

RESEARCH PAPER

Quantifying the cost of doing nothing

Identifying the cost benefits of a strategic approach to enterprise backup and recovery

December 2013

CONTENTS

Executive summary	p3
Nobody expects...	p4
The state of play	p5
Just in time?	p7
Under pressure	p8
A touch of class	p9
Must (and can) do better	p10
Conclusion	p11
About the sponsor, Dell	p11

This document is property of Incisive Media. Reproduction and publication of this document in any form without prior written permission is forbidden.

Executive summary

Working out the potential return on a new server or application is relatively straightforward – if it enables you to do more in less time, profits are likely to go up, making it worth buying. If there are no obvious benefits, however, and no immediate gains, then the computation is far less easy. This is the case with the long-time Cinderella of the IT world, backup and recovery, where the chief benefit is the ability to maintain rather than enhance the bottom line.

The trouble is that, although essential to business continuity, there is no easy way of quantifying the financial benefits of investing in backup and recovery. Not least because it involves predicting the future and estimating the financial impact should the worst happen; effectively thinking the unthinkable, if you like, and quantifying the potential cost of doing nothing about it.

Unfortunately that analysis is far from easy and, as a result, a backup and recovery solution is often only ever justified as part of another project. Projects such as an application migration or a datacentre upgrade with more easily definable financial benefits. Over time that, in turn, can lead to an ill-matched mix of legacy platform and application-specific solutions, each with its own dependencies, backup windows, recovery times, support requirements and so on.

More than that, once in place, backup and recovery tools are rarely reviewed, upgraded or consolidated, resulting in outdated silo solutions – solutions that are not only expensive to maintain and difficult to manage but which struggle to cope when it comes to recent developments such as datacentre virtualisation, cloud computing and Big Data.

This *Computing* whitepaper looks at the risks organisations need to consider when thinking what to do about backup and recovery. It also looks at the potential cost of doing nothing to address those risks and how to better justify and direct investment in modern, integrated and effective solutions to this business-critical issue.

Nobody expects...

Ask anyone to list potential disasters that might affect enterprise IT systems and flood, fire and pestilence are what tend to spring to mind. OK, maybe not pestilence, but you get the idea: most people think big, and rightly so as the worst can and will happen, as illustrated by the Fukushima nuclear meltdown and, more recently, typhoon Haiyan in the Philippines, to name just two.

Disasters, however, don't just follow in the wake of cataclysmic events. Software bugs can be just as damaging; malware and hackers are a constant concern and, seemingly, inconsequential actions can have a surprisingly big impact further down the line.

This is highlighted by a *Computing* survey of 120 datacentre professionals designed to find out how those charged with keeping enterprise IT systems operational quantify risk and, more importantly, how they use that information to justify spending on backup and recovery.

Unsurprisingly, all claimed to plan for obvious things such as fire and flood, along with power cuts, hardware and software failures, malware attacks and so on. When asked for specific instances, however, the responses were a lot more interesting and, at times, even comical – although the respondents may not have thought so when wrestling with some of the problems reported.

When asked to describe an unexpected issue leading to real data loss or downtime, for example, answers ranged from the mundane to the downright bizarre. Among the more commonplace was a contractor, for example, who accidentally leant on the “kill” switch, turning off an entire datacentre in moments. At the bizarre end of the scale, meanwhile, one respondent reported an operator vomiting over a rack of switches as a result of finding a decomposing and very smelly hedgehog in the air conditioning.

The former was clearly the result of bad planning and could have been avoided. Indeed, the respondent indicated that a cage had subsequently been fitted around the offending kill switch. But who would have expected a rotting hedgehog to bring down a whole network?

