




Data Storage Infrastructure

Multi-tier Storage You Always Wanted But Could Never Afford

Written by
Marc Staimer,
President, Dragon Slayer Consulting



Introduction

The history of data and its storage has always been a story of hot, cool, and cold. Starting with the hot. Users want fast application response times. There's good reasons for that. IBM research on application response times have shown when application response times exceed 2 seconds, user attention begins to wander and productivity slows. At 3 seconds or more productivity plummets. As application response times fall below 2 seconds, productivity increases. When the application response times drops below a second to approximately 400 milliseconds (ms) or less, productivity soars. Doubling productivity for the novice user and increasing by an order of magnitude for the expert user. This results in much higher morale, lower turnover, reduced costs, faster time-to-actionable-insights, faster time-to-action, faster time-to-market, and faster time-to-unique-revenues and profits. Faster application response times are good for the bottom line.


Getting those faster application response times requires the storage to have low latencies, high IOPS, and high throughput. That type of storage is usually referred to as primary storage, a.k.a. tier 0 or tier 1 storage. Primary storage is also the most expensive storage. The question becomes what happens when that hot data cools off? Usually, not much. The data tends to be sticky. It commonly stays where it originally landed in the expensive primary storage. It just stays there because IT organizations don't have to do anything. It's easy. They see data storage as cheap. It's simple for them to just buy more. Simple and cheap is not the same as low cost. Keeping that data on primary storage is in fact, quite costly.

As long as the application data is hot – i.e., used – and the fast response times are necessary, it makes sense for it to live on the fast, expensive, primary storage. But when it cools off – starts to cool after 3 days and becomes pretty cold typically after 30 days – that rationale becomes deeply flawed. Over time, the mismatch between the data value and it's consumption of costly primary storage gets wider. This seemingly simple cheap decision continues to grow to chasm like gaps as the storage is refreshed because that cold data is now migrated to the new expensive primary storage. Plus of course the cost of the data migration.

This created demand for secondary storage – cool, a.k.a. tier 2 – and tertiary storage – cold, a.k.a. tier 3. Both cost far less than primary storage at a performance level that is also considerably slower. Secondary and tertiary storage tend to use either HDDs or slower, low write-life, high capacity SSDs . Whereas primary storage mostly uses fast and high write-life SSDs . Secondary storage is also primary used for workloads requiring lower performance. Workloads such as websites, office applications, cad/cam, photos, audio, drawing, file sharing, and video editing. Secondary data storage systems commonly use NFS, SMB, or object S3 interface protocols versus block that is more common in primary data storage systems. Primary, secondary, and tertiary storage are primarily completely separate systems. Secondary and tertiary storage can also be tiers within a primary storage system. Or tertiary storage can be a tier within secondary storage. All three variations have known complex issues.

1. Generally they are 3D QLC SSDs, but are sometimes 3D TLC. Varies by vendor.
2. These are commonly a mix of SLC, MLC, TLC, or 3D TLC SSDs. Varies by vendor.

Premises




Aligning data value, data performance – the application response time – with storage performance and a cost that matches the data value, requires data migration between storage tiers. It is considerably easier to move data within a storage system than between storage systems. Especially when separate storage systems are different models or from different vendors.

Intrasystem data tiering movement tends to be automated based on policies or machine learning, depending on the vendor and the storage system. Some use HSM to move the data. That type of movement is extremely inefficient because of its stub architecture. When an application attempts to read or modify the data, it has to be moved back to its original tier first, whether it needs to or not. HSM adds a lot of additional latency and longer application response times, while rapidly reducing drive wear life. HSM stubs are known to have breakage issues that can orphan the data.

Other primary storage systems migrate data to lower tiers internally based on policies. Policies such as data age or time passed since the data was last accessed. That's albeit a better methodology, but too often, the accessed data still needs to be rehydrated – moved – back to its original tier to be read or modified because the system believes the data is hot again. So even this methodology often has similar issues of higher access latency, longer application response times, and reduced drive wear life.

Most secondary and tertiary storage within a primary data storage system come at a considerably higher cost than stand-alone secondary or tertiary storage systems. And that cost continues ongoing as the primary storage system is tech refreshed every three to five years.

