Unitrends Recovery Checklist
White Paper

Your "Must Do" Backup & Recovery Checklist

How to get your DR plans in shape for the new year

Your "Must Do" Backup & Recovery Checklist

How to get your DR plans in shape for the new year

How to Get Your DR Plans in Shape for the New Year

A new year or a new quarter is a natural time to take stock of business continuity and recovery assets and evaluate what areas need improvement. Data is the lifeblood of all enterprises and a valuable asset that requires having efficient processes in place to ensure the business can access critical systems in a timely fashion. It's time to review critical backup and recovery plans and determine if they meet your company's specific requirements for uptime Recovery Time Objectives (RTOs) and Recovery Point Objectives (RPOs) as well as your compliance Service Level Agreements (SLAs).

Your IT organizations is likely continuing to seek new ways to leverage virtualization and cloud computing services to better meet rapidly evolving business needs. While IT focuses on transforming the data center, updating data protection and recovery strategies become increasingly critical. Often-times, existing data protection and recovery strategies prove unsuitable or are not optimized for new production, data growth, and compliance requirements. An outdated disaster recovery solution could add layers of complexity and cost and make it frustrating to meet business continuity demands.

These and other important trends are converging to create chaos and great change for businesses across all industries. If you're not already, you'll soon be challenged to recover growing data assets faster while also managing applications running across hyper-converged systems, multiple operating systems, multiple hypervisors and in the cloud. All of this is a must while also attempting to assure 24/7 access to critical applications.

The Current State of Data Recovery

In review, a TechTarget survey[1] showed that nearly one-third (30%) of organizations were using cloud storage services and almost half (48%) were planning to add more or start incorporating cloud into data protection solution. Backup (62%) is still the primary application for those using cloud storage. The fastest growing use cases were disaster recovery (35%) and archiving/long term data retention (35%) to meet compliance and regulatory requirements.

However, concerns continue to challenge cloud storage adoption, including security, with 54% saying that they still worry about how their data is handled and managed in the cloud. The expectation that all company data and applications will be available at all times, from any location and any device, continues to grow. Substantial growth in cyber-attacks, data breaches, and new viruses increase the risk of data loss and continue to be of great concern. Combine these concerns with exponential data growth, new complex applications, increasing compliance and regulatory requirements, and the need to strengthen disaster recovery (DR) resources rises significantly. Add the threat of floods, fire, major storms, and other external forces and a shadow is cast over IT resiliency that can only be safeguarded by implementing a robust DR strategy.

With growing challenges and emerging technologies, sorting out the real trends from the noise can be confusing. Take a look at the current trends in DR and backup that may help reduce the stress of making the right decision to provide

1 http://searchcloudstorage.techtarget.com/feature/One-third-of-shops-using-cloud-storage-more-buying-in

cost-efficient disaster recovery and business continuity in the coming year.

Backup and DR Trends

Variations in cloud backup and recovery solutions will continue to grow.

There are wide variations in cloud options so IT leaders must do their homework before deciding which solutions are the best fit for their company's environment. At first glance, it may appear that the obvious solution is to simply move copies of backups to a hyperscale cloud vendor. With the cloud storage cost per GB as low as one cent per month with some hyper-scale cloud providers, that may seem like the most affordable option. For some companies, it's the right choice. But IT leaders must look at the total cost of each solution, including network egress fees, data retrieval costs, and the compounding storage of data for long-term retention. IT professionals may be shocked by the total costs for these low-cost options over a period of a few years. Additional factors such as the ability to get the data back quickly may not be so obvious and can be just as costly. There are also wide variations in the SLAs that the various solutions provide for uptime, guaranteed recovery times, etc.

There is also diverse functionality between vendors. Many cloud recovery solutions have simply taken existing backup techniques and bolted on a cloud extension. In other words, they have created a mechanism to move copies of backups to offsite cloud storage that previously would have been stored locally, or put on tapes and shipped to a vault. While this solves some problems by automating

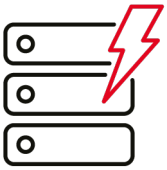

formerly manual processes of moving data offsite for basic archiving, those solutions may not provide true DR or operation in the cloud in the event of a full primary site failure. Therefore, due diligence is required when comparing options.

WAN bandwidth concerns persist.

Storage growth continues unabated. There is more data than ever to back up. However, while wide area network (WAN) bandwidth continues to grow, it remains a tremendous challenge moving data in hybrid cloud continuity architectures.

