




Inside Track Research Note

In association with



All-in-one Data Protection

Taking the pain out of
information management

January 2018

About this Inside Track

The insights presented in this document are derived from independent research conducted by Freeform Dynamics. Inputs into this include in-depth discussions with IT vendors and service providers on the latest technology developments, along with intelligence gathered from mainstream enterprises during broader market studies.

Tools and processes that were appropriate in the past may no longer be fully up to the job.

Legal and security concerns have become more prominent in data protection, whether due to regulatory demands or malware.

In a nutshell

Once upon a time, data protection simply meant backup and archiving. Now, it can encompass almost anything and everything, from compliance with national and international regulations such as the GDPR, to protecting data against security threats like ransomware. But implementing comprehensive software-based solutions to cover all these requirements can be quite daunting. Could the answer be to explore appliance-based options?

Data protection: not just backup

In today's fast-moving, information-intensive business environment, data protection and management is more complex and challenging than ever. New requirements, such as compliance with the EU's GDPR for any organisation processing the personal data of EU citizens, and new dangers like ransomware, are appearing all the time.

None of this can readily be dealt with using manual processes and scripts, or ad-hoc piecemeal automation. To make matters worse, you may well be relying on multiple tools merely to protect your data from loss or corruption, each with its own separate management system. This duplication of effort would seem like madness if you were setting it up today, but is hard to escape if you have historically had to deal with data protection requirements on a case-by-case basis as they arose.

Put all this together and it is very likely that, even though they were appropriate in the past, your data protection tools and processes are no longer adequate to meet your organisation's evolving needs. Meanwhile, data management and protection technologies have advanced significantly over recent years, so even if you do have working solutions in place, there may still be more efficient and effective ways of doing things.

The key factor here is that your data protection systems cover all of your data – or at least, they should do! That makes them a logical place to build a broad range of information management and governance capabilities, based upon the core file metadata that the systems have already gathered.

Assessing your current position

If it's been a while since you stood back and reviewed your data management and protection arrangements, then consider the list of issues and challenges below that frequently come up in our research. Do you recognise any of these?

- Legal and security concerns have become more and more prominent, whether it is dealing with regulatory demands to audit, edit and delete customers' personal information, or with threats that can delete, steal or encrypt data.
- Business data has become fragmented across application data stores, file shares, PC hard drives, your email system, and other locations, and you can't be certain everything is protected that needs to be.
- As users take more control, documents and datasets get downloaded, copied, forwarded, sync'd and otherwise replicated to the extent that the exact same information is often backed up many times from multiple locations.

Updating all the mechanisms that your backup and archiving processes depend on can require guru-like scripting skills.

Is it time to reconsider your options?

The ideal solution would manage everything – file stores, application databases, email repositories and more – from a central point.

Meta-data would allow your data protection officer to check for GDPR compliance.

- Static, historical data that seldom or never gets accessed clogs up live systems, occupying much-needed space on expensive premium storage devices and being backed up nightly, while you scabble around for budget to buy more capacity.
- Backup jobs take longer and longer to run as data volumes continue to grow, to the point where fitting everything into the available backup window is becoming more of a challenge.
- You periodically archive historical data, but according to relatively crude policies based primarily on age, then users complain when they need something that's been taken offline because it takes so long to retrieve.
- Keeping track of all the mechanisms and jobs that your backup and archiving processes depend on is hard, and some of the measures in place require guru-like scripting skills when modification or troubleshooting becomes necessary.
- You know the organisation is becoming more and more dependent on its virtualised infrastructure, but you've never quite figured out how best to handle the backup of virtual machines (VMs) and the data they might contain.
- You have disaster recovery (DR) policies and procedures in place, but you know that a lot has changed since they were devised, yet the thought of working through recovery time/point requirements and testing against these is daunting.

There are many other potential issues we could add, but if just three or more of the points mentioned above look familiar, it could be time to reconsider your options.

In an ideal world...

So let's imagine we have a magic wand, and could wave it to conjure up the perfect data management and protection environment. What would it look like?

For a start, it would enable everything to be managed from a central point – file stores, application databases, email repositories, virtual environments, and cloud storage. It would allow flexible policies to be set up through a simple point-and-click interface to deal with information backup and restore, archiving, and data disposal. Such policies could vary depending on the type of data, the source system and/or the owner or department involved (which might, for example, dictate regulatory requirements).

Our data protection system would create and maintain comprehensive metadata. This would enable files and objects to be quickly located on backup media or in the information archive. In addition, the metadata would allow your data protection officer to check that any processing or sharing of personal data has been compliant with the GDPR and other regulations.

Our ideal data protection system would also know what's needed to assure backup integrity, e.g. quiesce the databases before starting. It would compress data to save space, and be clever enough not to back up the same stuff more than once. When it came across the second, third or fourth copy of a document, for example, it would simply replace it with a pointer to the original, unless policy dictated otherwise. It would be secure too, to prevent a thief or snoopers from sidestepping the security on your primary systems to access equally valuable backups instead.

