

The Science of Storage

5 considerations to scrutinise your storage infrastructure

Deep dive into data lifecycle management, object storage, streaming data at the edge and how the Dell Technologies Power Series storage solutions offer extreme flexibility, scalability and security, to future-proof your storage needs.

Data#3

DELLTechnologies
TITANIUM PARTNER

Introduction

Data is everything; and today almost everything is data. Managing and leveraging the unprecedented volumes of information being produced is becoming increasingly complex.

Organisations grapple with the need to integrate storage across their entire IT stack – from networking and compute, to database and applications. Add to this the long-standing cloud versus on-premises debate which has dissolved to a de facto multi-cloud approach.

Now it's more a question of:

1. **How many clouds?**
2. **What workloads do I leave on-premises?**
3. **How do I split my storage between on-premises and cloud?**
4. **How can I centrally manage it all?**

Combining the security and control of an on-premises environment with connected access to cloud-based services and functionality, is how today's organisations unlock the best of modern storage solutions. Scalable limits, consumption-based pricing, simple disaster recovery and faster application deployment is waiting for you, with the Dell Technologies Power Series.



The Dell Technologies Cloud tackles the challenges of hybrid and multi-cloud environments, making it fast and simple to deploy and manage workloads across your infrastructure.

Data#3 is the first and only ANZ partner to attain the Dell Technologies Cloud Platform certification.

[Read the news story here](#)



Contents

1. Making sense of storage / **4**
2. Solving the data lifecycle problem / **5**
3. Object storage, welcome to the mainstream / **8**
4. Object Storage answering the cries of cloud native apps / **11**
5. Streaming data, the next challenge for storage infrastructure at the edge / **13**
6. Storage strategies under the spotlight. Learn from your peers at Hydro Tasmania, AngloGold Ashanti and the City of Kwinana / **15**

1. Making sense of storage

Storage remains a critical component of any business and it's becoming more complex than ever. You're juggling the need to harness the power of data through analytics, with consolidating old, new and emerging storage technologies, while also optimising your mix of storage mediums from on-premises to cloud.

Meanwhile, embracing the shift towards consumption-based pricing models with Everything as a Service changes the economics and is another unstoppable trend - but perhaps it shouldn't become a de facto standard across the board. You should be applying it in the right way, for the right workloads.

The core message is that the way we think about and utilise storage is forever changing and IT teams need to keep across these continual shifts.

Just as we've seen with the rise of data, analytics and security, the right storage strategy will underpin your next decade of infrastructure as new technologies, industries and applications continue to take centre stage.

2. Solving the data lifecycle problem

Collect, process, analyse, share, retain, destroy – commonly referred to as the data lifecycle. There are variations that split some of these groups into smaller segments, but essentially all organisations need to manage some form of this data cycle.

The problem?

The complexity of proactively managing mountains of data to this level. Knowing where your data is stored at any particular time – in the cloud or on-premises – still remains a difficult starting point. Is it trapped in application specific silos, or languishing in forgotten infrastructure? How much dark data is your organisation carrying?

Once you do know what data you have and where it lives, how will you manage data as it ages and becomes less relevant? When you know the information is no longer required, how do you mark it to be moved, archived or destroyed? Can these processes be automated for more efficient management?

Despite the best intentions, most organisations are focused on storing and backing up data and find themselves struggling to cope with the unrelenting growth of data and its impact on existing storage infrastructure.

BIG bytes of data

Byte is the unit of measurement for magnitude of data.

1 byte = a single character
10 bytes = a word

1 Petabyte (PB) = 1000^5 bytes

1 Exabyte (EB) = 1000^6 bytes

1 Zettabyte (ZB) = 1000^7 bytes

200,000,000,000,000,000,000,000 bytes or 200 ZB of data will be stored globally by 2025¹.