As a result, it's extremely important to use granular continuity approaches that can mix local archiving with advanced WAN acceleration techniques. More than ever, what matters is the backup and data retention solution's ability to efficiently deduplicate, compress, and otherwise optimize the amount of backup data that needs to be sent over the WAN.

There is an alternative method to get the initial set of data into the cloud without having to wait days or weeks to transfer many TBs over a WAN. This method is seeding. Seeding uses physical disks and overnight shipping to quickly create the first dataset in the cloud. Media is sent to the cloud provider to "seed" the initial full set of data and avoid the WAN challenge. However, not all vendors offer a seeding option.

Seeding can also be done in reverse, and this option may be even more important than the initial seeding of your data to get started. "Reverse Seeding," a.k.a. a Data Shipment Service Level Agreement, allows companies to get data back within 24 hours in the event of a disaster. If the customer has a disaster and loses all or a large amount of their data, the cloud vendor places their data onto disks or a new backup appliance and ships it to the customer. A data shipment SLA can be the difference in having a RTO of hours vs. weeks for large amounts of data.

Security continues to be a concern.

Consistently, security is the top concern of IT regarding the use of cloud technologies, and this trend is likely to continue. Despite the fact that most cloud facilities operate with far more extensive security measures that most enterprise data centers, there is natural concern among some IT and business leaders in moving data to anywhere outside of their control. In some cases, there may even be a regulatory requirement against it.

To get greater confidence, IT professionals must ask cloud vendors about SSAE 16 and SOC compliance certifications. SSAE 16 effectively replaced SAS 70 in 2011 and is useful in validating security and other financial reporting requirements. Additionally, data must be encrypted in flight and at rest on the cloud, preferably with AES 256-bit high grade encryption.

DR will continue to evolve to keep up with business demands.
IT strategies are changing as new software and cloud services are adopted to meet business goals. In this environment, it's not just about storing data but using that data for multiple workloads and environments. Therefore, we will see convergence fuel the modernization of backup in the coming year. This will involve streamlining how data is replicated and protected, and these new cloud-oriented architectures will encompass services and service-oriented architecture implementations. Whereas people once bought services and storage separately, there will be an expectation of convergence where a single application performs multiple functions including data protection, backup copy for long term retention, deduplication, replication, and security in-flight and at-rest. Cloud workloads and functions such as DR and testing instances will also converge. Alongside this, the management side will get joined up as well—the "single pane of glass" approach will allow IT to manage and view data from multiple sources in a single display.

A move beyond data protection to system and workload availability.
Once copies of backup data are in the cloud, numerous other possibilities open up. IT leaders can begin to think differently about data protection and start thinking about data and system availability. The dynamic nature of cloud resource provisioning and the mobility of data and resources within the cloud opens up possibilities that would have been far beyond the reach of most companies just a few years ago. Data stored in the cloud can be used to rapidly spin-up critical workloads. System availability and true DR can be done with a single click of a button or a phone call to declare a disaster.

Disaster Recovery-as-a-Service will continue to grow.
As cloud-based infrastructure continues to become the de facto standard for businesses, service offerings continue to grow in popularity and market share. Disaster Recovery-as-a-Service (DRaaS) will be a game changer in 2017 as it will enable businesses to meet and exceed customer and employee expectations around availability. As businesses place IT and availability at the center of their operations, expect to see SLAs with guaranteed backup and recovery times becoming the standard for the modern enterprise.

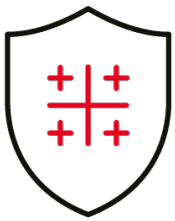

RTOs and RPOs for mission critical environments will continue to shrink.
IT professionals under pressure to maintain business continuity are increasingly working to meet more stringent RTOs and RPOs for mission critical systems such as databases, email, and business-specific data. Converged data protection solutions allow for the incorporation of higher performance recovery within existing data recovery infrastructure, therefore centralizing and reducing management complexity.

Increased demand for automated testing and validation of cloud backup and DR.
Backup and DR environments are important. However, what matters is whether the company can recover data and spin up critical systems when needed. Unfortunately, for too many companies, there are frequent doubts about the ability to recover. Today most DR environments are typically only tested every few months or once a year, if at all. These manual tests can be expensive and highly disruptive. True recovery assurance solution are now available that can solve these challenges. With recovery assurance, IT no longer has to wonder if the backups are viable and the applications will work properly in the DR environment when required. Recovery assurance ends the DR testing challenges that many enterprises face. Recovery assurance solutions automate the testing of backup and applications in the DR environment. By automating testing, IT professionals can test and certify, at the application level, that critical systems are functioning properly, providing IT with absolute confidence in the environments.

