

# THE CEO'S GUIDE TO INVESTING IN FLASH STORAGE



## **EXECUTIVE SUMMARY**

Flash storage is every data center's version of a supercharged sports car. Nothing beats it in speed, efficiency, and handling—though it will burn through your cash much faster than a standard consumer vehicle.

Despite the high cost of flash storage, most organizations can't afford to rely totally on traditional disk arrays. Today's markets are moving too fast, and data centers have to be able to keep up. But how much flash they need to meet demand without breaking their budget is a fine line to walk.

Most business leaders don't live and breathe IT topics every day—but the IT decisions they make will keep them up at night. This white paper gives those business decision makers everything they need to know (and maybe a little extra) to stay on top of the game and make the right decisions on the right storage arrays at the right time for their business.

If you have any questions, don't hesitate to call. At Red8 we live and breathe this stuff, and we are happy to give you a no-obligation consultation to get you headed in the right direction.

We hope you enjoy.

## **TABLE OF CONTENTS**

1. DEDUPLICATION: BRINGING FLASH TO THE MASSES OF APPLICATIONS	4
· Where deduplication plays into the field	4
· It comes down to input/output	5
2. HOW eMLCs SAVED FLASH AS WE KNOW IT	5
· The need for eMLCs	5
• eMLCs increase the life—and value—of AFAs	5
3. HYBRID FLASH ARRAYS: HOW TO SOLVE THE MONEY PROBLEM	6
· Skew: Meeting demand for data	6
· Hybrid flash arrays: For ultimate control	7
· Three ways to measure value	7
4. HOW TO MAKE THE RIGHT DECISION	8



## DEDUPLICATION: BRINGING FLASH TO THE MASSES ... OF APPLICATIONS

Flash arrays are everywhere—their performance is unmatched, and their size makes them scalable from thumb drives and smartphones to data centers of all sizes. But they're also an expensive initial investment with limited storage capacity. So why the craze over flash arrays?

The short answer: **deduplication**.

The long answer: Flash is closing the gaps between processor and storage performance across the IT landscape. But the only way to make flash cost effective for a wide variety of workloads is through deduplication. Let's take a look at the pros and cons of flash, how deduplication impacts SSDs, and see if flash still makes sense for your business.

### WHERE DEDUPLICATION PLAYS INTO THE FIELD

Deduplication isn't just file compression, which identifies redundant information within a single file. It analyzes redundant information in an entire volume of data, thus reducing storage requirements by a factor of 10 to 1, 100 to 1, or more. That's why deduplication is called a "force multiplier"—even a 10TB flash array may be able to hold over 100TB of data. Disk limitations suddenly become much more fluid. With that kind of multiplication, the cost per TB for deduplicated data shrinks substantially.

When used for a specific workload that combines performance as well as capacity requirements, flash arrays beat out traditional disks in several ways:

- They cost less with respect to initial cost of acquisition
- They provide better scalability and reliability over the lifecycle of the application
- They significantly reduce the cost for administration and performance tuning
- They provide better ability to manage unexpected peaks in workload or changes in application workload profile

Keep in mind that not all data can be deduplicated—some databases, images, and any highly compressed or already deduplicated data can't be reduced further. But when hosting hundreds or thousands of similar virtual machines or virtual desktop infrastructures, the savings from deduplication are, in a word, dramatic.

#### **IT COMES DOWN TO INPUT/OUTPUT**

Deduplication then opens the door to vastly improved I/O performance—more so output than input, since flash arrays write speeds are relatively unimpressive. But their RAM-like nature makes data access almost infinitely faster than what traditional disks can provide.

Deduplication allows you to store more of that information you need to access at random, which means magnets don't have to scan entire disks to find the right files. No matter how fast HDDs can scan, they can never match the submillisecond access flash arrays provide to critical data. "NO MATTER HOW FAST HDDs CAN SCAN, THEY CAN NEVER MATCH THE SUB-MILLISECOND ACCESS FLASH ARRAYS PROVIDE TO CRITICAL DATA."

