FLASH 1st

Right-sizing Tiered Storage using a simple Data Decay Model

Denis Vilfort
Sr. Director, USD Marketing

August 2011

EMC² where information lives

Today's Compute Hierarchy

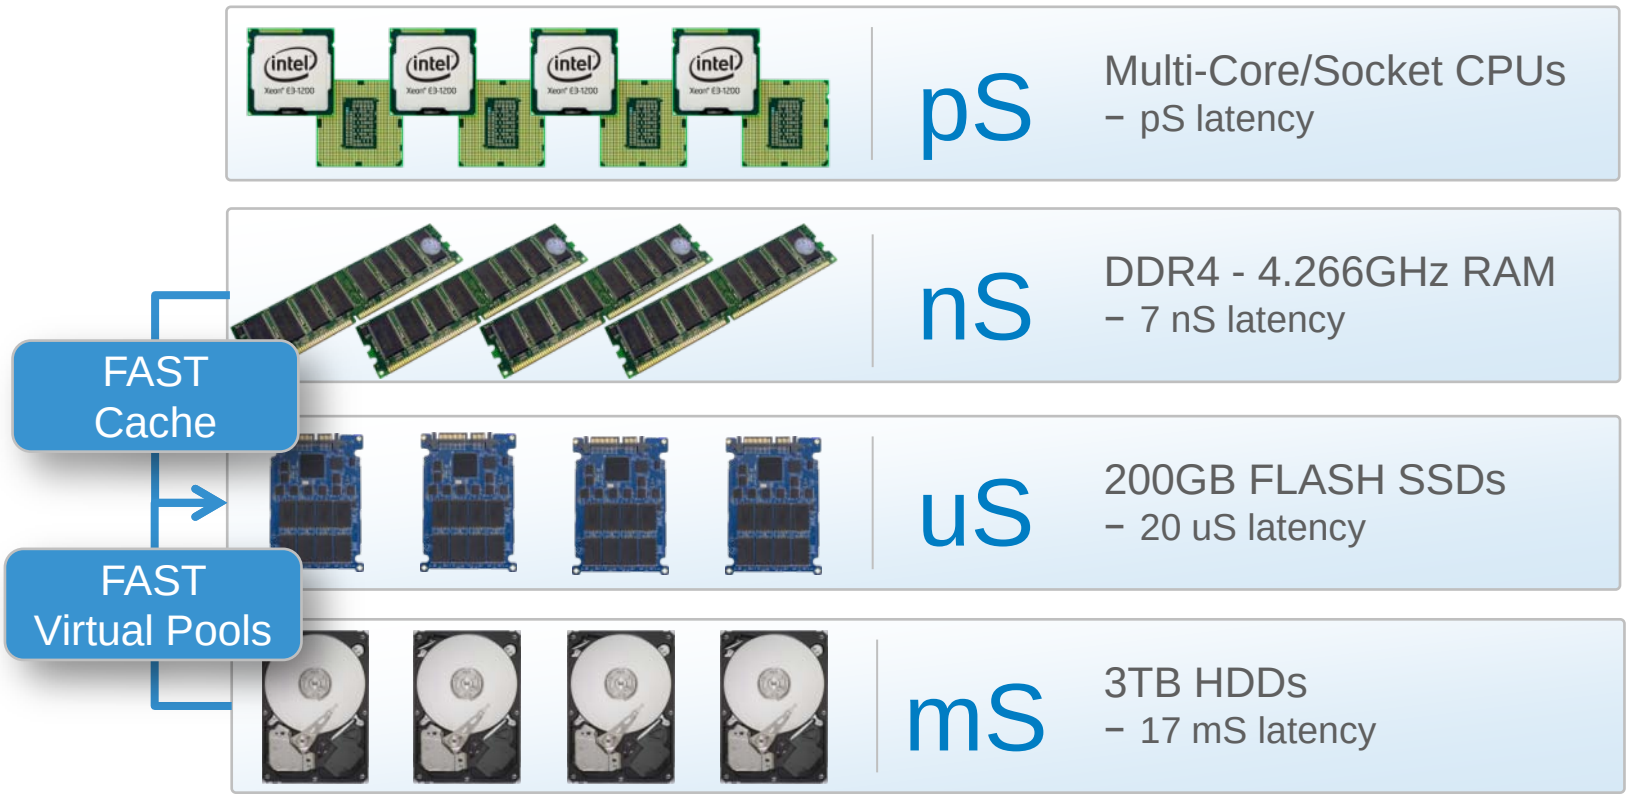

Increased Storage Efficiency

100 TB Usable Capacity Example

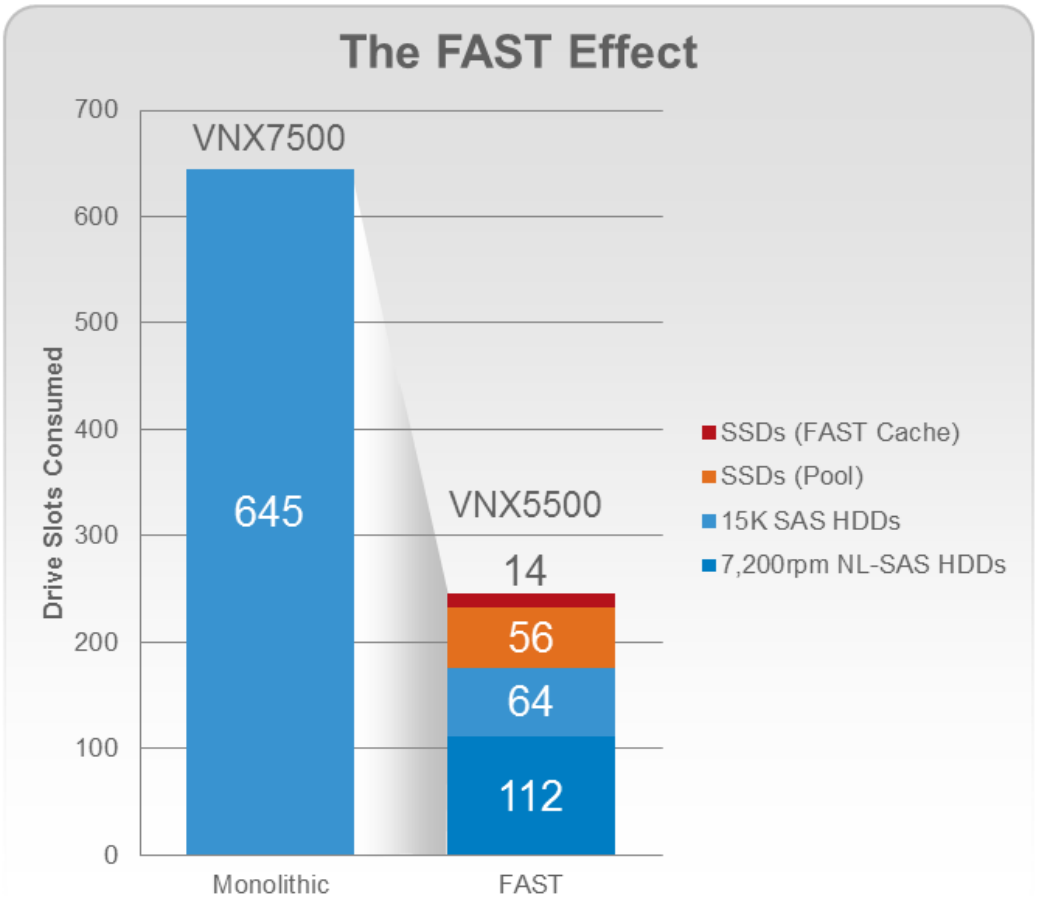

- 645 X 600GB 15K
- Replaced by 246 drives in a 6/10/84 blend:
  - 70 X 200GB SSDs
  - 64 X 600 GB 15K HDDs
  - 112 X 2TB NL-HDDs
- Smaller frame needed
  - 62% smaller footprint
  - 75% less power
  - 27% lower acquisition cost

