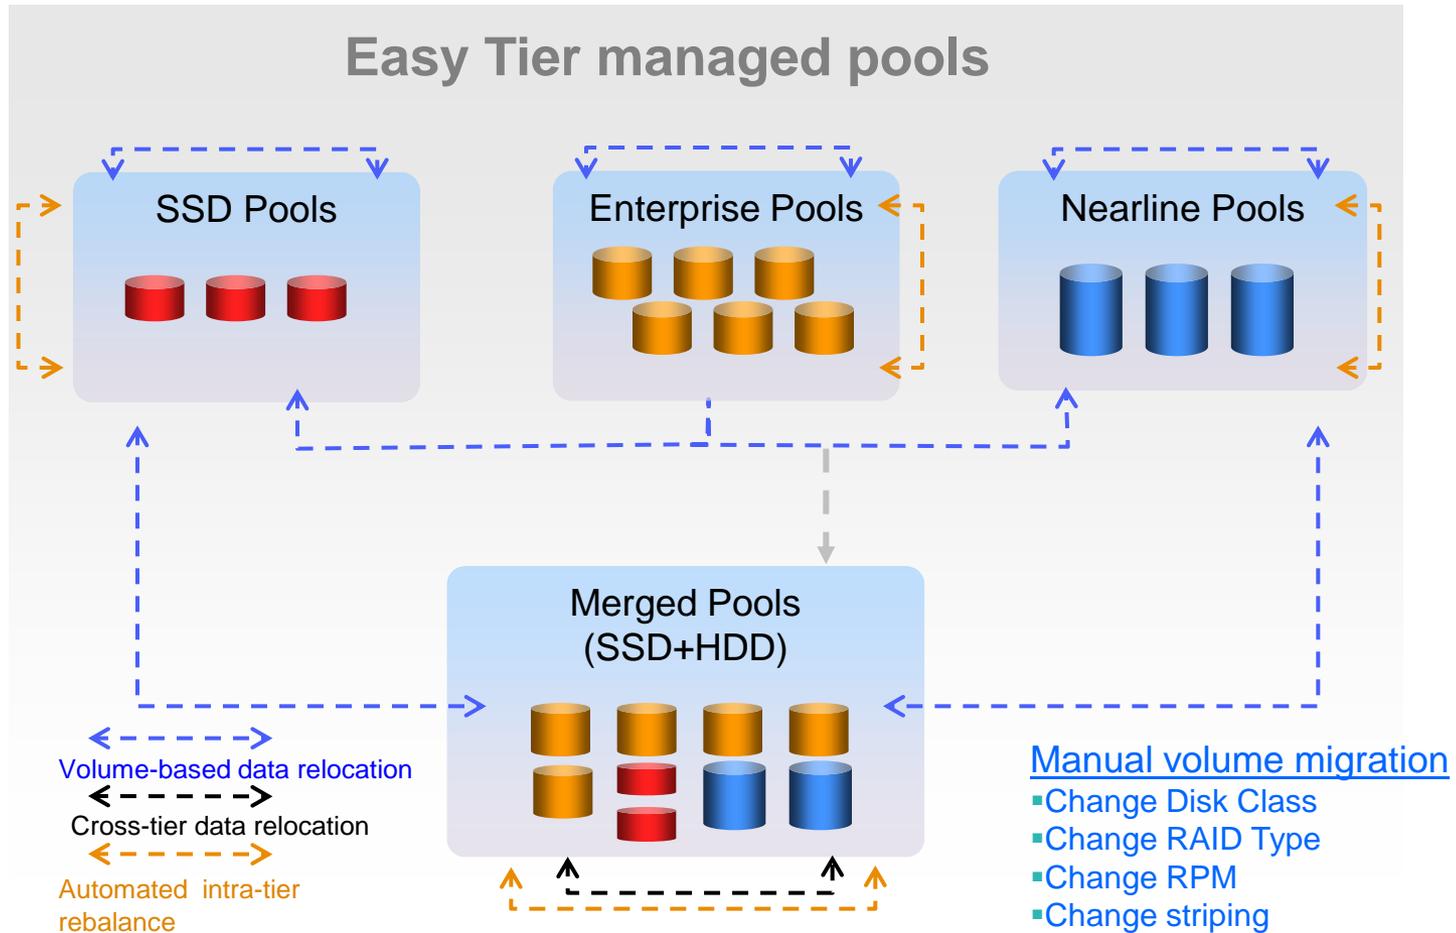
- Storage Pool Performance Statistics and Improvement Recommendation
 - *Existing Tier Status* showing rank utilization for each tier
 - *Recommended SSD, Enterprise, NL Configurations*
 - Expanding 1-tier HDD and any 2-tier configurations to 3-tier configurations
 - Expanding existing tiers when tiers are BW or IOPS overloaded
 - *Volume Heat Distribution*
 - Showing capacity and heat distribution by tier (hot data - warm data - cold data)

- Systemwide Recommendation
 - *Recommended SSD Configuration*
 - *Recommended Enterprise Configuration*
 - *Recommended NL Configuration*

Download STAT (no charge) and run on Windows

<https://www-304.ibm.com/support/docview.wss?uid=ssg1S4000876&wv=1>

Summary of Easy Tier Migration Capabilities



DS8000 Value

- **Easy Tier**
 - Enables the consolidation of more workloads onto the DS8000 by automating storage performance and economics of management across all 3 drive tiers
 - Security for data at rest with encrypting drives
 - Automatic rebalancing of extents for consistent performance
 - Manual Mode provides significant functional capabilities to storage administrators
- **Storage Tier Advisor Tool (STAT)**
 - Provides an understanding of what application data can benefit the most from relocation to SSDs, Enterprise (SAS/FC) drives or Nearline SAS drives
- **I/O Priority Manager (QoS)**
 - Provides tremendous scalability to manage more applications effectively and efficiently with a single DS8000 system
- **Ease of Use**
 - Common management across IBM storage portfolio
 - Simplified logical configuration setup and management



References

- **IBM System Storage DS8000 Easy Tier Redpaper**

REDP-4667

<http://www.redbooks.ibm.com/abstracts/redp4667.html?Open>

- **IBM System Storage DS8000: Architecture and Implementation**

SG24-8886

<http://www.redbooks.ibm.com/abstracts/sg248886.html?Open>

- **IBM System Storage DS8800 and DS8700 Performance with Easy Tier 3rd Generation**

WP102024

<http://tucgsa.ibm.com//gsa/home/s/s/sspadmin/web/public/docs/disk/white%20papers/Easy%20Tier%203rd%20Gen%20Whitepaper.pdf>

Trademarks and Disclaimers

IBM Corporation 1994-2012 All rights reserved.

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

Trademarks of International Business Machines Corporation in the United States, other countries, or both can be found on the World Wide Web at <http://www.ibm.com/legal/copytrade.shtml>.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Information is provided "AS IS" without warranty of any kind.

The customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Some information addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in IBM product announcements. The information is presented here to communicate IBM's current investment and development activities as a good faith effort to help with our customers' future planning.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Photographs shown may be engineering prototypes. Changes may be incorporated in production models.

