September 11, 2017

DXC Storage Provisioning

Performance Aware Storage Provisioning Application: Auto Tiering

Chuck Paridon
Master Storage Performance Architect
DXC Technology

DXC.technology

Implementing Auto Tiering storage configuration using PD

Topics

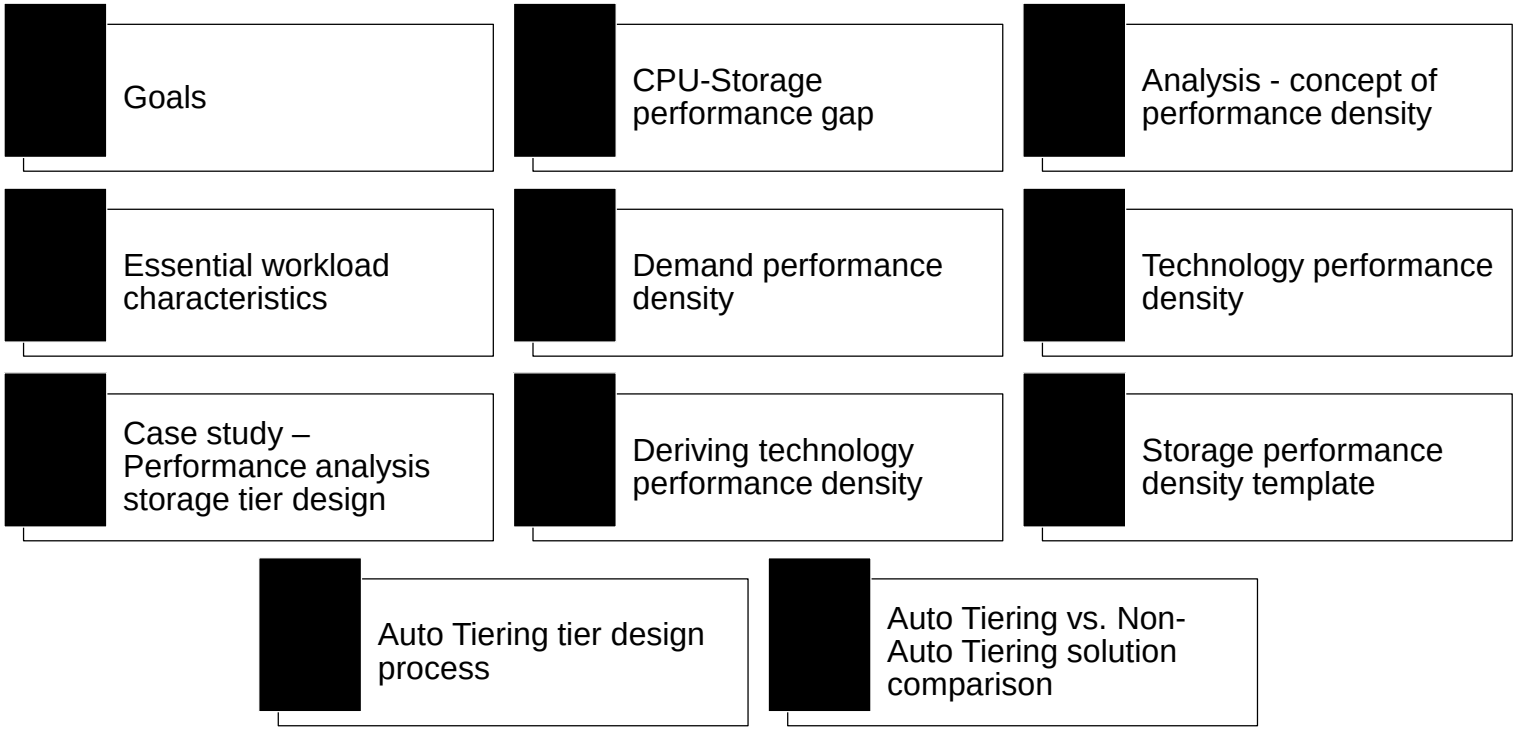

Goals

Monitor workloads to assess their fitness of optimization using Auto Tiering

- The two parameters to be derived via this process are the IO demand of the solution and the required capacity. These measurements are then used to derive the required performance density (IOPS/GB)

Selection of the appropriate tier technology to support the measured performance density

- Using a series of templates that are derived from a performance assessment of the various technologies.