Effective Cost Defined

Why utilization is a big deal

$$\frac{\text{Nominal Cost}}{\text{Utilization}} = \text{Effective Cost}$$

Low Utilization Example

$$\frac{\$10/GB}{10\%} = \$100/GB$$

High Utilization Example

$$\frac{\$10/GB}{80\%} = \$12.50/GB$$

Gaining Fast Application Response Time

Older approaches are increasingly inefficient

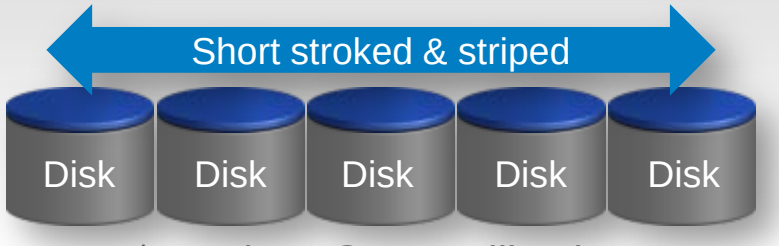

$1.53/GB @ 3% utilization

$54 per GB effective cost

- Mechanical drives are slow
- Utilization must be sacrificed in order to get the highest disk performance (short stroking)
- More and more drives must be striped as the need for business transactions increase

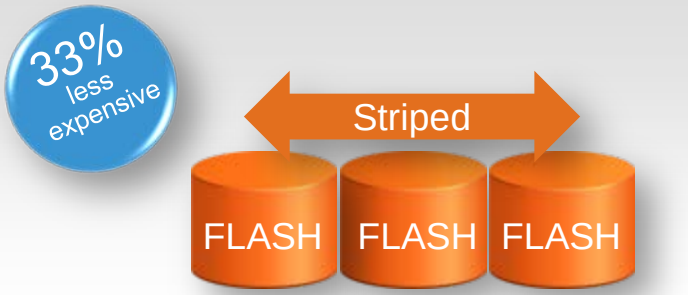

$34.80/GB @ 98% utilization

$36 per GB effective cost

- SDDs are up to 100 times faster than mechanical disks – but more expensive
- Almost the entire drive can be used, so fewer SSDs drives are needed
- Only the most critical data can justify the manual move to expensive FLASH SSDs

Solution: Automatic Data Optimization

The benefit of FLASH without the cost

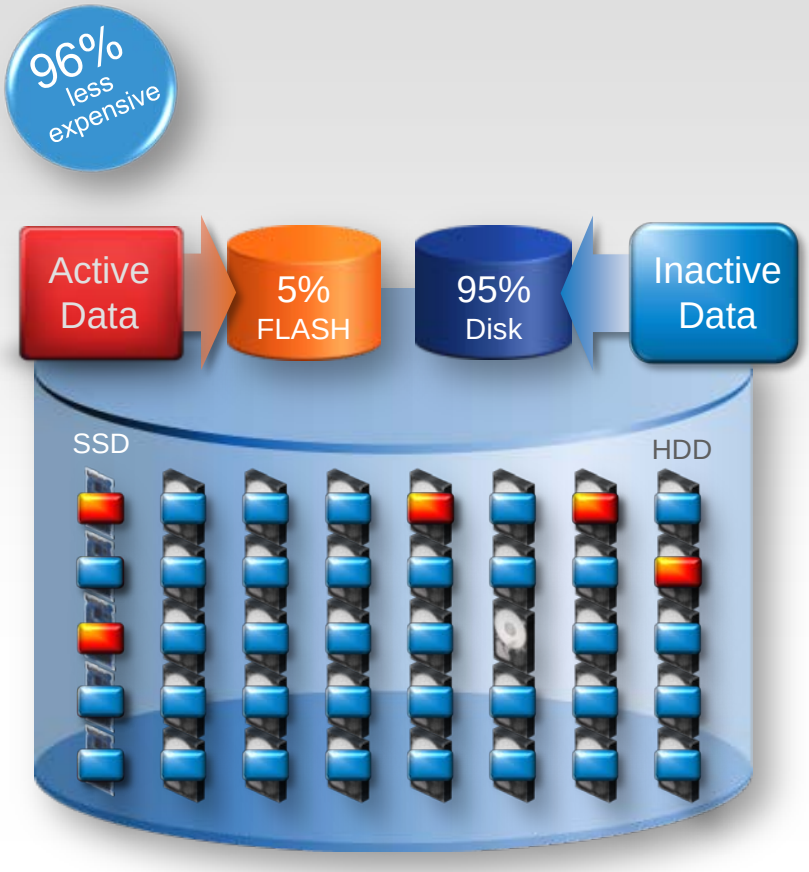

- All data activity is constantly tracked
- High activity data is automatically moved to FLASH SSDs
- Low activity data is automatically moved out of FLASH to low cost disk
- Applications now enjoy lowest response time possible for data that matters

90% utilization

$2.39 per GB effective cost

FLASH 1st Supercharges Vanderbilt

FAST Suite + SSDs help achieve IT excellence

VANDERBILT UNIVERSITY

- EMC Unified Storage Solution
  - VMware, SharePoint, SQL, Oracle
  - 108 TB's student data
  - Added 24 SSDs + FAST Suite
  - **1.6% FLASH serves 78% of all IO**

