

# Storage Through the Looking Glass

## An alternative business-centric view

**Tony Lock and Dale Vile**

Freeform Dynamics Ltd

January 2012

*The contents of this paper are based on extensive research into the evolution of storage requirements and associated technology conducted by Freeform Dynamics. The discussion presented incorporates insights derived from large scale surveys and in-depth conversations with IT and business professionals, as well as comprehensive briefings with technology vendors and service providers. The paper was sponsored by IBM.*



**Table of Contents**

Getting storage onto the business agenda..... 3

Five things to focus your mind..... 3

Technology and IT department related challenges ..... 5

The business impact of challenges and shortcomings..... 7

Ideas and principles for driving improvement..... 8

Engaging with business stakeholders ..... 11

Final thoughts and recommendations ..... 11

References and further reading..... 12

Appendix A – Overview of current technologies and techniques ..... 13

Appendix B – Organisational sources of challenge and pressure ..... 16

About Freeform Dynamics..... 19

About IBM..... 19

Terms of Use ..... 19

## Getting storage onto the business agenda

Storage has historically been one of the least glamorous of technology areas. If you start out with the serious matter of business processes at the top of the stack, then work your way down through critical applications, high tech middleware, servers that power everything, and database management systems that organise and retrieve our precious information, you eventually end up at the bottom with storage - essential but not something to get excited about.

Well at least that's the perception of many who are not directly involved in the area.

That is until recently, when some have started realising that it's not a case of 'storage, is storage, is storage'. How the organisation's valuable electronic assets are ultimately held and managed both physically and logically really does matter, and its importance is growing day by day. Use the right approach to storage and it becomes possible to do new things that have a direct and positive business impact, including potentially dramatic cost avoidance. Use the wrong approach and risks, costs and overheads can run out of control, possibly causing tangible business damage. Run out of storage or lose precious data, and the consequences can be disastrous.

But the changing nature of the storage game, whether looked at from a requirements perspective or from the viewpoint of technology and best practice evolution, is not broadly understood - especially by business people. Even those in IT responsible for storage often don't get the chance to step back from the day to day pressures and practicalities to look at the bigger picture.

With this in mind, our aim in this paper is to help you take that much needed step away from the detail and create some perspective. We reframe the storage discussion from one of technology operations to business enablement. So, if you are interested in raising the profile of the information storage discipline in your organisation, and making compelling business cases for investment in better tools and technology whilst showing cost savings to the business, then this is a paper for you.

## Five things to focus your mind

To kick things off, here are some points to get us away from the technology view of storage and tune us more into the business perspective - an essential starting point for convincing senior stakeholders to pay more attention.

### **Current approaches cannot effectively deal with the relentless growth in data**

If the rate at which data is being generated and accumulated in your environment is not high and accelerating, then your organisation is the exception rather than the rule. Whether it's transaction data residing in core business application databases, or unstructured information sitting in email servers and document sharing systems, the end result is a relentless increase in both the volume and diversity of business information that needs to be stored and managed. Together with all of the data held on PCs, mobile devices and so on, this can often escalate the problem beyond the level at which traditional tools and approaches can cope.

Whether you are at high stress levels now or not, the challenges are likely to catch up with you at some point if you don't prepare for them. The approach of just throwing more physical storage at the problem is not sustainable. The cost of continually adding capacity and the practical challenge of finding yet more physical space in computer rooms and data centres are obvious issues. Less obvious is the impact of increasing scale and complexity on the management overhead. Put simply, unless a different approach to storage management is taken, the only way of retaining control will be to grow storage administration teams; and how easy is that in the current economic climate?

### **Most organisations do not have a good handle on the data they hold**

Candid feedback from IT and business professionals (including senior execs) during our research frequently alludes to an uncomfortable truth - most organisations do not have a clear view of the information that exists across the business. As systems and databases have proliferated over the years, it has typically been no one's responsibility to keep track of all of the data generated or accumulated in any meaningful way. Certainly, some IT professionals may be able to look across

the storage landscape and provide a view of the data sets that exist and which application systems or business functions they support, but that's not the same as having a clear understanding of the nature of the information within those data sets from a business perspective.

The reality is that information too often lives in silos, and its nature or even existence is unknown outside of the business area, workgroup or even individual user that the silo serves. Think across your organisation about all of the data residing in departmental application databases, content shares, file servers and local PC file systems. Could you, or anyone else, provide a definitive or even good working view of the type, age and status of different records and documents that exist across that landscape? Without such an understanding, designing and implementing an effective storage, retention, security and access regime is hard, if not impossible.

### **Most of the information you store is unlikely to be accessed – ever again!**

An interesting question to consider is how much of the data that was created more than three months ago in your organisation has been accessed over the past week. What about data created two months ago, one month ago, a week ago? How much data, in fact, is never accessed again, even once, following its creation or capture? We will let you make your own judgement on such questions, but the chances are that even a rudimentary analysis will highlight that you have a lot of data sitting on spinning disks that is probably never going to be read again for business purposes.

An obvious follow on question is why so much information is sitting there doing nothing. Sometimes, the answer is because, for the reasons mentioned above, no clear view has been formed about whether or not it is safe to discard it, so it's kept by default. It could also, however, be that it has served its purpose from an immediate business value perspective, but needs to be kept 'just in case' for reference or increasingly to meet regulatory requirements. Either way, it translates to efficiency, cost and even environmental issues if value-less or essentially dormant data is being stored on the same fully managed high-performance devices as valuable and frequently accessed 'live' data.

