

Red Hat Virtualization Backup with Data Protector for Cloud Workloads

Data Protector for Cloud workloads supports backup and recovery for many hypervisors including Red Hat Virtualization. This enterprise virtualization platform supports key virtualization workloads including resource-intensive and critical applications. Built on Red Hat Enterprise Linux and KVM and fully supported by Red Hat.

Multiple Agentless Backup Strategies

- Full and incremental backups
- Tag and Name-based VM-to-policy auto-assignment
- Quiesced backup (File system freeze)
- VM disk exclusion
- Check available space for snapshot
- Snapshot management
- Power on VM after restore
- Individual file restore
- Instant restore

Virtualization is technology that **lets you create useful IT services using resources that are traditionally bound to hardware**. It allows you to use a physical machine's full capacity by distributing its capabilities among many users or environments. While VMware was the first significant hypervisor to provide and be used for virtualized applications, many more have followed and are now established, stable, and flexible technologies used extensively in the industry.

Data Protector for Cloud Workloads can integrate with multiple virtualization platforms and has the flexibility to store backups locally, using object storage, on S3 platforms, and in many more locations using the backup provider connection to Data Protector Premium. In addition:

- it is easy to use existing storage instances to keep backups of VMs according to the specific retention policy
- it has an easy-to-use web interface with the option to integrate into native administration consoles or other self-service solutions
- it allows virtualization admins to use tags or VM naming conventions to automatically assign policies to the VMs based on include/exclude rules
- it has native integration with OpenStack Horizon, and ready-to-use CloudForms integration

Red Hat Virtualization

Red Hat Virtualization is an enterprise virtualization platform that supports key virtualization workloads including resource-intensive and critical applications, built on Red Hat Enterprise Linux and KVM and fully supported by Red Hat. Resources, processes, and applications can be virtualized with a stable foundation for a cloud-native and containerized future.

Red Hat Virtualization Backup and Restore

Data Protector for Cloud Workloads Backup and Recovery is a stable, agent-less, Red Hat certified solution, providing backup and recovery services for virtual machines and container environments. RHV backup performance can be improved, and recovery tests automated with Data Protector for Cloud Workloads, which can significantly save your resources, time, and money.

Full and incremental backups with Changed Block Tracking provide flexibility and save resources, while tag-based VM auto-grouping and scheduling simplify and accelerate processes. Data Protector for Cloud Workloads performs generic application backup as well as application-consistent backups with guest file-system freeze. Full restores, with instant VM power-on after restore enable fast recovery. Individual file restore provides maximum performance for the greatest flexibility and granular operation. Metadata protection is provided for both Red Hat Virtualization and oVirt.

Contact us at:
www.microfocus.com

Like what you read? Share it.



Learn more at
www.microfocus.com/products/data-protector-for-cloud-workloads

oVirt Backup Solution

oVirt is the open-source version of Red Hat Enterprise Virtualization. Based on the same KVM hypervisor used with RHV, and it is built upon several other community projects, including libvirt, Gluster, PatternFly, and Ansible.

Data Protector for Cloud Workloads provides agentless, fast, and affordable data protection services for oVirt environments, with the easiest licensing model.

Snapshots

Keep the state of the virtual machine at any given time. One can quickly back up not only data, but also the configuration of the running virtual machines. This is an excellent supplement for data backup.

Open API

All communication between 3rd party systems goes via RESTful API exposed by the Server while tasks are being performed by the Node behind the scenes. End-user is going to use only a 3rd party system to invoke and monitor the status of the tasks.

Recovery Testing

With Data Protector for Cloud Workloads, the backup process can be easily automated, and this is also possible for the recovery processes. With these simple processes, the necessary verification tests can be performed quickly, saving time and effort. This ensures backup is performing correctly and a full recovery can be executed in case of failure. Recover your data with ease:

- to local storage
- to any remote location mountable on the Data Protector for Cloud Workloads Server
- directly to the virtualization environment

You can also restore individual files using mountable backups (directly from a web interface or node). And if one needs direct access to the original drive, there is also an option to share it from your backup over iSCSI. Especially if you need to preserve original file permissions or extended attributes.

Multiple Backup Destinations

Data Protector for Cloud Workloads provides support for multiple backup targets: Any mounted file-system, object storage (S3, SWIFT, Azure BlobStorage, GCS), and all Data Protector Premium backup targets using the backup provider. Multiple backup strategies are available to suit the customer's environment, IT infrastructure, and any specific requirements for backup. The backup strategies available include:

- **Proxy VM**—providing snapshot-disk attachment, and snapshot-chains (disk image transfer from hypervisor manager)
- **SSH transfer**—crash consistent snapshot using hypervisor's API, with optional FS freeze and optional pre/post snapshot command operation
- **Changed Block Tracking**—CBT in oVirt and RHV 4.4+ simplifies, speeds up, and improves robustness by backing up only changed blocks, and avoiding temporary snapshots.