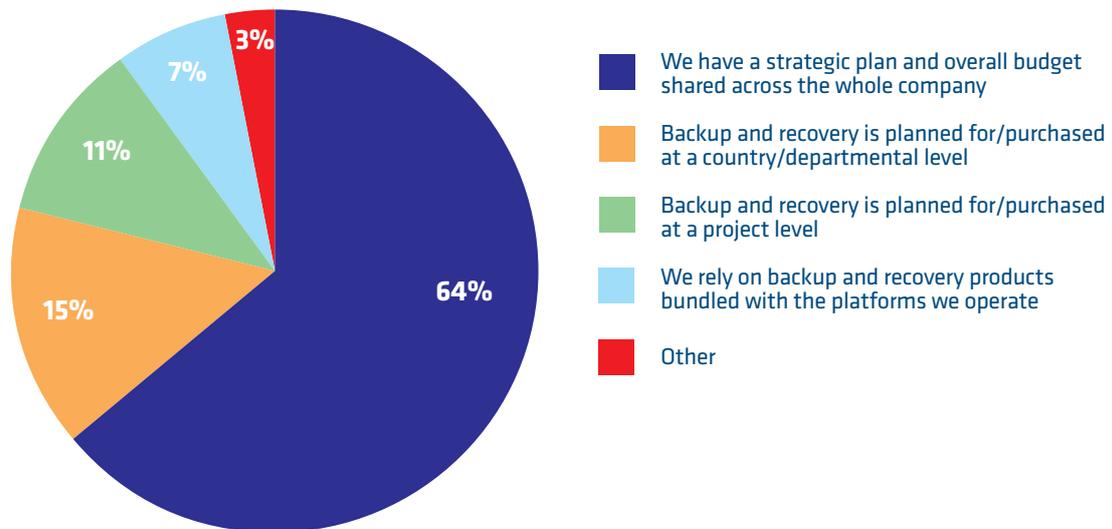
Some of the knock-on effects were equally unforeseeable. One company, for example, had prudently invested in a dedicated business resiliency centre but, when a nearby electricity sub-station exploded, found it was inside the area cordoned off by the police, preventing its use.

Another found that although tapes had been religiously loaded into the backup library every night, the backup routine itself had been descheduled. This was a simple error that wasn't picked up until staff found themselves faced with having to recover a crucial storage array in which two drives had failed. The need to verify backups after they have been taken suddenly took on a whole new meaning.

The state of play

There were plenty of other examples besides, with nearly all of those responding to the survey having an interesting anecdote to tell. Anecdotes which makes it all more surprising to find that, when asked, just 64 percent said their organisation had a strategic plan for backup and recovery (Fig. 1).

Fig. 1 : How do you plan and fund backup and recovery within your organisation?



More than that, 11 percent said that backup and recovery tools were purchased purely on a project-by-project basis while a further seven percent relied on the tools bundled with their platforms. This is a somewhat blinkered approach given the disasters and their consequences reported in the survey, but a clear reflection of the difficulties IT staff face when it comes to justifying investment in backup and recovery. In many cases it seems, backup and recovery can only be justified by either bundling the costs involved into another project with an easily quantifiable budget or by using something already available for free.

There are exceptions, of course. When asked for more detail the survey revealed that just under half (43%) of the companies surveyed were taking a more considered approach, with fully integrated backup and recovery systems able to protect all of their platforms and their line of business applications and to provide that protection across both real and virtualised infrastructures (Fig. 2).

