

Canonical OpenStack and TrilioVault

About Canonical OpenStack

Canonical OpenStack is the world's leading open cloud infrastructure platform. OpenStack was created as a solution that is fast to deploy, simple to develop and manage, whilst being massively scalable and free from vendor lock-in.

OpenStack consists of a series of interrelated projects delivering the necessary components of a cloud infrastructure.

- OpenStack Compute provisions and manages large numbers of virtual machines to offer on-demand computing resources
- OpenStack Block Storage or Object Storage allow for effective scale-out storage needs to be met, with a variety of options to meet the complexity of a project
- OpenStack Networking is a pluggable, scalable and API-driven system for managing networks and IP addresses
- Shared services span all three, addressing identity and image management, and integration with other OpenStack components and external systems
- The OpenStack Dashboard provides a web interface for administrators to control all of the above

Many organizations choose to build an OpenStack cloud alongside an existing, proprietary cloud. Typical reasons include: reducing vendor lock-in, being able to scale out without restrictions on capacity or proportional increases in pricing, testing the latest cloud technologies, which are almost always available on open platforms first, among others. The exact details of whether and how your existing cloud can work alongside an OpenStack cloud will vary depending on your existing technology and requirements.

TrilioVault for Canonical OpenStack

While OpenStack modernises cloud computing, few data protection technologies have kept up. Legacy solutions use decades-old code with proprietary backup schemas and formats, don't allow the tenant to have any control, and rely on clunky agents for backup. Without a reliable data protection solution, it's difficult to deploy an OpenStack cloud to production.

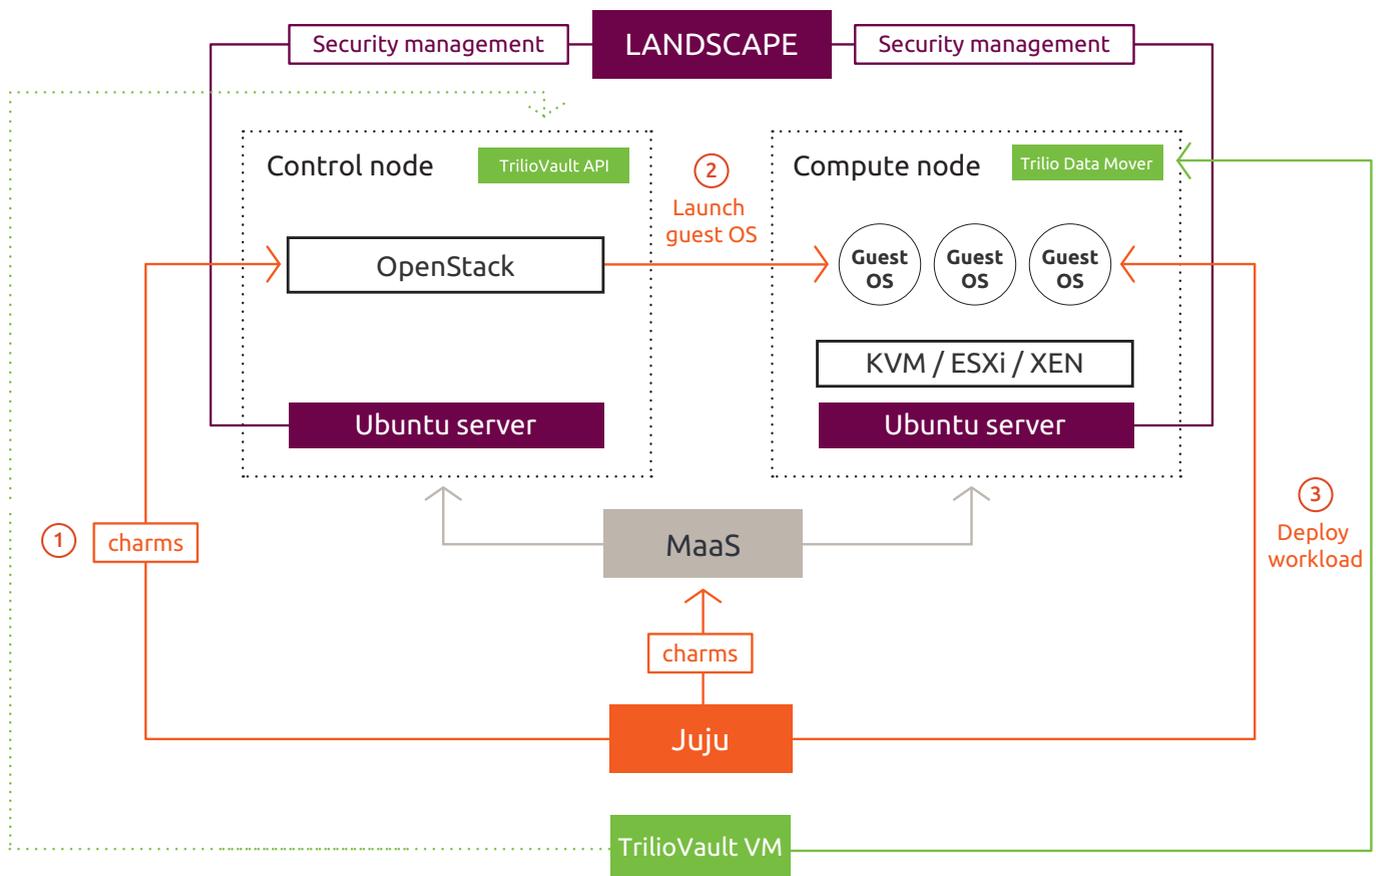
TrilioVault is the only OpenStack-native data backup and recovery solution that gives tenants and administrators the ability to restore entire workloads in one click.

TrilioVault's agentless, software-only solution provides tenant-level self-service protection and recovery of entire workloads. Architected for the cloud, TrilioVault is forever scalable with zero performance degradation.

TrilioVault helps Canonical OpenStack users protect their clouds and efficiently create, store, and manage point-in-time backups while providing for speed of recovery when required - a crucial element in the data protection continuum.

Plus, TrilioVault is agentless and non-disruptive by design, both during deployment and operation. You can even integrate it with your existing cloud lifecycle management to automate deployment via Juju Charms, Ansible, Puppet, Salt, and Chef.

With TrilioVault, businesses have complete control to backup their clouds in a way that's easily recoverable, requires little-to-no central IT administration, and reduces total cost of ownership.



TrilioVault features

- **Non-Disruptive Backup** - Capture incremental changed blocks of your workloads via a Data Mover
- **One-Click Restore** - Recover your whole environment or individual VMs to a point-in-time in place, or to other availability zones or clouds
- **Self-Service Management** - Tenants or administrators can restore individual items, VMs, or workloads on-demand

Contact us

If you are interested in building an OpenStack cloud with native data protection, please get in touch and our experts will advise you at ubuntu.com/cloud/contact-us.

You can also visit www.trilio.io or contact info@trilio.io to learn more.

With TrilioVault, you can

- Schedule automated incremental backups on a pre-defined or on-demand basis
- Easily test backups prior to recovery (stored in an open QCOW2 format)
- Quickly recover whole cloud environments in the event of a disaster
- Selectively restore virtual machines to target networks, availability zones, regions, and clouds