



Enterprise edition

Virtualization Data Protection Report 2013

Executive Summary

The Virtualization Data Protection Report 2013 builds upon the key findings from previous years to track the progress of enterprise-level organizations' data protection strategies for their virtual environments. Where previous reports looked at broad trends around data protection before focusing more closely on specific techniques such as server replication, this year's report investigates whether organizations are confident in data protection for their virtual environments. It identifies specific issues CIOs are having with their virtual environments and suggests some of the underlying causes behind this. In particular it highlights how capabilities, complexity and cost are still the key challenges for organizations wishing to consistently protect their most critical servers and the data they manage. Currently, 88% of CIOs are experiencing issues around capability, 84% around complexity and 87% around cost.

This report shows how attempts to "retrofit" data protection by applying physical world tools and techniques to the virtual environment continue to hold back the technology's true potential.



The concluding sections of the report outline some of the actions organizations are taking to deal with the challenges they face in implementing data protection in their virtual environments. For example, 58% of CIOs plan to change their tools for backing up virtual environments in the next 2 years.


The report is based on an online survey conducted in November and December 2012 by Vanson Bourne, an independent market research organization, of 500 CIOs from organizations across the United States, United Kingdom, Germany, and France that employ more than 1,000 people. The report is sponsored by Veeam Software.

Survey Background

As with previous surveys, respondents came from a cross-section of industries. While Manufacturing was the most well represented field (26% of respondents), it was closely followed by Retail, Distribution & Transport (23%); Financial Services (22%) and Business & Professional Services (20%). Other commercial sectors were responsible for 9% of the survey sample. This ensured that the responses gathered came from a wide range of enterprises across different sectors.

Chart A:
Types of organizations surveyed (%)





Part I | Capability Challenges for Organizations

1. Capability Challenges for Organizations

88% of CIOs identified capability-related challenges that are impacting their ability to backup and recover their virtual infrastructures. These challenges represent a failure to realize the full potential of virtualization-based data protection. If used correctly, virtualization enables much higher data protection capabilities than a traditionally managed physical environment. With the right tools, entire virtual servers or individual files and application items can be recovered in a matter of minutes, allowing IT departments to recover quickly from disasters both large and small. This rise in efficiency, coupled with the use of modern approaches to backup, means that organizations can further enhance data protection: for example, testing backups to ensure that they can be recovered when needed. By taking this approach, CIOs should be able to set and keep much more rigorous Recovery Point Objectives and Recovery Time Objectives, and therefore be better placed to meet SLAs for their organizations.

However, currently organizations are not taking full advantage of these benefits. While recovery of physical servers takes, on average, 6 hours, recovery of virtual servers is not significantly faster, at 5 hours (Chart 1). Worse still, this is no better than in the previous survey: while the difference of 1 hour was the same, recovery as a whole was faster at 5 hours for physical servers and 4 hours for virtual.




Chart 1: Time taken to recover backed up servers (hours)

2011
 ■ Time to recover virtual servers
 ■ Time to recover physical servers
2013
 ■ Time to recover virtual servers
 ■ Time to recover physical servers

This would seem to reinforce the concerns of 68% of organizations, who feel that their backup and recovery tools will become less effective as the amount of data and servers in their infrastructure continues to grow (Chart 2). With recovery times already increasing, it seems that this assertion is already coming true.




Chart 2: Organizations feeling backup and recovery will become less effective (%)

Financially, CIO respondents stated the cost per hour of downtime for their business critical servers that are not been protected by replication as \$324 793 (Chart 3). Coupled with a recovery time of 5 hours or more this means that, on average, each outage is costing organizations at least \$1.6 million. Unless data protection improves, these costs will remain high.




Chart 3: Cost-per-hour of critical servers being down (USD)

CIOs are also struggling with granular recovery from virtual servers, even though virtualization can make such recovery more straightforward. On average, recovering individual file or application items from virtual servers takes 3 hours (Chart 4): while faster than recovering a whole server, this is still a significant wait. Certain items can be even slower: recovering individual emails takes an average of 14 hours (Chart 4). These slow times are partly due to the fact that not all organizations have the capability for granular recovery despite their virtual environment: 71% often have to recover more than they need in order to reach specific files or application items (Chart 5).




Chart 4: Time to recover individual file, application items and individual emails (hours)

- Time to recover individual file
- Time to recover application item and individual emails




Chart 5: Organizations needing to recover more than desired to reach specific items (%)

- Organizations that often have to recover more than needed to reach specific items
- Organizations that can consistently recover individual items directly

Furthermore, CIOs show a lack of confidence in their ability to consistently recover from backups. On average, organizations experience problems when attempting to recover from backups 9 times per year (Chart 6). To place it in context, this accounts for 17,48% of all recoveries, meaning that the chances of an unsuccessful recovery are more than 1 in 6 (Chart 7).