¹ Arcserve. (2020) The 2020 Data Attack Surface Report. [Online] Available at: <https://1c7fab31m83f5gqjow2qqs2k-wpengine.netdna-ssl.com/wp-content/uploads/2020/12/ArcserveDataReport2020.pdf>

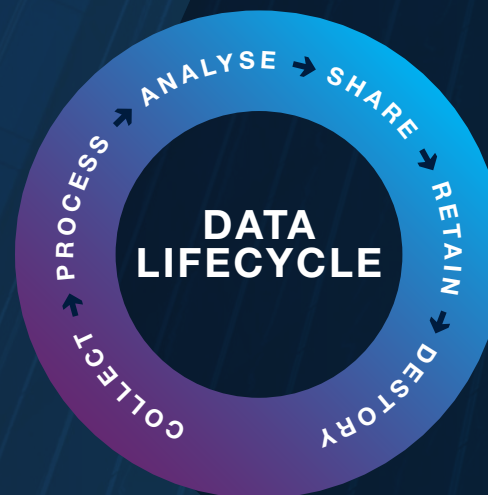
The concept of dataset management offers a solution

Dataset management is an evolving discipline that focuses on holistic data management throughout its lifecycle, rather than just storing and moving data. It takes storage from the old world of directories and files into the modern world of data and metadata.

Dataset management uses various approaches and technologies; like machine learning to add context to your data across infrastructure and data silos, whether on-premises or cloud-based. This not only provides the ability to instantly find data wherever it's located, but also makes it actionable, by leveraging that context for decisions on what should happen to the data next - allowing levels of automation to be introduced if required.

The better you can scrutinise your storage needs, the better decisions you can make about the infrastructure required to house and manage your data.

This proactive management of data according to the data lifecycle has significant financial benefits to your total cost of ownership. By uncovering optimisations in your storage usage, you can make informed decisions, such as reserving expensive high availability storage infrastructure for data that really needs it, or by moving less important data to cold storage options like Dell Technologies PowerVault or Azure Blob. This will allow you to maintain visibility, searchability and security over data whilst reducing costs.



Ask yourself:

- How much data do we have and how much growth are we expecting?
- Who needs access to this data?
- How often do they need access to this data?
- How quickly will I need to recover data from archive?
- Is there legislation or policies this data needs to comply with?
- How is the data secured?
- When do we delete data and how do we do so?



Manage massive amounts of unstructured data with the PowerScale Family

Delivers up to
80% storage utilisation

OneFS

The software-defined architecture that powers the entire PowerScale family.

Up to
16x faster
to identify HA problems

CloudIQ

Proactively monitor storage and use to reduce risk, plan ahead, and improve productivity.

Monitor up to
70 clusters and 2000 nodes

DataIQ

Intelligent insights and dataset management to speed up the data lifecycle and provide more value for the business.



Next-generation enterprise storage with PowerMax

An end-to-end non-volatile memory express (NVMe) multi-cloud storage solution, PowerMax architecture includes a true scale-out design with storage class memory (SCM) for persistent storage.



[Learn more about PowerMax](#)



3. Object storage, welcome to the mainstream

Demand for object storage is growing largely due to the rise in cloud-native app development and the significant impact of data-intensive analytics. Why? Flash memory and NVMe have reinvented what was once a slow and cold archiving system to a shiny and contemporary cloud-native storage solution.

Meanwhile, unstructured data – everything from email, images and documents, to videos, social media content and even application logs – have been exploding, to the extent that IDC predicts that **80% of data will be unstructured by 2025²**.

Finding a solution for managing all of this unstructured data has been a challenge. Next-generation applications that can handle rapidly growing unstructured data typically require the extreme performance of all-flash storage, but budget pressures have made it hard for organisations to commit to the new capital and operating expenses that these powerful systems require.

Object storage's flat structure and near limitless scale was designed to manage large amounts of unstructured data, with powerful metadata associated to broadly promote its use. This was a perfect environment for big data applications, however access and performance issues hindered its adoption in the marketplace. This has now changed and while object storage is definitely mainstream, it's still just one option in a growing sea of storage alternatives.

