VEEAM

Virtualization Data Protection Report 2011

Executive Summary

The Veeam Software Virtualization Data Protection Report 2011 builds upon the key findings from the previous year to track the progress of enterprise-level organizations' data protection strategies. Where the Virtualization Data Protection Report 2010 looked at the broader trends and issues related to data protection, this year's report focuses on the process of server replication, a critical component of data protection. It identifies how financial and resource constraints are leading organizations down a rocky path where tough decisions are being taken around data protection.

As with the previous year's survey, this report highlights how a 'physical world' mindset applied to virtualization continues to limit the true potential of this technology.

The concluding sections of the report outline key trends and issues that will need to be addressed before enterprises can further benefit from virtualization-enabled data protection.

Vanson Bourne, an independent market research organization, conducted an online survey in August 2011 of 500 CIOs from organizations across the United States, United Kingdom, Germany, and France that employ more than 1,000 people.

Virtualization Transforms Data Protection

A key insight from the 2010 Virtualization Data Protection report was that virtualization has the potential to transform data protection. For the 2011 study, CIOs were therefore directly asked whether they believed this to be true and an overwhelming 94% of agreed that virtualization can transform data protection strategies *(Chart 1)*.

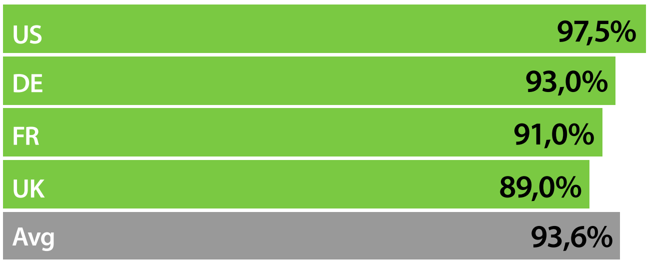

Chart 1: Virtualization Transforms Data Protection

Survey Background

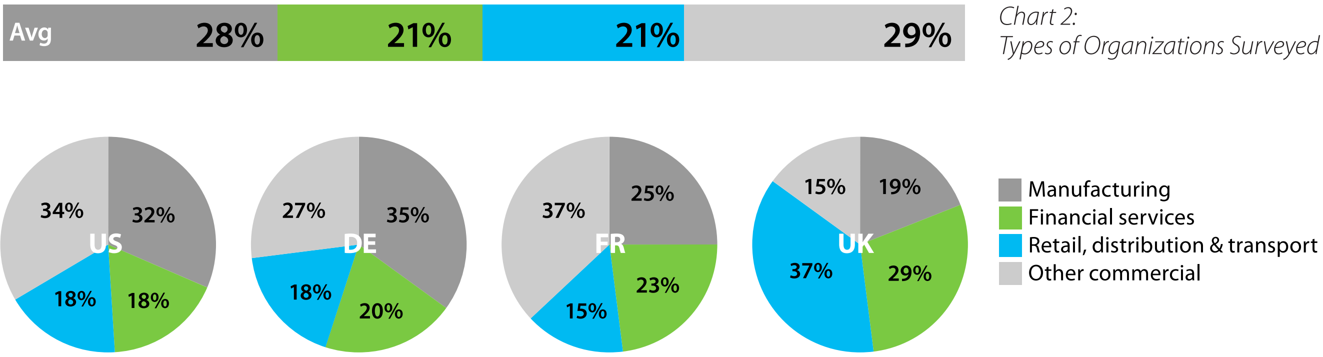

Chart 2: Types of Organizations Surveyed

As with the 2010 virtualization data protection report the organizations in this year's research survey *(Chart 2)* had a wide geographical spread of employees, with over half of the businesses in the sample having employees based overseas. Similarly, the companies were diverse, with the sample split almost equally between manufacturing (28%); financial services (21%); retail, distribution & transport (21%); and other commercial (29%). This ensured that the statistics generated were gathered from many different types of commercial organizations across different market sectors.