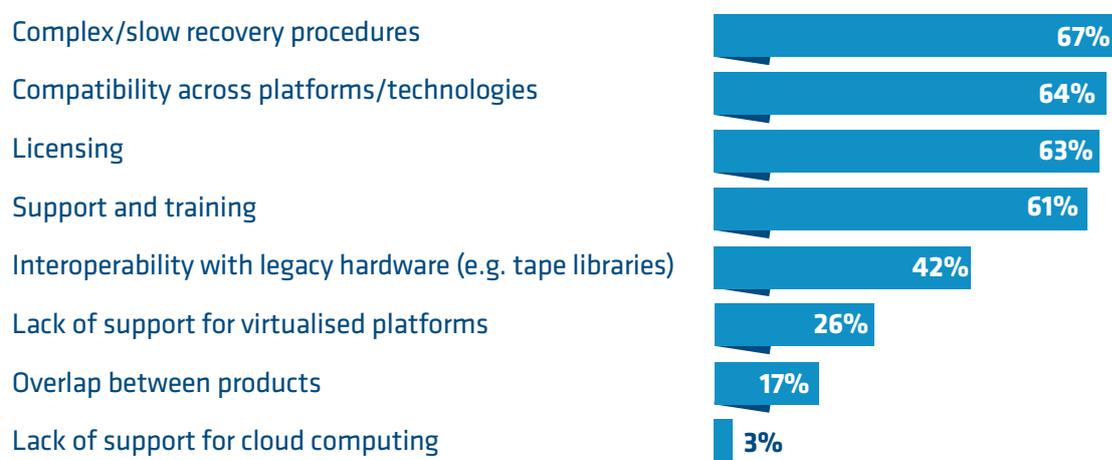
Fig. 2 : Which of the following best describes the way in which backup and recovery is handled within your organisation?

We have a backup and recovery solution to protect server and operating system platforms, but not individual applications	17%
We have a mix of legacy backup and recovery products for platforms/operating systems plus others to handle specific applications	28%
We have a mix of legacy backup and recovery products with bolt-ons/extensions to handle newer technologies and applications	4%
We have a fully integrated backup and recovery solution able to protect all of our platforms and applications, real and virtual	43%
We have migrated some applications to the cloud and, in addition to in-house products rely on our service provider to protect our cloud-based systems	6%
We have migrated most/all of our applications to the cloud and rely on the backup and recovery protection supplied by our service provider	2%

This is good news for the companies involved, but on the other side of the coin it reveals that more than half are less well prepared, with many struggling to manage with legacy products poorly equipped or, possibly, unable to cope with the latest technologies. In particular, virtualisation where tens of virtual machines are likely to share the same physical platform.

In this respect just under a fifth (17%) said they relied on being able to recover whole servers rather than being able to bring specific applications back online should problems arise. Unfortunately so-called bare metal recovery of complete server platforms can take a long time and its prevalence could be one reason why complex and slow recovery procedures came top of the list when respondents were a mixture of legacy products were asked to highlight issues with backup and recovery in their organisation (Fig. 3).

Fig. 3 : What issues arise because of the need to operate and support multiple legacy backup and recovery products?



* Answered by those with multiple legacy backup and recovery products; respondents could select multiple answers.

As the graph here shows, 67 percent pointed to recovery time and complexity as their major challenge, with compatibility across platforms and technologies trailing in second (64%). These were then closely followed by licensing, support and training of multiple backup products plus interoperability with legacy hardware such as tape libraries.

A quarter also said they had issues when it came to protecting virtualised platforms with the – mostly legacy – backup and recovery tools at their disposal.

Just in time?

When asked specifically about how quickly they thought they could recover from a “disaster”, the answers were equally revealing, especially given the general impression of companies having to compete in increasingly fast-moving, always-on and global marketplaces where every second counts.

Given that scenario, you might expect most companies to be equipped to recover lost data and even complete servers in seconds or at least minutes. As you can see from Fig. 4a, however, that expectation is widely optimistic and far from the case in a lot of organisations.

Fig. 4 : How long does it take to recover or roll back...

a) A particular file?

	Seconds	Minutes	Hours	Days	Other
Best case	31%	53%	13%	2%	1%
Worst case	4%	18%	38%	38%	2%

b) A full server?

	Seconds	Minutes	Hours	Days	Other
Best case	10%	36%	48%	4%	2%
Worst case	0%	8%	38%	52%	2%

When tasked with restoring or rolling back individual files, only 31 percent claimed that they would be able to recover data in seconds and only in the best of circumstances while, at worst, that figure fell to just four percent. Indeed, at worst nearly 40 percent thought it could take days to recover lost data or undo changes.

It was much the same when it came to estimating the time needed to recover an entire server (see Fig. 4b) where, again, the more pessimistic guesses were in the order of days (52%). One respondent from an oil company remarked that it could take days just to get to remote platforms before they could even begin the recovery process.