A third potential reason for data lying dormant is because business users either don't know it exists or cannot access it conveniently, even though it could be useful for analysis or business intelligence purposes. If it were a valuable piece of plant, machinery or some other tangible physical asset, or a stash of money lying there unused, then business management would be quick to highlight the opportunity cost. Curiously, the same thinking doesn't apply to unexploited information assets.

### **IT is often left to make decisions that business people should really be making**

Regulatory bodies are generally very clear on where the buck stops when it comes to protecting, securing, and tracking the use of business information. Ultimately, it's the officers of the company. At more of a day-to-day management and operations level, those responsible for departments, lines of business, operational processes, and so on, are the obvious 'owners' of business information.

Yet all too often, the problem of information management, security and compliance is thrown over the wall to the IT department, with business owners all but washing their hands of responsibility for key information related decisions. IT is left to figure out how information should best be organised, managed, secured, retained and disposed of, without necessarily knowing the full business context.

We can understand why this happens; business people can't be expected to know enough about the mechanics of handling electronic data to form a view on the storage technologies and techniques that are appropriate. However, while delegating the question of 'how' data should be stored and managed makes sense, problems arise when the 'what' and 'why' related questions are thrown across also. The upshot is that we see IT people making what are essentially business decisions on information related matters by default because no one else is willing to do so.

### **Lack of management interest and attention makes it hard to get funding for improvements**

When business stakeholders divorce themselves from the issues of information storage and management, it's also easy for them to become deaf to advice, warnings or even to acknowledge the risks faced by themselves and those in IT who are actually dealing with the challenges. We frequently hear stories in our research of budget holders being warned of storage related risks, but

failing to acknowledge the need for investment, often with disastrous consequences that were sadly very predictable. Whether it's to do with risk supervision and control or general optimisation, management disinterest is a common problem.

At its heart, investment in enhancing storage management provides an opportunity in many organisations to make much better use of existing storage assets, significantly delaying the requirement to buy more storage. The challenge for IT is to communicate the wide range of benefits available in new solutions without alienating business managers.

The chances are that at least some of the points we have raised so far will have resonated with you. After all, these are derived from our ongoing research, so we know that these frequently apply. It's also clear that a lot of what we have described translates directly to challenges and pressures that fall upon those responsible for data storage and management.

Before looking at strategies, tactics and solutions to improve things, however, we must first make sure we have a clear understanding of some of the fundamental factors that are making life more 'interesting' for IT executives, managers and operations professionals. Many of these are well recognised, but few are top of mind when it comes to looking at the storage and data management infrastructure as a whole. While the growth of data volumes, especially unstructured data, continues unabated, changing work patterns may slip under the radar. The highly political area of data ownership and governance is another matter that is beginning to come under the spotlight.

Appendix B pulls together many of the organisational sources of pressure that are now impacting storage.

## **Technology and IT department related challenges**

Switching our attention to storage technology and how it is deployed and managed, there are a number of additional challenges that it's important to acknowledge.

### **Hardware fragmentation and disjoints**

The systems deployed to store data vary enormously. Even in relatively small organisations it is by no means unusual to find a wide variety of storage platforms being utilised to hold data in different parts of the business. Many have server systems running with storage that is directly connected to the server itself. Beyond this they may also have systems connected to storage via a network. There will then typically be a number of tools being utilised to administer storage management.

Turning to the storage kit itself, this may well vary greatly in age, capacity and utilisation rates. Much of it is likely to be operated to support a single function or business application rather than run as a shared resource, available to support any business requirement. This siloed approach tends to result in much of the physical storage space available being wasted as it cannot be utilised by any other system which may be running low on capacity.

Beyond the computer room, some organisations may also find they have servers and storage that have been installed independently of IT to support specific business functions. It is by no means unusual for such systems to be protected inadequately against data loss, corruption or leakage.

And then there is all of the local data held on PCs, laptops, smart phones and tablet devices, much of which may not be protected well in terms of either securing it against loss or theft.

In truth, few organisations, regardless of their size, have an accurate appreciation of exactly what data they have, where it is and what device it is held on, never mind whether it is protected appropriately.

### **Patchy and inconsistent operational processes**

As has been mentioned above, the percentage of all storage capacity potentially available for use that is in productive use tends to be very low. In effect, most organisations are quite poor at exploiting their available data storage capacity. This is partly a result of the way IT systems have

functioned in the past, but it is also a consequence of the way storage systems are acquired and operated by IT, as well as the way in which users expect them to work.

In order to address perhaps the most challenging operational process in storage and its management, it is necessary to recognise a few characteristics of the data and files created by users.

Firstly, users often create multiple copies of files, many of which differ from each other only in minor details, if at all. Secondly, the majority of data files are rarely accessed more than 10 days after their creation, with considerable numbers never opened twice. There may well be compliance reasons making it essential to retain certain documents, but holding the vast majority of unchanging files in their original storage locations could prove to be costly over the years.

Today, the vast majority of this unchanging data stays online by default, often on expensive 'tier one' or primary storage platforms. The long term viability of this approach to data management is rapidly becoming untenable as volumes grow dramatically and as the overall costs of storage management rises to unsustainable levels, even before the additional costs associated with data protection and replication are considered. In essence, it can be argued that the time has come for IT to change the basics of how data is stored and managed over extended periods of time. But to achieve this transformation, business users need to be educated on how to use data storage systems efficiently and cost effectively.

