

# DBA's guide to Oracle data protection

How to integrate cloud backup with Oracle RMAN to reduce costs and complexity

## Introduction

Oracle's RDBMS has been the gold standard for managing structured data for decades, and today most major businesses rely on it for their mission-critical applications. Yet maintaining relational database integrity during a backup can be complex. It takes keeping physical parameters secure and database processes consistent as well as auditing data trails and performing risk-based validation. To reduce this complexity, Oracle introduced Recovery Manager (RMAN) as its standard tool to handle basic backup and restore functionality.

This white paper explains some fundamental backup concepts applicable to RMAN and describes how Druva works with an Oracle image copy, incremental merge, and backup set features to securely protect an Oracle database in the cloud. As an Oracle Backup Solutions Program (BSP) partner, Druva also gives additional control of data protection to your backup admins and teams while it provides:

- The most efficient protection for Oracle databases both on-premises and in the cloud
- On-demand scaling and storage without complexity
- Lower TCO for backup and recovery

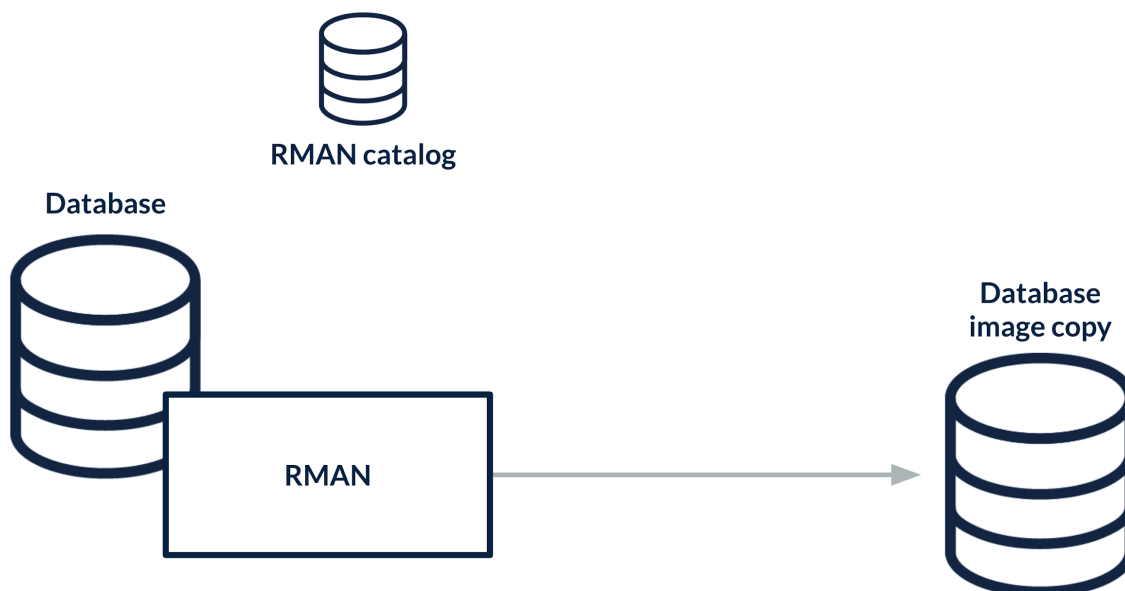
Today, Druva customers protect Oracle databases running on-premises and on major IaaS cloud providers including AWS and Oracle.

## Important Oracle data protection concepts

Oracle RMAN was launched with version 8.03, and it has been a standard Oracle component ever since. It is the intermediary between a database and a backup application that stores data on targets anywhere from physical and virtual tape, local disks, or remote clouds.

RMAN has two primary methods of compiling data for backup:

- **Backup set** – The most common way for RMAN to prepare data for backup has been to create an RMAN backup set: a logical structure that's the smallest unit of an RMAN backup. The backup set uses PL/SQL blocks and contains the data from one or more datafiles, archived redo logs, and control files or the server parameter file.
- **Image copy** – An alternate method is for RMAN to create an exact image copy of a single datafile and archived redo log file or control file. Image copies are not stored in an RMAN-specific format. They are identical to the results of copying a file using operating system commands.



## Incremental merge and block change tracking

An image copy lends itself to cloning the database to an alternate node host (copy data management) or refreshing it to make copies, for example, of *Dev*, *Test*, or *Staging* servers. There are other Oracle enhancements that make switching an image copy to an active copy relatively easy. Doing this with a backup set takes more time, and is a more laborious process.