1. Data Protection: More Data, Less Protection, Greater Risk

Whilst the 2010 survey looked at the general area of physical and virtual server backups another important aspect of data protection, particularly in the area of data recovery, is server replication. Replication is typically a process of copying or backing up data to production standard hardware that can be quickly brought back online quickly in the event of a data loss. 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, this must first be restored before it can be brought back online. Note: There are new virtualization-enabled technologies such as Veeam vPower that give the ability to start up a Virtual Machine (VM) directly from a compressed and de-duplicated backup file, without first restoring it. Learn more about Instant VM Recovery at: http://www.veeam.com/vmware-esx-backup.html

Server replication has traditionally been a cost- and resource-intensive process, especially as most replication solutions must be purchased separately to backup tools. In order to understand the value of virtualization in server replication, it is important to look at exactly where mission critical data resides and then the challenges enterprises currently face in protecting it.

This year's survey reveals that nearly half (45.64%) of all enterprise server estates both physical and virtual are viewed as mission critical *(Chart 3)*.

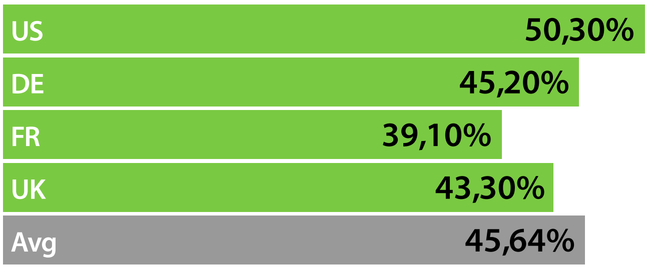

Chart 3:
Average Percentage of Mission Critical Servers

However, according to the survey, in enterprises that use server replication only 26.12% of business critical servers are replicated *(Chart 4)*.

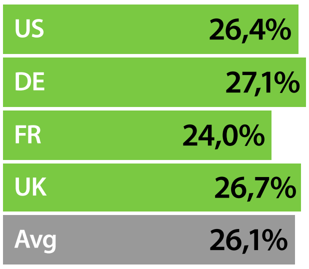

Chart 4:
Percentage of Mission Critical Servers Replicated

Unsurprisingly, the top reasons for server replication include: 1) Protection from Data loss (85%) 2) Protection from hardware failure (70%) 3) Protection from regular human error (49%) and Protection from data centre failure (49%) *(Chart 5)*.

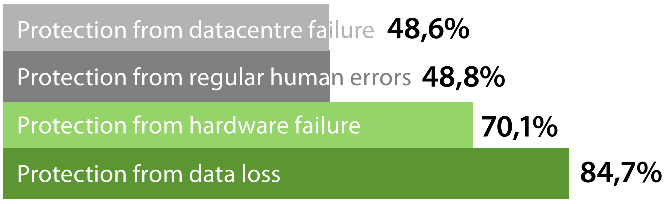

Chart 5:
Reasons for Server Replication

In terms of approaches to server replication, 45% of enterprises use hardware-based replication; 33% use software-based replication; 22% of enterprises do not use replication, only backup *(Chart 6)*.

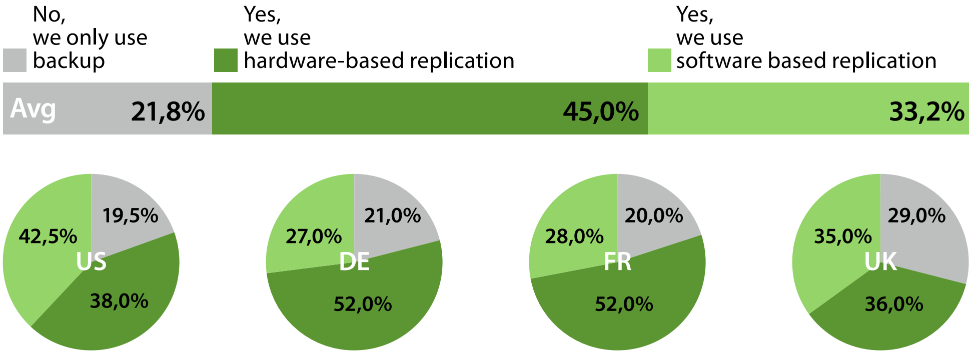

Chart 6:
Approaches for Replication

Drilling down further into the approaches for replication, the top three types of server replication include: 1) Replication of Physical servers, on-site (61%) 2) Virtual server level, on-site (43%) 3) Virtual server level, off site (40%) *(Chart 7)*.

