SURVEY REPORT

VIRTUALIZATION & Cloud Review

2018 VIRTUAL ENVIRONMENT DATA PROTECTION REPORT

By Nick Cavalancia

iland™

www.iland.com

Monitoring and managing network flows is a critical part of a secure and efficient approach to IT. Unfortunately, it's difficult to design networks to be monitored by the good guys without making it easy for the bad guys to do the same thing. One solution to this conundrum is packet brokers.

Even virtual environments require a data protection strategy – one that ensures the organization has both an ability to properly backup and recover data, applications, and VMs based on current business requirements. Many IT professionals think that the redundancy, durability, accessibility, and availability that comes with virtualization and the cloud is enough. But the reality is that, whether on-premises or in the cloud, even virtual environments can become victims of data corruption, ransomware attacks, data breaches, espionage, and more – all putting the state of your data in question.

So, how are organizations protecting their virtual environments?

With downtime propelling organizations into the headlines, costing as much as millions of dollars to rectify data loss, reputation concerns, loss of revenue, and more, it's necessary for every organization to have a data protection plan that ensures the business remains operational. It's reasonable to assume most organizations have some kind of plan or process in place. The question is what specific backup and recovery strategies, methods, and tactics are employed and how effective are they.

To find out, iland and Veeam partnered to survey 300 organizations about the current state of their virtual environment data protection, what kinds of outages are of concern, and what they are doing to remediate such scenarios.

In this report, we'll begin by looking at how organizations backup their virtual environments. We'll then take a look at the specific scenarios organizations are either worried about or have faced. Lastly, we'll cover how well organizations are able to recover their virtual environments based on business need.

ABOUT OUR RESPONDENTS

Nearly 300 organizations participated in this year's report, representing 25 countries, of which the US had the greatest representation (86%).

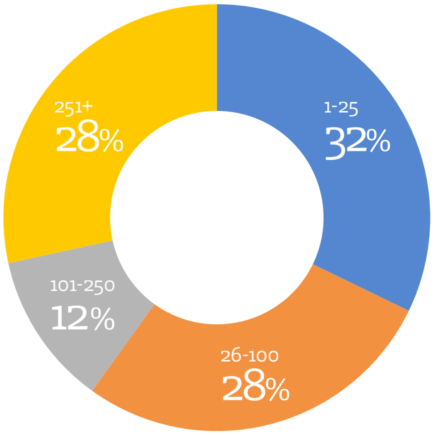

Valery Brozhinsky / Shutterstock.com

NUMBER OF VIRTUAL MACHINES

Response by size of virtual environment (# of VMs, shown at right) gave us a pretty equal level of visibility into what organizations with varying sizes of environments do to protect them.

INVESTED IN THE CLOUD?

Most of our respondents believe in the use of the cloud. A majority of them (80%) have some portion of their virtual environment in the cloud, with 29% of them hosting more than half of their environment there.

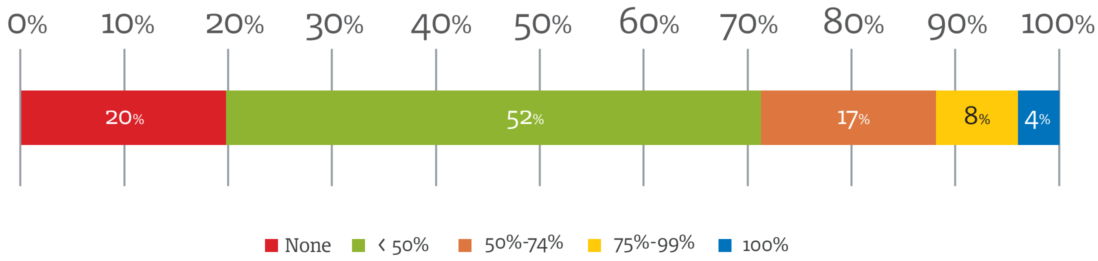

Percentage of virtual infrastructure that is cloud-based

BACKING UP THE VIRTUAL ENVIRONMENT

The first step in any data protection plan is to create backups of your virtual environment. Whether used as part of a simple recovery strategy or a complex automated disaster recovery effort, backups are the foundational element that drives your recoverability. We sought to find out exactly how organizations are backing up their virtual environment by asking a few questions.