Examples of IO Skew

Driven by Data Growth and Business Models

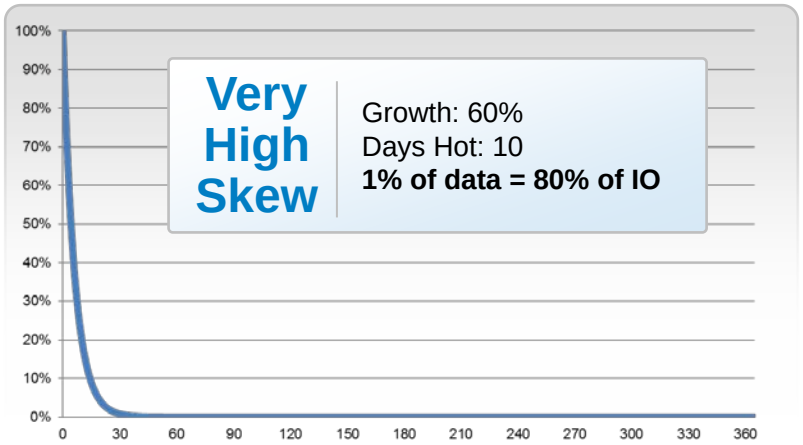

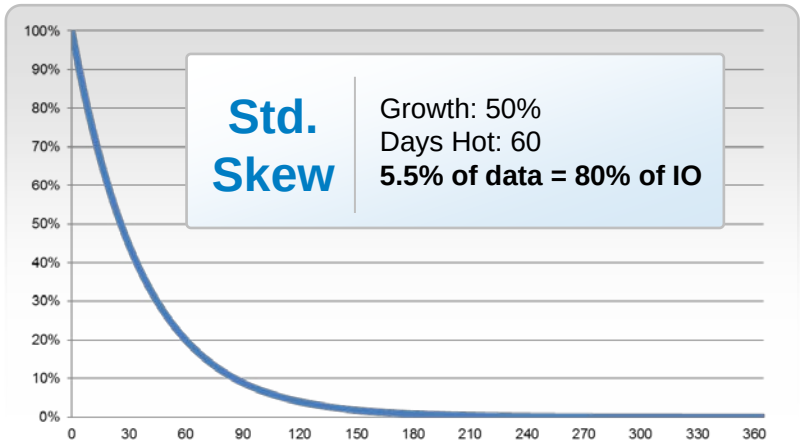

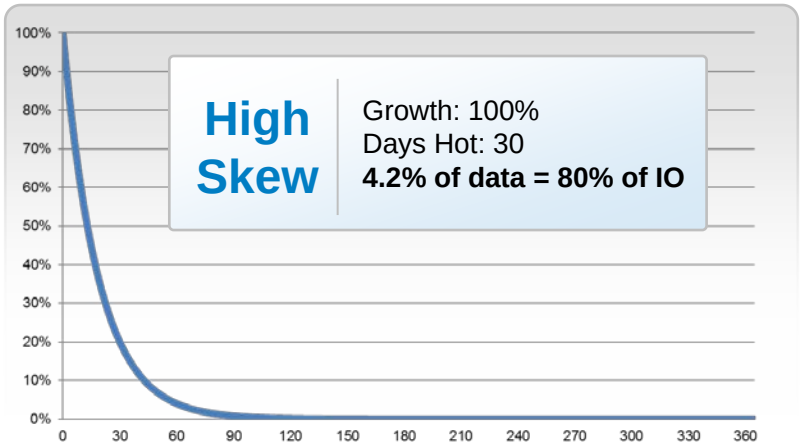

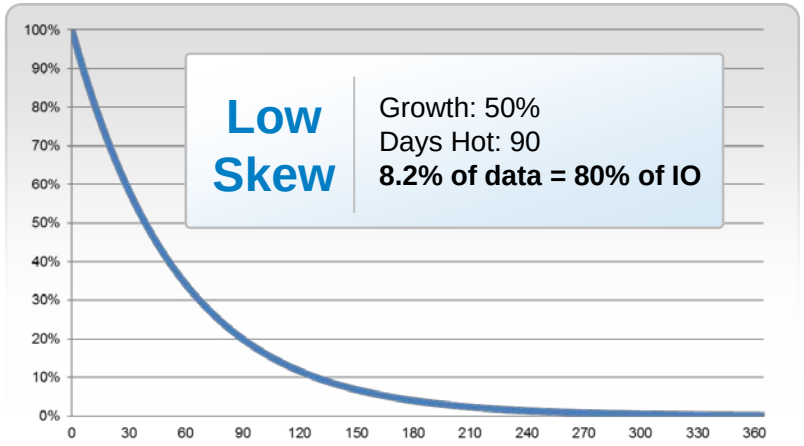

The FLASH 1st Data Strategy

Hot data on fast FLASH SSDs—cold data on dense disks

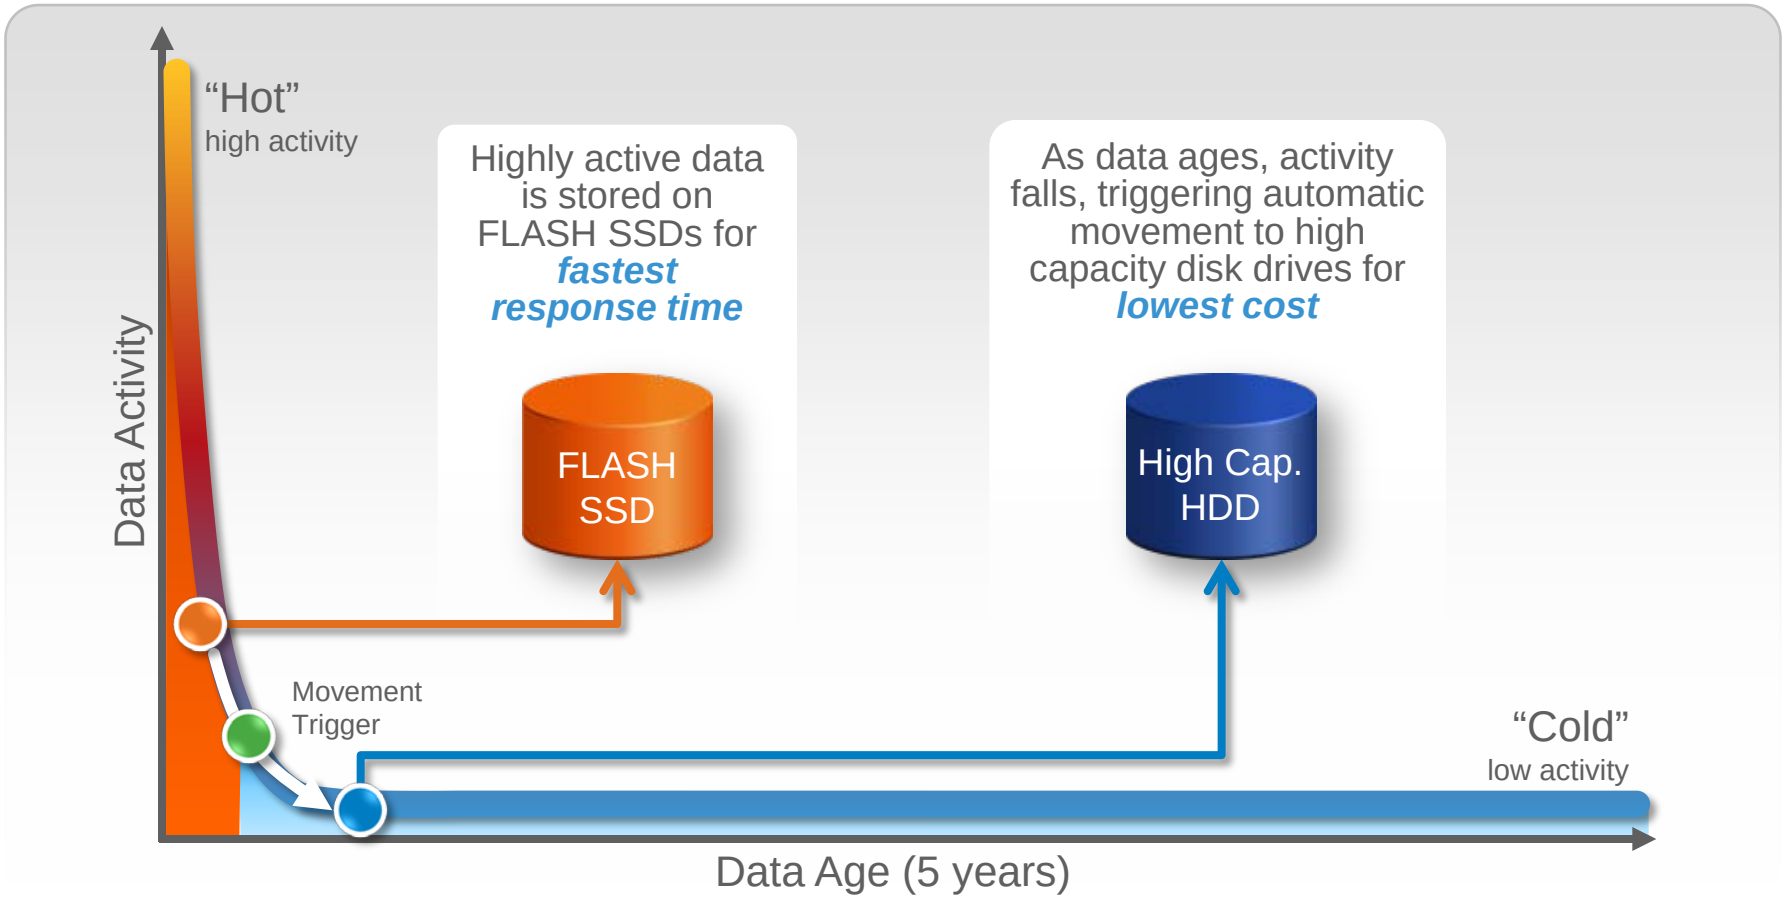

How Much FLASH?