² King, T. Solutions Review. (March 28, 2019). 80 Percent of Your Data Will Be Unstructured in Five Years. [Online] Available at: <https://solutionsreview.com/data-management/80-percent-of-your-data-will-be-unstructured-in-five-years/>

Flash-based object storage has overcome its initial teething problems with the ability to automate the process of tier to disk, creating an on-demand consumption model. The ability to move data from one tier to another when you need it, and into cold object storage when you're done, offers significant financial benefits. It has also removed previous capacity constraints, making it an interesting option for the development and analytics scenes. Machine learning, connecting limitless amounts of IoT data, coding and app development are not only highly achievable with flash-based object storage, but this technology provides a perfect tool for fostering innovation.

By allowing these data-intensive, time-critical priorities to take place in the flash tier, you essentially reserve that space for the most complex and sophisticated projects, with no competing accessibility issues. This takes the pressure off information sitting in cold object storage, which can be used to continue initiatives that don't have similar demands for performance and low latency.

Object storage has also proved itself a more efficient alternative to tape backup solutions, making it an ideal solution for both data retention and long term back up.

As object storage continues to evolve with flash and NVMe, the opportunities exist to bring in images, event data, logged information and other machine-generated inputs, which traditionally have been in file-access storage. For leading organisations, this has the potential to significantly increase delivery speed and performance and support greater analytics use cases - bringing it firmly into the mainstream.





High performance object storage with Dell Technologies unstructured data solutions

48%
Lower TCO
than Public Cloud

ECS Object Storage

Brings cloud scale, geo-distributed object storage to your enterprise, ensuring unmatched scalability, manageability, resilience, and economics to meet the demands of modern business.

12 GB/s
bandwidth
per chassis

Isilon H400 Hybrid NAS Storage

Provides operational flexibility with built-in multiprotocol support, supporting a wide variety of unstructured data workloads on a single platform.

Manage
petabytes
of data with a
single admin

PowerScale OneFS

Combines the benefits of traditional NAS storage and emerging object storage to provide an enhanced data lake capability and cost-effective unstructured data storage solution. For a PowerScale solution overview or virtual appliance tour [click here](#).



Powering your unique storage needs

Expanding capabilities for media and entertainment

Featuring secure collaboration, modular expansion, flexible consumption models and easy cloud integration, Dell Technologies PowerScale is an ideal choice for small VFX houses through to multi-national broadcasters.

Driving innovations in genome sequencing

Dell Technologies' portfolio of unstructured data solutions deliver the performance and ease of management needed to meet the large-scale processing challenges of gene sequencing and other biomedical research, helping to accelerate innovation in these fields.

4. Object Storage answering the cries of cloud native apps

With the rise of Kubernetes in powering containerisation and the rapid development of cloud-native apps, customers are looking to deploy object storage solutions - without having to procure and manage additional storage systems.

With the release of VMware Cloud Foundation with Tanzu and the integration of Kubernetes into vSphere, VMware has shown they are serious about providing optimum environments for running cloud-native apps in the enterprise. The result? A welcome opportunity for the next evolution of object storage.

One of the principles of object storage is that it is abstracted from the compute layer via APIs. That can be extended to take advantage

of Kubernetes native automated deployment, scaling and management capabilities, defining a next-generation of object storage. It will be lighter, faster and deployable anywhere that VMware Cloud Foundation with Tanzu is running – the data centre, the edge or public cloud.

This gives organisations the power to put data close to the applications they support, reducing latency and improving the user experience. In addition, object storage from disparate platforms will be able to cross-replicate for greater access, reliability and redundancy. You can utilise whatever object storage environment makes sense economically and with respect to the data lifecycle and how readily you need to access that data - be it on-premises or cloud.

ObjectScale is a cloud-native object storage solution that works with Tanzu to provision object storage via Kubernetes. It's a distributed, scale-out multi-tenant system with S3-compatible APIs and a single global namespace. Using ObjectScale, any organisation can deliver scalable cloud services with the reliability and control of private cloud infrastructure.