Designing, let alone implementing, a comprehensive data protection and archiving environment is not an easy task.

Much of the complexity and specialist skills requirement can be avoided by using a backup and archive appliance.

Advanced features may include policy-driven management to minimise administration overhead.

When procuring either an integrated or target appliance, the first step is to review your requirements and existing landscape.

From an operational perspective, it might be best to base the solution on a dedicated server and storage system, so backup and archiving jobs could be run against live systems without an undue impact on their performance.

You might have come across ideas like these before from software vendors. The chances are, though, that when you looked at how to make it work in practice, the integration effort required was prohibitive. That's because designing, let alone implementing, a comprehensive data protection and archiving environment is a complex process that demands specialist skills and knowledge. That's not just storage skills, it also requires expertise in areas such as regulatory compliance and virus resistance.

Making it easier

Much of that complexity can be avoided by using a data protection (backup and archive) appliance. This will most commonly come in one of two forms:

Integrated appliance: The idea here is to bring the right protection software together with the right hardware and management tools to create a self-contained single product that arrives in your datacentre or computer room in a fully functional and pre-optimised state.

Target appliance: This option provides a sophisticated, pre-optimised storage platform to underpin your existing backup and archiving software, and is particularly useful if you are using more than one backup software suite. Replication and deduplication features are sometimes included to reduce capacity demands and to ensure protection against system failures.

Apart from the speed of initial provisioning, a major advantage of the integrated approach is that the components will have been selected and assembled to work together seamlessly and optimally. This can be especially relevant if you have already taken a similar packaged approach with your compute layer, for example, by adopting converged or hyper-converged infrastructure. In this case, adding one integrated, out-of-the-box appliance to another can make for a simpler approach all round.

Both types of data protection appliance are likely to include functionalities such as data reduction, and replication to a remote site. More sophisticated examples will also feature at least a couple of additional advanced features. One example might be a full policy-driven management environment to minimise administration overhead. Another could be data discovery and classification tools to help with the task of locating personal data or information that's sensitive for other reasons, and protecting it appropriately.

Assessing the options available

When procuring either type of solution, the first step is to review your requirements and existing storage landscape to determine the appropriate size of appliance (bearing growth needs in mind), and the options to be included (e.g. mix of disks, incorporation of tape drives, etc).

Some manufacturers offer a range of solutions at different entry levels and price points, allowing them to populate the appliance with the amount of storage you need

Offerings in this space have been designed both technically and commercially so you only pay for what you initially need.

for the short to medium term, then add more capacity later. As an example, your appliance may be initially delivered with some of the available device slots unpopulated. As your requirements grow, you can then add more drives (of whatever type makes sense), swap out smaller capacity units for larger ones, or add a tier of tape or cloud storage.

You should look, therefore, for offerings that have been designed both technically and commercially so you only pay for the capability you initially need, yet give you the option to expand later as your needs evolve. The trick is to select a point in the manufacturer's product range that provides room for growth, without paying a premium for expansion capacity that you may never exploit.

Some specifics to consider

While solutions are generally built with capacity expansion in mind, the same is not always true when it comes to the breadth of their functionality. That is especially the case if you are considering an integrated appliance – some can be expanded with new software functions but others cannot. This becomes important once you move beyond the fundamentals of backup and restore to include newer elements of data protection, such as building a technical foundation to assist with GDPR compliance and ransomware protection, or even just active archiving.

It's important to do your due diligence and check that the box contains all the functionality you are likely to need or that relevant features can be added.

As you research this area you will come across options at all levels of capability, and it can be tempting to focus just on what you absolutely need in the short term. Remember though that with integrated appliances, unlike solutions assembled from individual components, you may not be able to make big changes to the functional scope down the line. It's important therefore to do your due diligence: check that either the box contains all the functionality you are likely to need in the longer term, or that the relevant features can be simply and cost-effectively added at a future date.

As an example, if your appliance provides straightforward file and email backup and recovery capability based on an open source software component, but does not deal with virtual machine environments or application databases intelligently, then you need to check if you can integrate this later. Similarly, if you decide in the future that you need flexible policy-based archiving, this may be difficult to add in.

And the same is true of the storage layer. If the appliance doesn't include self-management capability such as auto-migration of data between storage tiers based on real-life access statistics, then you probably won't be able to retrofit this later if, for example, you decide to go down the online rather than offline archiving route. Again, the flexibility to add different 'grades' (or even form factors) of disks, e.g. cheaper slower devices to handle longer-term online storage cost-effectively, is important from a future-proofing perspective.

If the appliance doesn't include self-management capability such as auto-migration of data, then you probably won't be able to add this later.

To address many of these issues, some manufacturers have constructed appliances based on state-of-the-art commercial data protection software. If it is done well, this means you may not be forced to compromise from a software perspective. Indeed, sometimes it's the opposite, in that the appliance manufacturer has made sure the right hardware configuration is in place for you to take full advantage of all the software features available.