### **Politics and storage funding – Legacy challenges**

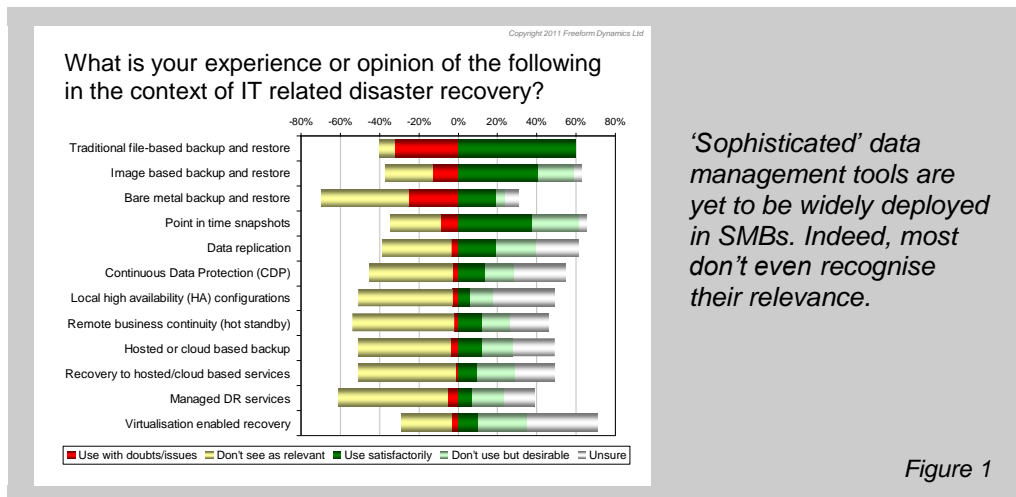
Our research tells us that in the majority of organisations, IT activity in the storage arena is mostly about keeping up with the pressures of managing growth in a reactive manner rather than planning and investing for the future. Indications are, however, that existing approaches are often inadequate for dealing with future requirements. New storage technologies and techniques are key to preventing costs and risks running out of control, and we will be looking at these shortly. Before this, however, we must consider challenges that are more to do with the organisation as a whole.

When looking at architectural changes and infrastructure investments, IT managers frequently run into cultural and political issues around funding and ownership. The traditional method of acquiring storage has been associated with particular projects, essentially as a core requirement to make things work. As a consequence nearly all storage has been directly funded by a specific project or cost centre budget, and has mostly been directly connected to the servers hosting the desired service. This acquisition model still tends to dominate, even when networked storage solutions have become available that allow multiple systems to access a single physical repository.

Typical consequences of this project funding model have been spending capital when storage resources remain unused on existing systems and a fragmentation / silo mentality. There now exist storage solutions that can mitigate many of the unused resource questions, even those that exist in direct attached storage. That said, we still hear reports of departmental budget holders stating that "we paid for the storage, we are unwilling to share it".

### **Storage technology awareness and perception**

A major challenge, especially in small and midsized businesses, is that few organisations have direct access to specialist storage experts and in the past little training has been focussed on storage as a discipline. As a consequence the level of understanding concerning new storage solutions tends to be quite low<sup>[1]</sup>. In addition, given how little time IT professionals have available to carry out 'blue sky research', the perception of solution suitability can drag some way behind actual solution development, especially in terms of their costs and the skills required to keep them operational in routine usage (Figure 1).



Improvements to service and storage costs will be enabled through a combination of taking some key principles on board, exploring and applying new technology where relevant as well as adopting best practice to keep operations running smoothly as data volumes continue to escalate.

## The business impact of challenges and shortcomings

We can discuss challenges and shortcomings to do with data storage and management from an IT perspective as much as we want, but it probably isn't going to radically change things. What's needed is to reframe the storage discussion in business terms. We need to be clear why business people should be interested in the discussion, so let's consider matters that concern business managers.

### The compliance and reputational risk of inadequate data management

The last few years have provided numerous, high visibility examples of the many risks associated with the long term storage of data. In addition, organisations of all shapes and sizes find themselves subject to a range of national laws, some of which are based on international agreements, covering numerous aspects of how they need to protect and secure various categories of data. Beyond generic national legislation, certain industries are additionally subject to explicit regulatory requirements concerning the types of information they hold, its retention and protection.

These regulatory and compliance requirements may detail for how long certain types of data, for example company financial records, health related details etc. must be stored. In addition there may also be specifications covering who can access certain classes of information. Following a spate of well publicised failures to protect data, it is now becoming more common for legislative directives to include considerable financial penalties for cases of non-compliance, sometimes backed up, at least in extremis, with the potential for directors to lose their liberty.

### The operational risk of poor data protection

Many organisations today rely on availability of their electronic data to function. Should an incident occur that takes critical data offline for an extended length of time – e.g. as a result of a fire, flood, inherent systems failure, or simply human error - tangible damage can result. How long can your sales related systems be down, for example, before customers start going elsewhere to place orders? And what about your customer services and logistics systems – how long before the backlog and associated cost and satisfaction issues build up to an unacceptable level?

And the problems can be even more acute if important information is permanently lost or corrupted. What would be the consequences, for example, of the day's sales or operational data being lost?

Such questions bring the potential impact of inadequate data protection – i.e. shortfalls in failover, backup, recovery measures, etc – into sharp focus. They also highlight the importance of business

and IT having a clear mutual understanding of the relative importance of different kinds of data from an operational perspective.

### The direct cost of keeping data in a chaotic manner

Quite apart from the external legislative and compliance driven factors enhancing the visibility of storage management, it should not be forgotten that failure to keep data effectively has real impacts, even when data is neither lost nor stolen.