Incremental merge is an Oracle technology, introduced in version 10, that lets you take a level-0 image baseline first and then incrementally restore copies of the level-1 backups on top of the level 0 to create a rolling copy of the database. In practical terms, it provides a level-1 incremental backup in perpetuity.

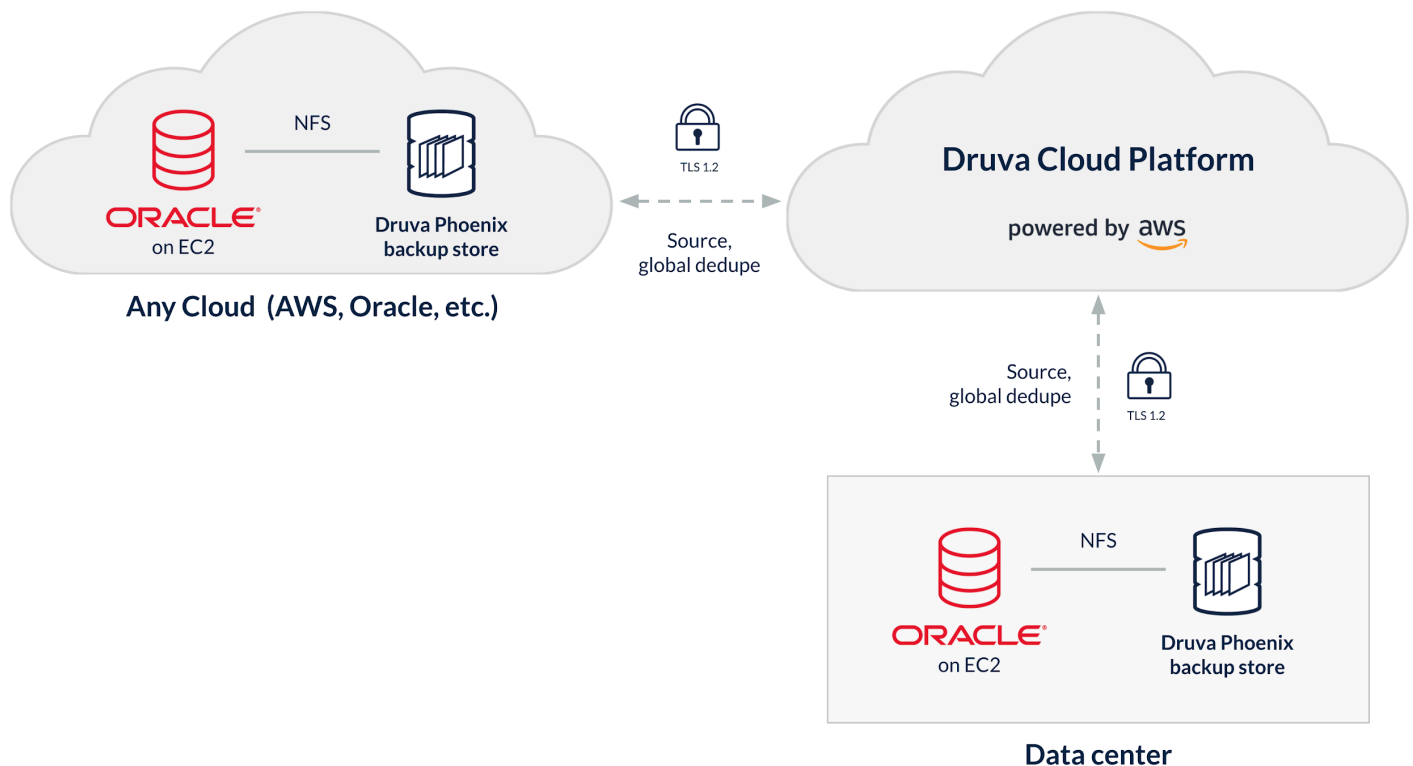
An incremental merge backup leveraging BCT consists only of changed blocks. Oracle recommends enabling block change tracking (BCT). This results in hyper-efficient, hyper-fast level-1 incrementals forever. You can back up a database with RMAN and only back up those changed blocks that improve your RPO. When using backup sets, BCT is typically not enabled and thus the RMAN binary has to scan through every block in a database to determine if it's changed or not changed and requires backup.

## Important implications for database and backup admins

Choosing to use backup sets or image copies affects database and backup admins differently. With backup sets, database admins must configure RMAN if they want to control backups and restores. This can be complex, particularly in UNIX environments. Backup sets are smaller, which is an attractive feature for backup admins. However, database admins are typically more concerned about their immediate access to backups and restores than they are about storage costs. Using image copies in the Druva for Oracle Backup Store, a DBA can control backups and restores with an easier-to-use data protection solution than just RMAN.