## HOW eMLCs SAVED FLASH AS WE KNOW IT

Flash storage provides more and better benefits over hard disk drives, including reduced physical space, lower maintenance costs, sub-millisecond latency, and more.

But SSD technologies are not without their issues, and they have evolved over the years to overcome numerous issues with respect to reliability, fabrication complexity, and cost. While there are emerging flash technologies such as 3D NAND and V-NAND flash that continue this evolution, the current enterprise-grade contender is eMLC flash.

#### THE NEED FOR eMLCs

Think about memory cells like very small bedrooms. You can fit one bed in, which is akin to one data point—that's a single-level cell (SLC). But when you need more space and you don't have more bedrooms, you put in bunk beds, meaning two data points in one memory cell—that's a multi-level cell (MLC). MLCs are at the core of SSDs for data centers.

MLCs give you far more memory capacity without significantly increasing fabrication costs or power consumption. On the other side of that coin, they also increase the number of program/erase (P/E) cycles on a single cell, which dramatically decreases the number of writes you have left on that cell before it starts to break down. For heavier workloads typical of enterprise data centers, the reduction of P/E cycles can be a real issue.

That's where enterprise MLCs (eMLCs) come in. eMLCs are programmed with specific commands and functions that give enterprises more room on the same space with less wear and tear, so data centers don't burn through SSDs in a period of months or even weeks.

#### eMLCs INCREASE THE LIFE-AND VALUE-OF AFAs

eMLCs have several tricks up their sleeves to extend the life of your flash array; for now, we'll take a look at four of the top methods for reducing wear and tear: wear leveling, bad-block mapping, garbage collection, and deduplication.

#### **1. WEAR LEVELING**

The first fail-safe is called wear leveling. Think about your SSD as a pair of shoes: patterns tend to show up in your soles, and you have to buy a new pair because of a hole in the heel, even though the rest of the shoe still looks great. Wear leveling is how an SSD evenly distributes where it is writing data to extend the life of the drive. It works by selling you 100TB (or any other amount) but keeping a secret stash of extra memory called overprovisioning, or OP. You can't technically access the OP so that the space can be used to evenly redistribute the load.

#### 2. BAD-BLOCK MAPPING

Wear leveling isn't perfect, and memory blocks still go bad—that's why bad-block mapping is crucial. Your SSD knows when a block has reached its P/E limit and uses its OP to copy data over to another block so that your data doesn't become corrupt.

#### **3. GARBAGE COLLECTION**

As data is distributed in wear leveling and then redistributed in bad block mapping, there's plenty of

opportunity for redundant data to gather in memory blocks across an SSD. That's called write amplification. To counter write amplification, garbage collection was built. Garbage collection is a function that identifies redundant, modified, or deleted data; writes the data that is still valid to a new block; and erases the original block, thereby freeing up space and keeping your SSD performing at peak efficiency.

#### 4. DEDUPLICATION

SSDs with eMLCs are constantly fighting against the effects of writing; but when you have deduplication, you actually avoid most redundancy issues before your SSD starts to write anything. Because deduplication is a force multiplier, you get far more data for far less wear and tear. With fewer P/E cycles per bit stored, deduplication thereby extends the life of your drive.

Though this is hardly an exhaustive list of all the precautions built into a flash array, these are differentiators that are currently in flux in the market. If you're hearing about something new that isn't on this list, give us a call and we'll be happy to answer any question you may have.

## HYBRID FLASH ARRAYS: HOW TO SOLVE THE MONEY PROBLEM

The cutting edge of processor technology is nothing short of amazing, and their capacities continue to increase according to Moore's law. But one area that has not made the same progress happens to be the one area that data centers desperately need to catch up: storage.

There is a performance gap between the capabilities of today's processors and the storage technologies that hold the data for those processors to do their work. But because all arrays don't have to access all their files all the time, there is a way to balance performance with cost.