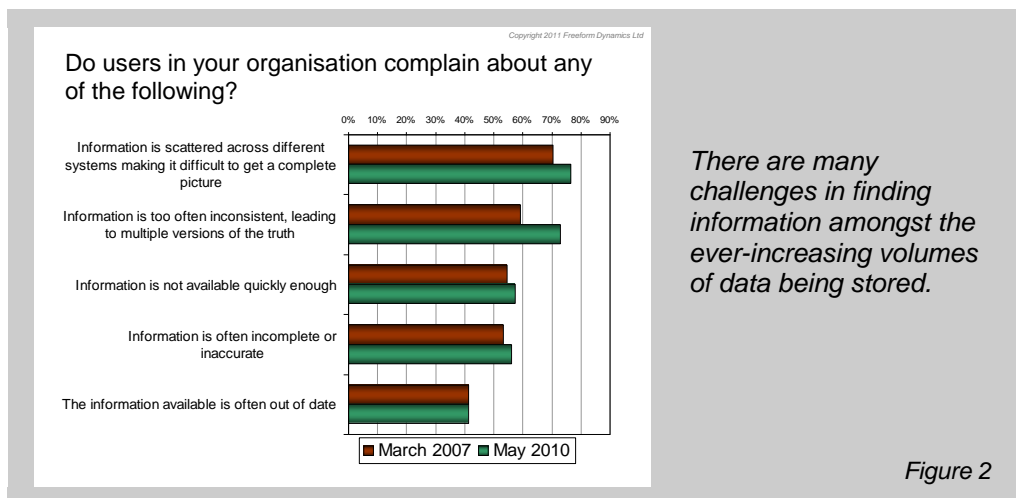
In addition, users now 'expect' that they will be able to open any file they have ever created, irrespective of how long ago it was last accessed. But keeping data costs money and consumes resources, no matter how well hidden these may be in operational budgets. And some estimates now calculate that the cost of storing data is one of the largest in most IT budgets, often exceeding anything other than staff costs.

With user expectation for holding data for ever longer periods of time being matched by similar regulatory pressures, the costs of data storage and its management are slowly becoming visible to business budget holders forcing them to consider long term solutions.

### The opportunity cost of failing to exploit valuable information assets

Quite apart from the growing cost of storing and managing data for long periods of time, there is another subtle challenge to be faced. How do you find data when it is stored haphazardly, in various data repositories, in a variety of shared file systems and on innumerable local disks attached directly to machines, whatever and wherever they might be?

Even more pertinently, how effectively can data be turned into information the business can exploit? There is little doubt that almost every organisation has data resources it could utilise more widely to enhance business success, if only it could find all the pieces and put them together quickly and easily enough. As shown in Figure 2, finding information is a pain, but not exploiting data is an opportunity cost<sup>[2]</sup>.



Many of the challenges discussed here are a direct consequence of data being inadequately classified at any stage of its lifetime, most notably at content creation. This lack of classification makes it awkward for users to find data. It also makes it very difficult for IT to manage data on behalf of the business with anything but the bluntest of tools.

## Ideas and principles for driving improvement

For lots of organisations, the way they think about and deal with storage and information management hasn't changed much over the years. However, things have moved on considerably in terms of both technology and best practice for optimising your storage and information management environment. With this in mind, here are some thoughts to help drive improvement.



## **Find out what you've got**

This is the starting point for everything else and without good knowledge of the data stores you have in the organisation, what's being held and what service metrics are needed to keep it over its lifetime, everything else will be a labour intensive struggle. This discovery process will involve two components. The initial step is to carry out some asset discovery in order to locate all the systems holding storage that need to be managed. After all, if you don't know where the data is being held, how can you look after it?

## **Take time to classify your data**

The second step is to perform some elementary data discovery and classification so that the access, protection and availability needs can be established for the lifetime of the data. This sounds straight forward but will require considerable planning and some effort, not least to establish data ownership and service requirements. It will be necessary to get business managers involved as well as possibly taking some external advice on securing data and its retention timescales.

Amongst the key points to establish will be:

- A need to assign data 'ownership' allowing policies to be established for each data category.
- The relevant retention/protection/performance requirements for each data category.
- Business owner input to establish likely access requirements over the lifetime for which the data is to be held.
- Legal/compliance input to cover security, retention angles etc.

After these are in place IT can look to work out long term storage options for each data category along with the management processes to be implemented as well as, at least to a certain degree, the cost / benefits / risk of each.

## **Understand the technology options available**

While technology is only part of the solution, it is important to appreciate that options have moved on in recent years so if it has been a while since you last looked at the approaches and solutions available, it is worth spending time getting up to speed.

The important areas to focus on when doing this are as follows:

- Backup to disk
- More sophisticated data protection solutions
- Storage tiering and the automation of data placement
- Storage virtualisation
- Storage management tools
- Deduplication and data compression
- Archiving solutions

In order to fill in any immediate knowledge gaps you might have, or to help you bring yourself up to speed in areas you haven't recently considered, we have presented an overview of each of the above in Appendix A. Appliances, virtual machine based solutions, and packaging/delivery in general are also discussed, as many technology vendors have spent time making sure their solutions are accessible to smaller organisations or at a departmental level in larger enterprises.

## **Put 'appropriate' availability measures in place**

It is fair to say that many users expect their data to be available, irrespective of its age whenever they desire access, without giving any thought at all to the cost of providing such availability. This makes the management of expectation a vital component of any storage management strategy.