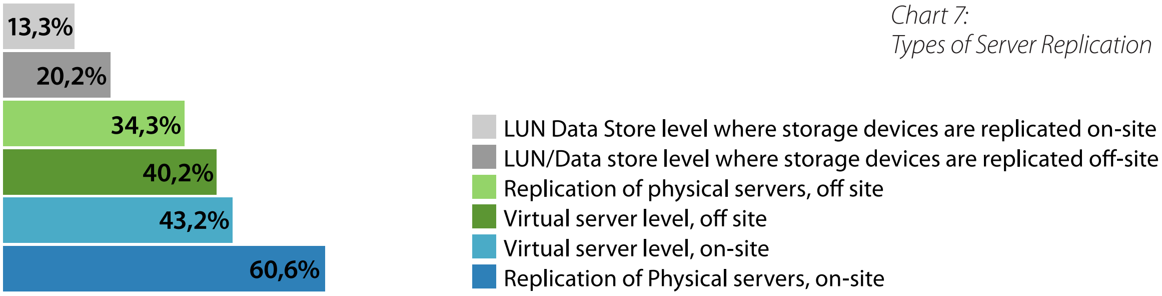

Chart 7:
Types of Server Replication

Cost implications

In order to understand the value of server replication, firstly CIOs from organizations that did not include replication as part of data protection were asked to estimate the cost per hour of outage to their business critical server estate. The average estimate here is over a quarter of a million dollars ($368,692) per hour *(Chart 8)*.

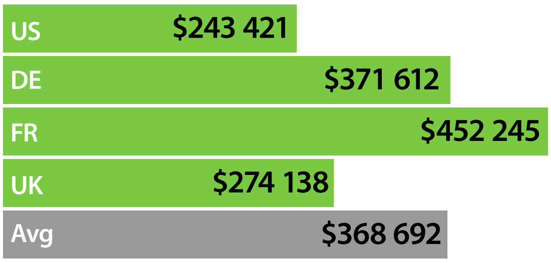

Chart 8:
Cost Per Hour of Critical Server Outage (In organisations that do not use replication), USD

Separately, CIOs in organisations that use replication for a portion of their business critical servers were asked to estimate 1) the cost of outage to replicated servers if they were not replicated and 2) the cost of outage to the remaining business critical servers that are currently not replicated. Taking into account the data revealing how on average only a quarter (26.12%) of business critical servers are replicated, these findings highlight further limitations around enterprise data protection strategies *(Chart 9)*.

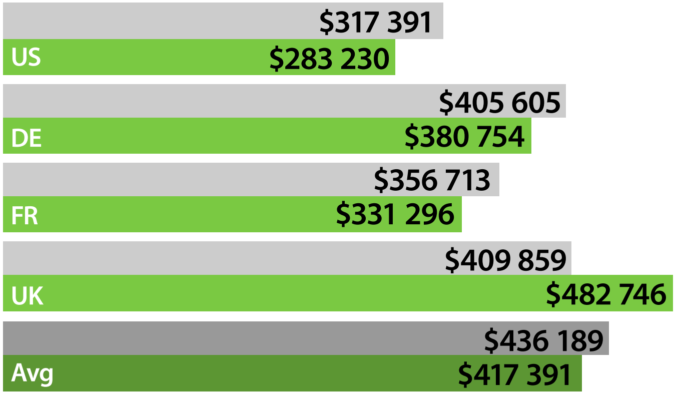

Chart 9:
Comparison of Cost per Hour of Critical Server Outage — Replicated Vs Non-Replicated (in organisations that use replication), USD

⬜ Non-Replicated Servers
🟩 Replicated Servers

Firstly, of those 26.12% of servers that are currently replicated, the cost per outage if they weren't is estimated at $417,391. Nonetheless, the cost per hour of outage to the remaining three quarters of business critical servers is $436,189. This suggests that enterprises today are taking calculated risks around the protection of business critical data. Moreover, having the ability to replicate a larger volume of business critical servers is critical. If this is not made possible then over the long term, as the volume of business critical data and servers increase, these risk decisions will inevitably become even more difficult. The goal for CIOs is therefore to improve data protection strategies by having the ability to replicate more data.

2. Replication: Addressing the cost conundrum

Increased replication will be clearly become a major requirement of a successful data protection strategy, but the reality for all enterprises is that this is currently an unachievable goal.