Used dynamically with FAST

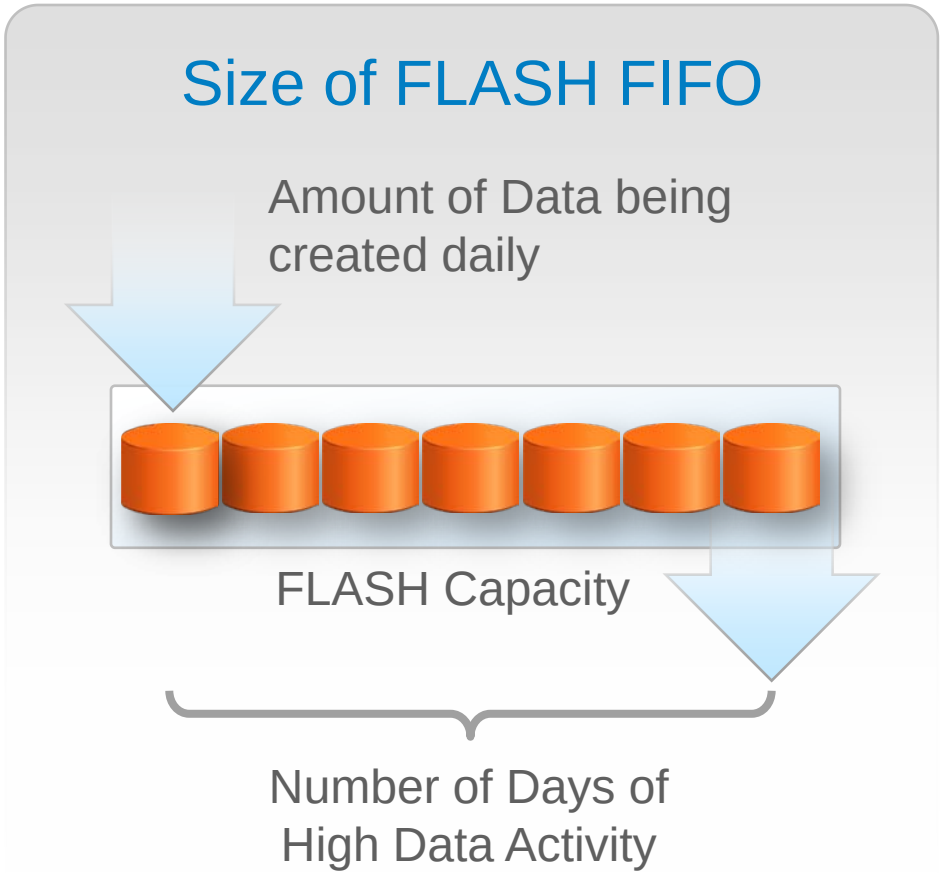

- Data follows a predictable decay in activity
- Older data is constantly being replaced by new highly active data
- The amount of FLASH required is determined by:
  - The amount of data created each day, and
  - The period of time it takes to cool

The FAST Suite

Dynamically optimizes IO for FLASH 1st at a 64K page size

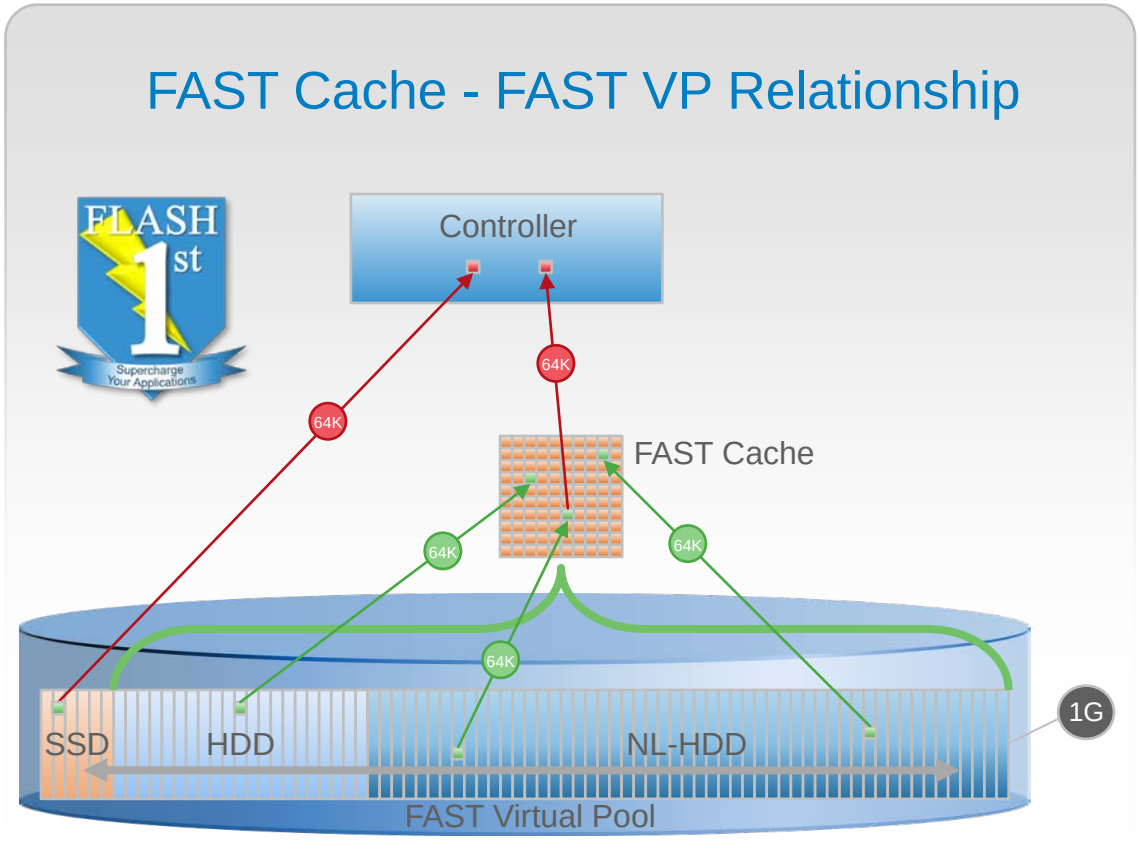

- FAST Cache
  - Caches data from the HDD or NL-HDD tiers in the pool
  - Operates at a page granularity of 64K
- FAST VP
  - Dynamically moves data between tiers in the storage pool
  - Operates at a slice granularity of 1GB
- Deploying both together ensures maximum IO granularity

How Long Does Data Stay Hot?

