

OpenStack Backup and Recovery

Data loss and data corruption happen. Hardware malfunctions during firmware updates, someone fat-fingers a configuration, your server room was dampened by a leaky pipe. The list goes on. No matter the reason, without a backup solution, you're risking losing your data forever.

THOUSANDS OF OPENSTACK CLOUDS LACK DATA PROTECTION

The OpenStack market is alive and thriving. According to 451 Research, OpenStack revenue is expected to grow from \$1.8 billion in 2016 to **more than \$6 billion** by 2021. That's an astounding CAGR of 30%.¹

Although many companies have shifted their data from traditional data centers to public and private clouds like OpenStack (or are in the process), the risk of data loss is more concerning than ever. Plus, it's costly: a Juniper Research study revealed the estimated cost of data breaches and data loss will top \$8 trillion by 2022. That's a staggering figure.²

But most of the companies implementing OpenStack haven't invested significant time and resources into a disaster recovery plan. According to a survey by EMC, 51% percent of organizations do not have a disaster recovery plan prepared or process documented for their "emerging workloads" — which includes those on hybrid cloud, big data, and mobile platforms. Should trouble befall their storage environments, 71% said that they are not fully confident in their ability to recover any lost data.³

That lack of confidence equates to real money. Without a documented process and the appropriate resources, it becomes incredibly difficult to recover an environment in time to meet your customers' or company's Recovery Time Objective (RTO). It happens all the time: 79% of businesses have experienced a major IT failure within the past two years, and only 7% were sure that they could recover operations within two hours.⁴

This whitepaper outlines common data loss and data corruption scenarios for application workloads running on OpenStack clouds and presents an overview of existing backup and recovery options. Specifically, we will discuss the challenges with existing backup solutions and demonstrate how TrilioVault can help businesses reliably recover workloads with significantly improved Recovery Time Objective (RTO) and Recovery Point Objective (RPO).

¹ *OpenStack Pulse 2017*, 451 Research, October 2017.

² *Cybercrime & the Internet of Threats 2017*, Juniper Research, May 2017.

³ *Over \$1.7 Trillion Lost Per Year from Data Loss and Downtime According to Global IT Study*, EMC, December 2014.

⁴ *New Research Reveals Major Challenges with Traditional Backup and Disaster Recovery*, Dimensional Research, November 2014.

The estimated cost of data breaches and data loss is expected to top **\$8 trillion** by 2022

LEGACY WORKLOADS IN OPENSTACK CLOUDS

Companies spanning the Fortune 1000 — in industries like media, telecommunications, financial services, logistics, transportation, and more — are moving toward open architectures that are based on open source technologies and backed by vibrant development communities. Cloud-based infrastructure is now a central component to modern IT strategies, but public clouds often leave IT managers without the appropriate control; consequently, OpenStack has become a top choice for private and managed clouds.

Over the past decade, OpenStack has matured significantly and gained traction among small and large organizations alike. OpenStack was architected to be one of the first true cloud platforms with ephemeral virtual machines (VMs). However, most IT workloads today are mixed — part cloud-native and part legacy applications. IT administrators have grown comfortable running legacy applications in OpenStack environments, often without consideration of the cloud platform's key differences.

As more enterprises move their OpenStack projects into production, they continue to be challenged by an important missing component: data protection. IT teams need backup and recovery solutions to quickly restore data and applications in the event of a total outage, data corruption, data loss, or event for version control (so they can roll back during upgrades). Without it, they aren't able to bring the systems back into operation quickly enough to meet strict internal service-level agreements or corporate compliance requirements for recovery.

LACK OF CLOUD-NATIVE BACKUP AND RECOVERY

With increasingly complex and critical IT environments, companies are looking for ways to fully protect their business while providing easier, faster, and more reliable recovery. This is a huge challenge for OpenStack administrators today: there are few tools on the market, and most are legacy hardware-based or VM-based backup tools that burden the cloud with clunky agents. So companies are left choosing between these costly legacy tools or attempting a do-it-yourself solution like scripting.

The lack of a **cloud-native backup and recovery solution** continues to be one of the biggest barriers to bringing OpenStack deployments into production.

Cinder does provide API support for taking full and incremental snapshots, but they're disruptive and require the application to be taken offline in order to take the snapshot.