Even today, one of the major issues around data protection is how regularly a server is backed-up or replicated. In short, the stronger the data protection process the smaller the chance of data loss. Yet the survey reveals that when it comes to the frequency of replication of business critical servers, the timing ranges from Continuous Data Protection (CDP) to just once a day (Daily) — *(Chart 10).*

Chart 10: Frequency of Server Replication

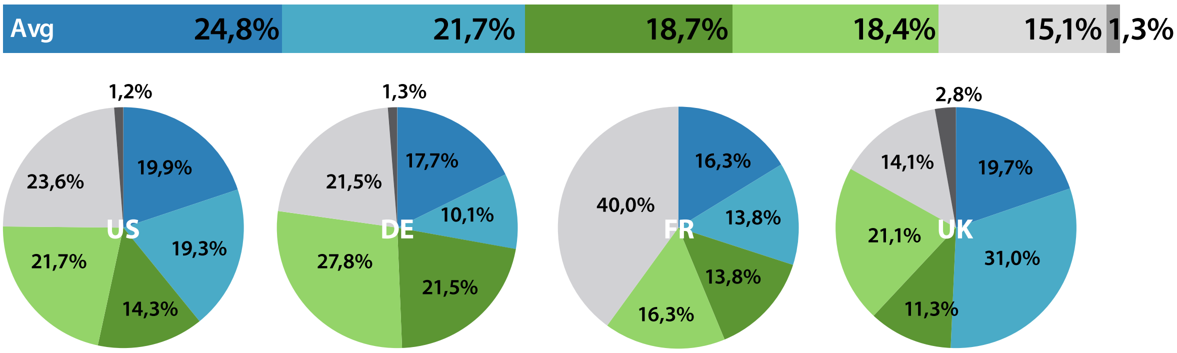

🟦	Daily	25%
🟦	Several times a day	22%
🟩	CDP	19%
🟩	Near-CDP (Incremental data replication or backup is performed every 15 minutes or less)	18%
⬜	Hourly	15%
⬛	Less regularly than daily	1%

As data is backed-up on production level hardware, one assumption with server replication is that the data will indeed be recoverable in the event of loss. This is only true with certain solutions, such as Continuous Data Protection (CDP). However, CDP only constitutes 19% of server replication approaches in enterprises that use replication. Furthermore, according to the survey 80% of CIOs say high cost is the key barrier preventing adoption of this solution. As with general back-ups, testing the recoverability of replicas is also critical to ensuring a strong data protection strategy. However, the survey reveals that on average enterprises only test the recoverability of replicated servers every 9 weeks *(Chart 11).*

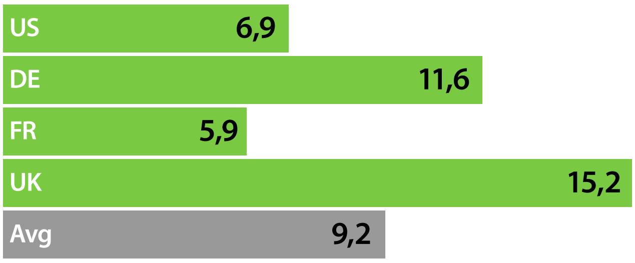

Chart 11: Frequency of Testing Replicas for Disaster Recovery, weeks

When CIOs from all organisations (i.e. including those that both use and don't use replication) were asked about the key barriers preventing them from using replication, the top three answers were exactly the same. The top three barriers preventing increased replication were 1) Cost of hardware (60%) 2) Cost of replication software (52%) and 3) Complexity (42%). Likewise the key factors preventing organisations from adopting replication were 1) Cost of hardware (55%), 2) Cost of replication software (44%) and 3) Complexity (39%) — *(Charts 12 and 13).*