Given the breadth and consistency of the answers, slow recovery times are clearly the bane of IT departments across companies of all types and size. It also begs the question as to why those same departments don’t do some relatively simple maths and use those very same lengthy recovery estimates to quantify the financial impact on the business of not investing in solutions to address the issue.

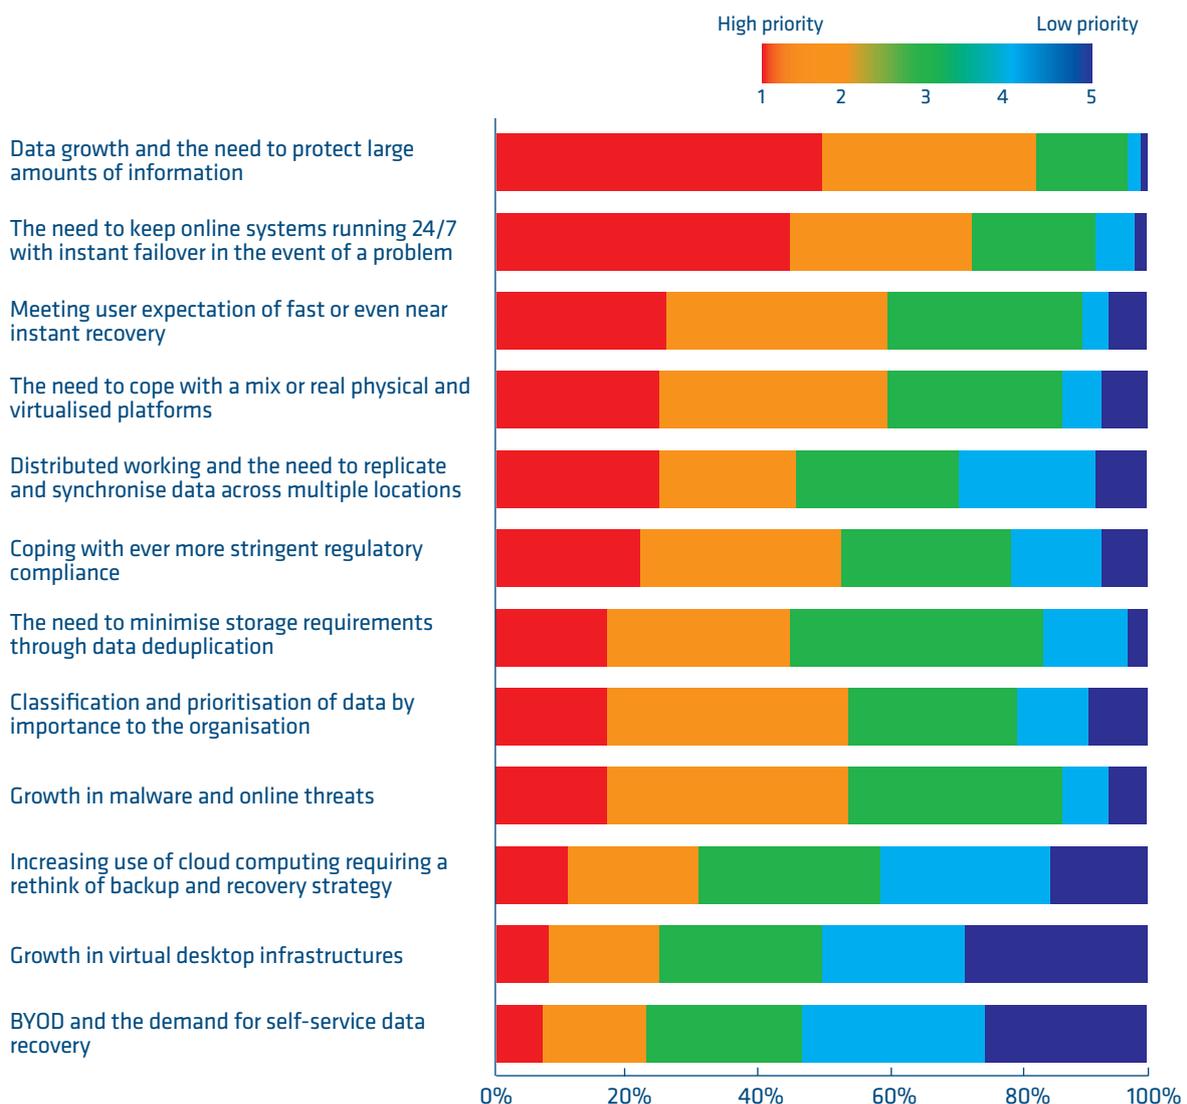
Quantifying the cost of doing nothing

OK, it's not something you can do in a few minutes, but it is not rocket science. A server down or an application offline means lost business and the longer it takes to recover, the more business will be lost. The cost of that lost business can, at the very least, be guessed at, if not accurately estimated and used when bidding for the purchase of new backup and recovery solutions or when updating an existing setup to cope with technological and business changes.

Under pressure

One reason why this may not be happening could be the sheer volume of data, applications and host systems that IT department have to protect these days. The *Computing* survey asked respondents to rate a number of pressures they faced when it came to backup and recovery. Data growth and the need to protect large amounts of information came top of their priority list (Fig. 5).

Fig. 5 : Changes in technology, business practices and user expectations are placing new pressures on IT teams when it comes to back up and recovery. Rate the following in terms of priority



