BRIEF

TRILIO

TrilioVault Cloud Native Data Protection

The Advantages of TrilioVault Over Legacy Backup Solutions

Data loss happens. 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 cloud native backup solution, you're risking losing your data forever.

Legacy data protection solutions for the old world of bare metal servers or purpose-optimized virtualized environments cannot keep up with the demands of OpenStack. By force-fitting an outdated approach onto a modern cloud, companies realizing the importance of backup capabilities must settle for disruptive, agent-based solutions that lack self-service, multi-tenancy, and only store data in proprietary formats that require their tools to unlock. Often, normal day-to-day changes in the cloud environment require extensive, ongoing manual intervention just to keep records current and functional.

You need to be able to go back to a specific point in time, and to quickly and reliably restore your entire workload to the last-best-known state. A cloud-native backup solution that has been built specifically for OpenStack clouds not only enhance performance, but also reduce time spent on management activities. This makes data easier to backup, and thus restore.

So why is TrilioVault's OpenStack-native data protection the clear winner when it comes to data protection in OpenStack?

AGENTLESS

Legacy data protection solutions rely on agents that are installed on the machines that need to be protected. These agents have to reach and be reached by their central data protection server in order for backups to occur. This agent-based approach is contrary to one of the key elements of OpenStack clouds: the usage of private networks inside tenants that use floating IPs to provide access only to the machines that really need to be accessed directly. Agents that have to be reached by a media server automatically demand that every machine with an agent has a floating IP. This raises the question: how many floating IPs does the environment have, and do you really want to create a floating IP for each tenant environment?

Agents also lock the service provider from knowing which legacy data protection solutions are in use by tenants and prevent preinstallation of the agent in the image used to create the machine. This means that the tenants' administrators have to install the agents afterwards, which is often time-consuming. Plus, once recovery is required, the agent must be reinstalled on the machines, which of course equates to more manual processes.

TrilioVault is agentless and does not rely on any agent inside the machine. Instead, it uses native OpenStack APIs and TrilioVault Data Movers as part of the compute node, which do not require additional servers or resources. Therefore, TrilioVault reduces the number of floating IPs needed for data protection to zero, while also reducing the load for the administrators to provide data protection to the machine.

MULTI-TENANT

OpenStack is multi-tenant, which gives every tenant their own space to work within, without the ability to see or access other tenant spaces. A typical cloud supports hundreds or thousands of tenants who deploy and manage their applications on an ad hoc basis. They have total autonomy (within a given role) to provision the resources they need as they need them, including VMs, network, and storage from their portal.

Backup, recovery, and management shouldn't be any different. Each tenant should be able to set the data protection policies that their applications require and manage data retention based on individual application SLAs. Rather than limiting these activities to backup experts (and watching the help desk tickets pile up in the process), OpenStack tenants should be able to log into Horizon and administer their own backup policies at their leisure.

Legacy data protection solutions don't provide multi-tenancy. To give every tenant full data protection capabilities without the ability to see backups from a different tenant, the OpenStack provider or the tenant themselves need to contain the complete data protection architecture (including the media server) inside the tenant — utilizing expensive OpenStack resources while they do so. Additionally, all of these media servers need to be managed, which is a burden on the backup administrators. This quickly becomes unmanageable.

TrilioVault allows tenants to take full control of their backup policies, dynamically testing and reusing their copies all the way to restoring production — completely on their own. We provide this multi-tenancy by connecting a data protection activity to a specific tenant, allowing tenants only to see and work with their own data protection activities, without the need of any special installation inside the tenant itself. Each tenant has the ability to create, configure, and restore backups of their own workloads without relying on a central IT administrator.

SELF-SERVICE

OpenStack provides all services to tenants as self-service, and administrators have usually no access to a tenant space. Legacy data protection solutions rely on a central backup administrator who is informed when a machine needs to be protected/restored. The backup administrator then owns and performs all of the necessary tasks to create a backup/restore job. Typically, those backup administrators have a lot on their plate, so it takes days to weeks before a task is complete. Compare this to the lifespan of the machine inside OpenStack: there is a high chance that the machine is already gone before the backup job has been created. So, the OpenStack user might have to wait until the job is created before moving the workload into production, tying up important resources and costing the company money.

Beyond that, the OpenStack administrator is rarely up-to-date on each individual workload status, and the OpenStack user usually knows best what needs to be protected and what doesn't. The user is forced to relay this information through the backup administrator as to get the needed protection rather than doing it himself or herself.

TrilioVault provides all its capabilities in an easy-to-use self-service interface, whether it's via Horizon, CLI, or OpenStack APIs. The TrilioVault dashboard is a plugin to the OpenStack Horizon dashboard to provide a management GUI. In line with the seamless OpenStack design, tenants and cloud administrators automatically view different built-in tabs, providing them with full control of the data protection of his or her tenant space, without the need for a backup administrator in the middle. This takes the burden off of the backup administrators, giving them back time that they can use to further improve the data protection service as a whole.