The fact that technology developments have made the delivery of highly available storage solutions both more affordable and more manageable should not detract from establishing whether each type

of data held would actually benefit the business by its being available with guaranteed continuity. And whilst highly available storage solutions are now more affordable, they do still come at a price premium, a matter that business users must take into account when they decide on service level requirements and economic viability.

It should also be borne in mind that considering each storage service request on its own and with its own project funding leads, inevitably, to siloed solutions that run at lower cost effectiveness than can be obtained compared to building a storage infrastructure capable of supporting the needs of the entire business. Employing an architectural approach to building storage systems makes it more affordable to deliver premium storage to meet the requirements of business critical information and data, and usually also provides much greater flexibility to modify service level delivery as the sensitivity of the data changes with age.

### **Think in terms of architecture**

Amongst the options for discussion and evaluation following the discovery processes established above, it might become necessary to consider using new approaches to store data over many years. Some options are likely to be relatively easy to assess, especially the physical consolidation of storage platforms onto a smaller number of systems.

Others may be somewhat more political unless storage funding models are changed to encourage shared storage platforms rather than individual project silos. High profile options falling into this category could include the virtualisation and pooling of storage platforms to create a shared storage infrastructure. Beyond the use of storage virtualisation systems themselves, it is important that the impact on managerial processes is taken into account up front and that appropriate management tools and training are procured.

All of the potential architectural evolution of storage platforms should make better use of storage capacities whilst limiting wasted space. At the same time they might also improve availability through the use of modern approaches, potentially extending higher availability to larger classes of data beyond the traditional 'mission critical'.

Another couple of areas may also manifest themselves during such an investigation for changing storage management, one covering the models by which such systems are funded and the other looking at the entire domain of data governance.

### **Sort out your governance and funding**

In order to optimise IT services there are certain architectural changes that make economic sense as well as potentially enhancing operational service quality. The use of shared IT infrastructure platforms capable of meeting a variety of service needs from common 'pools' of resources holds great potential. But to employ such systems requires a fundamental shift from a traditional project based view of the world to a more 'infrastructure' based alignment. And to get this far will require a change in IT financing models to support such investments and operations.

This is a pre-requisite for everything else and is probably the single biggest challenge faced by IT and business users as they learn new ways to work together effectively. In most organisations it will require substantial alterations to internal political processes and business / IT funding / budget structures. Decision making and governance processes will also need to evolve to allow new patterns of service requests and service delivery to take shape, potentially bringing forward matters such as internal transfer charging based on resource usage metrics.

Questions of 'ownership' and 'responsibility' will need clear, well understood answers. In particular the old model whereby business managers just assume that the problems of security, compliance and resilience, will be handled by IT without any management involvement is not sustainable.

Beyond this, there has been, and continues to be, considerable discussion amongst IT vendors concerning the pressure that organisations and businesses are experiencing to change the way in which they fund IT procurement. At its base the argument runs that organisations are under pressure to reduce capital expenditure and instead link spending to business use and therefore, hopefully, to their income.

In truth though, our survey results tell us that there is yet to be any major change in how organisations spend money. Recent research we have carried out shows that IT departments expect to increase their use of non-capital spending to acquire goods and services only slowly in the coming year<sup>[3]</sup>. There are great potential benefits to be obtained by organisations should they fund IT projects using alternative models, be they leasing, cloud etc., but understanding of the model options, their impacts and benefits are limited.

With fundamental changes to be considered, how should business stakeholders be engaged?

## Engaging with business stakeholders

The most important facet to try to smooth the acquisition of new storage solutions, or indeed any other IT related investment, is to educate line of business managers and budget holders / project approvers of the benefits achievable using language and ideas they recognise. The idea is to focus on the kind of things it makes sense to talk about and emphasise when trying to drive discussion with the business, including things that IT managers should be asking of them.

As we discussed earlier, one of the fundamental ideas to get across to business is the need for them to become actively involved in data and information management, especially when it comes to taking responsibility for data ownership, rather than just throwing the problem at IT. Getting business people to take 'responsibility' for data ownership should help business managers appreciate the long term cost of keeping data when volumes grow rapidly, as well as focussing IT on the identification of suitable storage solutions.

Once appropriate solutions are identified from the host of possibilities now suited for use in even small and mid-sized businesses, it becomes possible for IT to move discussions with business users away from the challenges / issues / risks of data storage towards recognising the benefits hopefully followed by the unlocking of new opportunities.

Many of the benefits available with new storage solutions have been described already, but here we summarise them in terms that resonate with line managers:

- Better data protection with a consequent reduction in operational, regulatory, commercial and reputational risk.
- Improved cost control, not only in IT, through helping users find, access, manipulate and protect data.
- Greater flexibility to respond to new requirements thereby supporting rapidly changing business demands through the quick provisioning of storage.
- Improved ability to handle unexpected events and developments, without the need for constraining quotas or the disruption of having to deal with unplanned storage expenditure.
- Future proofing the storage environment, through platform re-use, repurposing etc., enhancing capabilities to deal with increasing data volumes and complexity.
- Better exploitation of information assets through improved access for key business users.
- Green and environmental benefits arising from more energy efficient platforms and architectures.

## Final thoughts and recommendations

New storage solutions offer great potential to limit costs whilst enhancing data protection, but perhaps more importantly they also offer a platform which users can employ to better exploit data and turn it into business value. To do so IT managers must first understand their existing usage of storage systems, a matter that few organisations can show they do well, as well as finding the time to understand what solutions are now available to them whilst evaluating the pros and cons of each.