And scripting is no easy feat: consider a simple workload. In order to perform a regular backup of a workload using existing OpenStack APIs, one has to perform the following steps:

- Pause VM 1 and VM 2
- Detach Storage Volume 1 and Storage Volume 2 from respective VMs
- Snapshot VM 1 and VM 2 and store on Glance
- Call Cinder backup APIs to back up Storage Volume 1 and Storage Volume 2 to Swift
- Keep track of these copies in an Excel sheet
- Attach Storage Volume 1 and Volume 2 back to VM 1 and VM 2
- Resume VM 1 and VM 2

These manual processes consume significant time and resources. While manual efforts may successfully back up the data, it will be nearly impossible to recover it when needed.

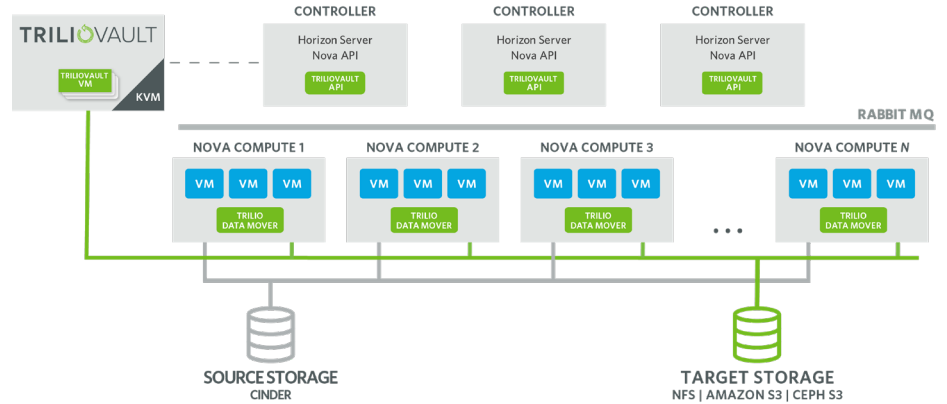
Legacy backup solutions, by contrast, were built for the old world of bare metal servers and struggle to keep pace with the demands of OpenStack. By force-fitting an outdated approach onto a modern cloud, companies must settle for disruptive, agent-based solutions that lack self-service, multi-tenancy, and store backup data in proprietary formats that require their tools (and licenses!) to unlock. Often, normal day-to-day changes in the cloud environment require extensive, ongoing manual intervention just to keep records current and functional. It quickly becomes impossible to quickly and reliably restore your cloud workloads to their last-best-known state.

TRILIOVAULT: THE ONLY OPENSTACK-NATIVE BACKUP SOLUTION

Ideally, your backup solution would act like any other OpenStack service that a tenant consumes, providing them with self-service capabilities to back up and recover workloads according to your company's policies. Just as importantly, the backup process must not disrupt running workloads, respecting required availability and performance.

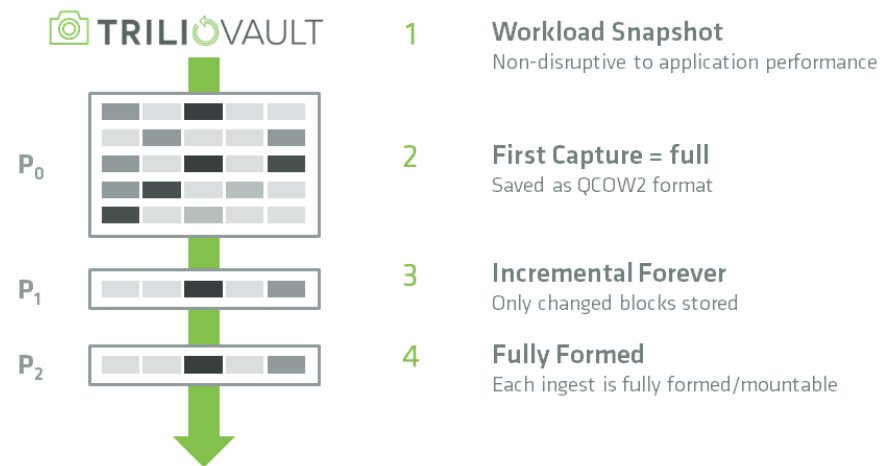
Only TrilioVault makes this possible. TrilioVault is agentless, API-driven software that natively integrates into your OpenStack cloud. Plus, it's entirely non-disruptive, capturing only changed blocks and effectively alleviating burdens on the backup storage appliances.

