Buyers' Guide
IT Professionals

FREEFORM DYNAMICS

and

SITUATION PUBLISHING

in association with

komprise

Analytics-Driven Storage Management

An evaluation and decision guide for IT leaders and professionals

Freeform Dynamics Ltd, 2020

Management at the storage level is no longer enough

Traditional storage management was about predicting needs, then implementing, optimizing and administering hardware to meet performance, service level and data protection requirements. As part of this, the mantra was often to err on the side of stability, and only change things when absolutely necessary.

The world of storage has moved on, however, becoming far more dynamic and data-rich. In response, storage vendors have added new capabilities and automations, but all too often they have simply expanded the old ways of doing things. If any of the following challenges seem familiar, you'll know that the old ways are no longer adequate.

The desire to keep everything

Whether it's for its potential future value or a regulatory demand, the business increasingly wants to keep all the data it collects, forever.

Data tiering that brings lock-ins

Many advanced storage systems can transparently move older data to a cheaper storage tier, but only if the data stays within that storage system.

Storage platform proliferation

Devolved decision making and the use of newer technologies and cloud services to meet new demands has led to greater storage landscape diversity.

Today's Challenges

Resource, time & skills limitations

Managing storage as demand escalates and complexity grows is hard enough for specialists, let alone the non-specialists now often involved.

Relentless data growth

As well as all the usual business files, videos and so on, almost every process and electronic device now generates a data trail or log which may also be of value.

Costly storage is used inefficiently

Rarely-used data fills up expensive primary storage because it is too complicated or awkward to move it somewhere cheaper and more appropriate.

Data management: your data is now in the driver's seat

The storage infrastructure still matters in our digital world, of course, but now the management focus needs to move to what's really driving storage growth, and that's your organization's data - the files and objects that sit on top of that storage.

In the past, it's been acceptable to manage this data relatively loosely, but as volumes grow so do the physical and administrative costs of buying and managing storage, ensuring backups, etc. A new, more granular and analytical, approach is required.

Let's take a look at some of the key issues and opportunities that you need to consider as you plan for this new and more efficient approach: managing your data directly.

Visibility & insight

Clarity of vision is vital: If you cannot clearly see and understand what you are really managing - i.e. data, not storage - you can only manage it crudely and indistinctly.

Seamless transparency

It is easy to move data, but hard to do it without users noticing. Data needs to flow to the most appropriate or efficient storage location without appearing to move.

Openness & flexibility

Storage lock-in can be a trap. Taking a more open route can reduce costs, allow a greater choice of technologies and services, and minimise end-of-life migration issues.

Emerging Imperatives

TAKE ACTION

New storage technologies

From non-volatile memory and Flash to object storage and cloud archiving, there is great flexibility to create new balances of cost versus performance.

Fit the service level to the data

There's many storage types and services available, offering widely differing cost/performance, so it's increasingly important to have data on the right class of storage.

Efficient use of resources

Your main NAS systems or storage arrays are designed to perform well for active data. Keeping inactive data on them is wasteful and inefficient, but can be hard to avoid.

Addressing the cold data challenge

Clearly it is a waste to have inactive 'cold' data sitting on an expensive primary NAS system. Moving it to online archive storage, say, could keep it available in case it's needed, while saving costs not just on the storage itself but also on the backup workload.

Conversely, most organizations want to keep their active 'hot' data on their primary storage systems. Moving it to cheaper storage, on-site or in the cloud, is likely to introduce latency and may very well have an unacceptable effect on application service levels.

And in all this, the decision-making and data movement has to be as seamless and automatic as possible. If you have to manage manually, you'll soon run up more costs than savings! Here's some of the key technical issues and options to consider.

The hardware route	Many modern storage systems can host multiple types of storage and transparently move data between them. **Most can only tier within the box, however.**
Software for data management	Identifies hot/cold data, moving it up and down through the storage tiers as appropriate. **It typically needs an agent on each connected server, which may not be ideal.**
File stubs for relocation	When data is moved, it can be replaced by a small file called a stub that tells applications and users where to look. **Stubs are a simple approach, but non-standard.**
Symbolic links for redirection	Symbolic links are shortcuts, transparently redirecting file access to a new location. **They are standard NFS & SMB elements, supported by most modern storage.**

Understanding your data: the need for analytics

While seamless data movement is essential, a complete data management solution requires more than that. The movement must be transparent to your users and applications, and as automatic as possible. It also needs to be governed and driven by deep analytics that understand not only how your data is really being used, but also the costs and performance characteristics of those various storage tiers.

Ideally it should also be able to model different solutions for you. And you will probably not want that freedom of movement and user transparency to come at the expense of storage lock-in. Here are some practical questions and issues to consider.

Visibility & insight	Can you see how much data you have, what kind it is, who's using it and how fast it's growing? And can you model the effects of a particular data movement plan?
Openness & flexibility	How free will your choice of storage hardware and services be? Can movement between them be fully transparent to users, file systems and applications?

Ask questions, act on the answers

Data management is much more than just the ability to move data between storage tiers. So as you evaluate data management options, we recommend you ask yourself a series of questions, starting with the list below. For completeness, you could first assess how confident you are in your ability to answer these questions positively today. Then move on to answer them for the data management solutions you're evaluating.

Solution assessment: key considerations

	Level of confidence		
	High	Medium	Low
Will I have full visibility across all my storage to know the "who/what/when" of data access?	☐	☐	☐
Can I see how data growth varies across the organization and its applications?	☐	☐	☐
Can I see data ownership and usage, to help build migration policies and do charge-back?	☐	☐	☐
Can I model multiple data migration scenarios for cost-effectiveness and performance?	☐	☐	☐
Will I need to install software agents or additional applications on my servers?	☐	☐	☐
Once I identify cold data, can it move without affecting users, applications, or hot data?	☐	☐	☐
Once data has been moved, is it still freely accessible via normal file protocols?	☐	☐	☐
How many storage tiers are supported, and can I add more as new options appear?	☐	☐	☐
Is cloud storage fully supported, both for data migration and analysis?	☐	☐	☐
Can I tailor data migration policies to the needs of different users, data-sets, etc.?	☐	☐	☐
Can the same tools replicate data, e.g. for disaster recovery, as well as move or migrate?	☐	☐	☐

About Freeform Dynamics

Freeform Dynamics is an IT industry analyst firm. Through our research and insights, we aim to help busy IT and business professionals get up to speed on the latest technology developments, and make better-informed investment decisions.

For more information, and access to our research library, visit www.freeformdynamics.com.

About Situation Publishing

Situation Publishing brings you the latest news and views on the global tech scene from its offices in London, New York, San Francisco, Sydney — and through the Anseatic Alliance — Europe.

We deliver an unrivalled mix of fiercely independent, analytical, and original stories and events to technology professionals worldwide, giving our readers what they want.

For more information, visit www.situationpublishing.com.

About Komprise

At Komprise, our mission is to help businesses handle the incoming deluge of data with analytics-driven data management that helps customers know, move and control their data to save costs and extract more business value—without disrupting access.

For more information, visit www.komprise.com.