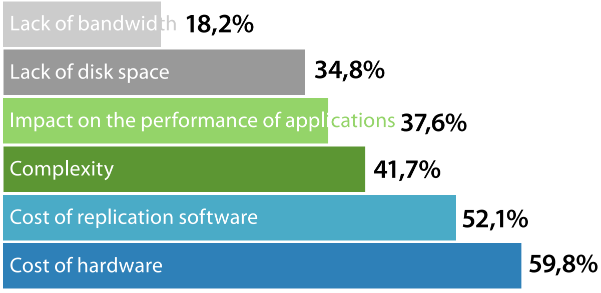

Chart 12:
Barriers to Increased Server Replication

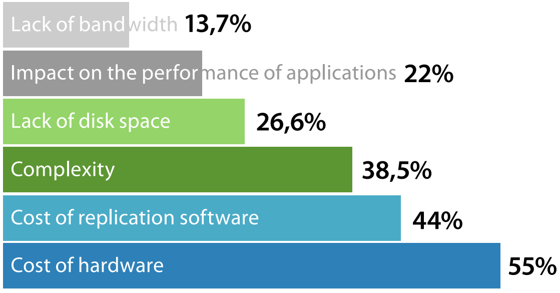

Chart 13:
Factors Preventing Server Replication

The relatively low cost and hardware-agnostic benefits of virtualization should offer enterprises the ability to potentially address the budget challenges faced with replication. As more virtual machines can be replicated to a recovery site it is possible to at least reduce the cost of hardware. Yet 80% of CIOs said that due to the agent-based approach of traditional replication solutions, there is minimal difference between physical and virtual machines when it comes to the actual volume of data that can be replicated. Currently, even virtualization is not improving the ability to increase data replication *(Chart 14)*.

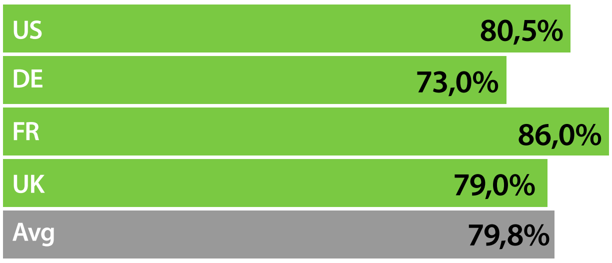

Chart 14:
Virtualization Does Not Improve Ability to Increase Server Replication

3. Disaster Recovery: Pushing the boundaries of acceptable risk

Similar to the 2010 report, many of the issues raised in this latest 2011 virtualization data protection report are a major consequence of a physical IT mindset. Consequently, enterprises are continually pushing the boundaries of acceptable risk at a time when this is no longer necessary. The 2011 report continues to see evidence of this mindset, particularly in the area of disaster recovery.

Escalating risks from disaster

One of the central components of any data protection strategy is the strength of enterprise disaster recovery processes. Here the 2011 survey reveals a number of insights:

Looking at the rate of growth in business critical data within their businesses, CIOs raise a major concern around disaster recovery. More specifically, 87% of CIOs say recovery times from large-scale disaster continue to grow as the volume of business critical servers (physical and virtual) within the enterprise increases *(Chart 15).*

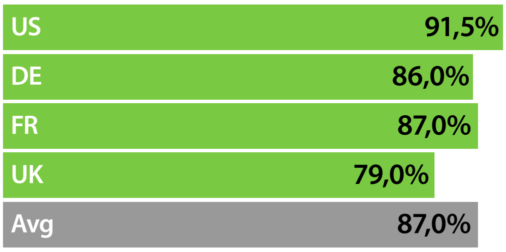

Chart 15:
Disaster Recovery Times Growing as Volume of Business Critical Servers Increases

According to the 2011 survey, while recovery from physical server loss is currently reported at 5 hours on average, there is only a 1 hour difference from the average virtual server recovery time of 4 hours *(Chart 16)*. As virtual server deployments begin to exceed those of physical servers, this minor difference in recovery time is both unsustainable and highly risky in the event of a disaster.

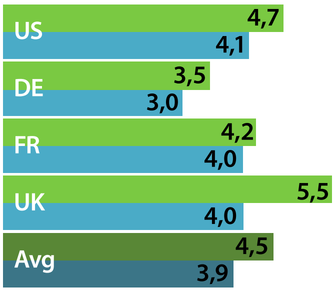

Chart 16:
Average Server Recovery Time — 2011, hours

🟩 Physical Servers
🟦 Virtual Servers

Going deeper into specific aspects of data recovery reveals that enterprises also incur significant delays in granular-level recovery from virtual machines. More specifically, the average recovery time for an individual application item is 1 hour and recovery of an application or file 2 hours *(Charts 17 and 18)*.