This also means developers and hyperscalers can access object storage data from anywhere, even if it exists in a globally distributed infrastructure, because architecture design principles allow it to be managed under one namespace, even with interconnecting platforms. This enables hybrid cloud strategies that apply cloud capabilities to high performing on-premises storage infrastructure, making it easier to manage by giving IT departments the control and support needed to scale cloud applications.

When deploying object storage for cloud-native apps and services, consider the following:

1. Architecture layering is important for scalability and flexibility. When operating in this distributed platform, it's important to make sure your layers are independently available for scaling at all nodes.
2. Ability to manage globally distributed storage infrastructure under one namespace without compromising on access to data, saving time and effort.
3. Easily accessible object storage that is tightly governed and securely connected to containers and virtualisation platforms.

ECS Object Storage is an enterprise-grade object storage platform that allows organisations to flexibly capture, store, protect and manage unstructured data behind the enterprise firewall, with scalability that rivals the public cloud.


5. Streaming data, the next challenge for storage infrastructure at the edge

There is a growing argument to rethink the treatment of data as streams, rather than individual files or objects. This means log or event data can be consumed in the same way as watching a live football game on an app, giving you the option to pause, play and rewind so you don't miss critical moments.

However, IT infrastructure wasn't designed to cope with this type of on-demand, real-time playback function. Ingesting and processing stream data has been limited by traditional storage solutions, although improvements in CPU performance and machine learning can now provide instant autoscaling to manage streaming workloads fluctuating throughout the day.

In order to deliver this capability, streaming data needs to be processed chronologically, logically, cumulatively and one record at a time so organisations can respond to situations almost immediately. An example is digital advertising; the moment you've browsed a product, or scrolled through property listings, recommendations soon appear on social media feeds. That type of immediacy can open up a new world of possibilities for your marketing department - yet managing and storing that data in your current infrastructure will require some clever new approaches.





Today, overcoming the challenges of streaming data remains critical for industries processing a lot of data at the edge, such as universities streaming lectures, media and broadcast publishers, security companies managing mountains of CCTV or organisations overseeing large venues like stadiums, shopping centres and airports. As data stores begin to reach exabyte scale across industries, IT departments will be challenged with capacity and compute issues like never before, demanding a solution that is truly agile.

SDI storage options are a strong candidate to manage high-performance workloads originating from streaming data, while **container-based architectures** can also address performance and stability risks. Your ideal infrastructure may be a mix of both, but what will be critical is a storage solution that's smart enough to scale automatically as soon as the need arises.

With the ability to leverage data as streams for capture, playback and archive, a whole new era of industry-based applications will be possible.

6. Storage strategies under the spotlight

Hydro Tasmania

With their existing data centre nearing end of life, Hydro Tasmania was looking to reduce risk and increase the flexibility and room for growth in their storage solution.

The team turned to a hyper-converged solution, implementing 40 Dell EMC VxRail storage nodes across four data centres. The agility of the new solution meant that as new workloads are introduced, additional nodes can simply be added, with minimal impact to business.

“ The hyper-converged data centre infrastructure that Data#3 offered from Dell Technologies gave us lots of benefits in terms of flexibility, scalability and ease of operation.”

David Ovington, Program Manager,
Hydro Tasmania.

Dive into Hydro Tasmania's complete storage story



AngloGold Ashanti

AngloGold Ashanti Ltd needed to replace complex legacy architecture at their Tropicana gold mine in Western Australia, to improve application performance, minimise costs and strengthen their disaster recovery.

Leveraging the Dell Technologies portfolio, Data#3 stepped in with a customised solution including PowerEdge servers, Unity Storage, a Data Domain upgrade, PowerVault TLO Tape Library and VMware vSphere.

The outcome, a robust, scalable and secure infrastructure that provides greater capacity, is easier to manage, and helps reduce risk through lower RPO and RTO. Improved application performance and reduce complexity to streamline operations has ultimately delivered a more efficient and productive gold mine.