## Druva for Oracle data protection

Seamlessly integrated with RMAN, Druva provides simple and secure cloud backup and recovery for Oracle databases located both on-premises and in the cloud. It efficiently manages backup sets and image copies as needed with alacrity.



Druva uses cloud- and client-side components to interact, manage, and control backups in your infrastructure. These components include the:

1. **Phoenix Backup Store (PBS)** – An NFS software appliance that can be deployed on a physical or virtual system on-premises or in the cloud which serves as the target for backups (and restores). A mount on the store is a shared directory that serves as a local target location for RMAN operations. The Phoenix Backup Store creates a snapshot of the backup and uploads it to the Druva cloud. The backup remains available for immediate restores or clones in an NFS target managed by the Phoenix service on the PBS.
2. **Druva Management Console** – A web-based app that provides complete visibility and management of all Phoenix Backup Stores and their mounts.

## Configuring Oracle backups with Druva

Setting up Druva is straightforward and consists of three steps:

- Deploy a Phoenix Backup Store (PBS)
- Create and map a mount
- Define RMAN backup schedules

The basic backup workflow is:

1. An admin defines a backup schedule in CronJob, OEM Cloud Control, or the Windows Task Scheduler to execute Druva-provided RMAN script templates.
2. The copy is deduplicated, compressed significantly (up to 4:1) with ZFS, and encrypted.
3. The deduplicated, compressed copy is uploaded to an air-gapped location in the Druva cloud.
4. It is compressed further and is globally deduplicated.

For restores up to seven days, database admins can access the backup copy on the local PBS per RMAN recovery window. Older backups (as old as needed for legal compliance requirements) that have been snapshotted to the Druva cloud are accessed via the Druva console and are mounted to the PBS or alternate destinations as desired.

### Deploy a Phoenix Backup Store

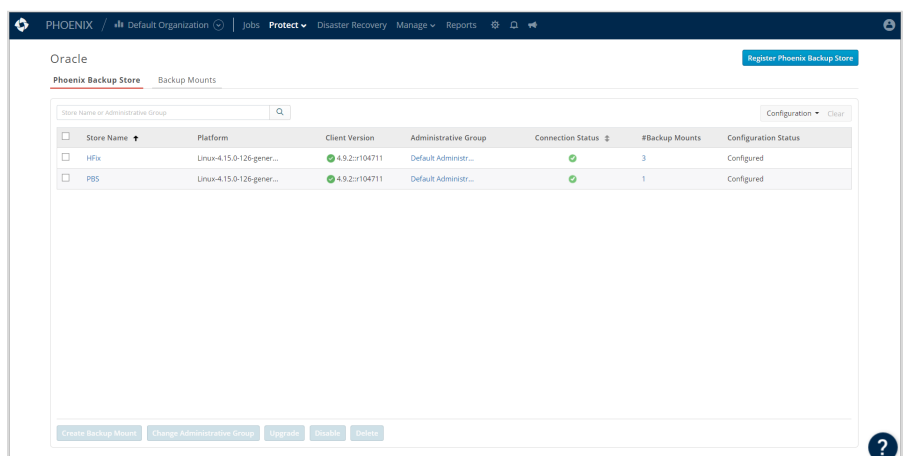
You can deploy the Phoenix Backup Store on an Ubuntu server or on an Ubuntu VM on-premises. In AWS, an EC2 instance will suffice. The Phoenix Backup Store works as an NFS server and as a ZFS server. Druva provides the following deployment packages:

- A Debian package is available to install on an Ubuntu server to work as the Phoenix Backup Store. Or, the package can be deployed on an Amazon EC2 instance (with Ubuntu as its operating system) to work as the PBS.
- An Open Virtual Appliance image is also available for download to deploy on a VMware VM to create a virtual machine as the Phoenix Backup Store by simply importing it into an Ubuntu VM allocated from the vCenter. Storage for the ZFS pool on the PBS is also allocated from a vCenter datastore.

### Create and map a mount

After the Phoenix Backup Store is deployed and activated, log in to the PBS.

1. From the top menu, select the *Organization* where you want to configure the PBS. On the *Organization* page, click *Protect* > *Oracle* on the top menu.
2. On the *Phoenix Backup Stores* tab, select your PBS, click *Create Backup Mount*, and follow the instructions in the wizard.
3. Map the PBS to the Linux or Windows server that hosts the RMAN. Detailed instructions are available in the [support section](#) of [druva.com](#).