Storage capacity growth over time

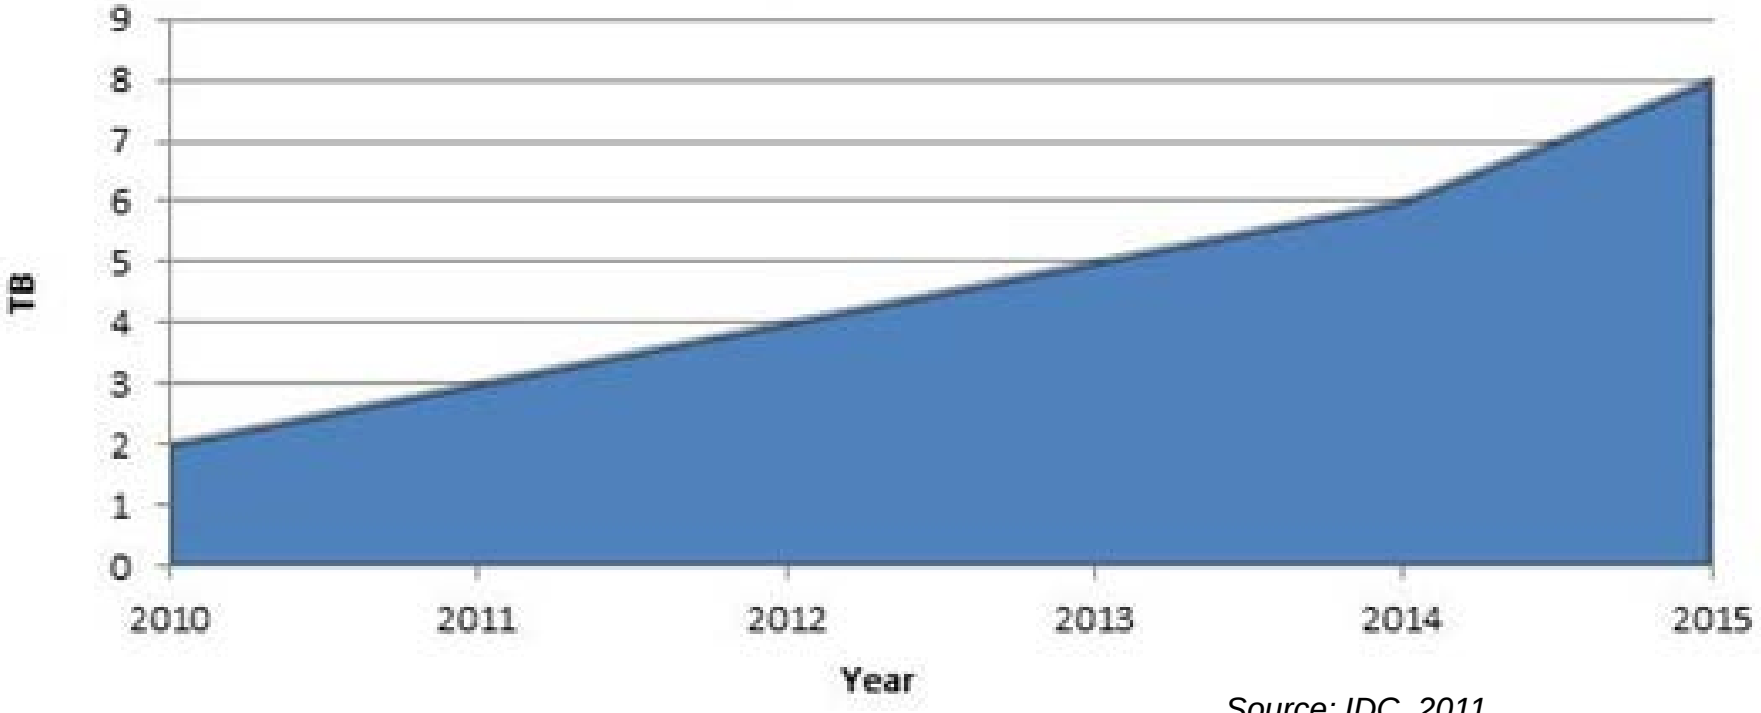

Source: IDC, 2011

Storage performance growth over time

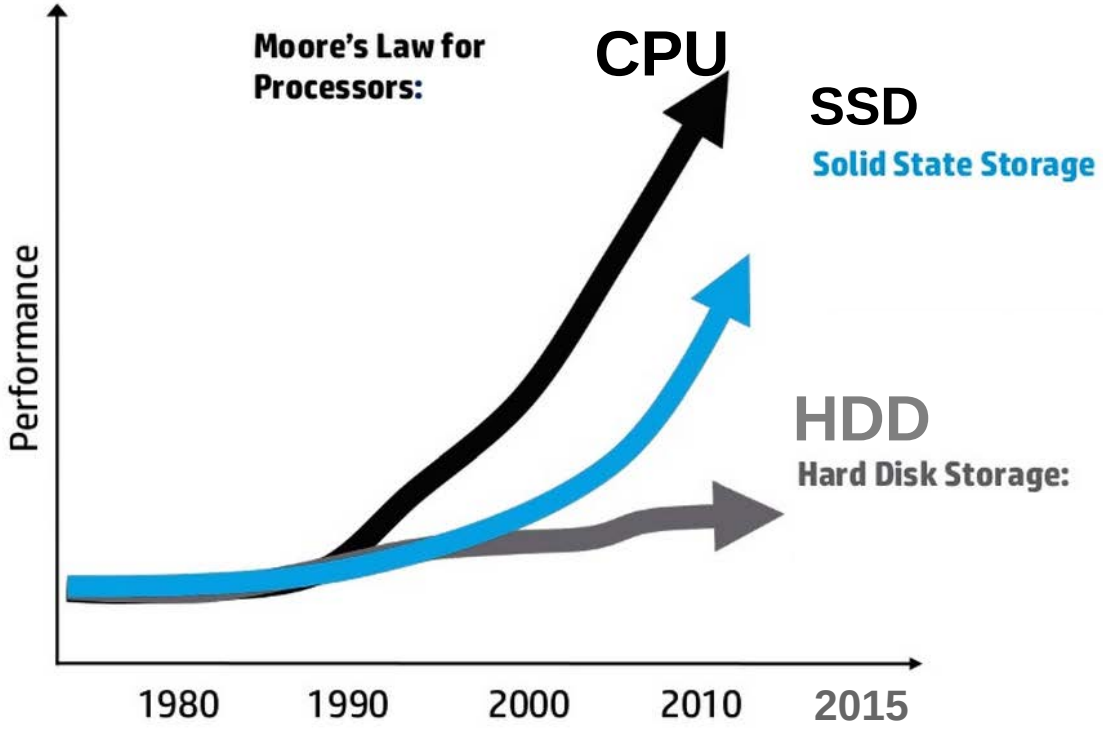

Auto Tiering storage design concept: Performance density

- Past practice has been to design the configuration of the storage based on capacity.
- As storage devices (HDD/SSD) have grown in capacity, but not performance, the performance per GB has dropped dramatically. This has resulted in considerable performance shortfalls for many applications.
- We must adopt an alternate method for sizing that takes performance considerations into account
- Performance Density (IOPS/GB) allows us to size storage systems based on both Capacity and Performance.
- Many tiered storage performance tool statistics are expressed using Performance Density So, why not use the same metric to select the capacity and technology of disks?