“ The simplified technology stack helps us to streamline application performance, increase capacity, reduce risk, and minimise costs.”

Adam Wahby, Systems & Infrastructure Specialist, AngloGold Ashanti.

Dig deeper into AngloGold Ashanti's storage success



City of Kwinana

The City of Kwinana south of Perth, knew their cloud-based storage wasn't working. The solution was slow, patchy and didn't support mobile access. Download speeds during peak times could slow to 10Mb/s and server outages were common.

"We had almost no control over our own data", reported Mathew Smith, IT Manager at the local council.

The technology overhaul saw the city deploy an on-premises cloud-connected solution built on four Dell PowerEdge Servers, alongside management and backup servers. Paired with VMware wireless networks and virtualisation software, as well as Dell Endpoint Management for devices such as Dell Latitude Rugged tablets, the city was able to regain control of their data. Smith estimates the new solution is saving the city \$300,000 annually.

" Dell Technologies actually exceeded the specifications we wanted – well within our budget."

Matthew Smith, IT Manager, City of Kwinana.

Discover how the City of Kwinana saved on storage



Next Steps

The way organisations design and deploy storage infrastructure is fundamentally changing. As the number of cloud-native applications increase and the economics of cloud versus on-premises storage evolves, the reality of Everything as a Service matures.

Separate object-based storage, for the huge amounts of unstructured data, and file-based storage, for database and critical business applications, has meant data has become siloed and fragmented. A focus on deploying the right storage infrastructure for specific workloads has sometimes come at the cost of an overarching, cohesive storage strategy.

In this eBook, we've explored five considerations for storage infrastructure, but there are many more questions to be asked if you want to develop a comprehensive storage strategy. You need to better understand your data and how that data should be (or could be) supporting applications and business outcomes.



Drill into questions on workload, capacity, performance and availability and ask yourself:

- How big should it be, how fast does it need to be, how often is it accessed and how do you protect it?
- What corresponding demands are there on compute and memory resources?
- Where can you consolidate and how will that affect the type of storage infrastructure you deploy, as well as the rest of your infrastructure?

Complimentary Data Centre Assessment

Traditionally, the process of gathering actionable workload characteristics and performance details to support investigations like this has been complex and resource-intensive, with the research and evaluation phase typically taking months to complete. Often taking so long the data isn't entirely accurate anymore.

A complimentary **Data#3 Data Centre Assessment** can help fast track the research and review phase. Using Live Optics – a lightweight, remote and agentless software – we can quickly collect, visualise and share in-depth data about your IT environment and workloads. This drastically simplifies the once clunky discovery process because we can analyse your data to gain a deeper, real-time understanding of workload performance. Learn more about the discoveries of a data centre assessment [in this blog](#).

As the Dell Technologies [2020 Transformational Partner of the Year](#) for the APJ region, as well as the **ANZ Server Partner of the Year**, we understand how the right storage strategy underpins your overall infrastructure approach, so ultimately, we can create a mutual data-led understanding of optimal infrastructure requirements.



Book your free Data Centre Assessment today

To discover more Dell Technologies storage solutions from Data#3, experience virtual product tours or to connect with a storage specialist, [click here](#).

Data#3 and Dell Technologies

Data#3 and Dell Technologies have helped our customers build better businesses since 2012. Across the network, data centre, workplace and cloud, we help customers maximise their IT by combining skilled resources with robust and secure technologies from across the Dell Technologies group.

A Titanium Solutions Provider Partner with the largest number of technical certifications, including the **first and only partner in ANZ to attain the Dell Technologies Cloud Platform certification**, Data#3 is uniquely positioned to provide end-to-end infrastructure design and services that enable our customers to optimise budgets and outcomes.

Data#3

DELLTechnologies
TITANIUM PARTNER



1300 23 28 23



facebook.com/Data3Limited



twitter.com/Data3Limited



linkedin.com/company/data3/



youtube.com/user/data3limited