## Configure RMAN for schedules

To configure RMAN for a scheduled backup task:

1. Download the [Druva Phoenix RMAN template scripts](#) to the Oracle Server:
  - oracle\_rman\_data\_backup.sh
  - oracle\_rman\_archivelog\_backup.sh
2. Define a CronJob schedule to accommodate RMAN scripts as follows:
  - ./oracle\_rman\_data\_backup.sh /home/test-usr/rman-log-directory 192.0.2.1 testmount
  - ./oracle\_rman\_archivelog\_backup.sh /home/test-usr/rman-log-directory 192.0.2.1 testmount

When the script runs, the RMAN creates an Oracle RMAN backup and stores it on the backup mount of the PBS. All the logs of the Oracle RMAN backup job are created automatically and stored in the NFS target the DBA provides at the time of running the script.

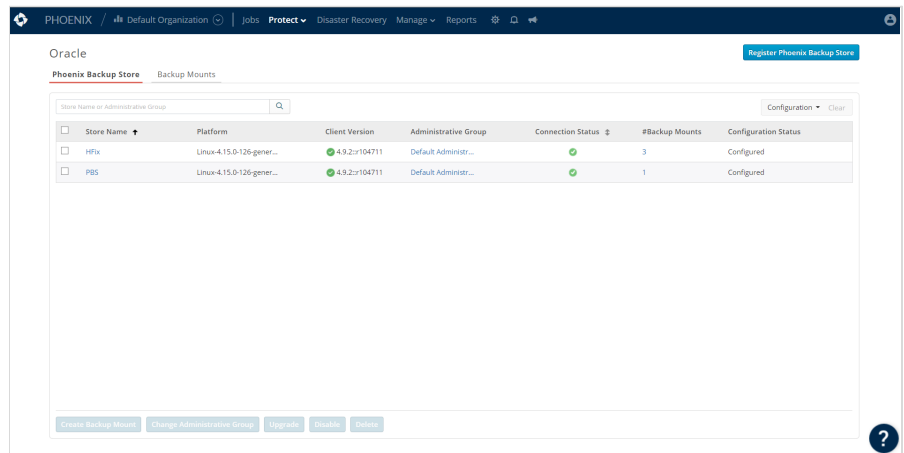
Detailed instructions are available in the [Backup and Restore Oracle Databases](#) section of Druva's documentation.

## Druva for Oracle management console

An intuitive management console lets database and backup admins control and monitor the PBS, mounts, and all backup and restore activities. The following are some of the most-used tabs.

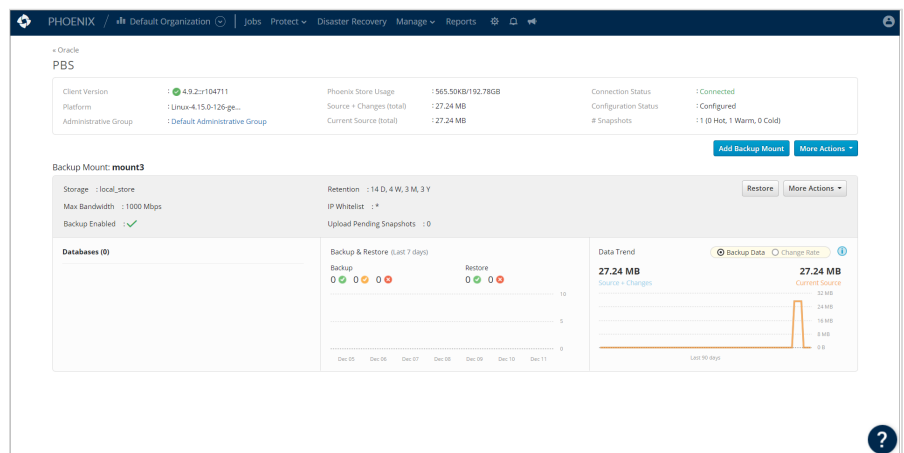
### Protect

This tab lets you manage every aspect of a PBS and its mounts.