HOW OFTEN?

We asked what percentage of the virtual environment was backed up at least daily. As shown below, the largest response in each VM environment segment was that of organizations backing up 100% of the virtual environment, with an average of 37.5% of organizations.

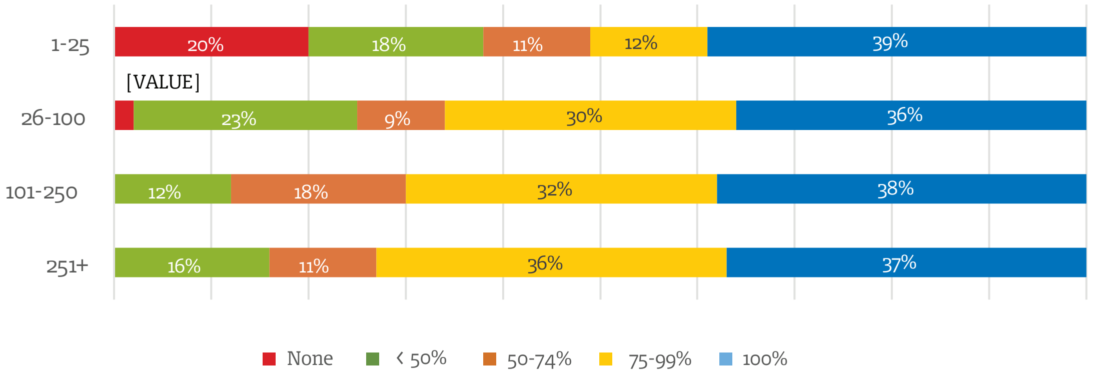

PERCENTAGE OF THE VIRTUAL ENVIRONMENT BACKED UP

It was surprising to note that an average of nearly 23% of organizations backup half or less of their VMs daily, with the largest representation by those with 1-25 VMs.

HOW?

Just as not every organization backs up the same amount of their virtual environment, organizations don't necessarily back up their VMs the same way either. We asked how VMs are being backed up. As shown below, organizations use an average of two methods, with the use of on-premises disk-to-disk as the primary method.

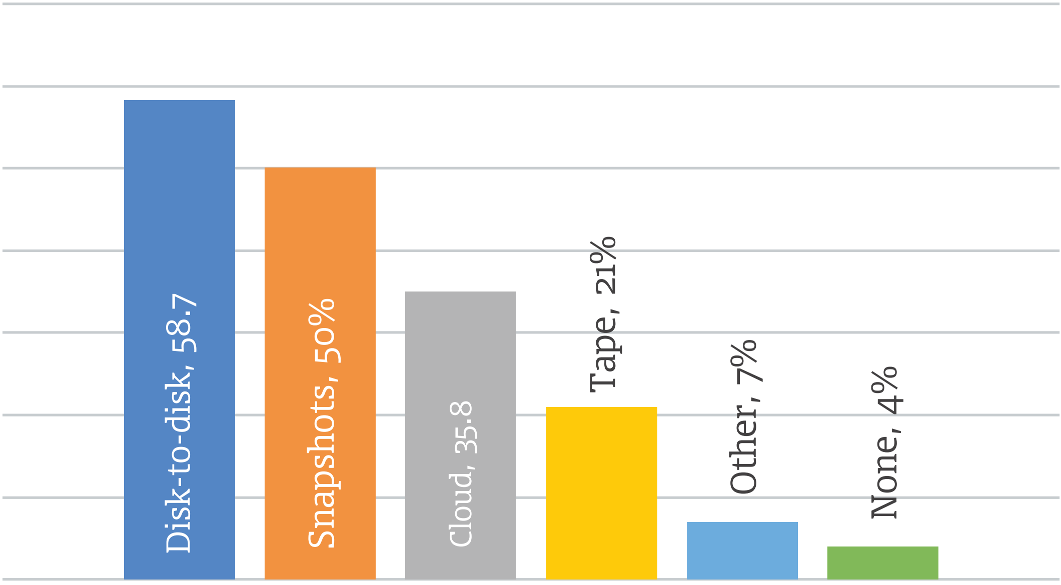

The technology used to backup VMs

Interestingly, we found that only 43% of those organizations with some portion of their VMs in the cloud, use the cloud as one of their backup technologies.