Chart 6: Number of times per year organizations experience problems attempting to recover from backups





Chart 7: Percentage of recoveries that present problems (%)

Chart 8: Frequency of backup testing (months)

One reason for this relatively high failure rate may be a lack of opportunities to test backups. Currently, organizations test their backups for recoverability every 3 months (Chart 8). However, on these occasions they only test on average 7,43% of all backups (Chart 9). As a result, the vast majority of backups remain untested. While organizations can concentrate on the most critical areas, there is still a large proportion that must essentially be left to chance. In turn this makes it harder for CIOs to guarantee their availability SLAs.




Chart 9: Percentage of server backups tested when testing for recoverability (%)

As shown, there are significant capability-based challenges facing CIOs. However, another important question is whether CIOs recognize all of the challenges they face. While 88% admitted to facing capability-based issues, 45% stated that backup takes too long, while recovery taking too long was an issue for 38%. This suggests that 62% do not believe they have an issue with their recovery times. Given the average recovery time of 5 hours, it may be that CIOs do not recognize that backup and recovery capabilities should be better. Alternatively, they may have come to accept this level of performance as the norm thanks to using tools poorly suited to the virtual environment. Similarly, 26% stated that file- and application-level recovery was too difficult, 23% said that SLAs were missed and 16% that backup or recovery often fails (Chart 10). This again suggests either that CIOs may not be aware that they even have an issue with these tasks or that they have become used to the relatively poor performance of their existing backup and recovery tools, given the recovery times and frequency of backup and recovery failure referenced above.





Chart 10: Capability-related challenges identified (%)

- Backup or recovery often fails
- SLAs missed
- File- and application-level recovery too difficult
- Recovery takes too long
- Backup takes too long



Part II | Complexity Challenges for Organizations

2. Complexity Challenges for Organizations

While capability provides the most obvious issues for organizations, 84% of CIOs said that they are experiencing complexity-related challenges with backup and recovery of virtual environments. Among these are backups needing ongoing management (experienced by 57% of CIOs); backup tools being difficult to configure (33%); too many virtual servers to backup (32%); and difficulty backing up to tape (23%) (Chart 11). Modern data protection tools that can simplify management, configuration and scheduling of backups can be a huge help with these issues.




Chart 11: Complexity-related challenges identified (%)

- Backing up to tape is difficult
- Too many virtual servers to backup
- Backup tools difficult to configure and use
- Backups need ongoing management

A significant difference between physical backup tools and more modern approaches is the use of software agents on protected machines. By requiring agents to be installed, monitored and updated, agent-based backup adds an extra layer of complexity to data protection: in turn making it easier to miss SLAs. While this is the way data protection has traditionally been performed, agents can be done away with in a virtualized environment. This in turn removes the extra layer of management, making the process faster and less complex for the IT department.

Currently, 76% of CIOs surveyed say that their backup tool requires agents inside virtual servers (Chart 12). In turn, 77% of these experience problems or management issues due to agents: this represents 58% of all organizations surveyed.




Chart 12: Organizations experiencing issues with agent-based backup (%)

- Organizations experiencing issues with agent-based backup
- Organizations with no issues around agent-based backup
- Organizations not using agent-based backup

Common problems and challenges for CIOs using agent-based backup and recovery include: agent management, for example installation, upgrading, updating and managing conflicts, which affects 43% of those CIOs; backups failing far too often (32%); restores failing too often (28%); the expense of agent-based backup (20%); and agents slowing down system performance (14%) (Chart 13). Based on these statistics there will be a large number of organizations that are experiencing multiple issues with agent-based data protection, in turn increasing the cost and complexity and reducing the capability of their backup and recovery strategy while increasing the chances of missing SLAs.





Chart 13: Agent-based backup issues identified (%)

- Agents slowing down system performance
- Expense of agent-based backup
- Restores failing too often
- Backups failing too often
- Agent management

The survey also suggests that many organizations are unaware that agent-based backup tools have been superseded by more modern alternatives. While the majority of CIOs surveyed are experiencing issues that can significantly affect performance, 53% believe that it is better if a backup tool uses agents to aid backup and recovery. A further 19% believe that there is no difference between agent-based and agentless tools (Chart 14). As organizations become more familiar with the potential that modern data protection presents, we would expect this to change: essentially, organizations will recognize the issues that are caused by agent-based data protection and instead begin to favour modern, agentless data protection tools that enable far more rigorous SLAs to be met.