Premises



Many primary storage systems additionally are single access protocol type (block, file, or object) focused. Even those that are “unified”, utilize one primary access protocol and convert the others to that protocol. A good example is one well known file storage vendor that supports block. However, the block access is converted to file inside the storage system so that the system’s storage services can be applied to the block data as well as the file data. That conversion adds latency. Higher latency increases application response time. Higher application response times result in reduced productivity.

But then again, all is not milk and honey with separate storage systems either. Intersystem data tiering between different storage systems is far more complicated, time consuming, and error-prone. It generally requires costly 3rd party software products to automate the data movement on high bandwidth, low latency networks. Or it will require extensive internal use of scripting and/or open source products requiring extensive modifications, testing, quality assurance, documentation, and manual labor.

Remember that the reason those stand-alone lower tier storage systems are cheaper is because they have reduced system performance and drive performance with less primary storage functionality. Functionality such as extensive snapshots, cyber resilience, replication, analytics, and more.

Premises

That generally leaves IT with three unsavory options:

1. Pick a primary storage system with multiple tiers, but expect:

- ▶ Higher cost.
- ▶ Higher latency.
- ▶ Lower performance from the secondary and tertiary tiers.

2. Choose separate secondary and/or tertiary storage, but expect:

- ▶ Complicated, time-consuming, error-prone, even data loss in the data movement.
- ▶ Multiple systems to manage, operate, troubleshoot, tech refresh, etc.
- ▶ Reduced functionality and performance.

3. Keep cool and cold data on expensive primary storage media in the tier they landed on.

- ▶ i.e., do nothing.

The obvious ideal solution would be a high performance multi-tiered primary storage with Enterprise functionality, availability, and resilience at the cost of tier-2 storage. Unfortunately, that has simply not been an option...until now with the StorONE Data Storage Platform.

3. Hierarchical Storage Management



Data Storage Infrastructure

StorONE Data Storage Platform


The StorONE Data Storage Platform is unique storage software. It converts two common off the shelf (COTS) servers into extremely high performance, multi-tier, Enterprise class storage, with functionality that goes beyond Enterprise primary storage, at secondary storage costs. The question becomes then, how did they do it? It's a good question. A deeper dive provides insight.

The StorONE Vision

Uniquely capable storage software that solves the difficult data storage system problems user are commonly resigned to live with. But do so at an affordable cost – not just price.

The StorONE Data Storage Platform Heart

StorONE's Data Storage Platform is a remarkable storage engine. It's highly robust and efficient code empowers exceptional performance while exceeding the performance and storage services of many Enterprise multi-tier storage. It does this by consuming the absolute bare minimum of CPU and memory resources. Accomplishing that magic trick required StorONE to develop hyper-efficient cacheless software. This is a non-trivial task taking many years of diligent development.



The results are astonishing. The StorONE Data Platform produces the lowest latency, maximum IOPS, and highest throughput possible from two Intel/AMD servers. It does this from the fewest number of NVMe SSDs. That's just the beginning.

Too many primary and secondary storage systems see their performance begin to degrade at 50% capacity utilization. The storage systems become increasingly unstable especially when capacity utilization approaches 80%. The StorONE Data Storage Platform sees no decrease in performance or instability as capacity utilization exceeds 90%. This is a huge cost savings. What about those storage services?

IT organizations expect a broad sweep of storage functionality today. The StorONE Data Storage Platform delivers in spades by exceeding those expectations. These capabilities are the soul of the StorONE Data Storage Platform.

4. Cacheless improves reliability as it eliminates data loss risk or corruption from cache coherency or cache hit failures. By directly accessing and storing data on the underlying media achieves a much higher level of data integrity and application consistency is achieved.

Data Storage Infrastructure

StorONE Data Storage Platform



The StorONE Data Storage Platform Soul

Hyper-efficient code not only frees up resources for performance, it also enables the StorONE Data Storage Platform to provide storage services not available in most Enterprise primary storage systems let alone secondary storage systems. These are the soul of the StorONE Data Storage Platform.