INTERESTINGLY, WE FOUND THAT ONLY 43% OF THOSE ORGANIZATIONS WITH SOME PORTION OF THEIR VMS IN THE CLOUD, USE THE CLOUD AS ONE OF THEIR BACKUP TECHNOLOGIES.

GOING OFFSITE?

As part of the 3-2-1 Backup Rule, every organization should have at least one copy offsite (the 1 in the rule). Only 14% of organizations (shown below) have no backups stored offsite, with another 14% considering data mirroring as an offsite backup (although susceptible to data corruption)

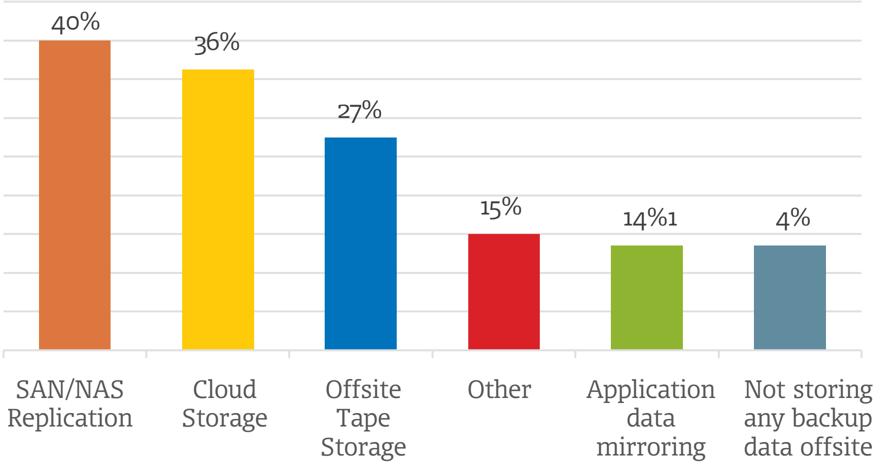

Where offsite backups are stored

SAN/NAS, offsite tape, and application data mirroring were most seen in the larger environment sizes. Use of the cloud storage was most prominent in the smaller VM environments, higher than even the largest environments use of the cloud.

CONCERNED ABOUT OUTAGES

Organizations face a number of possible threats to their business continuity – each having an impact on production, people, and profits. Organizations are generally most concerned with the direct impact an outage has on their data, as shown below. Productivity came in a distant second, with other impacts to the business coming in last.

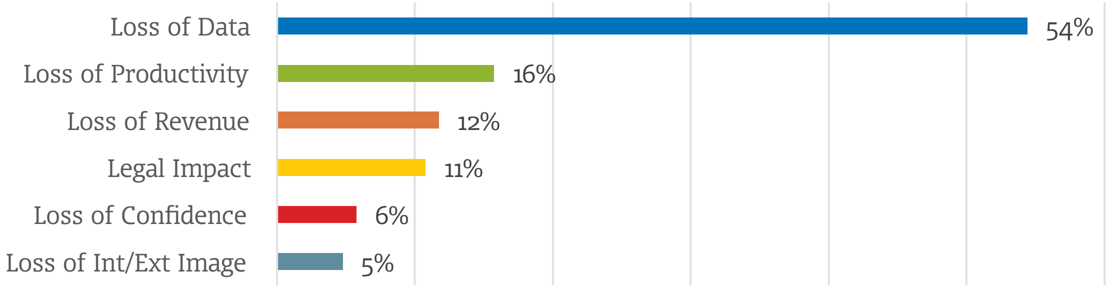

What worries organizations the most about a data interruption event

So, what happens when an outage actually occurs?

A little more than three-quarters of organizations (as shown below) had one or more issues resulting in an outage. Hardware Failures and Human Error were experienced most by the largest VM environments.