To that end here are a few potential steps to help move things forwards. Some may apply to you, others may not, but they will help to frame your thoughts.

- Asset discovery
  - Find out what storage you have
  - What data is being held?
- Classify the data
- Classify business needs
  - Not all data has the same storage and management requirements
- Understand the technology and management options available
- Frame a storage architecture that can be built step by step
  - Stop creating storage silos
- Address data ownership, governance and funding
  - Business managers must take responsibility for creating and using data management policies
- Put in place good data management and reporting processes

But in order to be able to evolve storage platforms and their associated management tools from where they are today to something capable of supporting the business through the next few years requires IT to make some subtle changes in how it communicates with business management and budget holders. To do so IT managers will have to recognise that budget holders do not care about new technology for its own sake, whatever the wiz-bang and flashing lights. Instead, new storage management offerings must be discussed in terms of the benefits they will deliver directly to the business.

## References and further reading

1. **Enabling rapid and effective IT recovery**  
*DR insights and tips for small and mid-sized businesses*  
<http://www.freeformdynamics.com/fullarticle.asp?aid=1383>
2. **BI 2010 – the state of play**  
*Balancing information management with delivery*  
<http://www.freeformdynamics.com/fullarticle.asp?aid=1150>
3. **X86 Server virtualisation check point**  
*How dynamics should your data centre be?*  
<http://www.freeformdynamics.com/fullarticle.asp?aid=1241>

All of these are available for free download from [www.freeformdynamics.com](http://www.freeformdynamics.com).

## Appendix A – Overview of current technologies and techniques

In this appendix we will take a high level look at what changes have occurred in storage solutions over the last few years. We will not go into great detail on any particular offering, but will rather consider why they matter.

One area that has been noted by disk vendors is that the cost per unit for disk storage (£ / Tb, € / Tb, \$ / Tb etc.) continues to fall, but few organisations are seeing any positive bottom line impact on the cost of storage solutions. The primary reason for this discrepancy is easy to see – data volumes are growing just as rapidly, if not even faster, than the cost of raw storage is falling.

Let's take a look at some storage solutions and where they might deliver benefits. We will start with something close to every IT manager's heart, backup and data recovery, still a source of frustration and anxiety in many organisations.

### Backup to tape and disk

Backup and archiving have been with us since the very first days of IT, yet today our surveys report that large numbers of organisations, large and small, still struggle to get their backup and recovery systems to function effectively. Traditional backup to tape is by no means at the end of its life, and indeed remains the predominant data protection strategy in most organisations, but we are seeing a steady growth in the number of organisations now employing backup to disk as an option in their data protection strategies.

Indeed, it is interesting to record that our surveys show that the use of disk as a primary backup medium is even higher in very small businesses, whilst the use of disk to disk to tape backup solutions are well represented. Backup to disk is primarily employed for its speed, both during backup operations and for data recovery. The combination with tape is usually utilised to provide either long term archiving capabilities or for offsite storage of backed up data at a secondary location for resilience.

### Sophisticated backup and data protection solutions

A survey<sup>[1]</sup> carried out by Freeform Dynamics earlier this year indicated that a significant proportion of small businesses do not recognise that some recent developments in data protection solutions might be suitable for their use, whilst IT generalists in larger organisation admit to not having a good understanding of modern storage solutions.

In the last few years, the area of data protection and recovery has witnessed the development of many solutions of considerably more sophistication than the traditional backup to tape that has been the backbone approach almost since commercial computing hit the mainstream. Solutions such as data deduplication, image based backups, data replication / point in time copies and snapshots have now become well established in large enterprises, but have yet to penetrate far into small businesses.

Even solutions such as those targeted to provide continuous data protection (CDP) have yet to be perceived by many smaller organisations as being suitable to their business requirements or as being affordable. Perhaps of more concern yet is that many smaller organisations perceive sophisticated solutions as incapable of being operated in organisations of their size and are only suitable for large enterprises.

In fact there are now many examples of data protection solutions based on more sophisticated technologies than backup to tape that have developed, in terms of scalability and price, making them suitable for smaller organisations to acquire. More importantly, many of them have also developed relatively straightforward management tools making them usable in daily operations by small teams of generalist IT staff rather than being restricted to storage specialists.

Such solutions provide the potential to provide faster backup times for larger volumes of data and, crucially, to make it possible to recover information more rapidly and reliably than from tape systems alone, something usually demanded by business users as a right. Tape still has a vital role

to play although its job is changing. Tape is no longer the only game in town, but it is certainly far from dead.

### **Storage tiering**

Using different types of storage to hold different data sets is well established in storage management, but until recently the storage performance / reliability characteristics had to be chosen at the time of new storage platform acquisition. In essence, storage characteristics were defined up front in a project when the system hardware was specified. Once data was stored on the system, it was there for life.

Thus business critical systems were usually installed with top of the range, and hence expensive, storage platforms. But as data ages it is accessed less and less frequently and at some point usually reaches a stage where it is considered to be far less mission critical than when it was created. But moving such data off mission critical storage platforms to something cheaper was not easy to manage.

The drop in price of flash storage is now making it economically viable for a growing range of uses, but it is the development of storage management tools capable of automating the movement of data between storage platforms of different performance characteristics that is really making it possible for organisations to look at storage tiering. Until very recently, storage systems with automatic tiering capabilities were only available at the large enterprise end of the market but such solutions, sometimes also employed with storage virtualisation, are now being delivered in systems suitable for the use of small and mid-sized organisations.