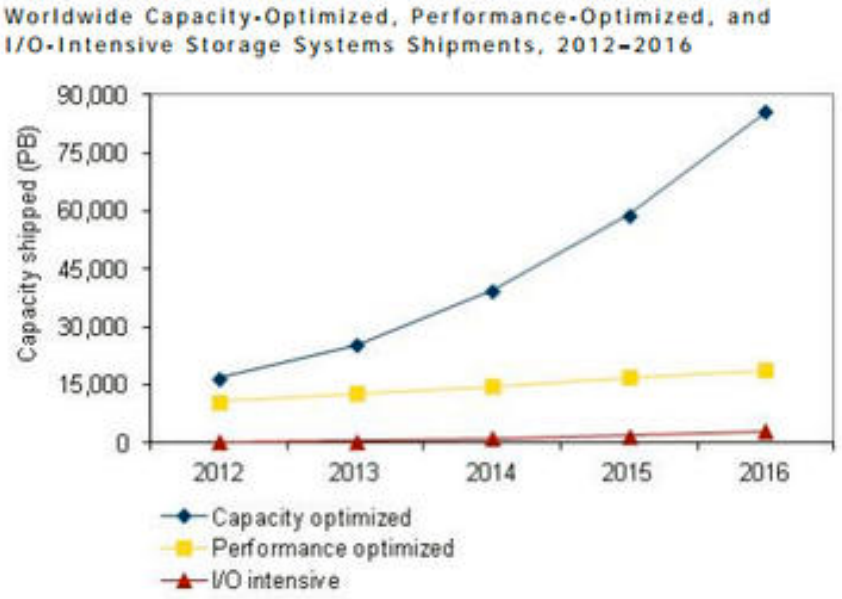

Source: IDC, 2013

Characteristics of workloads enabling the benefit from Auto Tiering

- Asymmetric (non-uniform) IO demand intensity A.K.A Hot Spots
- Location of the Hot Spots
- Relative demand regions
- Periods of reduced IO demand
- Repetitive nature of changing workloads

IO demand performance density

This metric is dependent solely on the characteristics of the workload

It is simply the total number of IOs per second divided by the total space

- For uniform access patterns, this metric applies to the entire space
- For non-uniform access patterns, different subsets of the space will have different performance densities

This metric can be used to size both Auto Tiering and traditional configurations

- The tools for measuring Performance density are different for Auto Tiering vs. Non-Auto Tiering environments.
- Host based tools may be used for uniform access environments
- Non-uniform access applications require storage product specific tools

Technology performance density

This metric is dependent solely on the technology and number of storage devices used

It is the total number of IOPS that can be delivered divided by the total space

- For uniform access patterns, this metric applies to the entire space
- For non-uniform access patterns, different subsets of the space will have different performance densities and will therefore required different devices (tiers)

By using the performance density as the sizing metric, we ensure that both the capacity and performance criteria will be optimally met

- Wasted capacity is minimized
- Maximum value for a given device technology is achieved as full performance is delivered