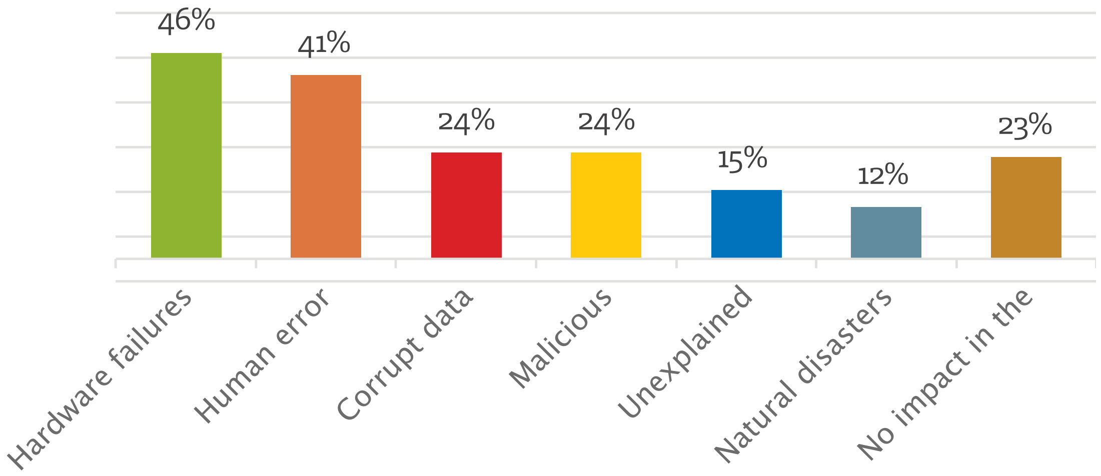

We specifically inquired about any occurrences of one of the most damaging and widespread sources of downtime – ransomware. Of those organizations impacted by ransomware, only 4% paid the ransom, with 71% leveraging backups and/or a formal disaster recovery plan to retrieve the ransomed data.

With organizations concerned about the impact on data, we asked how much downtime was experienced with each outage incident. As shown at right, of those organizations experiencing outages, nearly two-thirds of incidents (66%) caused four hours of downtime or less. The larger VM environments tended toward longer downtime durations, with half of the "less than 15 minutes" responses belonging to the smallest VM environments.

With so much potential downtime possible, it's imperative for organizations to have a plan that ensures the ability to recover as quickly as possible. In the next section, we'll look at just how prepared organizations today are to quickly recover their virtual environments.

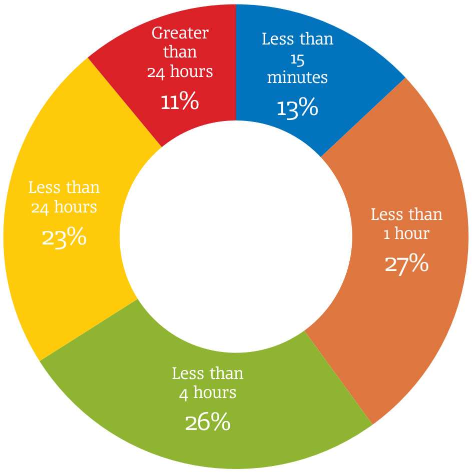

Average downtime per incident

RECOVERING THE VIRTUAL ENVIRONMENT

It's not enough to simply have backup copies of your VMs. What's necessary, at a minimum, is a recovery methodology that has undergone planning and testing to ensure it will work in your organization's time of need.

As shown at right, over two-thirds of organizations have needed to fully recover an application or virtual machine due to an outage.

In this section, we'll cover several factors that will both demonstrate how organizations are addressing recovery and determine whether that intent is enough to see a successful recovery.

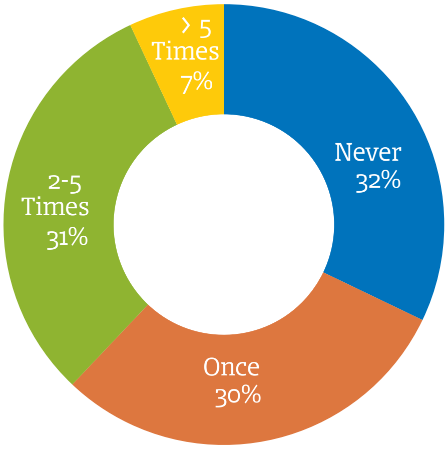

Numer of times orgs recover annually

BACKUP TESTING

One of the cardinal rules of backups is to test your backups before a recovery event to ensure the backup data set is viable. As shown below, there are some loose consistencies around the backup test frequencies based on VM environment size. In the larger three segments, we see how Monthly, Quarterly, Annually, and Only at Recovery have similar results.