Given almost as much attention, the need to keep online systems running 24/7 and provide instant failover in the event of a problem came second. This is yet another indication of the kind of pressure IT departments are coming under as they attempt to grapple with changes in not just technology, but business and trading practices, the global marketplace and users/customer expectations.

So much so that a comprehensive backup and recovery solution to protect all of the IT resources within an organisation may still be hard to justify. Moreover, it may still be viewed as prohibitively expensive even when measured against the cost of losing business for days in the event of what might be a relatively minor “disaster”.

A touch of class

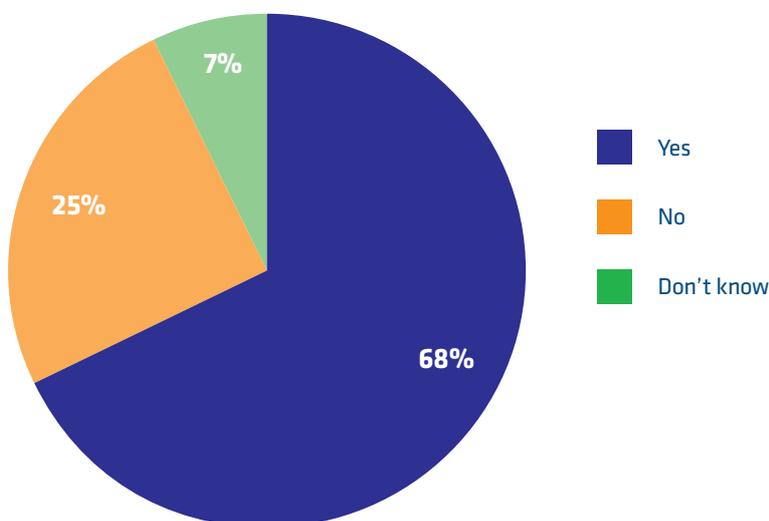
So what more can be done? Take another look at Fig. 5 and you will see that, when it comes to their impact on business resiliency planning, companies simply don't see things like BYOD or even cloud as major concerns. Rather, they are having to cope with the same issues of storage growth, compliance and security which have plagued them for decades. So instead of trying to protect everything in the same way, why not take a different tack and identify the resources of most importance to the business? Resources such as servers, applications and data that would cause the biggest operational harm should they become unavailable for any length of time.