Compliance concerns become more critical.
Enterprises are increasingly getting hit with new compliance requirements. For companies in most industries, compliance regulations require reliable and provable business continuity plans, including the implementation of offsite DR protection processes. As a result, DR is no longer an option, but a necessity so that the organization is well prepared to pass its next compliance audit. Looking ahead, DR in the cloud will become a key initiative

for most organizations looking to assure a compliant enterprise as it addresses the need for business continuity without heavy investments in additional physical resources or the need to build an offsite data center. Hybrid cloud backup and recovery solutions are emerging as a key solution to enable business continuity 24x7x365. The mix of local backup solutions, combined with copies of backups and archive in the cloud, plus DRaaS and recovery assurance can empower a company with absolute confidence to withstand almost any downtime event or disaster. To determine the most appropriate plan and course of action, we first need to understand the real cost of doing nothing in today's always on environment.

The High Cost of Not Having a DR Plan

When information and communication systems are disrupted, it's much more than an inconvenience. Digital records and legal documents can be lost, employee and customer trust can be weakened, and productivity and revenue can be threatened severely. As we have seen during events such as natural disasters, situations that shut down business-critical systems and applications for any length of time (or wipe them out completely) can have devastating direct and indirect costs to the business—costs that make it absolutely vital to have a solid DR plan. But even in the wake of some of the most severe disasters on record, it appears that many CIOs aren't preparing their companies for the next one. It's crucial for IT to prioritize customer-facing and other business-critical systems for immediate recovery and business continuity, while certain file data and backed-up data (user files and archives) can be put off until later.

The first step of disaster preparation is to assess the impact of downtime. It's imperative to understand the effect of downtime cost on the business. Unexpected IT outages can unleash a series of direct and indirect consequences both short-term and far-reaching. The dollar amount that is assigned to each hour of downtime varies widely depending upon the nature of the business, the size of the company, and the criticality of the failed IT system to primary revenue generating processes.

The following chart includes both direct productivity and revenue loss as well as indirect cost such as damage to reputation, loss of customer opportunity, and damage to the company brand. When calculating loss, each of these factors should be included in downtime cost calculations. An average estimate, according to studies and surveys performed by numerous IT analyst firms, downtime cost businesses between $90,000 and $300,000 (US dollars) for every hour of IT system downtime. Loss to large financial institutions, telecommunications, transportation, manufacturing, and energy companies can be significantly higher.

Beyond loss of revenue, productivity, and reputation, businesses today cannot function without computer access and functionality. FEMA reports[2] that most businesses that suffer catastrophic data loss or an extended IT outage go out of business within two years of the disaster.

No matter what the cause, downtime impacts more than day-to-day interactions. It can impact the integrity of databases as well as the applications that use them. Some businesses can survive some data loss, while others are dependent on electronic data interchange, or are required to