Depends on business model, applications and workload

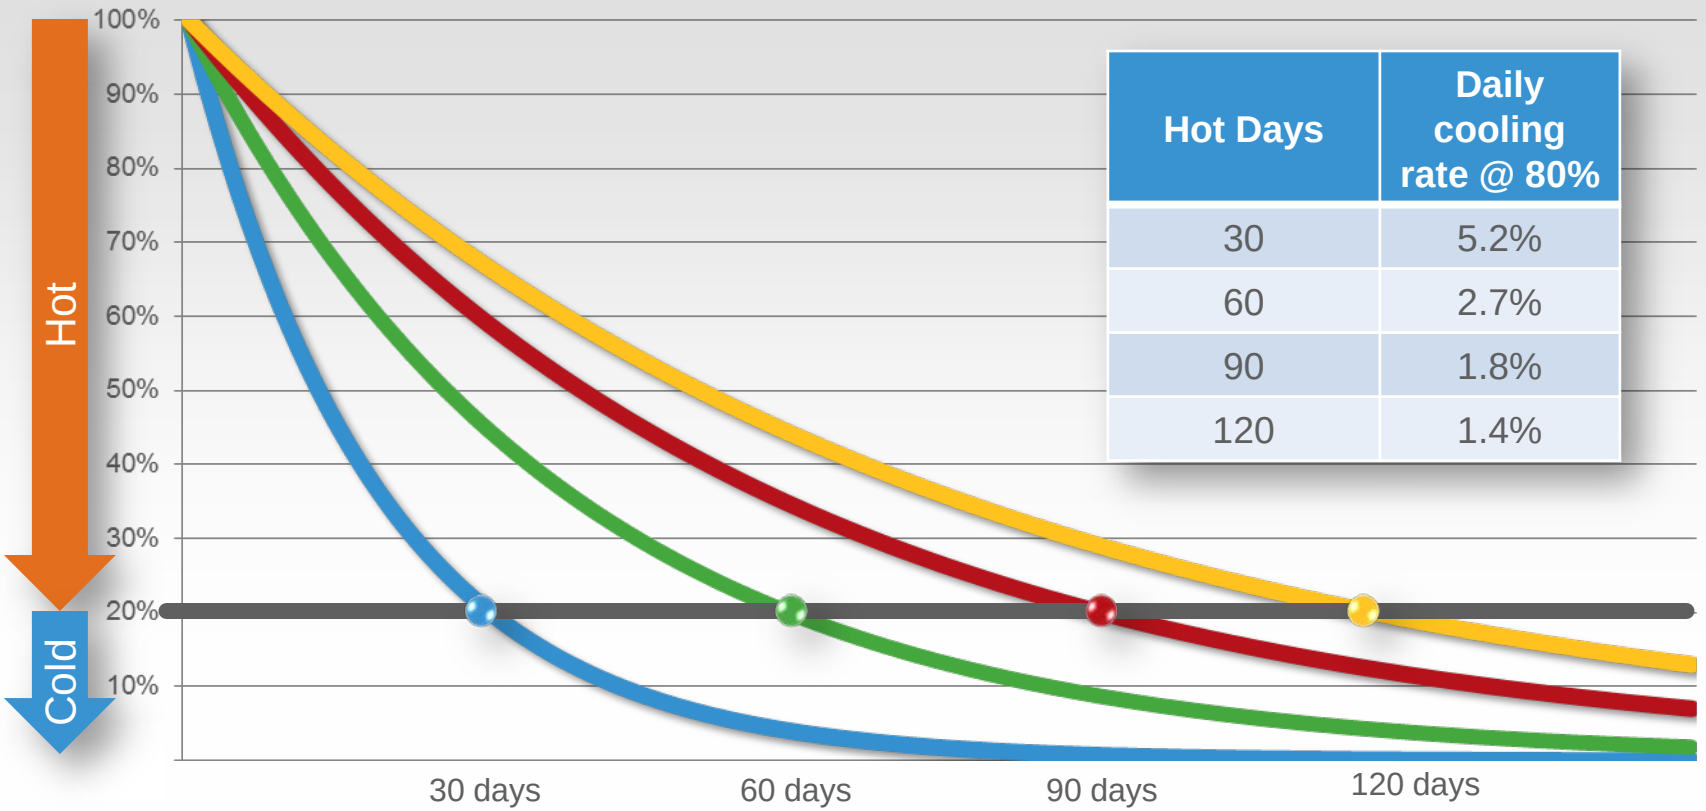

Hot Days	Daily cooling rate @ 80%
30	5.2%
60	2.7%
90	1.8%
120	1.4%

Calculating How Much FLASH

FLASH Portion

$$\text{FLASH \%} = \frac{\text{Yearly Growth Rate\% X Number of Hot Days X 100}}{\text{365 X (Yearly Growth Rate\% + 100\%)}}$$

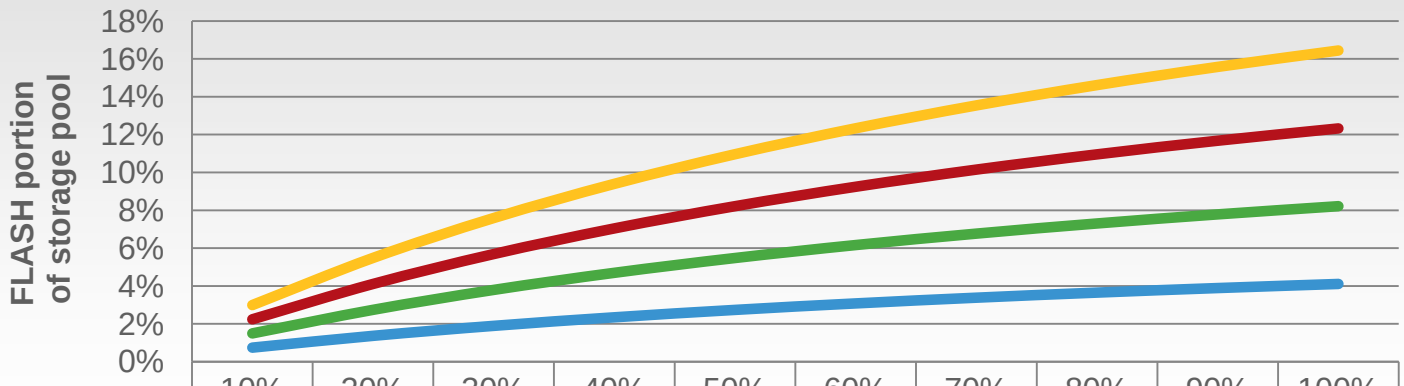

	10%	20%	30%	40%	50%	60%	70%	80%	90%	100%
30 days Hot	1%	1%	2%	2%	3%	3%	3%	4%	4%	4%
60 days Hot	1%	3%	4%	5%	5%	6%	7%	7%	8%	8%
90 days Hot	2%	4%	6%	7%	8%	9%	10%	11%	12%	12%
120 days Hot	3%	5%	8%	9%	11%	12%	14%	15%	16%	16%

How Much FLASH?

3 fundamental business questions

- How much data is under management today?

✓50 TBs

- How much is your data growing each year?

✓50% YoY

- How long does your data stay hot?

✓60 days

Calculating Net New Data

Net new is a function of amount of “starting data” and “growth”

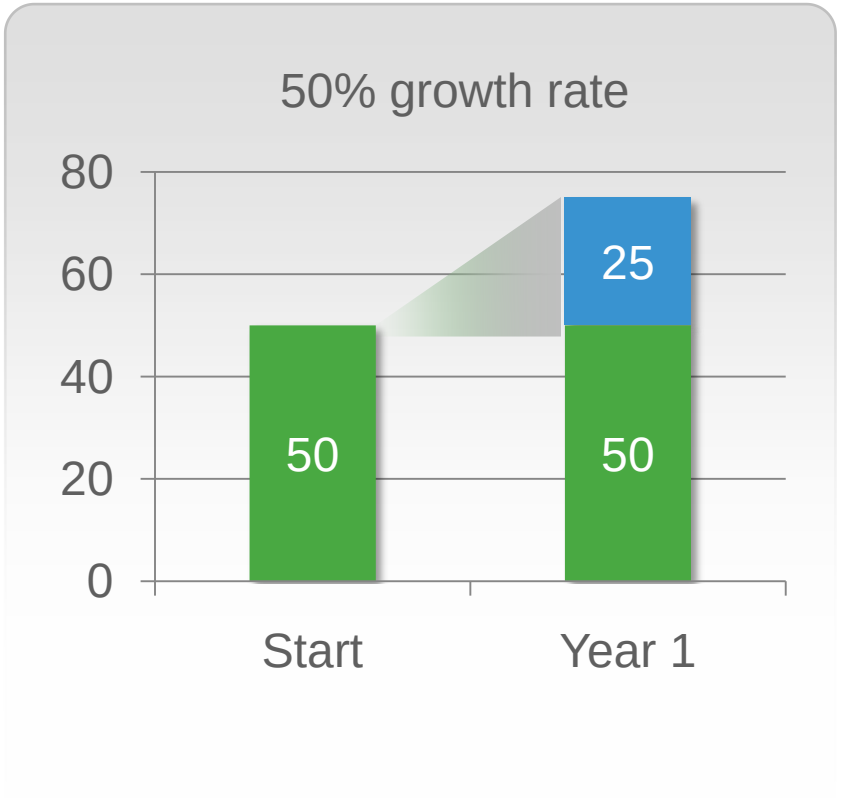

- 50% of 50TB is 25TB
- The average amount of data generated each day:
  - 25 X 1024 GB / 365 = **70 GB per day**