Example: Non-uniform IO demand density

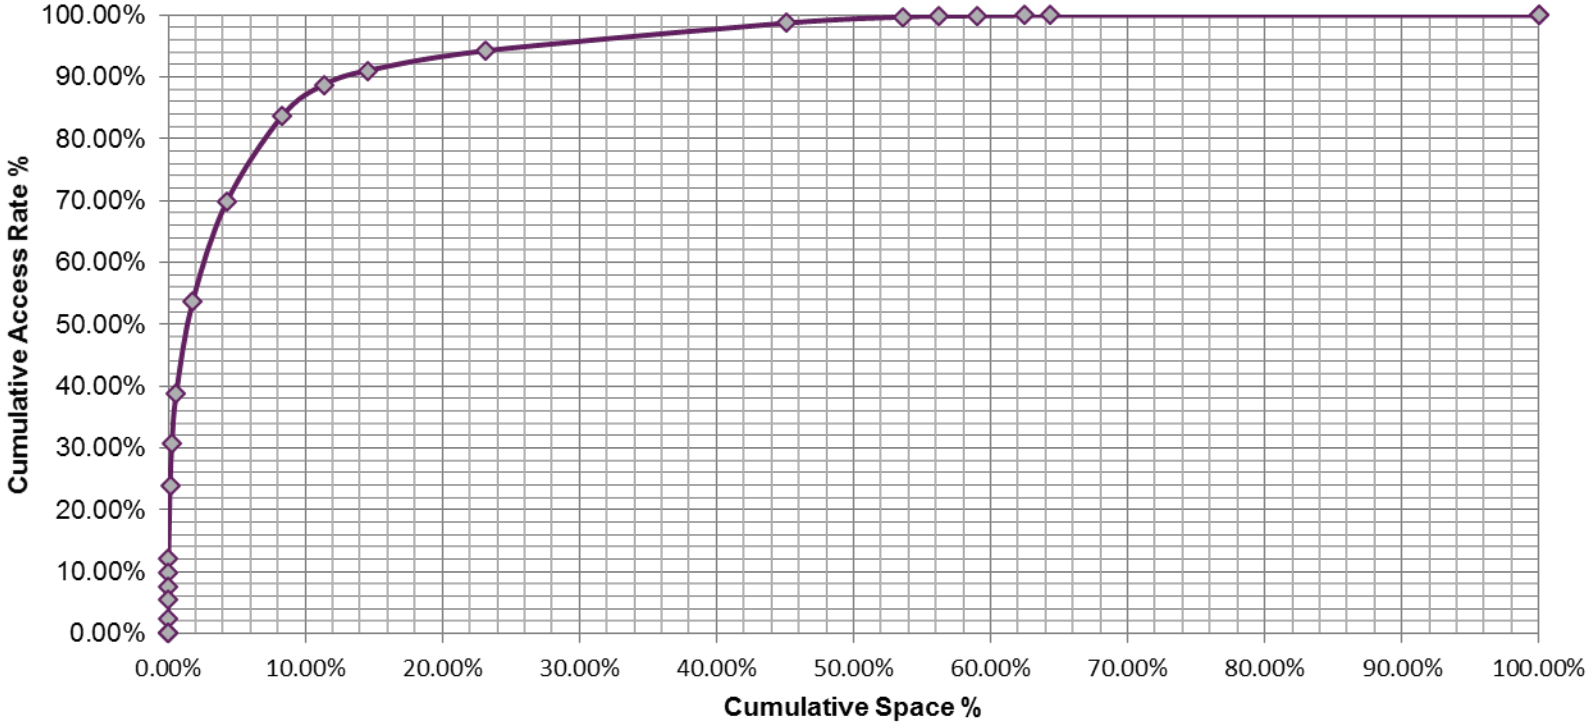

Note that for uniform access, the curve is a straight line corner to corner

Storage performance analysis (Auto Tiering tier design)

Performance Density Measurement Analysis

From the preceding chart we can glean the following:

- Based on relative IOPS, we see that **90%** of the IO demand can be contained within **12%** of the total space.
- By knowing the absolute proposed IO demand, and the total space used, we can calculate the required performance density of the tiers.
- Consider the space required to be **81TB** and the total proposed IO demand to be **75,500 IOPS**
- Therefore the required performance density of the tier is **6.99 IOPS/GB (75,500 *.9 /0.12*81000)**
- Workload is assumed to be **8 KB 70/30 R/W** random mix
- If this were a uniform access workload, we'd need a tier capable of delivering **75,500/81,000 = 0.93 IOPS** per GB
- Next, we take the remainder of the IO and space demand to determine the technology and capacity of the second tier....

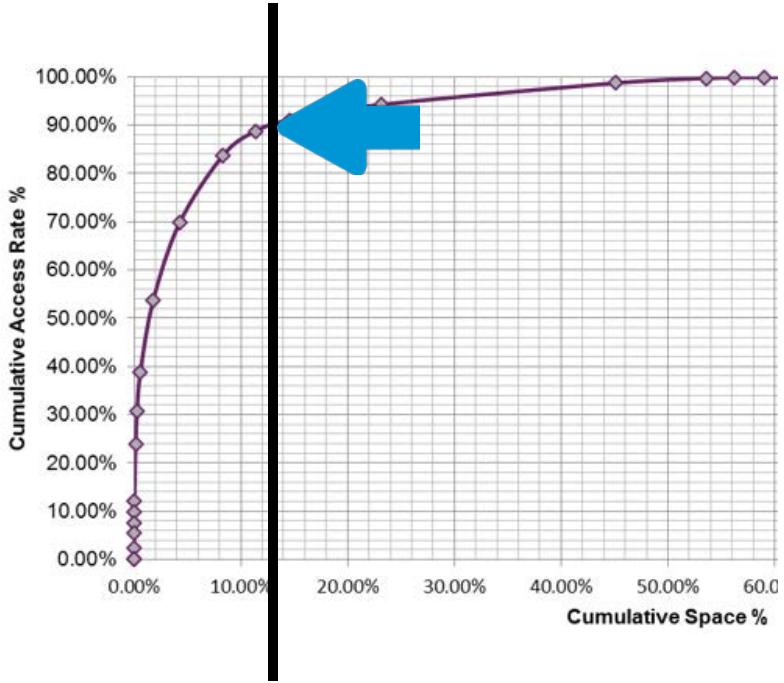

Storage performance analysis (Auto Tiering tier design)

Performance Density Measurement Analysis

For the next Auto Tiering tier, from the preceding chart we can glean the following:

- Based on relative IOPS, we see that **8%** of the IO demand can be contained within **28%** of the total space.
- Again, using the absolute proposed IO demand, and the total space used, we can calculate the required performance density of the second tier.
- Restating, the space required to be **81TB** and the total proposed IO demand to be **75,500 IOPS**
- Therefore the required performance density of the second tier is **0.266 IOPS/GB (75,500 *.08 /0.28*81000)**
- Workload is assumed to be **8 KB 70/30 R/W** random mix
- Next, we take the remainder of the IO and space demand to determine the technology and capacity of the third tier....