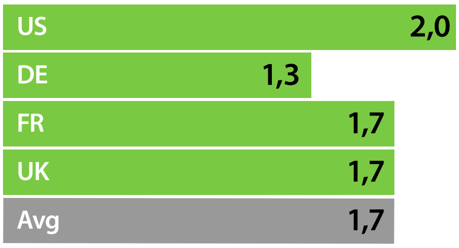

Chart 17:
Application Recovery Time
from Virtual Server, hours

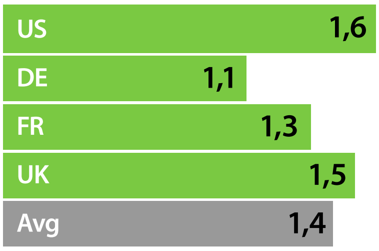

Chart 18:
File Recovery Time
from Virtual Server, hours

These continued challenges around virtual and physical server recovery explain why 79% of CIOs today believe that as the growth of data increases, the current tools used for disaster recovery will be less effective *(Chart 19)*.

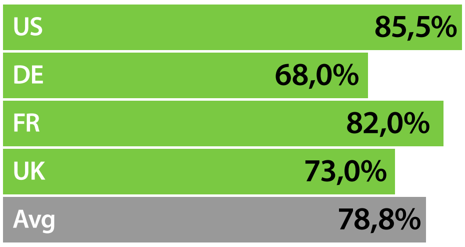

Chart 19:
Existing Tools for Disaster Recovery
will Become Less Effective

Despite the issues challenges with server replication, as noted in Chart 1 there is a strong sense that virtualization does have a role to play in improving disaster recovery. Again, 94% of CIOs say virtualization can transform data protection strategies.

Summary and Conclusions

This year's annual report further demonstrates not only the continued limitations of physical world-based data protection strategies but also how the true potential of virtualization is not effectively being extended to disaster recovery. More specifically, it highlights the following key trends and issues that will need to be addressed before enterprises can further benefit from virtualization-enabled data protection:

- As replication is currently only possible on a quarter of business critical servers, a vast proportion of IT infrastructure is exposed to lengthy outages following a disaster
- Until replication can be extended to a wider proportion of server estates, enterprises will continue stretch the boundaries of acceptable risk following an IT disaster
- Despite the growth of virtual infrastructure, disaster recovery times will continue to grow due to ineffective tools

About Veeam Software

Veeam Software, innovative provider of VMware data protection, disaster recovery and virtualization management solutions for virtual datacenter environments. This second annual Virtualization Data Protection report reveals a particular flaw in scalability of enterprise disaster recovery processes, Veeam acknowledges that these limitations are likely based on a lack of knowledge around the true potential of vPower™ technology for Virtualization-Powered Data Protection™.

In November 2011, Veeam released Veeam Backup & Replication™ v6. This new version of Veeam's award-winning solution improves upon the ground-breaking vPower™ technology for Virtualization-Powered Data Protection™. Veeam Backup & Replication v6, a product which combines Backup & Replication in a single solution at no additional cost, allows businesses to overcome the limitations of traditional disaster recovery tools to enable increased server replication at a fraction of the cost and complexity. The latest features include:

- **Enterprise scalability:** Enhanced distributed architecture streamlines deployment and maintenance of remote office/branch office (ROBO) and large installations. Also speeds up backup, replication and restore over WANs.
- **Advanced replication:** Accelerates replication by 10X, streamlines failover and provides real failback with delta sync.
- **Multi-hypervisor support:** Reduces the cost and complexity of managing multi-hypervisor environments with new support for Windows Server Hyper-V and Microsoft Hyper-V Server within Veeam's existing data protection infrastructure (one install and one console).
- **Numerous enhancements, including 1-Click File Restore:** Extends Veeam's existing file-level recovery with delegated, web-based restore directly to the original virtual machine (VM), without requiring a direct network connection or in-guest agent.

In addition, Veeam Backup & Replication v6 includes the key Veeam innovation vPower™ — that runs a VM directly from a compressed and deduplicated backup file in production or in an isolated lab, enabling these five industry firsts:

1. **Instant VM Recovery:** Recover an entire VM from a backup file in minutes.
2. **U-AIR™ (Universal Application-Item Recovery)**: Recover individual items from any virtualised application, on any operating system, without additional backups, agents or software tools.
3. **SureBackup™**: Automatically verify the recoverability of every backup, of every VM, every time.
4. **On-Demand Sandbox**: Create test VMs from any point in time to troubleshoot problems or test workarounds, software patches and new application code.
5. **Instant File-Level Recovery for any OS and file system**: Recover an entire VM or an individual file from the same image-level backup.

© 2011
Veeam Software
All rights reserved.
All trademarks are the property of their respective owners.