### PBS

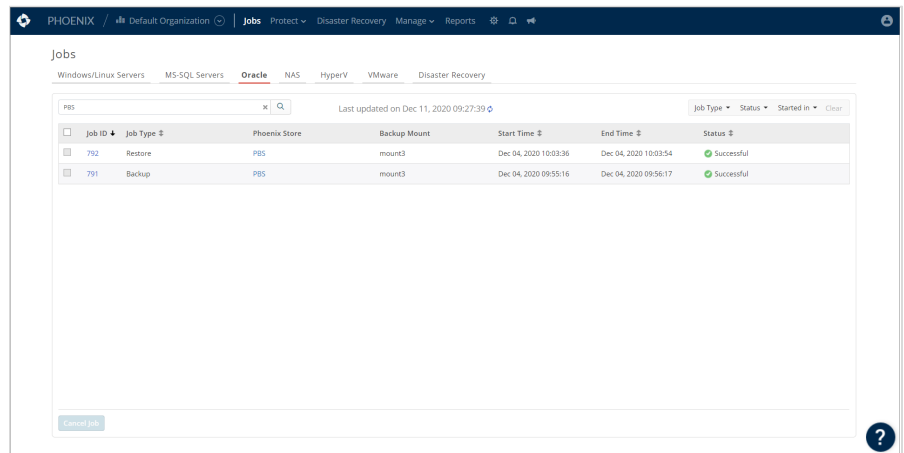
This tab shows every detail of an active PBS.





## Jobs

This tab lists every active job with time and status details.

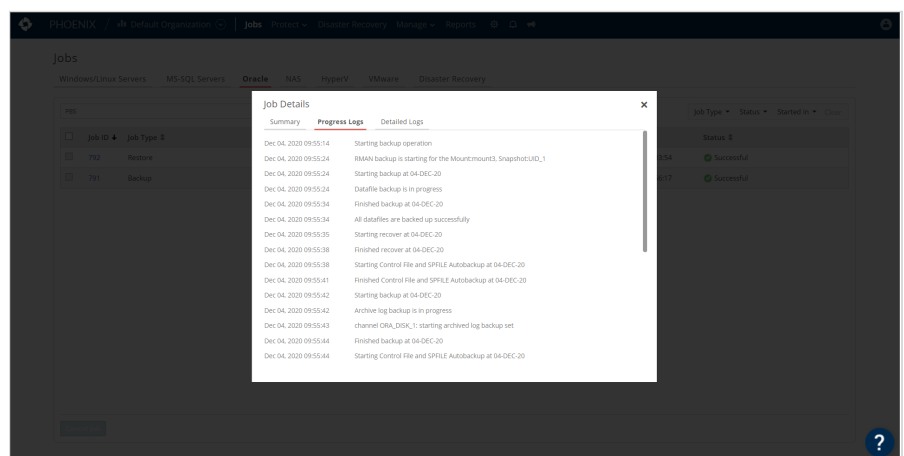


The screenshot shows the 'Jobs' tab in the Druva console. The table lists two jobs:

Job ID	Job Type	Phoenix Store	Backup Mount	Start Time	End Time	Status
792	Restore	PBS	mount3	Dec 04, 2020 10:03:36	Dec 04, 2020 10:03:54	Successful
791	Backup	PBS	mount3	Dec 04, 2020 09:55:16	Dec 04, 2020 09:56:17	Successful

## Jobs/Progress logs

This screen shows job progress details.



The screenshot shows the 'Jobs' tab with a 'Job Details' modal open. The modal displays the following progress logs:

Summary	Progress Logs	Detailed Logs
Dec 04, 2020 09:55:14	Starting backup operation	
Dec 04, 2020 09:55:24	RMAN backup is starting for the Mount:mount3, SnapshotID:1	
Dec 04, 2020 09:55:24	Starting backup at 04-DEC-20	
Dec 04, 2020 09:55:24	Datafile backup is in progress	
Dec 04, 2020 09:55:34	Finished backup at 04-DEC-20	
Dec 04, 2020 09:55:34	All datafiles are backed up successfully	
Dec 04, 2020 09:55:35	Starting recover at 04-DEC-20	
Dec 04, 2020 09:55:38	Finished recover at 04-DEC-20	
Dec 04, 2020 09:55:38	Starting Control File and SPFILE Autobackup at 04-DEC-20	
Dec 04, 2020 09:55:41	Finished Control File and SPFILE Autobackup at 04-DEC-20	
Dec 04, 2020 09:55:42	Starting backup at 04-DEC-20	
Dec 04, 2020 09:55:42	Archive log backup is in progress	
Dec 04, 2020 09:55:43	channel ORA_DISK_1: starting archived log backup set	
Dec 04, 2020 09:55:44	Finished backup at 04-DEC-20	
Dec 04, 2020 09:55:44	Starting Control File and SPFILE Autobackup at 04-DEC-20	

## Druva's Oracle direct-to-cloud solution

Druva's Oracle direct-to-cloud (DTC) solution is a novel implementation of the SBT\_TAPE library Oracle DBAs have been using for decades. Druva's innovative approach utilizes RMAN to leverage the SBT\_TAPE library, no longer using RMAN to stream data to tape. Data is now streamed directly to the highly optimized and performant Druva filesystem, part of the Druva Cloud Platform, running on AWS.