SCALABLE

OpenStack as a software solution is designed to use the scale-out architecture. To have more resources available, you can simply add a new server to the OpenStack environment rather than increasing a single server.

Legacy data protection solutions are designed to scale up. To protect 100 more machines, you must upgrade to the next bigger version or the protected environment will need to be divided into multiple protection islands. This means that the growth of your OpenStack environment will be limited by architecture of legacy data protection solutions. You'd also need to plan for future growth upfront, so you have the appropriate resources available.

TrilioVault is infinitely scalable with zero performance degradation. It utilizes scale-out architecture by adding service components that match the growth of the OpenStack environment that needs to be protected. This allows the environment to grow as needed, without having to consider data protection limits while planning your deployment. It further allows you to add resources as you need them, rather than planning for them upfront.

NON-DISRUPTIVE

It's common for backup software to cause infrastructure or business disruption during deployment and execution. No company today has the luxury of providing big backup windows for their applications, as any downtime usually results in a loss of revenue.

Legacy data protection solutions often require significant downtime in order to perform updates or take snapshots of the OpenStack storage volumes.

TrilioVault 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. There are no manual configuration changes and no restarts at either administrator or tenant level. Similarly, operations are executed outside of tenants' machines, so copies can be taken at any time with no impact to production. Resource consumption by TrilioVault is easily accommodated through proper sizing as part of normal operations by cloud administrators.

CLOUD NATIVE

OpenStack is designed to do everything by API calls, since each service is provided by a project that focus on a specific service, like Nova for compute and Neutron for network. This design results in a lot of metadata that is needed to define any machine running inside OpenStack.

Legacy data protection solutions are not designed to capture any information beyond the actual volume data of a machine, meaning those solutions are not able to identify which nova flavor was used to spin up the machine, to which network it was connected, and so on. As a result, the restore of any machine is a tedious task; all of the information needed about a machine inside OpenStack must be gathered and pieced together before the restore can be done, typically using spreadsheets and manual tracking.

TrilioVault was designed as a native OpenStack service to provide self-service, policy-based cloud data protection. It captures point-in-time application data from block storage (Cinder), along with complete OpenStack configuration data. This allows tenants to seamlessly recover, reuse, and relocate entire workloads with just one click. With all the needed information available, TrilioVault is then able to easily restore any backup, again utilizing the OpenStack APIs, to natively create the base machine where the volume data will be restored.

DEVOPS INTEGRATED

These days, any serious cloud implementation includes a robust DevOps process. DevOps and Infrastructure-as-Code (IAC) have became synonymous with cloud life cycle management: DevOps tools help you with configuration management and IAC lets you version control your configuration. Essentially, the life cycle management of cloud is fully automated, without ever needing manual intervention.

Re-architecting your cloud is easily one of the most challenging tasks in IT. When you choose a backup solution, look for one that adheres to your established cloud management practice. Whether you are deploying a new cloud, scaling a cloud, or upgrading the cloud, your backup solution should be managed by the same processes you implemented for your cloud. TrilioVault integrates with your existing cloud lifecycle management to automate deployment via Ansible, Puppet, Salt, and Chef.

FUTURE PROOF

Since many OpenStack backup projects are relatively immature, we've seen a lot of folks in this space cobble together a custom solution by choosing different technologies from different vendors. In most cases, they end up selecting a backup engine from one vendor, a dedupe hardware appliance from another, and backing up the incompatible solutions separately in their proprietary formats.

Unsurprisingly, this doesn't typically end favorably. It's easy enough to tuck them away on a remote server as a literal "back up plan," but if you need to recover from proprietary formats or just manage them, you will need to have valid license(s). Forever. That's astonishingly restrictive. By contrast, TrilioVault backups are all stored in an open QCOW2 format, so you are not reliant on vendor solutions to use and restore backups. TrilioVault is also storage target agnostic, so you can back up via NFS, Ceph, or S3 regardless of whether it's on-premises or in a public cloud.

WORKLOAD CAPTURE

Clouds simply aren't tidy. They tend to have thousands upon thousands of workloads spanning multiple VMs that are using a few dozen variations of network and storage configurations. Backups from legacy solutions are incomplete as they only capture files, databases, VMs, or data volumes. So when you apply this legacy backup solution to your cloud — with applications distributed among multiple VMs — you simply won't be able to capture all of the critical configuration data associated with those VMs.

TrilioVault snapshots capture entire workloads, including applications, VMs, operating system, networks, and more.

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.

CONTACT US FOR A DEMO

TRILIO

WWW.TRILIO.IO