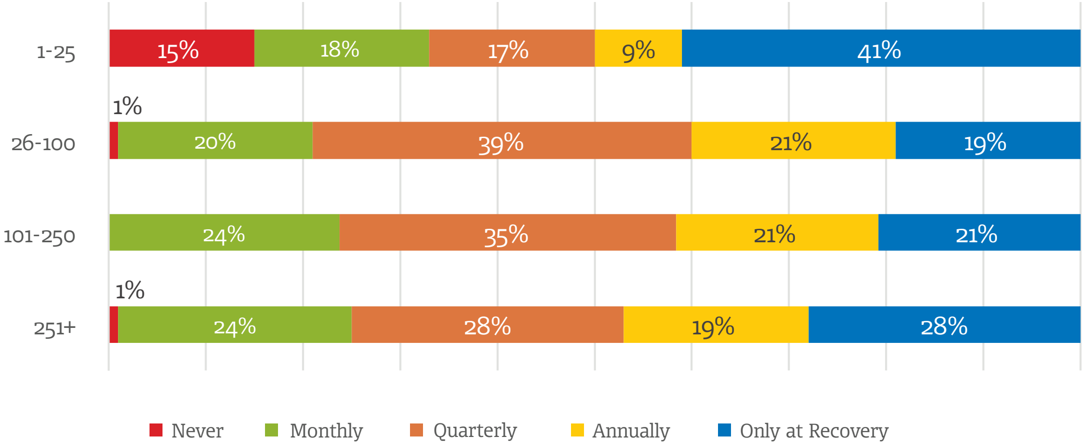

How often orgs test recovery from a backup

What's surprising is the material portion of each segment that only test annually or at recovery. With a combined average of nearly 45% of organizations, this lack of testing may reveal issues when it comes to recovering during a disaster scenario.

PRESENCE OF A DR PLAN

Looking beyond the basic backup and recovery of VMs, the presence of a disaster recovery plan helps to define how the process of getting the business back up and running (rather than, say, a single VM) should occur. A DR plan defines levels of workload criticality, associated recovery objectives (which can differ based on workload), testing schedules, recovery processes, contingency plans, and more.

Not every VM workload necessarily needs to be a part of a DR plan – a temporary development VM as part of a DevOps initiative, for instance. But all critical workloads should be considered when building the plan. But, as shown below, there is nearly 50% representation by organizations that either have less than half their VMs protected by their DR plan, or have chosen to have none of their virtual environment protected by that plan. The lack of a DR plan is greatest in organizations with less than 25 VMs (39%). It's also surprising to such a material percentage of organizations (21%) in the None category.

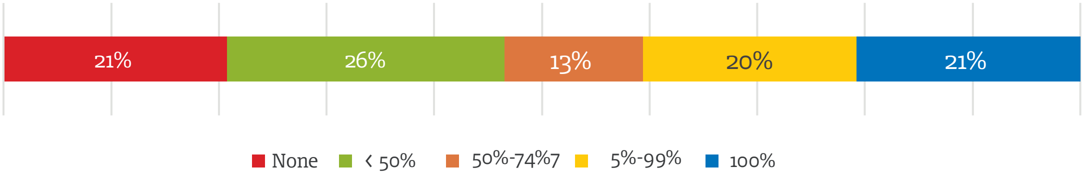

Percentage of VMs that are protected as part of a disaster recovery plan

Organizations utilize a variety of technologies to execute on a DR plan. As shown below, most are using replication methods (at the VM and data levels). The challenge with replication alone is data durability. If the source is encrypted from ransomware, for example, so is the replicated data.

