The Very Large Database Problem

How to Backup & Recover 30–100 TB Databases

actifio

Radically Simple

16-140 E161007

The Very Large Database Problem

How to Backup & Recover 30–100 TB Databases

We can't solve our problems with the same thinking we used when we created them. – Albert Einstein

100+ TB databases do exist, and surprisingly, there are quite a few such Very Large Databases (VLDB). Most are Oracle and some are Microsoft SQL Server based. Applications with such large databases are mission critical to the enterprise and easily attract 60% to 80% of IT budget. VLDBs of 50 TB or 100 TB incur significant compute, storage, and DB license costs. But the majority of the costs are taken up by copies of production databases that are created for use cases such as 1) Backup (onsite and offsite), 2) Replication and Disaster Recovery, 3) Application Development, QA testing, UAT, and 4) Data Warehousing and Analytics.

This document will focus specifically on the backup and recovery, and disaster recovery use cases, including a look at the requirements, prevailing approaches, and challenges. The second part of this document will look at how these challenges can be resolved.

VLDB PROTECTION EXAMPLE

For a large healthcare enterprise, Actifio worked with a solution provider partner, Initio, to design an architecture to protect a single-instance Oracle database that was over 100TB in size. With Actifio's innovative copy data virtualization technology, Initio was able to deliver a comprehensive, standardized and auditable data protection solution that saved the client $13M compared with an alternate approach comprised of EMC and Veritas NetBackup software, hardware, and extensive consulting and customization fees. In terms of ROI, Initio recorded payback in less than seven months for their client.

Backup and Recovery

Requirements

The basic requirement is that IT administrators should have backup copies of their applications data onsite and offsite in order to recover from disruption of local services, application failure, user error, virus or any other mishap, from a prior point in time.

Figure 1: **Backup Copy On-Site and Off-Site**

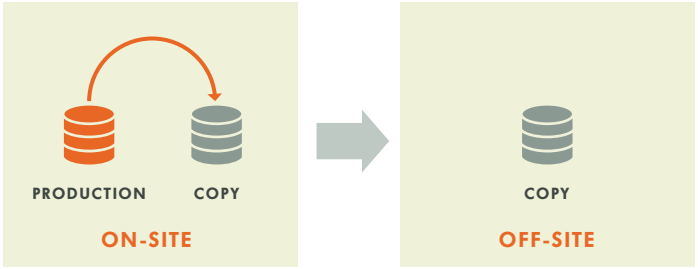

Retention of point-in-time backup images, whether of a database, virtual machine or file system, can be for only a few days, weeks, months, or even years.

Current Approach

Today's typical approach to backup has IT teams deploying media servers with backup software, along with deduplication appliances on-site and at off-site locations. Most traditional backup products do not have incremental forever backups for databases. As a result, for VLDBs, such products just can't get a full backup done on a daily basis. The compromise then is to do full backups over the weekend, and also do incremental backups daily.

Figure 2: **Weekly Full and Daily Incremental Backup Example**

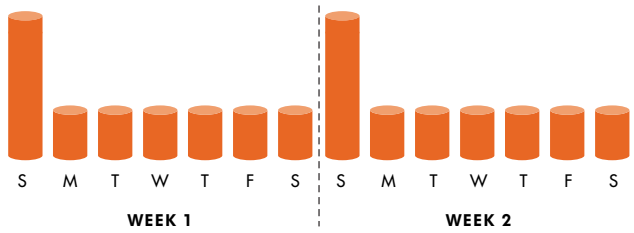

Challenges

- **Long RTO–Data conversion and reconstituting full backups**
  During restores, the entire VLDB has to be restored from wherever it resides, potentially a deduplication appliance, a tape library, or from an off-site tape storage service. Restoring a full VLDB from a dedup appliance can be a very long process for a number of reasons. First the user selects the point in time –usually an incremental backup, and the backup server will have to restore that incremental, as well as the prior full backup. Then both the full and incremental database backup images will be rehydrated from their deduplicated state so the backup software can read the data in its proprietary backup format, combine them, and then converting and writing the database back to the target host in the native application format. If tape is the storage media instead of a disk appliance, then the tapes may need to be recalled from their off-site location before the restoration process can begin–adding even more time.