TrilioVault's OpenStack-Native Architecture



TrilioVault is completely agentless, working behind-the-scenes using native OpenStack APIs and Trilio Data Movers as part of the compute node — so your backup and recovery operations can execute without interruption.

TrilioVault’s forever-scalable OpenStack backup and recovery platform can recover any point-in-time with a single click using the Horizon dashboard or leveraging Trilio’s rich set of APIs. These snapshots help protect application workloads from data corruption and data loss through seamless workload recovery in a single click.



Plus, unlike legacy vendors that take a snapshot of the application data running on a single compute node alone, Trilio takes a non-disruptive, point-in-time snapshot of the entire workload. That snapshot can be scheduled or on-demand, and consists of the compute resources, network configurations, and storage data as a whole. The benefits are a faster and reliable recovery, easier migration of a workload between cloud platforms and simplified virtual cloning of the workload in its entirety.

THE DIFFERENCE OF TRILIOVAULT'S CLOUD-NATIVE APPROACH



NON-DISRUPTIVE WORKLOAD CAPTURES. Trilio is non-disruptive by design, both during deployment and execution of software. TrilioVault deployment is performed live with no impact to cloud operations. Tenants are unaware of the process and can start using TrilioVault as soon as it is deployed.

TrilioVault is also completely agentless, using native OpenStack APIs and Trilio Data Movers as part of the compute node — so your backup and recovery operations can execute without additional servers or resources. These backups can be stored in a destination of your choosing through NFS, Amazon S3, and Ceph S3.



ONE-CLICK WORKLOAD RECOVERY. Any backup copy, irrespective of its complexity, can be restored with a single click. This one-click feature evaluates the target platform and restores the copy, once the target platform passes the validation successfully.

On the other hand, Trilio's Selective Restore feature provides enormous flexibility, allowing users to restore to a new geolocation, new public or private cloud, or simply a different environment. Tenants simply choose the target platform and TrilioVault fills in the rest, regardless of backup image resources, hypervisor flavors, availability zones, networks, storage volumes, etc. Based on policy, a tenant can back up a workload (scheduled) and replicate that data to an offsite destination. This provides a copy to restore a workload in case of an outage at one of the geographical locations.

Full restores are not required, though. Trilio enables you to recover not only the entire workload but also individual files. Individual files can be from a point-in-time snapshot via an easy-to-use file browser. Trilio thus provides end-to-end recovery, all the way from workload to individual virtual machine to individual file, providing flexibility to the end user.



MULTI-TENANT, SELF-SERVICE WORKFLOW. Trilio empowers backup administrators, developers, and testers to create a copy of an application workload to restore an application in the event of an outage and for test and development purposes via an easy-to-use, intuitive GUI. With Trilio, an administrator can test and restore an application during an outage rather than pulling a team of administrators to orchestrate the recovery process. Similarly, test and development teams can create a copy, refresh, or retire an application — without waiting for lengthy approval cycles.

The backup and recovery solution is also integrated with the OpenStack Horizon dashboard, so administrators and tenants can use the familiar, self-service portal to schedule and take on-demand backups of their workloads.

MEET COMPLIANCE REQUIREMENTS WITH TRILIOVAULT

Companies deploying critical workloads on OpenStack know the importance of backup and recovery and the ability to meet SLA and compliance requirements. With Trilio's technology, tenants and backup administrators can recover the entire workload, irrespective of the size and the complexity of your deployment, without requiring a team of administrators to orchestrate the recovery process — accelerating RTO. Your organization can also significantly improve your RPO with Trilio's non-disruptive and efficient backups running in the background and reducing your backup window.

For these critical workloads, Trilio provides a comprehensive, multi-tenant backup and recovery solution that is non-disruptive with efficient backups. With TrilioVault, you can reliably recover workloads from data corruption or data loss scenarios, providing organizations significantly improved RTO and RPO.

For more information about how Trilio can help you, visit www.trilio.io today.

© Trilio Data 2018. All rights reserved. This document or any portion thereof may not be reproduced or used in any manner whatsoever without the express written permission of the publisher.

About Trilio

Trilio is a leader in data protection for OpenStack and KVM environments, and the only provider of OpenStack-native backup and recovery solutions. Since 2013, Trilio has been on a mission to give tenants more control over their ever-changing, growing, complex, and scalable cloud-based architectures. Today, Trilio is trusted by businesses all around the world to protect their clouds in a way that's easily recoverable, and requires little-to-no central IT administration.

