

SOLUTION BRIEF

What's Lurking in Your Database Backup Plan? The Risks of Relying on the Dump and Sweep Method.

What does the dump and sweep method of protecting databases cost you?

Today's databases are larger, faster, and more complex than ever. Businesses rely on these increasingly powerful databases to store valuable information, making it essential to protect data against loss, corruption, or attack.

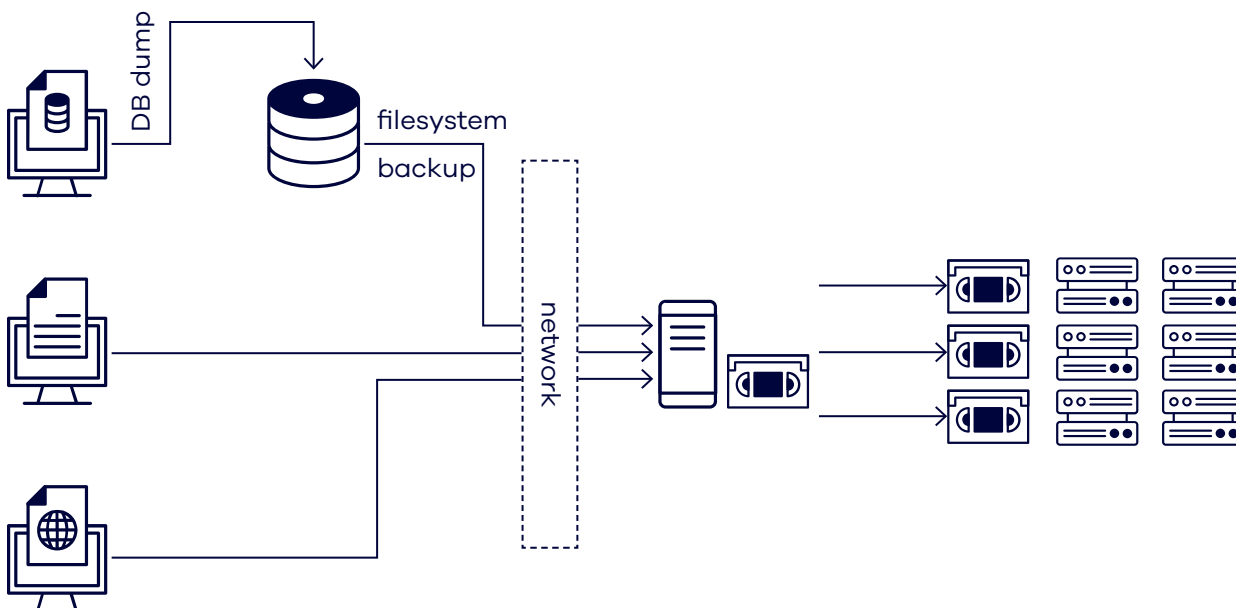
A semi-common way of database protection is the dump and sweep method that consists of two distinct and separate tasks: a database dump is triggered by Database Administrators to local storage, generating a copy of the database, usually as a single large flat file. A file-based backup agent installed on the same host then sweeps through the filesystem and writes the database dump to the backup server's protection storage. However, this inefficient method can leave your data vulnerable to attacks while producing less-than-desirable side effects, such as:

High costs—Maintaining storage for backup of all database data when exporting a live, full-size database is expensive. Not only must you pay for extra personnel to manage the extra space, but the disk landing zone is typically a Tier 1 expensive disk and much more costly than if it were intended for backups in the first place. The more backups required, the more disk space you consume, which can easily grow to 3-7x the database size. Furthermore, the double movement of data from source to target leads to higher CPU utilization and storage and network traffic, thus increasing costs.

Increased risk—This approach increases security risk due to its reliance on manual processes and the proliferation of multiple copies of data across various locations. As a result, malicious actors have more opportunities to access data, and there is no way to audit the process if something goes wrong, making it difficult to trace where and why the issue occurred.

Resource inefficiencies—Relying on the dump and sweep method is a laborious and time-consuming process. Identifying databases to back up and scheduling backups are manual tasks. Furthermore, databases require manual configuration for each instance, adding complexity. This lack of automation can increase data sprawl and costs due to longer restore times when needed.

Limited environment portability—A different solution is often required each time a database migrates to a new platform, whether it be physical, virtual, on-premises, or cloud IaaS/PaaS. This is often further compounded by the need for more administrative resources to facilitate the backups, adding even more cost in terms of human capital.