#### SKEW: MEETING DEMAND FOR DATA

Some data is accessed more than other data—that differential is called "skew," and it describes how unevenly the data in an array is used. The specific details vary for each array, but generally speaking, 20% of the data stored in an array is accessed by 80% (or more) of the I/O requests.

In a nutshell, skew means that you can dramatically improve the performance of a given workload by speeding up the performance of the data with the highest access density. With smart software within the array that can detect data access patterns and redistribute data accordingly, you have the ability to significantly improve performance without needing to move all of the data onto solid-state devices. That brings us to hybrid flash arrays (HFAs).

HFAs are a simple enough concept: blending spinning disks with SSDs in a single array to boost an array performance without paying the premium for an all-flash array. Consider it like a supercharged minivan that gets off the starting line quickly but still has all the capacity you need for people and groceries. But as is always the case with flash, the reality is a little more complicated.

## "CONSIDER IT LIKE A SUPERCHARGED MINIVAN THAT GETS OFF THE STARTING LINE QUICKLY BUT STILL HAS ALL THE CAPACITY YOU NEED FOR PEOPLE AND GROCERIES."

### HYBRID FLASH ARRAYS: FOR ULTIMATE CONTROL

HFAs can use a mix of SSD, high-speed SAS, and/or slower nearline SAS (NL-SAS) drives within the same array, creating a multi-tier storage system. Getting the right mix is a highly custom process—your workload skew shows what you need out of your array, so you can know exactly what to put in it. In addition to the I/O requirements, you also need to consider how data moves across tiers and the speed with which the array needs to move data. It all has an effect on the right drive mix.

Whatever goes into an HFA, its success or failure depends upon processing power and software capability within the array to relocate data based upon access patterns. In this way, multi-tier arrays can accommodate a wide variety of workloads, scale to evolving data access patterns, and provide the best price/performance ratio.

#### THREE WAYS TO MEASURE VALUE

There are three ways to measure the value of HFA, but none of them are perfect. All are dependent, of course, on your workload analysis and what your customers need from your server capacity; but these four points will at least give you a reference point to help decide on which array to use and how to measure its value once it is installed.

The first measure is \$/GB. It is perhaps the least effective measure because of the processor/data storage performance gap we described. While cost per capacity is a consideration for every decision, it should only be one of the factors in evaluating a solution unless the performance of the storage solution is unimportant.

The second measure is \$/IOPS. This is a better way to value your array, because it says more about what kind of performance your array delivers rather than simply the capacity available to store data.

The third measure can be more difficult to calculate, but it's a much more accurate representation of how two or more solutions may differ in the long run: total cost of ownership (TCO). What you know from the start is that form, energy output, and maintenance all go down with flash, putting more money in your organization's pocket. The most difficult thing to measure is the new business opportunity gained because your personnel now have more time for other projects. Even a rough estimate of TCO is worth the effort to obtain, because it shows your over/ under for the life of the array.

## HOW TO MAKE THE RIGHT DECISION

You're constantly in the business of upgrading your servers as old servers reach the end of their life. But knowing what to invest in to future-proof your data center is tough. It depends almost entirely on analyzing your server's workload and mapping your current and future needs to the arrays best suited to your needs.

We've discussed flash storage at length here and in other places on our site, but all-flash arrays are not the end-allbe-all solution for every organization. Hybrid flash arrays may be exactly what you need to solve your business needs. Whatever decision you make, you need to know what you're buying and why.

This eBook is a quick rundown of what we consider to be the most salient points of how flash impacts your data centers and your business. But it's hardly a conclusive résumé of all things flash—if you have more questions, please give us a call. We're happy to give you a consultation and help you find the right vendors with the right solutions to keep your data centers running at peak efficiency and keep your bottom line growing every day.

#### We'll talk to you soon.



if you have questions or would like more information on any of Red8's solutions or services please contact us.

611 Anton Blvd Costa Mesa, CA 92626 www.redeight.com

info@redeight.com sales@redeight.com