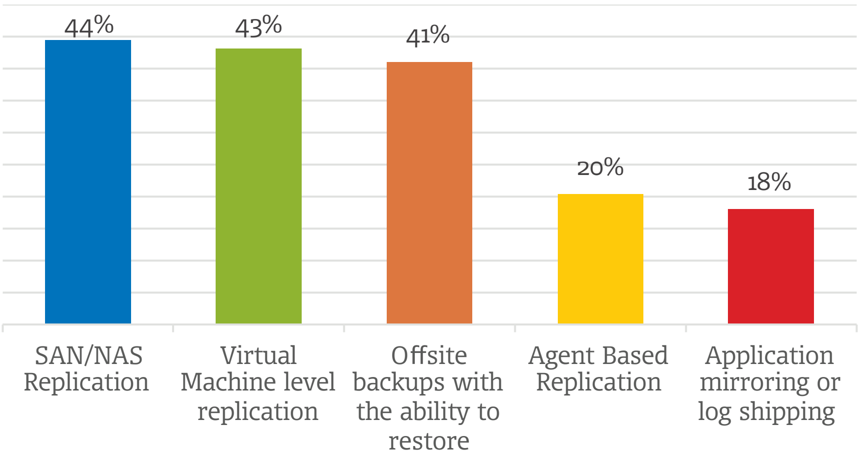

Types of disaster recovery technology used as part of a DR plan

The choice to restore from offsite backups is a close third choice, aligns more closely with the 3-2-1 backup rule, and facilitates a more flexible recovery should there be a loss of location or network connectivity.

DR TESTING

More critical than with backups, the testing of the DR plan is a critical exercise. Simulating real disaster conditions, the plan involves validating the infrastructure, backups, people, and processes all work properly, resulting in a fully recovered mock environment.

As shown at right, over half of organizations test either annually or during recovery itself (which, by definition isn't a test; it's the recovery!). Even the largest VM environment segment were guilty – those with more than 251 VMs represented the largest contributor to annual testing and second largest to testing during recovery.

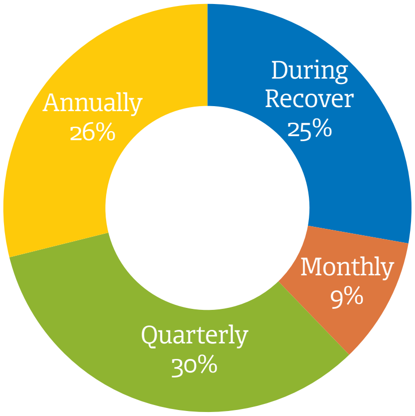

DR Plan Testing Frequency

RECOVERY OBJECTIVES

One of the foundational definitions of a recovery strategy is the recovery objectives. These are often established based on the criticality of a given data set, application, VM, or workload. Both the Recovery Time Objective (which defines how long you have to recover) and the Recovery Point Objective (which defines how much data can be lost) establish the parameters by which IT must perform their backups and recovery.

While most mature recovery strategies have multiple RTOs and RPOs (each based on the business requirements of a given workload), we asked organizations to characterize their recovery time and point requirements.

As shown below, the majority of organizations require either a few hours or as long as 24 hours to recover critical applications. While these aren't RTOs, they do demonstrate the current capability of organizations to recover their virtual environment.

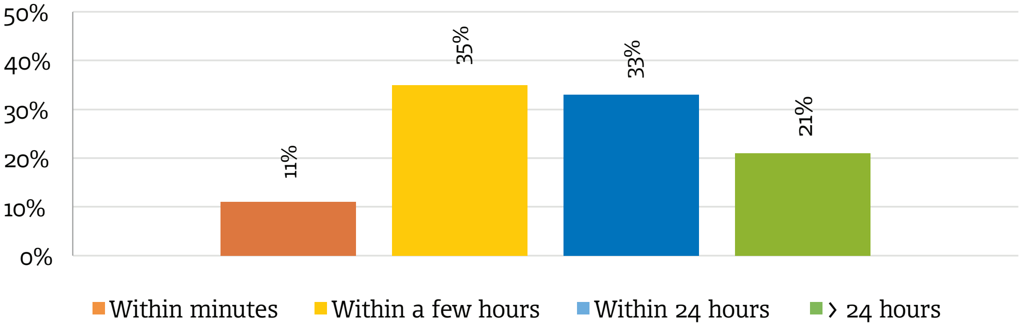

Time necessary to bring critical applications back online

To get a better idea of whether these are acceptable to the business, we can look at what is in essence an RPO. As shown below, 25% of organizations will accept zero data loss. Given that only 11% of organizations (shown above) can recover within minutes, it's highly unlikely that an adequate DR infrastructure and plan is in place to facilitate zero loss of data and near-100% uptime.