Having said that, there are situations in which it can be difficult to pin down your longer-term requirements for specific backup, recovery and archiving features. You

Moving to a modern, integrated all-disk based platform can simplify management and allow you to keep more data online.

Implementing an effective data protection and archiving strategy needs to be driven as a business initiative.

Focus on key applications and data, setting default policies for the remainder.

may also have perfectly legitimate reasons to start simple but leave the door open to enhance functionality over time. If either of these use-cases apply to you, an easy option may be to go down the target appliance route. You can then upgrade or extend your existing software applications when the time is right.

Target appliances are also useful to strengthen and enhance your existing software environment. Moving from a storage layer made up of generic disk and/or tape technology to a modern, integrated all-disk or disk/flash hybrid platform can simplify management and allow you to keep more data online to speed access and recovery. And if you simply must use multiple software suites, the consolidated data reduction process may lead to higher deduplication rates, further lowering storage capacity demands.

The right solution may even emulate aspects of your old platform, e.g. by presenting virtual tape devices so existing software and processes are not unduly disrupted as you migrate from your current environment. Similarly, some modern appliances can create an additional copy on cloud storage, as a last line of defence, without the backup processes being aware that this is happening in the background.

But the technology only gets you so far

While an appropriately selected modern appliance removes much of the pain and provides many options from a technology perspective, implementing an effective data protection and archiving strategy needs to be driven as a business initiative. It's beyond the scope of this paper to go into details, but as a minimum you will need to:

- Review your current application portfolio and associated databases or file stores, and categorise documents and data-sets according to business requirements. Considerations here include disaster recovery basics such as recovery point objectives (RPOs) and recovery time objectives (RTOs), along with retention, archiving and disposal policies in line with both regulatory and business needs. It is perfectly legitimate to focus on key applications and data here, setting default policies for the remainder: don't risk your project stalling due to excessive demands on time and resources.
- Step back and take a wider view of your information management needs. Your data protection system should cover all your data and generate metadata on the files it stores. As mentioned, that makes it a logical base for a wide range of other information-related tasks, from archiving to regulatory compliance. Of course, there will be use-cases where this is not appropriate, but they are likely to be the exceptions.
- Don't just focus on the initial setup and implementation of the solution. Once a policy is established, whether at the document lifecycle or physical storage level, it will generally be applied automatically, but policies are often not static. The right data protection and archiving setup therefore will allow you to create a 'living' policy environment, i.e. one in which you can optimise storage, retention, migration and disposal on an ongoing basis. For this reason, you need to think about who is responsible for which policies, and make sure procedures are in place to deal with reviews, escalations, exception handling, etc.

This last point illustrates that in some ways, modern technology in this area can actually create work – but importantly, it's the right kind of work. Rather than crossing

your fingers and hoping on the data protection front, and dealing with storage and retention policy on an ad hoc, manual basis, you can now develop the visibility and an automation capability to deal with such matters in a robust and business-like manner.

Requirements for effective and efficient data protection and management are only going to increase in the future.

The bottom line

Requirements for effective and efficient data protection and management are only going to increase in the future. Transaction systems are producing higher volumes of data at faster rates, and then we have the growing use of business analytics and big data. Meanwhile, users have a seemingly-limitless capacity to create more and different types of content as they collaborate electronically, which creates additional pressures.

If these factors weren't enough, regulators continuously ask for more in the area of compliance, whether to do with personal data privacy, financial integrity or industrial safety. And criminals seek to profit in any way they can, including fooling your staff with infected emails or stealing your data.

Reduce risks, save money, and sleep more soundly.

Against this background, a piecemeal approach relying on old, fragmented and often manual solutions to protect and manage data is most probably not sustainable, even in the short term. Whether or not you go down the appliance route that we have discussed, you still need to get yourself onto a firmer technical foundation. Do this well and you'll reduce risks immediately, save money over the longer term, and generally sleep more soundly.

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 library of free research, please visit www.freeformdynamics.com.

About Fujitsu

Fujitsu is the leading Japanese information and communication technology (ICT) company offering a full range of technology products, solutions and services. Approximately 162,000 Fujitsu people support customers in more than 100 countries. We use our experience and the power of ICT to shape the future of society with our customers.

This includes a strong portfolio of data protection solutions helping customers to backup, recover and archive data in a simple and efficient way.

For more information, please visit www.fujitsu.com/fts/products/computing/storage/data-protection/.

Terms of Use

This document is Copyright 2018 Freeform Dynamics Ltd. It may be freely duplicated and distributed in its entirety on an individual one to one basis, either electronically or in hard copy form. It may not, however, be disassembled or modified in any way as part of the duplication process. Hosting of the entire document for download and/or mass distribution by any means is prohibited unless express permission is obtained from Freeform Dynamics Ltd or Fujitsu. The contents contained herein are provided for your general information and use only, and neither Freeform Dynamics Ltd nor any third party provide any warranty or guarantee as to its suitability for any particular purpose.