Such systems can allow organisations to keep frequently accessed or business critical data available on highly performant platforms aiding its use in value generation and business operations and then moving it to less speedy storage when either its importance drops or its access rates decline.

### **Storage virtualisation**

Nearly every IT professional has some understanding of the value of server virtualisation solutions. They have been widely used to consolidate applications previously running on poorly utilised servers onto more scalable systems that operate at higher efficiencies. This approach can save money in a range of ways including lowering system maintenance fees, power consumption, physical space and sometimes even in application and middleware licensing charges.

Storage virtualisation works in a similar vein and also holds the potential to greatly increase storage platform utilisation rates ultimately requiring far less storage to be acquired than in standalone physical systems. In addition, virtual storage solutions can often extend to making a range of performance options available at a variety of price points, each with different service characteristics. Overall, storage virtualisation can help control costs as well as providing greater flexibility, possibly combined with higher availability for a greater range of data. Such solutions are now considered to be mature and are becoming suitable for routine use in smaller organisations, although deployment has thus far been limited.

### **Storage management tools**

Storage management has traditionally been the domain of a small number of IT specialists, few of whom worked in small or medium size businesses. Indeed, in most SMBs, storage was only really considered at system implementation time, usually as a corollary to server platform acquisition. The slow, but steady growth of networked storage solutions (SAN and NAS) in organisations of all sizes has been enabled by a continuous improvement in the capabilities of storage management tools.

The usability of such tools has now reached a stage where they have become suitable for use in an expanding range of scenarios making it possible for data and storage management to now be handled more dynamically as business requirements vary. As organisations pick up on such management tools, the protection of data will improve and ultimately data should be available to end users to generate business value, whilst security and integrity will be enhanced.

The understanding of such management tools remains low in most organisations as few IT generalists have time to investigate the options that may be available to them, even though potential benefits could be high.

### **Data deduplication and archiving**

As data volumes continue to grow rapidly, it is likely that solutions that seek to limit the impact of data growth will become more visible and, ultimately more widely deployed. Two areas that offer great potential to ease the increasing cost burden associated with rapidly expanding data volumes are the well-established, though not so widely deployed, data archiving and newer solutions seeking to deduplicate data.

Data and content archiving solutions have been around for some time and are used to move less frequently accessed information from expensive, online, storage systems to cheaper long term systems whilst allowing the user to still search the system as if the data were still present. Once the domain of large enterprises, and even there only in specific business contexts, it is likely that economic pressure will encourage the development, and adoption, of archiving solutions capable of supporting wider use across lines of business.

Data deduplication systems are less advanced in their development. A number of different architectural approaches are available from suppliers that seek to remove any data from storage systems that is an exact copy of existing material. Such solutions offer the promise of reducing overall storage needs as it is well established that many copies of documents with little or no variation exist in nearly every company as users work collaboratively generating reports or other materials.

### **Data compression**

Twenty years ago when floppy drives were becoming central to movement of information inside businesses, data compression was of critical importance due to the low capacities of such diskettes. Today, many people take for granted that they can store and move huge files with little thought to resource usage or costs. Yet many of the data formats commonly used in everyday business, including popular office productivity tools, do not store data as effectively as they might. Data compression solutions provide another means by which to limit the impact of the escalating data growth being experienced today. It's another case of 'yesterday once more'.

### **Appliances, virtual machine based solutions, packaging/delivery in general**

In recent years storage vendors have begun to produce a variety of appliances and, more latterly, virtual machines that package up sophisticated functionality and hardware into an appliance format. In this way it is possible for small and mid-sized businesses to obtain systems targeted at delivering capabilities previously only available to very large enterprises blessed with sophisticated specialist staff. The use of appliances can open up new opportunities.

## Appendix B – Organisational sources of challenge and pressure

In this appendix we highlight some of the non-technological factors organisations report as having considerable impacts on how they operate today. The likely impact on the storage infrastructure resulting, directly or indirectly, from these evolving business work practice changes is discussed.

### Proliferating sources of data and associated risks

Ask any long-term IT professional and they will tell you that data proliferation is nothing new. The cycle where new storage technology comes along freeing 'restrictive' limits has always been followed by users and software finding new ways to eat up any storage headroom available. But today, the scale and speed of the growth in data being created is at a level that would have been unimaginable just a few years ago.

No single development has triggered the staggering growth of data that most organisations are experiencing, rather it is a combination of many, subtly interconnected, changes. Whilst any individual may well recognise one or more of the 'symptoms' if they are pointed out to them, few business users are in a position to visualise the whole picture.

So where is data growth coming from? The table below highlights some of the typical areas that inevitably, but not always visibly, result in escalating data volumes and complexity. Also highlighted are some the frequently overlooked risk factors arising from commonly observed activity.

Area	Example	Exposure / risks	Storage impact
Growth in unstructured data	Word processing documents, presentations, spreadsheets, PC data bases etc.	Where is the data? Is it 'protected' and backed up?	Fastest growing area of data being created. Need to backup data and / or centralise
Workforce flexibility	Use of laptops	Possible unsecured / unprotected data	Backup / data replication to central storage?
	Use of smart phones	Should company data be held here? Possible unsecured / unprotected data	Backup / data replication to central storage?
	Use of personal tablets	Should company data be held here? Possible unsecured / unprotected data	Backup / data replication to central storage? Desktop virtualisation?
	Use of home PC	Should company data be held here? Is the data properly protected and backed up?	Possible need to replicate data to central systems
SaaS / cloud services	Hosted email, Web mail and Software as a Service (SaaS) based business applications	Where is the data? Should company data be held here? Is it 'protected' and backed up? Do we need to replicate valuable data back onsite?	Data replication to central storage?