DITCH THE DUMP AND SWEEP: STREAMLINE DATABASE BACKUPS WITH COMMVAULT® CLOUD

Commvault Cloud provides database and IT admins with streamlined, easy-to-use data protection that grants the visibility they need to minimize the risk of data loss for critical applications that depend on a broad range of databases. Our solution provides other significant benefits as well:

Cost savings—Reduce costs and save time with a more efficient way to process and move data via a one-step vs. two-step process.

- Built-in deduplication and compression optimize performance and storage
- Reduce network traffic with client-side deduplication
- License metering is limited to the last full backup – logs and incremental backups are always free from the metering.*

*This is only available when we can index a backup as a database backup with the iDataAgent.

Ultimate security—Protect business-critical databases with built-in security protocols to secure data, prevent unwanted access, and drive compliance.

- Encryption of data in-flight and at-rest
- Role-based, SSO, SAML authentication controls
- WORM lock to prevent possible accidental deletion
- Ability to catalog a backup file as an application-specific data set allows for an auditable solution
- Anomaly detection and air-gapped copies for ransomware protection

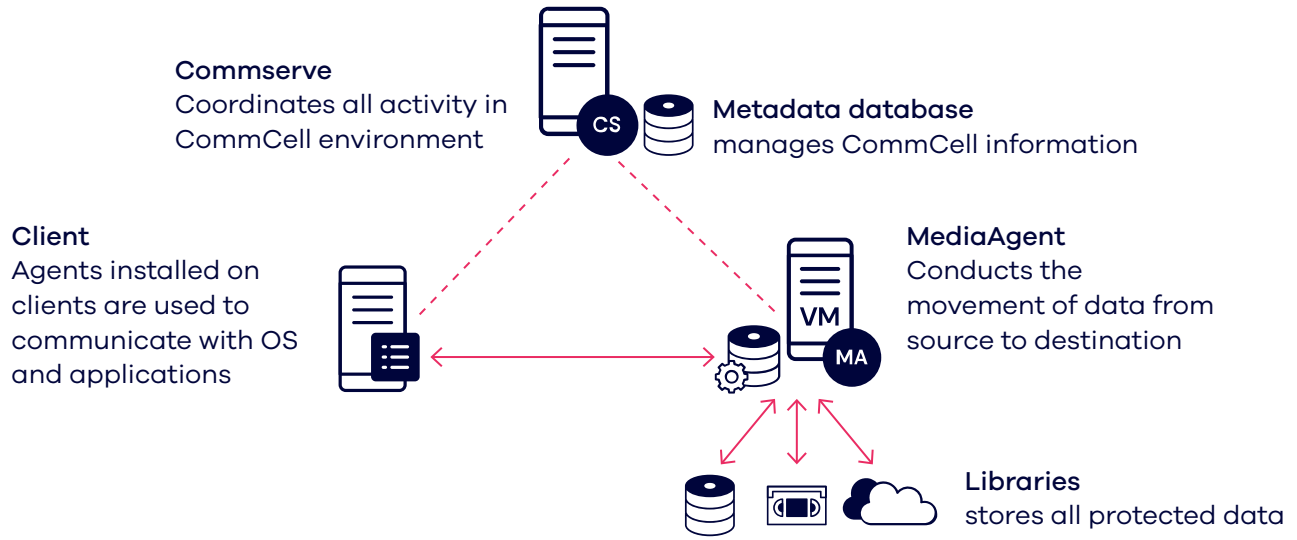
Ease of configuration—Take back time with powerful automation. Avoid multiple point products and writing scripts with a simple, integrated solution for databases and other structured data.

- API-first support provides simplified and unified control with cloud-native integration to automatically discover, protect, and manage the database copy lifecycle
- Automated log backup processes can monitor disk usage and issue backups upon reaching a desired threshold, while disk cache logging allows for RPOs down to 1 minute for most database types
- Replication support for most databases for disaster recovery, migration, and reporting
- Automated discovery of new databases

Broad workload, cloud, and multi-cloud portability support—Commvault has the largest cloud, database workload, and multi-cloud portability support compared to any backup vendor in the market today.

- Single UI supports the most popular databases, such as Oracle, Microsoft SQL Server, SAP HANA, DB2, MySQL, PostgreSQL, Sybase, Informix, Cloud PaaS DBs, and more
- Ability to define uniform RPO policies across hybrid cloud and multi-cloud environments with a blend of snap and backup
- Ability to protect instances, complete regions, or multiple regions

See the full list of [Commvault Cloud-supported databases](#).



6x reduction

in database restore times

Baptist Health reduces the time to restore a 50G database by about six times, giving users access to critical information hours faster—enhancing patient care.

30 Minutes

to complete backup for core databases

Yantai Shinho saves time and costs with automated, centralized backups.



To learn more, visit [commvault.com](https://www.commvault.com)