Calculating Needed FLASH

SSD tier only needs to be large enough to hold hot data

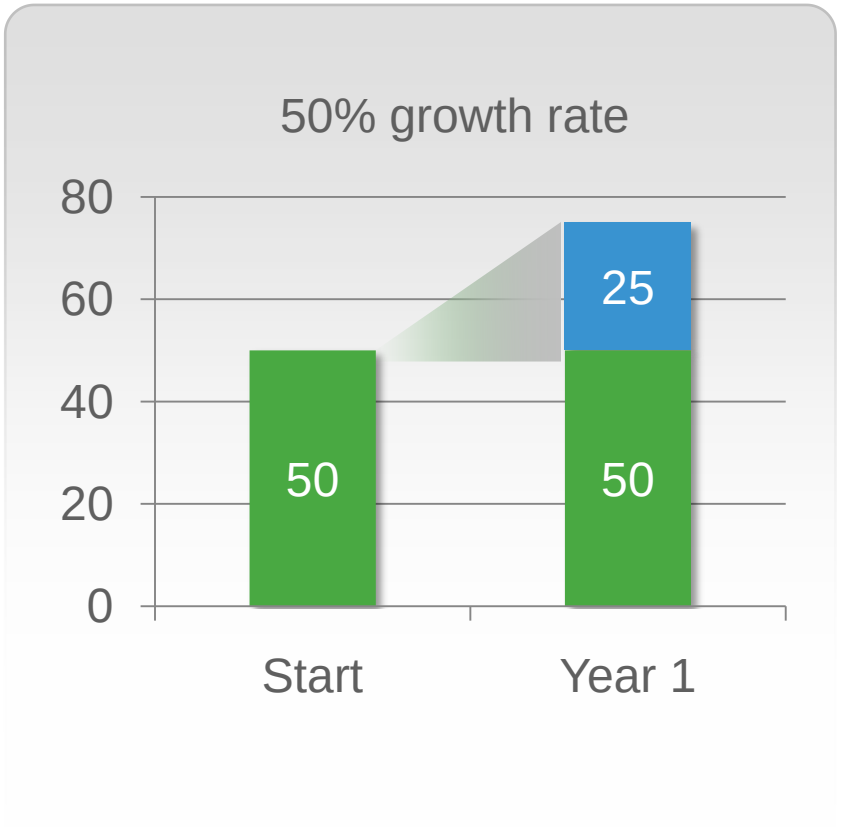

- FLASH Capacity:
  - 60 days X 70 GB = **4,200 GB**
- FLASH Percentage:
  - FLASH Capacity/Total Capacity
  - 4,200 GB/(75 X 1024) GB X 100: **6%**

FLASH Percentage

A function of data growth and service level

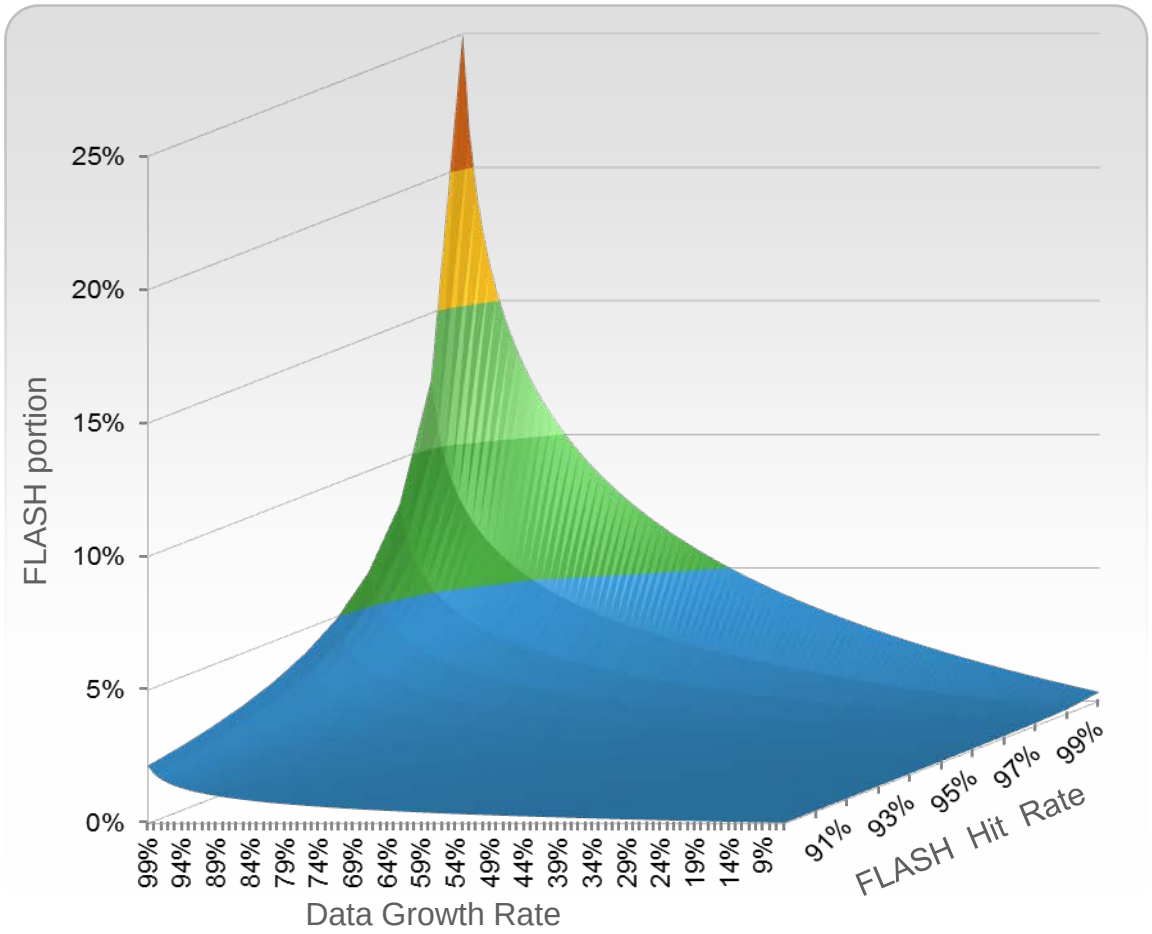

- Relative more FLASH is needed when:
  - Service level (FLASH hit rate) is elected high, and
  - Data growth is high
- More than 25% FLASH is *highly unlikely!*

Modeling "Data Decay"