Chart 14: Opinions on use of agents in backup (%)



- It is better to use agent-based backup
- It is better to use agentless backup
- Either way will work
- Don't know

Part III | Cost Challenges for Organizations

3. Cost Challenges for Organizations

Finally, 87% of CIOs are also facing cost-related challenges with backup and recovery of virtual servers beyond those costs already incurred by lengthy downtime. These cover three distinct areas: high ongoing management costs, affecting 53% of CIOs; expensive licensing models (50%); and backups either requiring or using too much storage (42%) (Chart 15). Reducing management costs; easy-to-understand and low-cost licensing; and making backups as storage-friendly as possible will be vital to addressing these challenges.




Chart 15: Cost-related challenges identified (%)

- Backups require too much storage
- Expensive licensing models
- High ongoing management costs

Part IV | Potential Upheaval in the Data Protection

4. Potential Upheaval in the Data Protection Market

As we have seen, there are still issues with data protection in the virtual environment. The capabilities offered by many backup and recovery tools, especially physical tools retrofitted to work with virtual infrastructures, are still not at the level that should be expected of the technology, resulting in missed or unambitious SLAs. Cost and complexity challenges are further adding to the difficulty of implementing a suitable data protection strategy. At the same time agent-based tools are presenting a range of issues that make it harder for CIOs to do their jobs: while this is not fully recognized, retro-fitting physical tools onto new virtual environments does not look like solving organizations' problems. These issues can be clearly seen with techniques such as replication: while virtualization has minimized hardware-based barriers, uptake has not greatly increased since 2011.

However, there are signs that organizations are recognizing this and planning to change the way in which they work. Currently, 58% of organizations are planning to change their backup tool for virtual servers in the next 24 months (Chart 16): on average, the expected time until a change is in fact only 10 months (Chart 17). This suggests that by 2014 a large proportion of organizations will have fresh backup and recovery tools in place that may improve on a number of the issues above.




Chart 16: Organizations' plans to change backup & recovery tools (%)



Chart 17: Average timescale for changing backup tool (months)

Those CIOs planning to change their current tools cite a number of reasons around cost, complexity and capability. The top reason given is Total Cost of Ownership, including management and maintenance (51%). Also popular is complexity (47%), while 42% are planning to change due to the hardware and software costs of their current tool. Lack of capability provides other reasons for changing: failure to meet Recovery Time Objectives (32%) and Recovery Point Objectives (24%) are both given (Chart 18).




Chart 18: Reasons for changing current backup tool (%)

- Failure to meet Recovery Point Objectives
- Failure to meet Recovery Time Objectives
- Software and hardware costs
- Complexity
- Total Cost of Ownership

As we can see, cost, complexity and capability challenges are becoming increasingly important to organizations: so important that they are driving their data protection strategies for the next 2 years.

Appendix

Appendix 1: The State of Virtualization Data Protection

As server virtualization continues to grow in popularity, it is becoming an ever-more important part of the IT infrastructure. Indeed, virtualization is now the dominant means of providing IT services. This in turn provides new opportunities and challenges for enterprises. To begin, there is still the decision of what infrastructure to protect, and to what extent. A virtual environment is potentially much easier to backup and recover than a physical one: as a result, organizations have the capability to protect more of their infrastructure than ever before, while using fewer resources and at an ultimately lower cost. At the same time, since IT infrastructure is far less limited by physical constraints, enterprises can make greater use of more advanced techniques such as replication.

Currently, on average 51,42% of the production server estate is virtualized in organizations: virtualization has to some extent already proved itself as the primary IT infrastructure. Within the next 2 years, this proportion is expected to grow steadily to 63,48% (Chart 19).




Chart 19: Current and predicted virtual estates (%)

- Percentage of production server estate predicted to be virtualized in 2 years
- Percentage of production server estate currently virtualized

When backing up their virtual environments, IT departments generally use one of three approaches. The first is to use native tools that are part of their applications or operating system: these have the advantage of no additional licensing costs beyond the application or OS itself, but tend to have limited awareness of virtualization compared to specialist data protection tools. Second, organizations use a single third-party tool to backup both their physical and virtual environments: this allows an organization to continue using their legacy backup solution. However, those legacy solutions have not been designed to backup and recover virtual environments; meaning the cost-saving will be offset by a lack of performance. Last, organizations can use separate specialist tools to backup their virtual environments. While such tools are relatively new, they can exploit the nature of virtualization to provide the best possible performance and capabilities.

While virtualization is growing in popularity, the majority of organizations are not backing up every virtual server. 74% of organizations do not backup all of their virtual servers (Chart 20): on average, all organizations surveyed backup 68,28% of their virtual environment (Chart 21).