Storage services such as:

- ▶ AI Smart Auto-Tiering.
- ▶ Virtual Storage Containers™ (VSC).
- ▶ vRAID Enhanced Redundancy.
- ▶ Robust Data Protection, Rapid Recovery, and Granular Management.

5. The StorONE Data Storage Platform supports most standard storage protocols simultaneously. Protocols include iSCSI, Fibre Channel (FC), Infiniband, NVMe-oF (RDMA and TCP), NFSv3, NFSv4.2, SMBv2, SMBv3, SMBv3.1, and object S3 compatible.

Data Storage Infrastructure




**Multi-tier Storage You Always Wanted
But Could Never Afford**

StorONE Data Storage Platform

AI Smart Auto-Tiering

StorONE's AI Smart Auto-Tiering is highly advanced data placement. It intelligently and dynamically puts data across multiple storage tiers based on sophisticated autonomous code. That code eliminates the need for customer intervention in setting up policies or parameters on when data should be moved from higher tiers to lower tiers. Data is automatically shifted between high-performance and high-capacity media producing the ideal balance of performance, cost and data value. It utilizes real-time data usage patterns to automate that movement. Data identified with lower value is placed on higher density, lower performance, and lower cost media. Hot and warm data – high value data – are placed on lower density, higher performance, and higher cost media.

This optimized auto-tiering meaningfully reduces storage costs and administrative overhead without user intervention, scheduled downtime, or outages of any kind. An impressive aspect of StorONE's optimized auto-tiering is how it dynamically adjusts the upper tier size per-volume based on real-time workload requirements. It also boosts hardware utilization by combining both random and sequential writes and normalizes them in primary tier, then sequentializes the writes to the lower tier for efficient and high-performance migration. This accelerates the entire process while extending media wearlife regardless of the type – QLC SSDs or HDDs – in the lower tier.



Data Storage Infrastructure

StorONE Data Storage Platform

Virtual Storage Containers (VSC)

VSCs isolate logical storage functionality from the physical hardware. This is analogous to bare metal servers with containers isolating logical server functionality from the physical hardware. VSCs simplify storage services, access protocols (block, file, object) per volume, isolate each volume from one another, while sharing from the same media pools. Protocol support easily extends to emerging storage protocols and technologies, ensuring future-proof connectivity.

The hardware abstraction additionally enables use case-specific media pools (flash/HDD or combinations) based on application requirements. It can also eliminate the noisy neighbor through shared nothing pools so workloads can be isolated.


This allows normally resource-competitive workloads to co-exist within the same StorONE Data Storage Platform without impacting one another. This eases administrative overhead, simplifies deployment, while reducing support costs.

VSCs moreover permit mixing of storage media, expansion of existing storage pools, or integration of new storage media into existing systems.

That ensures ample capacity and capability for current and future data storage requirements while allowing capacity on demand.

Capacity utilization is intelligently managed adaptively via thin provisioning. The dynamic storage allocation enables resource allocation on-demand both within and between VSC volumes. It also allows users to consume $\geq 90\%$ of the StorONE Data Storage Platform capacity without performance loss or instability.

Keep in mind VSCs are managed and operated as a single system. One platform for primary, secondary, and tertiary storage.



Data Storage Infrastructure

StorONE Data Storage Platform

vRAID Enhanced Redundancy

StorONE Data Storage Platform virtual RAID or vRAID is a major step up from traditional RAID. Traditional RAID has problems with significant overhead consumption, long rebuild times, reduced system performance during a single drive rebuild, and plummeting performance with two concurrent drive rebuilds. Erasure coding was supposed to be and heavily promoted as the next generation of RAID. Unfortunately, there was a significant performance issue. Erasure coding added considerable latency to writes and reads making it a bad fit for primary storage workloads. That high latency has relegated it primarily to secondary and tertiary storage.

StorONE has changed the game by taking erasure coding to the next level. With hyper-efficient softwares, StorONE Data Storage Platform has brought erasure code latencies down to equal or better than RAID. This means vRAID erasure coding is appropriate for all workloads.