2 Protecting Your Business" June 15, 2015 https://www.fema.gov/protecting-your-businesses

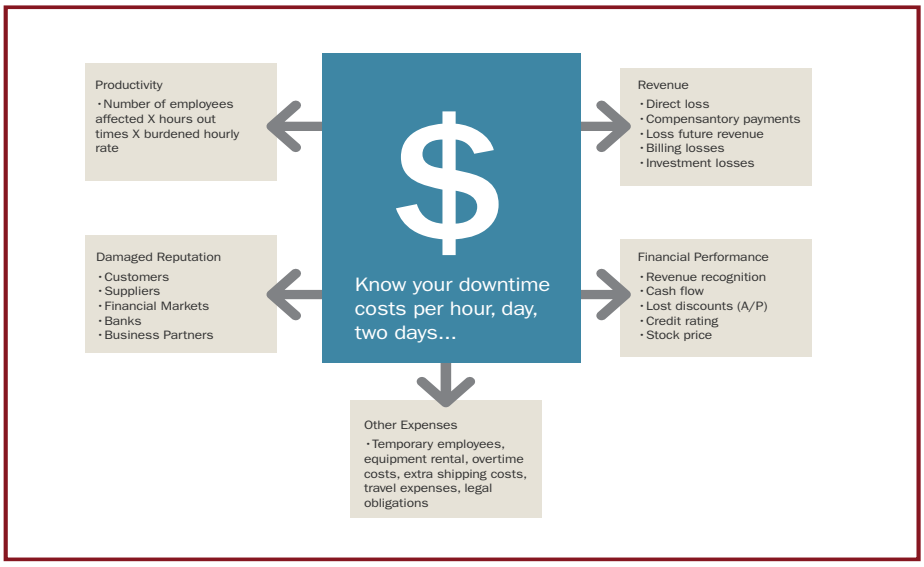

Figure 1: Know Your Downtime Costs

archive information to meet stringent audit, regulatory, and compliance requirements. Businesses that employ a global workforce that collaborate around the clock, or provide eCommerce to make sales and deliver customer service 24/7, also cannot afford data loss. While all organizations feel the impact of IT system downtime, organizations that rely on real-time data dependency will suffer most from the effects of any downtime, both immediately after the data loss and well into the future.

As an IT professional, it's inevitable that you will one day experience a system failure, outage, or complete site disaster. Your organization may already have a full or partial DR plan in place, but are you certain you certain you can recover your critical applications and data in a timely manner that meets corporate business continuity requirements and customer expectations? Now is the right time to review your priorities and ensure you are on the right path to recovering data and applications to meet your business needs.

Priorities: A Checklist for Starting on the Right Path

A disaster just isn't what it used to be. In years gone by, most companies defined a disaster as an act of nature—a hurricane, tornado, flood or fire that ravaged a building and wiped out a company's ability to conduct business. Today, with worldwide networks, Web apps, and 24/7 call centers, even a common electrical failure could spell disaster if it brings communications and online transactions to a screeching halt.

When establishing priorities, Business Continuity needs to be understood. Sometimes, even while certain IT systems are down, the business needs to continue to operate critical applications via alternative processes and or locations for reasonable periods of time.

The most important step in disaster recovery planning is to understand how an unplanned outage can affect an organization. Without the ability to determine impact of an unplanned outage in a meaningful way, it becomes very difficult to determine the type of disaster recovery strategy needed. The greater the potential impact, the more important it is to immediately restore an application or service. Include the following in your impact analysis:

- What applications and data does your operation use?
- What is your tolerance for downtime for each application?
- What is your tolerance for data loss for each application?
- What dependencies (databases, DNS, Active Directory, Web Services, etc.) are needed for each application?
- What are the regulatory and compliance requirements (internal and external) for data retention, and for how long?

Each application and service has a specific function and needs to be prioritized. Some are more time-sensitive than others in the face of an outage or disaster. Time-sensitive applications require fast recovery times based on the latest copy of data, or else significant damage to the operation can occur. Time-critical applications may cause extra work and lead to increased cost or lost productivity, but the organization can continue to operate, so these application may be of lower recovery priority in the middle of a disaster. Evaluating potential risk and the impact on each application or service is key to understanding the organization's needs and providing a differentiated level of service availability based on business priority. With more data and services on premises and in the cloud, businesses need to ensure they have strategies to backup, protect, and restore their data on all fronts. What follows are some important steps to ensure your backup and recovery plan is in shape for the coming year.

- Take an inventory. Ensure all backups are automated and completing within allotted timeframes required to meet business recovery and operational requirements. Be certain that the entire environment (all physical and all virtual) are protected by a single application and with a single management console, eliminating cost and complexity.
- Review and rank the most critical applications. Which applications are essential to keeping business operations running? How do the applications differentiate the business from others in the segment? Next, look at the application's data volatility. Those with the highest rate of churn or change (e.g., online transaction processing databases, inventory data, financial transactions, etc.) should be ranked at the top, while those whose change is more gradual (e.g., file data), should be ranked toward the bottom.

- Look at how applications interact. Applications continue to grow in complexity, including dependencies on other applications to operate correctly. Identify which applications are n-tier and have interdependencies, and then group them together to ensure data stays in sync. Backup, replicate, and retain the entire group with the same priority assigned to the most loss-sensitive application within that group. Coordinate and test the entire group together to ensure that the latest data is in sync, and that dependencies work together and still meet recovery objectives.
- Review or Define RPO and RTO. For each application or application group, set a recovery point objective (RPO) and a recovery time objective (RTO). An RPO is a measure of how much data a company can afford to lose in an outage scenario. Operationally, the RPO translates into how frequently application data must be backed up. An RTO is a measure of how long a company can go without an application. In practice, the RTO dictates how soon after a failure an application or application group needs to be able to recover.
- Review security and WAN Bandwidth. Data continues to grow unabated, leading to excessive demands on network, storage, and management resources. IT's best response is to incorporate global deduplication, which compares and eliminates data copies across all sources. This significantly reduces both local storage and WAN requirements while retaining and replicating only changed data. Additionally, security in-flight across

open WAN networks and at-rest on a cloud are critical concerns. Incorporate high grade encryption (such as AES 256) both in-flight and at-rest to ensure data is securely transmitted and protected in the cloud.