When using Druva a DBAs no longer require a disk landing zone, tape drive, or any other type of storage hardware. The data that RMAN streams is encrypted, compressed, air-gapped, and globally deduplicated by the library itself. All data streamed has its backup metadata sharded from itself, with multiple copies of the data created to preserve the data's integrity. The air gapping, compression, and encryption make backups immutable and resistant to corruption or ransomware.

### Configuring Oracle Direct-to-cloud

DTC is an agent-based methodology, meaning the DBA works with either sudo-privileges or the lead system administrator (SA) to install the agent. Once the agent is started, it auto-discovers all databases and places them in the Druva console. The DBA then uses the discovered databases to assign a mechanism of either OS for Oracle or database credential authentication. The next step in the process is to configure the database for backups. Configuring a database is a simple process, with options for choosing cloud storage and creating a backup policy if one does not yet exist.

Druva provides backup policy templates to further expedite the process, and the DBA/user has the opportunity to customize as they see fit. The DBA may set up a schedule for full, incremental, and/or archive log backups, and choose the intervals at which these take place. The backup policy also allows the DBA to customize the retention of the backups, RMAN configuration, and the settings that RMAN uses.

## Restore and recovery with Oracle direct-to-cloud

DTC also affords great latitude in restore and recovery operations. The DBA may use an automated form of recovery via the UI to restore or recover the original database on its host. Alternatively, one may specify a snapshot restore of RMAN backup sets or a point-in-time restore of the database via the RMAN backup sets. This provides the ability to leverage a set of RMAN snapshots, for example, those downloaded to an alternate node host from the Druva console. Once the snapshots are on the alternate node host, they are essentially copied to the host via download from the Druva Cloud in a restore job.

The DBA then designates the appropriate Druva restore job directory and, by invoking RMAN, clones a copy of the database on the alternate node host. This affords the DBA the option of leveraging his or her database backups for cloning production to development, test, and/or UAT. DTC offers great flexibility via this methodology as the RMAN backup sets in a snapshot format are capable of being utilized in a manner that RMAN dictates. Hence, the DBA may use any type of RMAN operation with snapshots to provide a wide variety of recovery and restoration options.

## Conclusion


Druva for Oracle brings DBAs, backup admins, and cloud teams together with a secure platform that offers the visibility, control, and automation needed to meet SLAs. It also increases predictability, the transparency of costs, and agility. It enables:

- Automated, single-step, point-in-time, and granular recoveries either from a local target or a secure snapshot in the cloud.
- Incremental merge and level-1 backups for as long a duration as needed for legal compliance.
- Reduced bandwidth and storage requirements with global, source-based deduplication.

When you choose Druva, you're getting a cloud-based data protection service for backup and recovery as well as resilience.

Druva provides all-inclusive services with no need to manage hardware or software, reducing the costs and complexities of backup infrastructure and data management.

To learn more, visit [druva.com/solutions/oracle](https://druva.com/solutions/oracle)



**druva** Sales: +1 888-248-4976 | [sales@druva.com](mailto:sales@druva.com)

Americas: +1 888-248-4976  
Europe: +44 (0) 20-3750-9440  
India: +91 (0) 20 6726-3300

Japan: +81-3-6890-8667  
Singapore: +65 3158-4985  
Australia: +61 1300-312-729

Druva® delivers Data Protection and Management for the cloud era. Druva Cloud Platform is built on AWS and offered as-a-Service; customers drive down costs by up to 50 percent by freeing themselves from the burden of unnecessary hardware, capacity planning, and software management. Druva is trusted by thousands of companies worldwide, including over 50 of the Fortune 500. Druva is a privately held company headquartered in Sunnyvale, California, and is funded by Sequoia Capital, Viking Global Investors, CDPQ, Neuberger Berman, Tenaya Capital, Riverwood Capital, and Nexus Partners. Visit [druva.com](https://druva.com) and follow us on [LinkedIn](#), [Twitter](#), and [Facebook](#).