By distributing vRAID data across all drives in a pool, data resilience is markedly improved and the benefits are huge. The first is a major reduction in storage consumption overhead. Instead of 50-200% RAID consumption overhead, it's reduced to 25%. That translates into much greater usable capacity.

Better yet, when one or more drives fail, only the data is rebuilt, not the entire drive. Rebuilds are orders of magnitude faster because many more drives are utilized in rebuilding the data. And best of all is no negative impact on performance for writes, reads, and even rebuilds.



Data Storage Infrastructure




**Multi-tier Storage You Always Wanted
But Could Never Afford**

StorONE Data Storage Platform

Robust Data Protection, Cyber Resilience, Rapid Recovery, and Granular Management

StorONE Data Storage Platform has incomparable data protection. It starts with vSNAP. vSNAP is so efficient it allows unlimited snapshots. That means it can achieve a recovery point objective (RPO) every 60 seconds with nominal capacity utilization and no effect on application workload or storage performance. Ongoing snapshots are highly granular and differentiated by the changed data. vSNAP specifically enables rapid recovery with the ability to roll back to specific points in time, minimizing downtime and ensuring business continuity.

vSNAP snapshots are immutable naturally. This prevents any changes, deletions, or unauthorized encryption to the snapshots – a must today for cyber resilience. By capturing immutable snapshots at the granular level, critical data remains unaltered and safeguarded against changes. This advanced level of data protection provides an extra layer of defense against data loss, malicious actors, and ransomware attacks.



Data Storage Infrastructure

StorONE Data Storage Platform

When combined with multi-factor authentication, multi-administrative authorization, electronic air-gapping, anomaly detection – that continuously monitors system behavior for any abnormal patterns that could signal a security threat or system, real-time logging, monitoring,

and alerting of potential security anomalies within the StorONE Data Storage Platform, StorONE is delivering a very robust cyber resilience indeed.



Replication is another must for disaster recovery (DR) today and vReplicate delivers. It simplifies and centralizes replication management with an intuitive interface for efficient replication control and management. It comes with flexible replication topologies adjustable to the application workloads' specific requirements. This allows data to be replicated between diverse storage systems, data centers, or cloud environments.

Replicating data copies in different locations ensures enhanced data availability and utilization, minimized downtime and faster business continuity in the face of outages or disasters. Of course vReplicate can also facilitate more than just Business Continuity and Disaster Recovery.

It enable cloud bursting, data sharing, data replication to data lakes, and more. All without impacting production workload performance.



Data Storage Infrastructure

Conclusion

For years, IT organizations have been promised high performance multi-tiered storage that would align data value with storage costs. And for years they've been left with miserable and unsavory choices.

They could:

1. Pick a primary storage system with multiple tiers, but expect higher cost, higher latency, and significantly lower performance from the secondary and tertiary tiers.

2. Choose separate secondary and/or tertiary storage, but expect complicated, time-consuming, error-prone, even data loss in the data movement with multiple systems to manage, operate, troubleshoot, tech refresh, etc.

3. Keep cool and cold data on expensive primary storage media in the tier they landed on; i.e., do nothing.

None of these were or are distasteful choices and all cost way too much.

StorONE is the first to solve this problem with the StorONE Data Storage Platform. It's an affordable high performance multi-tiered storage.

The StorONE Data Storage Platform is uniquely:

- ▶ **Hyper-efficient and high performance.**
- ▶ **Automated multi-tiered storage – primary, secondary, and tertiary storage.**
 - ▶ All flash.
 - ▶ Mix of SSD types.
 - ▶ Mix of SSD and HDD.
- ▶ **Multi-protocol – block, file, and object – via Virtual Storage Containers.**
- ▶ **Robust data protection, cyber resilience, and fast recoveries.**

In other words, the StorONE Data Storage Platform is the ideal storage for all workloads, all tiers, at an affordable cost.

For More Information

Go to ▶ www.storone.com

6. The StorONE Data Storage Platform automatically detects, alerts, and notifies immediately upon discovering any suspicious activities.

Data Storage Infrastructure



**Multi-tier Storage You Always Wanted
But Could Never Afford**