A simple GB-Day state model describes typical data behavior

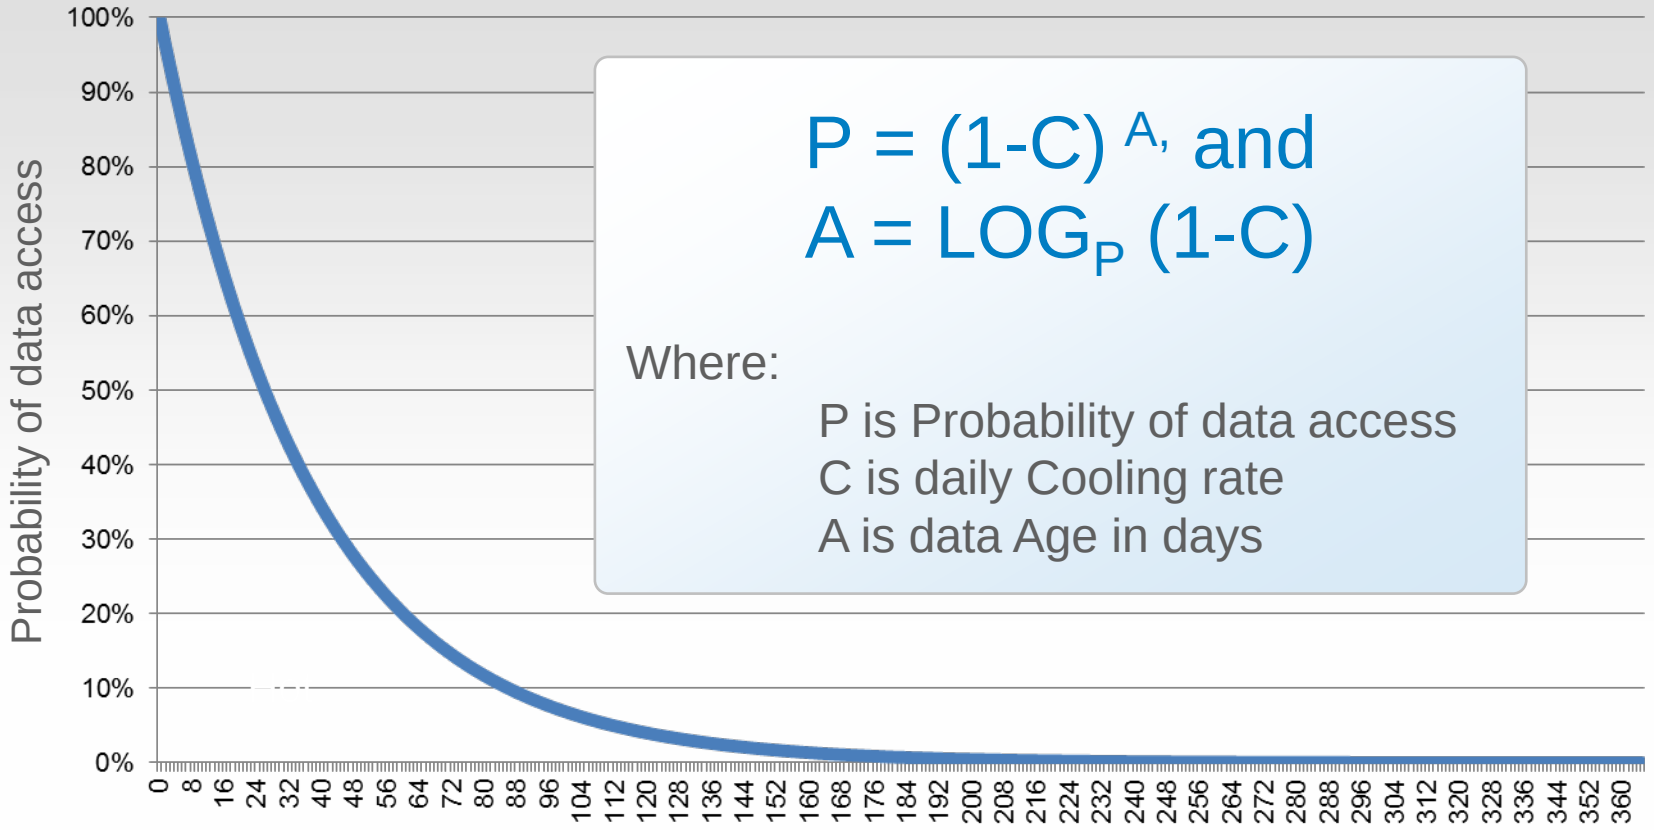

Modeling "Data Decay"

A simple GB-Day state model describes typical data behavior

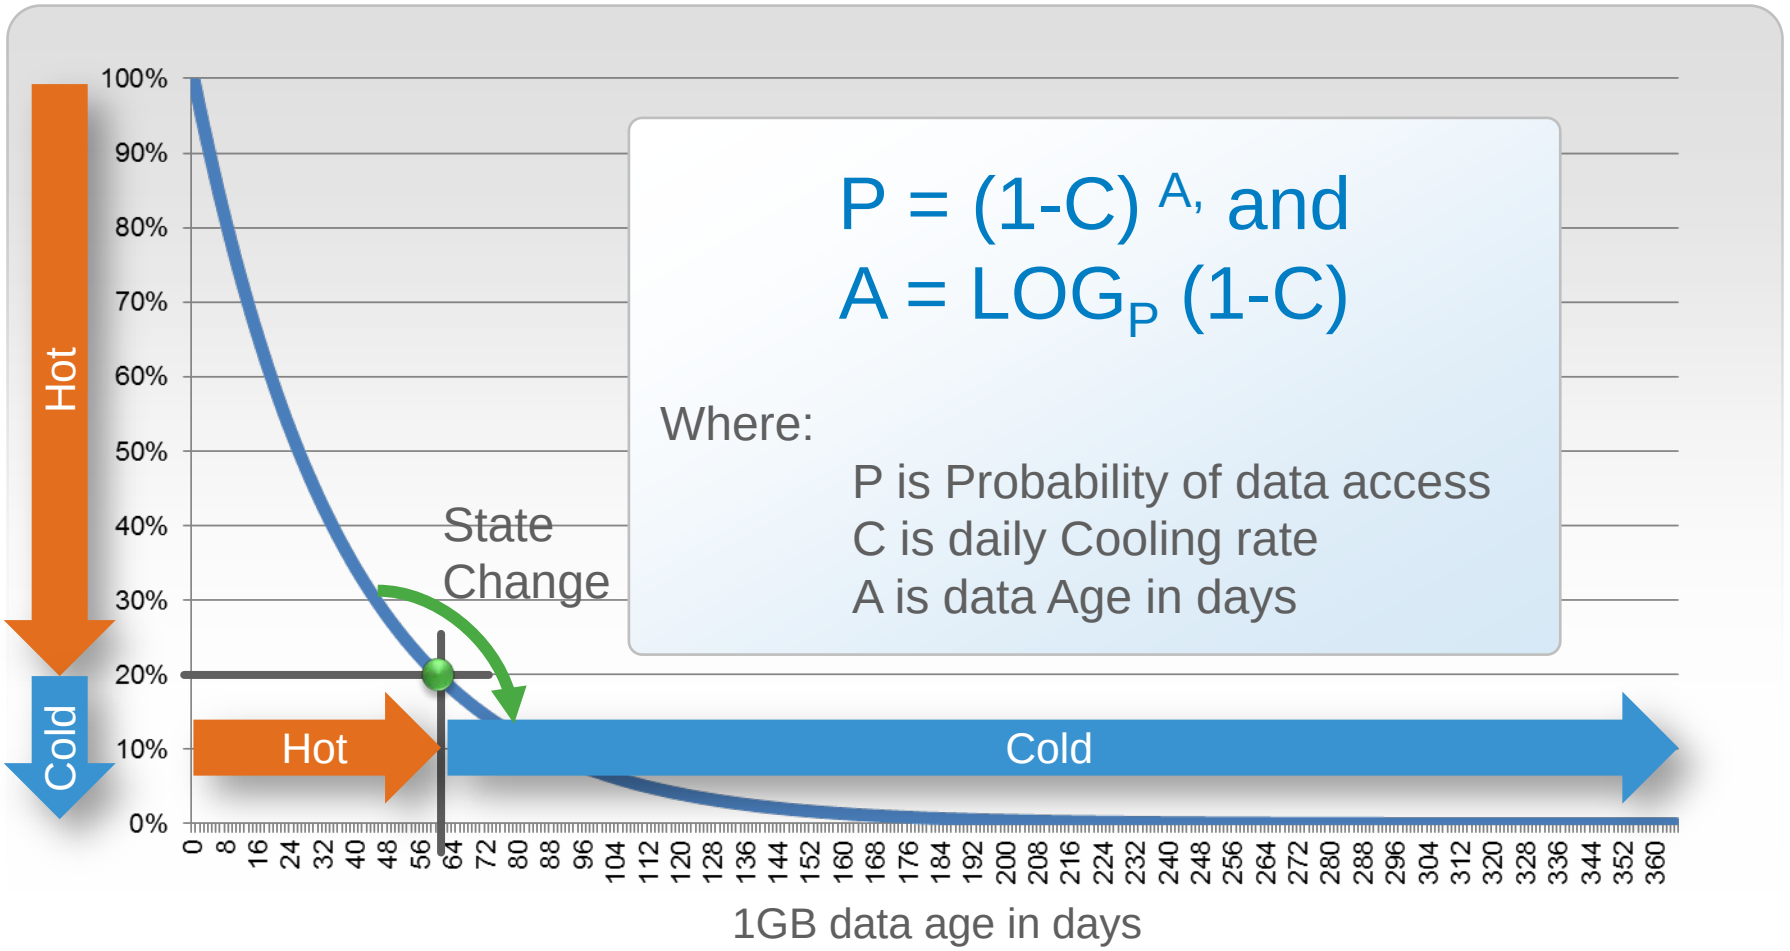

$P = (1-C)^{A,}$ and

$A = LOG_P (1-C)$

Where:

P is Probability of data access
C is daily Cooling rate
A is data Age in days

Using Business Metrics to Configure

With no trace data, the “decay model” can be a useful guide

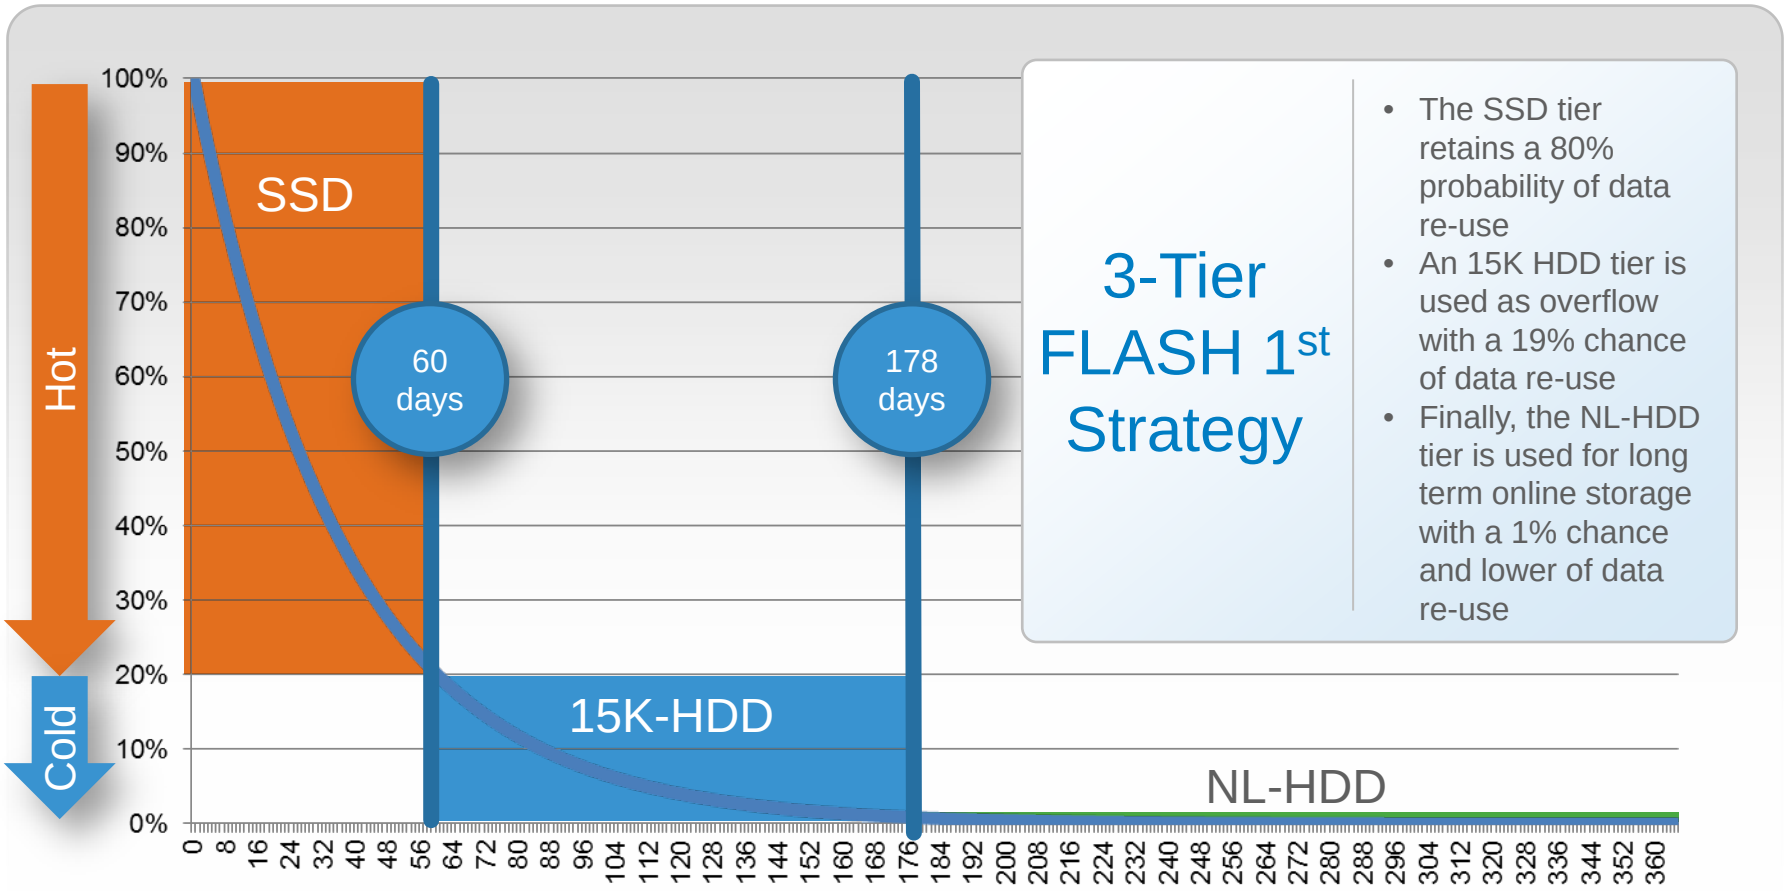

3-Tier FLASH 1st Strategy

- The SSD tier retains a 80% probability of data re-use
- An 15K HDD tier is used as overflow with a 19% chance of data re-use
- Finally, the NL-HDD tier is used for long term online storage with a 1% chance and lower of data re-use

Configuration Example

Using the model for 50TB, 50% growth, 60 days hot @ 80%

VNX5700

Large Foot Print

325 SAS 15K HDDs

$481,360

Monolithic

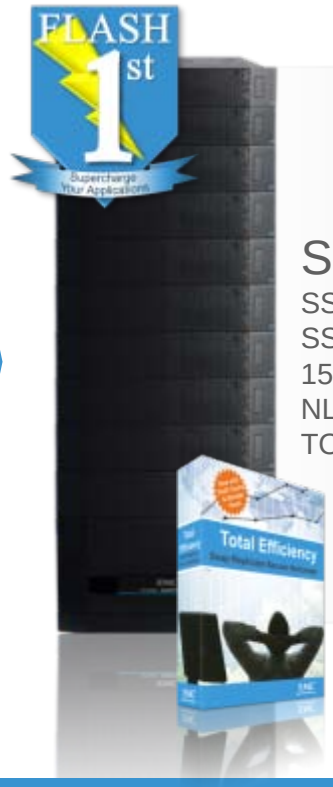

VNX5500

Smaller Foot Print

SSDs for FAST Cache	8
SSDs for FAST VP	32
15K SAS HDDs for FAST VP	32
NL-SAS HDDs for FAST VP	56
TOTAL	128

$434,691

Note: pricing only an example, NOT a quote. Individual cases may vary.

Clear Benefits of FLASH 1st

- Faster Response Times
- Smaller Footprint
- Less Power
- Lower Cost

THANK YOU