The table lists just a few of the sources of data creation likely to be found inside any organisation. At the heart of the storage challenge today, unlike twenty years ago, the majority of data in small and mid-sized business, as well as larger ones, now no longer resides in centrally managed business applications and their associated databases. Indeed, many studies have instead highlighted that it is the world of 'unstructured' data, mostly the files generated by individual workers, that today accounts for the fastest rate of data creation.

Such files are now routinely megabytes in size and contain everything from complex data sets to email messages with embedded video files. And usually multiple copies of the same file may exist on different machines, or there could be dozens of file variants with only slight differences from each other. Everyone is generating data, and not in small amounts. Furthermore, the context in which they are creating it is evolving.



## **Changing working patterns and people related issues**

The increasingly broad landscape of devices used in business – including home computers, laptops, tablets and smartphones of various flavours - is encouraging more movement and distribution of data. Today, for example, it is by no means uncommon for the finance director or MD to take a few documents home to work on over the weekend, irrespective of policies that might say this should not happen. As many senior managers work more flexibly using portable or mobile devices, even personal equipment, for work purposes, the task of keeping track of, securing and protecting company data poses considerable challenges.

While technology based solutions exist that can help with some of these challenges, it is important to recognise that the 'people factor' must also be dealt with. Will a senior manager, for example, be happy for you to encrypt part of their personal device, and/or to take backups of their local file system onto the corporate network? Probably not, at least until they have had the embarrassment of leaving a device containing sensitive data in a taxicab, or the pain of losing data through a hard disk crash or a tablet spontaneously resetting itself to factory defaults.

But educating senior managers to account for such technology and business risks is not easy, let alone the broader workforce. The response of many users to discussions of how to manage and protect the data they create is often: "It's not my problem – it's something for IT to deal with" – which is ironic given that much of the activity has been initiated independently of IT.

It's also extremely hard to resolve the employee tendency to abdicate responsibility with another commonly observed trait. Experience tells us that users tend to be very possessive of their laptops, PCs and other 'personal' IT equipment, even when it is supplied and supported by the company, to the extent they resist any form of IT management of kit. Equally, many will attempt to work around any processes or systems if they perceive them to be either limiting their 'freedom' or to be impacting the performance of the machines.

Whether it's corporate data held in central systems, the departmental Access database, the files held on personal devices, or the data copied to USB sticks or other removable media, one of the trickiest questions when it comes to data storage and management is the one of ownership and responsibility.

### **The tricky question of ownership and responsibility**

With some of the dynamics we have been discussing, a key question is who should be responsible for the data once it is created, and do they know how to protect it appropriately? Easy to ask, but not so easy to answer.

The question of responsibility is an issue fundamental to the management of information, but it's one that is rarely tackled directly. The reality is that, outside of organisations where there are severe regulatory or legislative pressures, little effort is expended putting policies in place governing information-related responsibilities. This sounds quite harsh when you say it in this way, but the reasons for ownership and responsibility related issues are very understandable, as they often 'creep up' on the organisation over time.

When IT systems are first commissioned, for example, business managers are usually involved helping to decide on matters of security, availability and recoverability. IT can then put appropriate measures and solutions in place to deal with the requirements. So far so good.

But as time goes by systems may alter in the use that is made of them, and sometimes the importance of such changes may not be notified to the IT team, so the security, compliance and protection measures in place are not updated accordingly. Furthermore, those involved in the initial discussions may leave the company or move to another part of the business, so IT is sometimes left in the position of having no one available from the business with whom an informed discussion about data management can be had.

Adding to these challenges are issues associated with 'unofficial' or 'temporary' solutions that ultimately become permanently embedded in the business, that were never really 'owned' by anyone in any formal sense. Then, of course, there is the question of ownership and responsibility

when it comes to the aforementioned data held on PCs, laptops and now smart phones and tablets, whether such devices are used officially or unofficially. IT can take steps to help manage and secure such data, but unless the questions of ownership and responsibility are addressed, only so much can be achieved.

## About Freeform Dynamics



Freeform Dynamics is a research and analysis firm. We track and report on the business impact of developments in the IT and communications sectors.

As part of this, we use an innovative research methodology to gather feedback directly from those involved in IT strategy, planning, procurement and implementation. Our output is therefore grounded in real-world practicality for use by mainstream IT professionals.

For further information or to subscribe to the Freeform Dynamics free research service, please visit [www.freeformdynamics.com](http://www.freeformdynamics.com) or contact us via [info@freeformdynamics.com](mailto:info@freeformdynamics.com).

## About IBM



At IBM, we strive to lead in the invention, development and manufacture of the industry's most advanced information technologies, including computer systems, software, storage systems and microelectronics.

We translate these advanced technologies into value for our customers through our professional solutions, services and consulting businesses worldwide.

For more information on IBM, please visit [www.ibm.com/storage](http://www.ibm.com/storage).

## Terms of Use

This document is Copyright 2012 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 report for download and/or mass distribution of the report by any means is prohibited unless express permission is obtained from Freeform Dynamics Ltd.

This report is provided for your general information and use only. Neither Freeform Dynamics Ltd nor any third parties provide any warranty or guarantee as to the suitability of the information provided within it for any particular purpose.