By classifying data and other resources this way, IT managers can then concentrate any investment in backup and recovery on making sure those resources are adequately protected. Firstly by making sure that extra redundancy is built into host platforms, networking and WAN services to keep business critical systems running. And just as importantly, by choosing backup and recovery products expressly designed to handle the mix of platforms and applications employed, ensuring that, should the worst happen, those business-critical resources can be recovered quickly with minimal impact on the bottom line.

The simple truth is that IT systems are not all made equal and are not of equal importance to the operation of the business. Some servers, applications and data sources need to be accessed all the time and will have to be recovered in seconds to keep trading. At the other end of the spectrum, however, there will be servers that are only used occasionally, applications that only run from time to time and data that is rarely, if ever, accessed at all. If it were to take a few hours or a couple of days to recover such systems it really wouldn't matter.

Classification of data to better concentrate backup and resources would seem to be something all companies should do. Unfortunately that doesn't appear to be universally understood as, when asked whether they classified data resources by importance to the business, only 68 percent of respondents to the *Computing* survey answered “Yes” (Fig. 6).

Fig. 6 : Do you classify data by importance to your business?



Those that did, however, clearly saw faster recovery as the main advantage of taking this, more considered, approach to business resiliency planning. In fact, quicker recovery was ranked top of the list of benefits by 68 percent of respondents, well ahead of reduced storage overheads when taking backups (41%), reduced backup costs (33%) and shorter backup times (27%), although these benefits are all worth having and can be achieved simply by concentrating on the resources that are key to the business.

Must (and can) do better

Classification of resources by importance to the business is important, and worth doing. At the same time, however, it can only be part of the solution when it comes to coping with business resiliency planning ills. Equally important is the need to understand the requirements and limitations of the technologies companies increasingly rely on to run their business. Technologies such as virtualisation which brings many benefits when it comes to hardware consolidation, speed of deployment and operational flexibility but which makes systems more vulnerable by enabling multiple virtual machines to be hosted on a single platform.

Backup and recovery solutions need to understand this and address the needs of virtual as well as physical resources. Ideally they also need to be able to cope with legacy systems as well as new platforms, be easy to manage and meet the recovery expectations of the organisation, its employees and customers alike.

It is a tall order but it can be done.

Conclusion

As shown by the *Computing* online survey on which this whitepaper is based, enterprises are very much aware of the potential disasters that can befall IT systems, and that those threats can come in a variety of, seemingly, innocuous guises.

The majority are also aware of the inadequacies of legacy backup and recovery systems to protect against such threats, leading many to invest in modern integrated solutions better able to cope with changing technologies, business practices and user expectations. A significant number, however, continue to adopt a piecemeal approach finding it hard to justify investment in new or upgraded solutions except when part of a project with clearly identifiable cost benefits.

The end result is a mixed picture of best practice together with continued and widespread reliance on legacy products leaving many businesses ill-equipped when it comes to the specific demands of virtualisation and other new technologies. Some companies are able to cope but many cannot, with slow recovery times the biggest issue faced when problems do arise.

As well as being an issue, however, slow recovery times can also be used to quantify the cost benefit of investment in improved backup and recovery solutions. Equally, by classifying data resources by importance to the business, that investment can be better directed to deliver the maximum benefit at the minimum cost, even in companies faced with escalating data growth and the need to keep business-critical systems running 24/7.

It is possible that with the right planning, the right backup and recovery solutions targeted at the most important systems, businesses can be confident of meeting and coping with disaster no matter how it comes along.

About the sponsor, Dell

Dell is the only vendor to offer a portfolio of proven backup and recovery solutions that give you the power to match your backup to your business – now and in the future – saving time, reducing risk and nearly eliminating data loss.

Through products such as Dell AppAssure™, Dell NetVault™ Backup, and the DR and DL families of purpose-built backup and recovery appliances, organisations of all sizes can build fully-optimised data protection environments quickly and easily, without changing their overall data protection strategy.

Leverage the versatile RTO/RPO potential of the Dell Data Protection portfolio to restore your most critical data in seconds and your entire infrastructure in minutes, from simple to complex.

Visit www.dell.co.uk



Software