- Test. Test backups, applications, and recovery regularly, both on premises and in the cloud to ensure RPO and RTO objectives. Testing should be automated and performed locally with every new data backup and at least once per month in the cloud to ensure both applications and data can be recovered in the event of an outage or complete site disaster. Many organizations who have a DR/ BC strategy in place fail to test that strategy on at least a monthly basis. Oftentimes, this failure to test results in strategy failure in the event of a disaster or disruption. Having a plan in place is only half the battle. Testing it ensures your organization will survive in times of crisis.
- Document. Formalize the business continuity plan into a written document and then communicate it to all of the businesses' key employees and key stakeholders.
- Review. Revisit and test out the business continuity plan as often as needed to ensure the plan continues to meet business, compliance, and regulatory needs as conditions evolve.

How to Stay on Track and Minimize the Cost of Change

Every business has a different tolerance level when it comes to application downtime. Some can endure a few hours, while others will start losing revenue, and even customers, within minutes of a failure. At the end of the day, though, the total impact of a failure really comes down to the types of applications the business is running.

As we move into the future, applications continue to increase in complexity and variety, some of which are more critical to the operation of the business than others. As such, it has become increasingly important that the business gain a clear understanding of how the failure of one or more of these applications will affect their business, and how much downtime they can afford if one were to go offline for a few minutes, a few hours, or even a few days. As you review your backup and recovery strategy, long term retention needs and DR alternatives, here are some thoughts on how to stay on track while reducing cost:

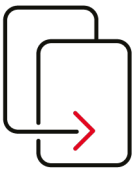

Simplify storage with deduplication. Global data deduplication can substantially decrease the need for storage and reduce WAN bandwidth requirements by only moving change data to a DR site. Daily, weekly, and monthly backups can be retained longer and restored in a more granular fashion to a particular point in time.

Move away from point solutions.
As the types and locations of the data and applications protected has grown, the number of backup tools has proliferated, with many designed to address only one particular type of data or environment—the classic example being physical and virtual servers. These point products have been deployed in silos over time, creating complexity and higher cost in management, training, operation, and storage.
Covering all the bases for backup and DR doesn't have to be complex. Whether you're running strictly physical servers and workstations, operating in any of the popular virtualization environments, or taking a hybrid approach, creating backup policies for every machine on a network should be as simple as pointing and clicking from the console. This strategy should include the process of setting up agentless and agent-based (where needed for more granularity) backups for virtual hosts, backups for physical systems, and local and cloud-based long term data retention. Recovery should be equally as easy, whether instant recovery or bare metal recovery to quickly recover a failed application or server, recover from a virus or software error, or continue operation of critical application in the cloud in the event of a total site outage.

Implement a hybrid cloud architecture.
Hybrid cloud backup and continuity solutions, in particular, have a variety of benefits that can greatly enhance the depth and breadth of protection for businesses that need to stay up and running through all sorts of downtime challenges. It's now possible for a single backup and recovery solution to protect data locally and have multiple backup copies stored in various cloud locations for business continuity in the event of a primary site outage. Local protection allows companies to have the most recent backups cached near IT for instant recovery of minor issues, such as an accidentally deleted email, server failure, or roll-back in the event of a software bug or virus. Storing multiple copies in the cloud provides geo-redundant disaster protection and failsafe long-term retention. While this was possible previously, the most powerful cloud recovery solutions today do this automatically and at a cost far cheaper than most companies can do it themselves.

Reduce cost with a unified management console.
The IT user experience of backup and recovery solutions is not strictly features-based—it is workflow-based. IT users are scheduling, adding instances, monitoring, troubleshooting, replicating, coordinating archive copies, and creating and monitoring DR policies and processes for thousands of machines, instances, and applications. Their activities are action-oriented. They are less concerned with the behind-the-scenes functionality and more concerned with whether or not the backup and recovery solution is helping them be more productive in achieving their main goal—protecting corporate data and ensuring business continuity in the event of an outage. Having a unified management console application that is consistent across all functions makes it possible to quickly identify risk and problems. An effective console displays status information about individual physical and virtual machines as well as reports job success and failure in real-time. This gives users the power to quickly identify risks that may exist across an organization's entire infrastructure, local and remote. With a single pane-of-glass management console, business users can get insight from anywhere in real-time, with a high-level of granularity. This enables planning, execution, prediction, and simulation, and ultimately enables informed decisions that can be made on the fly for faster business impact.