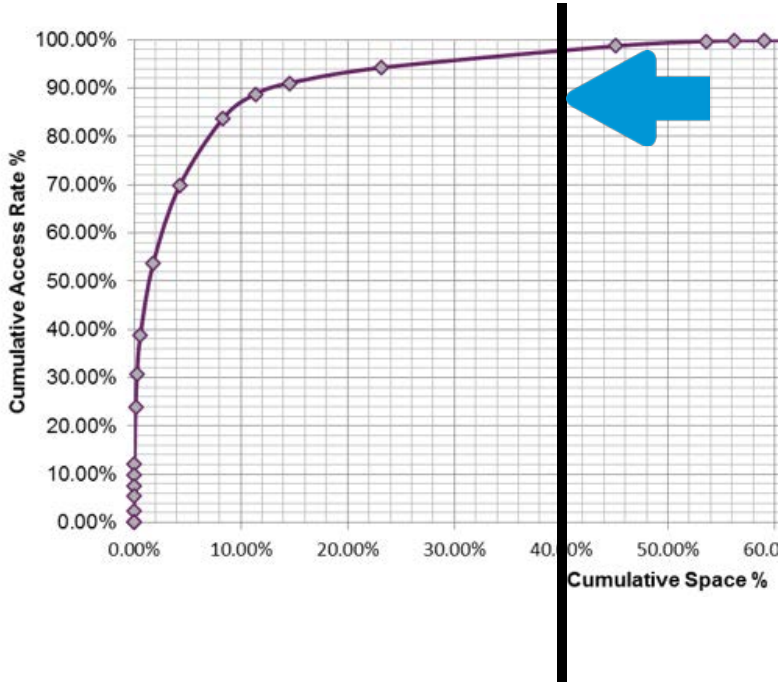

Storage performance analysis (Auto Tiering tier design)

Performance Density Measurement Analysis

For the third Auto Tiering tier, from the preceding chart we can glean the following:

- Based on relative IOPS, we see that **2%** of the IO demand can is contained within **60%** of the total space.
- Again, using the absolute proposed IO demand, and the total space used, we can calculate the required performance density of the third tier.
- Therefore the required performance density of the second tier is **0.031 IOPS/GB (75,500 *.02 /0.6*81000)**
- Next, we index into a Tier technology performance template to determine the appropriate devices to use for the three tiers.... But first, how are these tables derived?

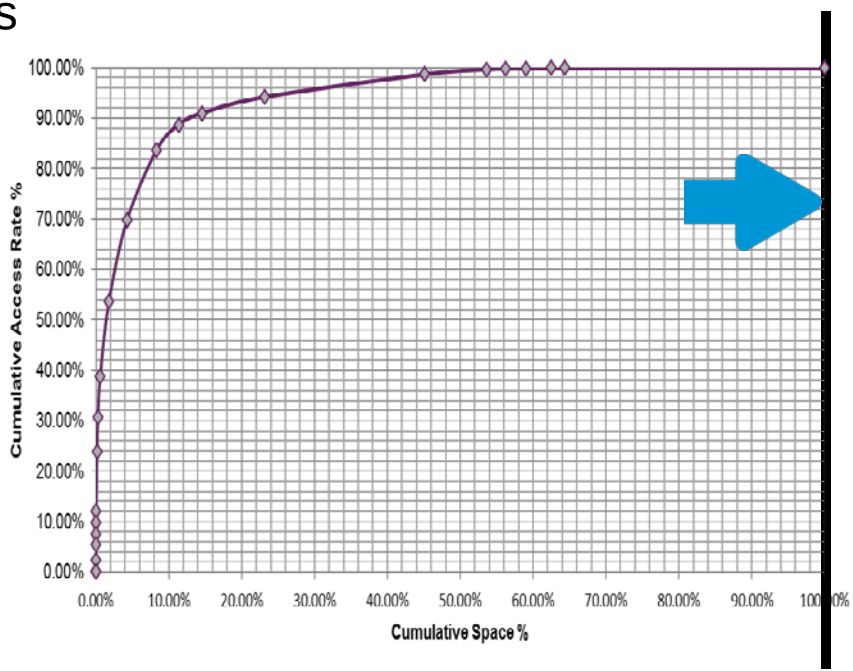

Deriving technology performance density

Technology Performance Density Calculations

Previously: It is the total number if IOPS that can be delivered divided by the total space