Chart 20: Percentage of Organizations protecting 100% of virtual servers (%)

- Protect 100% of virtual servers
- Do not protect 100% of virtual servers





Chart 21: Percentage of virtual infrastructure protected (%)


- Percentage of virtual infrastructure protected
- Percentage of virtual infrastructure unprotected



Organizations are divided on the tools they use to backup their virtual environments: 7% use native tools to backup their virtual servers, 54% use a single third-party tool to backup both physical and virtual servers, and 38% use a specialized tool for their virtual environments. Less than 1% of respondents use another method to protect their data, such as SAN replication (Chart 22). This suggests that organizations are still most comfortable using legacy tools to protect their virtual environments, rather than investigating more modern solutions.



Chart 22: Tools used to protect virtual servers (%)



Appendix 2: The Evolution of Replication

One data protection technique that has been greatly improved by the use of virtualization is replication. Replication is typically a process of copying data to production standard hardware that can be quickly brought back online in the event of a server or site failure. This differs from the process of "backup", whereby data is basically compressed and then stored on relatively inexpensive hardware. In the event of data loss or a server or site failure, the backup must first be restored before the data can be brought back online.

Server replication has traditionally been a cost- and resource-intensive process, especially as most replication solutions must be purchased separately to backup tools. While virtualization can help make replication less costly, for example by enabling more efficient creation of the required infrastructure, its use is not yet universal. Currently, 76% of organizations replicate at least a few servers: 51% using hardware-based replication and 25% using software-based replication (Chart 23).





Chart 23: Use of replication by organizations (%)

- Organizations using hardware-based replication
- Organizations using software-based replication
- Organizations not using replication

However, replication has a clear cost benefit for those organizations that use it. Replicated servers would cost \$409 531 per hour of downtime if they were not so protected (Chart 24). Given the average time to recover a server of at least 5 hours, we can see that those organizations using replication are essentially saving themselves over \$2 million each time they need to make use of their replicated servers.



Avg	\$409 531
DE	\$328 409
FR	\$419 270
UK	\$393 038
US	\$446 296

Chart 24: Cost-per-hour of replicated servers being down (USD)

Organizations are also experiencing almost the exact same barriers to increased use of server replication as they did in 2011 (Chart 25). The top 3 barriers are, again, the cost of replication software (53%), cost of hardware (51%) and complexity (44%). Interestingly, the positions of hardware and software costs have swapped since 2011, indicating that hardware costs may be becoming less of an issue as virtualization grows in use. For those organizations that do not use replication at all, the top barriers have changed a little since 2011, with 53% citing the cost of replication software, 45% the cost of hardware, 36% a lack of disk space and 35% the complexity (Chart 26). Again, the barriers presented by hardware and software have swapped importance while a lack of disk space is now held to be almost exactly as much of an issue as complexity, whereas in 2011 it was some way behind.




Chart 25: Issues presenting greater use of replication, 2013 & 2011 (%)

- Complexity 2011
- Complexity 2013
- Cost of hardware 2011
- Cost of hardware 2013
- Cost of replication software 2011
- Cost of replication software 2013




Chart 26: Issues preventing adoption of replication, 2013 & 2011 (%)

- Complexity 2011
- Complexity 2013
- Lack of disk space 2011
- Lack of disk space 2013
- Cost of hardware 2011
- Cost of hardware 2013
- Cost of replication software 2011
- Cost of replication software 2013

About Veeam Software

Veeam® Software develops innovative solutions for [VMware backup](#), [Hyper-V backup](#), and [virtualization management](#).

Veeam Backup & Replication™ is the [#1 VM Backup](#) solution. Veeam ONE™ is a single solution for real-time monitoring, resource optimization, documentation and management reporting for VMware and Hyper-V. Veeam extends deep VMware monitoring to Microsoft System Center with Veeam [Management Pack™](#) (MP), and to HP Operations Manager with Veeam [Smart Plug-In™](#) (SPI). Veeam also provides [free virtualization tools](#). Learn more by visiting www.veeam.com.

Virtualization changes everything – especially backup. If you've virtualized on **VMware** or **Hyper-V**, now is the time to move up to the backup solution Built for Virtualization: Veeam.

Unlike traditional backup that suffers from the “3C” **problem** (missing capabilities, complexity and cost), Veeam is:

- **Powerful:** Restore an entire virtual machine (VM) or an individual file, email or database record in 2 minutes
- **Easy-to-Use:** It just works!
- **Affordable:** No agents to license or maintain, works with your existing storage, and includes deduplication, VM replication, Microsoft Exchange recovery, and more

Join the 60,000 organizations who have already modernized their data protection with Veeam.
Download Veeam Backup & Replication today! To learn more visit [our website](#).