Prioritize disaster recovery applications
One of the primary challenges with disaster recovery is dealing with its inherent interdependencies. The coordination required between the various functional areas of IT requires an end-to-end perspective—disaster recovery is only as effective as its weakest link. It does no good to replicate storage or recover data if servers and applications are not available, and nothing gets done without a functioning network and the availability of necessary functions like name and directory services (e.g., DNS, LDAP, and Active Directory) as well as databases and appropriate Web front-ends.
The fast pace of a digitally-dependent world where information rules will continue throughout the coming year. Pressure will increase for IT to efficiently manage the data created,

safeguard it, and locate it quickly. Determining how best to use and exploit data will be key to business vitality. As seen throughout this paper, data will remain the crucial ingredient that organizations cannot do without, while managing this data throughout its life-cycle will rise to meet the demands of fast recovery, business continuity, and compliance.

Data will continue to be generated at such a velocity that it will be difficult for IT to keep pace. This puts greater pressure on IT executives to create infrastructures that are dynamic, agile and scalable. Effective strategies will require a high degree of intelligence and analytics to handle the diversity of data that must be protected.

The Unitrends hybrid cloud approach is that architecture needed to meet the growing demands of mixed use data protection, recovery, and compliance and regulatory requirements. With a convergence of backup appliances and cloud solutions, Unitrends gives you an integrated, single-point solution that protects the dynamic, high data growth environment that is becoming the norm for most businesses. Unitrends' end-to-end protection recovery assurance architecture guarantees that application specific RTO and RTO metrics are always met, even for the most complex multi-tier applications. Protection is available locally for immediate recovery and as continued operation in the cloud in the event of a primary site outage. All this is managed by a single-pane of glass management console for automating the protection of data across diverse operating systems and virtual environments. Plus, Unitrends provides the ability to manage all backup, backup copy, and recovery operations across an organization's entire infrastructure, include local and remote sites as well as public and private clouds.

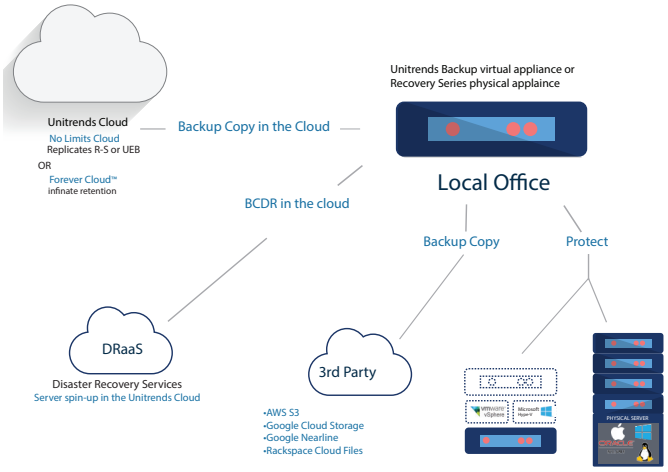

Figure 2: The Unitrends Cloud Architecture

Conclusion

New DR trends, converging architectures, and increased demands for immediate access after an outage do not need to create chaos. With proper DR planning and tuning of your data protection and recovery environment, you will be able to recover data assets faster, manage applications running across complex infrastructures, and ensure that users and customers will have 24/7 access to critical applications required to keep your business operating without damage to revenue, reputation or data. The Unitrends hybrid cloud approach converges data protection archive, instant recovery, deduplication, WAN optimized replication, P-V transformation, and management of multiple cloud services with DR all in a single solution—everything you need to complete your Backup and Recovery Checklist.

About Unitrends

Unitrends increases uptime and confidence in a world in which IT professionals must do more with less. Unitrends leverages high-availability hardware and software engineering, cloud economics, enterprise power with consumer-grade design, and customer-obsessed support to natively provide all-in-one enterprise backup and continuity. The result is a "one throat to choke" set of offerings that allow our customers to focus on their business rather than backup. Learn more at unitrends.com. Follow us on Spiceworks, LinkedIn, Facebook, or Twitter.

Need help with your backup and recovery checklist this year? Contact us.