We'll use the measured data from lab tests to determine the IO rate

We'll use configured raid capacity to determine the space

- For this example,tier 0, lets consider **400 GB** SSDS, Raid 5 in an Enterprise Array (70/30 Read Write Workload)
- The following is a performance table for several device types
- We see that each SSD is capable of **2100 IOPS** for the 70/30 workload.
- We also know that in a Raid 5 configuration, there is **25%** capacity overhead (set size 4)
- So the capacity of each SSD is **400 GB * 0.75 = 300 GB**
- Finally, per device we can expect **2100 IOPS / 300 GB = 7.0 IOPS/GB**
- This appears to be a good fit for Tier0 as it’s performance density is slightly higher than required **6.99 IOPS/GB**

Component performance

HDD (Disk)

Disk Type/Speed	IOP/s	MB/s
15K FC	200	45
10K FC	150	45
7.2K NL	75	30

SSD

Workload	IOPs Aligned	IOPs Unaligned
100% Read	3950	3900
70% Read/30% Write	2100	1500
50% Read/50% Write	1800	1150
30% Read/70% Write	1500	1000
100% Write	1600	1150

Storage performance template: Enterprise Array (Raid5 Set size 4)

Performance Density

Workload	Array Limits	200 GB SSD	400 GB SSD	800 GB SSD	15k 300 GB	15k 450 GB	15k 600 GB	10k 300 GB	10k 450 GB	10k 600 GB	7.2k 1 TB	7.2k 2 TB
Random Reads (IOPS)	365,500	26.33 IOPs/GB	13.16 IOPs/GB	6.58 IOPs/GB	0.89 IOPs/GB	0.50 IOPs/GB	0.44 IOPs/GB	0.67 IOPs/GB	0.44 IOPs/GB	0.33 IOPs/GB	75 IOPs/TB	38 IOPs/TB
Random Writes (IOPS)	139,500	10.67 IOPs/GB	5.33 IOPs/GB	2.67 IOPs/GB	0.89 IOPs/GB	0.50 IOPs/GB	0.44 IOPs/GB	0.67 IOPs/GB	0.44 IOPs/GB	0.33 IOPs/GB	75 IOPs/TB	38 IOPs/TB
Rnd 70/30 Mix (IOPS)	256,040	14.00 IOPs/GB	**7.00 IOPs/GB**	3.50 IOPs/GB	0.89 IOPs/GB	0.50 IOPs/GB	0.44 IOPs/GB	0.67 IOPs/GB	0.44 IOPs/GB	**0.33 IOPs/GB**	75 IOPs/TB	38 IOPs/TB
Seq. Reads (MB/s)	14,660	3.67 MB/s/GB	1.83 MB/s/GB	0.92 MB/s/GB	0.20 MB/s/GB	0.13 MB/s/GB	0.10 MB/s/GB	0.20 MB/s/GB	0.13 MB/s/GB	0.10 MB/s/GB	30 MB/s/TB	15 MB/s/TB
Seq. Writes (MB/s)	10310	3.09 MB/s/GB	1.55 MB/s/GB	0.77 MB/s/GB	0.20 MB/s/GB	0.13 MB/s/GB	0.10 MB/s/GB	0.20 MB/s/GB	0.13 MB/s/GB	0.10 MB/s/GB	30 MB/s/TB	15 MB/s/TB

Deriving technology performance density

In a like manner, we'll examine the table to determine the technology of the second tier

- Consider **600 GB** 10 k RPM HDDs, Raid 5 in an Enterprise Array (70/30 Read Write Workload)
- Again, referring to the performance table for the device types
- We see that the performance density of 600 GB 10k RPM HDDs is **0.33 IOPS/GB** for the 70/30 workload.
- This appears to be a good fit for Tier 1 as it’s performance density is slightly higher than required **0.266 IOPS/GB**

Storage performance Auto Tiering tier design process

Implementing the Tier(s)