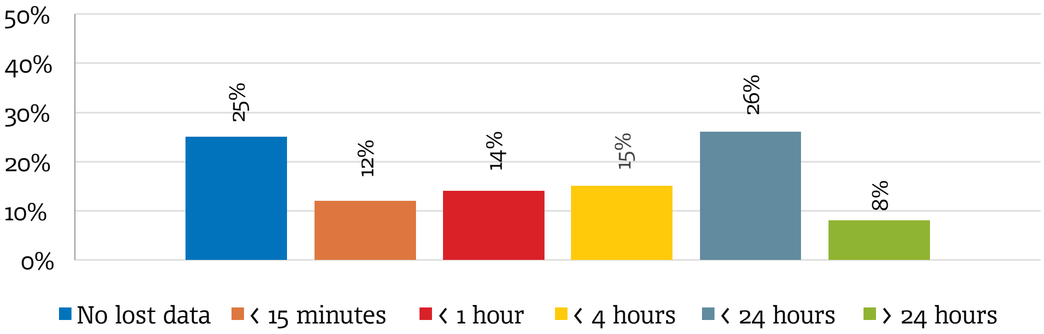

Tolerable amount of data loss seen as acceptable by the business

As shown below, the vast majority of organizations are either Confident or Very Confident that they have an ability to meet the established recovery time and point objectives.

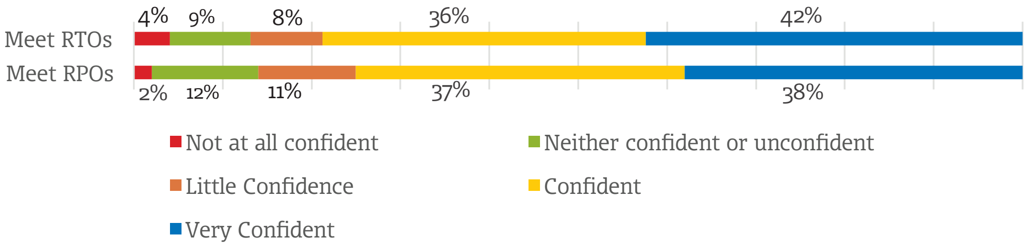

Organizational confidence in meeting recovery objectives

A FRACTION OF VMS ARE BEING BACKED UP. THERE'S A CLEAR RELIANCE UPON REPLICATION THAT POTENTIALLY HAS NO RECOURSE IN THE CASE OF CORRUPTION OR RANSOM.

RECOVERING THE VIRTUAL ENVIRONMENT: SUCCESSFUL OR NOT?

We've covered a lot of information around how organizations are currently protecting their virtual environments. But the question needs to be answered are they doing enough to successfully recover?

Let's recap a few facts uncovered by this report:

- An average of nearly 23% of organizations backup half or less of their VMs daily
- 44% of orgs rely or some form of replication as part of their backup and/or DR strategy
- 24% of orgs have experienced data corruption or a ransomware attack
- Nearly 50% of orgs protect less than half their VMs with a DR plan, or have no DR plan
- Only 39% of orgs test their DR plans monthly or quarterly
- 25% of orgs have a zero data loss standard
- 89% of orgs have no ability to recover their critical applications within minutes

What you have here is a recipe for a recovery disaster.

A fraction of VMs are being backed up. There's a clear reliance upon replication that potentially has no recourse in the case of corruption or ransom. DR plans are tested infrequently. And, a vast majority of orgs can't recover quickly enough, despite a requirement – in some cases – for zero loss of data.

What all this means is, if you are in a similar boat to the data depicted in this report, you have a serious deficiency between the business requirements surrounding recovery, and the backup and recoverability measures in place that would make those recovery requirements a reality. Based on the data in this report, some organizations definitely have their act together. But, if your organization feels a lot more like this last bullet list, it's time to rethink how you approach protecting your virtual environment.

FIND OUT MORE AT:
WWW.ILAND.COM

Nick Cavalancia is founder & chief techvangelist at Techvangelism. Nick has more than 20 years of enterprise IT experience, and is an accomplished consultant, speaker, trainer, writer, and columnist. He has several certifications including MCSE, MCT, Master CNE and Master CNI. He has authored, co-authored and contributed to over a dozen books.