- From the previous template, we see that **400 GB** SSDs, Raid 5 are capable of delivering **7.0 IOPS/GB**
  – We also know we need **12%** of **81TB**, or roughly **9.72 TB** of high speed capacity.
  – This translates to **9.72 /0.75 = 12.96 TB raw = ~ 32 SSDs**

- For this part of the workload in terms of IO demand and required space, **32-400 GB** devices using Raid 5 would deliver **67,200 IOPS** which is very close to the required demand

- For the second tier, we agreed on 10k RPM, 600 GB HDDs
  – Space required: **22.68 TB**

  – Raw Space required: **22.68/0.75 = 32.4 TB -> 50 10k RPM 600 GB HDDs**

Storage performance Auto Tiering tier design process

Implementing the Tier(s)

- For Tier 2, the required performance density is **0.031 IOPS/GB (75,500 *.02 /0.6*81000) = 31 IOPS/TB**

- Again, consulting the Performance density Table we see that 7.2k RPM, 2 TB HDDs can deliver **38 IOPS/TB**

  – Space required: **60 % of 81 TB = 48.6 TB**

  – Raw Space required: **48.6/0.75 = 64TB -> 32 - 7.2k RPM HDDs**

Storage performance template: Enterprise Array (Raid5 Set size 4)

Performance Density

Workload	Array Limits	200 GB SSD	400 GB SSD	800 GB SSD	15k 300 GB	15k 450 GB	15k 600 GB	10k 300 GB	10k 450 GB	10k 600 GB	7.2k 1 TB	7.2k 2 TB
Random Reads (IOPS)	365,500	26.33 IOPs/GB	13.16 IOPs/GB	6.58 IOPs/GB	0.89 IOPs/GB	0.50 IOPs/GB	0.44 IOPs/GB	0.67 IOPs/GB	0.44 IOPs/GB	0.33 IOPs/GB	75 IOPs/TB	38 IOPs/TB
Random Writes (IOPS)	139,500	10.67 IOPs/GB	5.33 IOPs/GB	2.67 IOPs/GB	0.89 IOPs/GB	0.50 IOPs/GB	0.44 IOPs/GB	0.67 IOPs/GB	0.44 IOPs/GB	0.33 IOPs/GB	75 IOPs/TB	38 IOPs/TB
Rnd 70/30 Mix (IOPS)	256,040	14.00 IOPs/GB	**7.00 IOPs/GB**	3.50 IOPs/GB	0.89 IOPs/GB	0.50 IOPs/GB	0.44 IOPs/GB	0.67 IOPs/GB	0.44 IOPs/GB	**0.33 IOPs/GB**	75 IOPs/TB	**38 IOPs/TB**
Seq. Reads (MB/s)	14,660	3.67 MB/s/GB	1.83 MB/s/GB	0.92 MB/s/GB	0.20 MB/s/GB	0.13 MB/s/GB	0.10 MB/s/GB	0.20 MB/s/GB	0.13 MB/s/GB	0.10 MB/s/GB	30 MB/s/TB	15 MB/s/TB
Seq. Writes (MB/s)	10310	3.09 MB/s/GB	1.55 MB/s/GB	0.77 MB/s/GB	0.20 MB/s/GB	0.13 MB/s/GB	0.10 MB/s/GB	0.20 MB/s/GB	0.13 MB/s/GB	0.10 MB/s/GB	30 MB/s/TB	15 MB/s/TB

Auto Tiering vs. non-Auto Tiering solution comparison

Auto Tiering Solution Device Type and Count

Fast Tier: **(Tier 0) 32 – 400 GB SSD**

Intermediate Tier **(Tier 1) 50 – 600 GB 10K rpm HDDs**

Bulk Tier **(Tier 2) 32 – 2 TB 7.2k RPM NL HDDs**

Total Auto Tiering Solution Device Count: **114**

Non-Auto Tiering Solution Device Type and Count

Single Tier: **0.93 IOPS/GB -> 300 GB FC Drives**

Total Non-Auto Tiering Solution Device Count: **360 (81 TB / 0.75) / 300 GB**

Total Device Count Reduction: 68%

Hardware Savings (list price): $18,570 or ~6.7%

Improved Performance and Same Capacity ☺

September 11, 2017

DXC Storage Provisioning

Performance Aware Storage Provisioning Application: Auto Tiering

Thank You

DXC.technology