Figure 3: **Slow Traditional Restore Process for VLDBs**

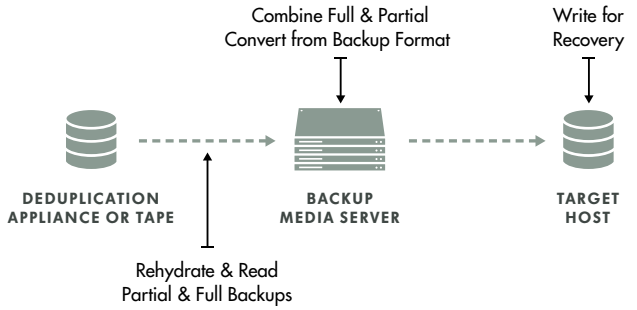

- **Production System Impact**
  During a full backup of a VLDB (every week), high storage I/O, heavy network loads, and high CPU and memory utilization combine to have huge performance impacts on the infrastructure supporting production databases. This is increasingly problematic for global enterprises whose customers and workforce use applications 24x7x365.

Costs

There are at least 6 cost components that combine to drive up the TCO of traditional backup approaches:

1. Backup software licensing and maintenance
2. Media server hardware
3. Storage for on-site backup copies
4. Bandwidth charges for replication of backups
5. Off-site media server costs, and
6. Off-site backup storage.

If there is a final step of backing up to tape and potentially using a tape management service, the costs pile up even more.

Figure 4: **Various Cost Components**

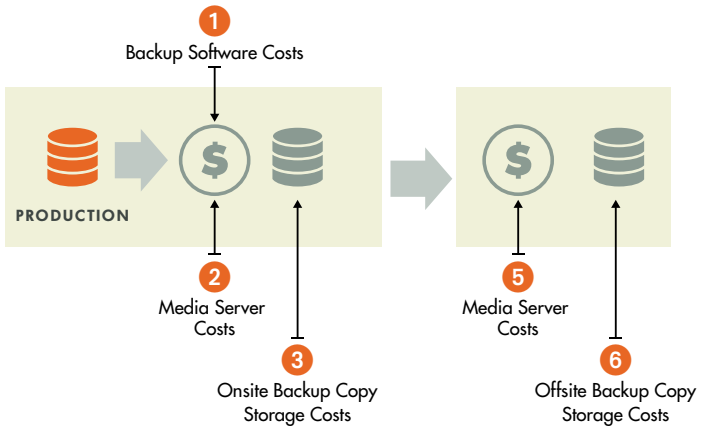

Despite these mounting costs, to solve the problem of recovering large databases quickly particularly off-site, an additional set of technologies, e.g. replication, is typically deployed for workable DR solution, and that introduces a whole new set of costs and challenges.

Disaster Recovery with Replication

Requirements

The typical requirements are, first, to recover data quickly with low RTO including off-site, which leads to the second requirement, a way to replicate data from production to remote site.

Current Approach

Some of the technologies available to help drive down recovery times include

- **Storage Snapshots:** A snapshot of storage LUNS that are being used by critical production database can be taken on a periodic basis and mounted for recovery purposes.

  SNAPSHOT CHALLENGES:

  - The snapshots happen on production storage and that impacts production performance.
  - If the production storage goes down, so do the snapshots.
  - Snapshots alone do not address off-site disaster recovery needs.
  - Snapshots exist on production arrays, which are typically costly, as their primary goal for enterprises is high-performance to meet business SLAs. Additionally, many vendors charge for snapshots, or lump them into upsell packages, driving up costs.

- **Storage Array Replication:** Storage replication products can keep track of changed blocks and replicate to a storage array from the same vendor at the remote site. Some can also create CDP-like (continuous data protection) journals that can be used to do point-in-time recoveries.
  CHALLENGES:
  - **Vendor Lock-in:** Enterprises have to use production storage (to host production database), storage replication software, and target storage at remote site - all from the same vendor. Plus, maintaining compatible firmware versions on both the source and target arrays can create even more headaches.
  - **Cloud Mobility:** Most organizations want to begin using cloud infrastructures, often test the waters with IT operational tasks, such as backups, or to use on-demand cloud resources for on-demand activities such as disaster recovery, so they don't have to maintain duplicate on-premises infrastructure. The problem is that storage replication software typically doesn't replicate to VMs running in a public cloud such as AWS, Oracle, or Azure. This often delays cloud adoption plans for organizations.

FIGURE 5: **Challenges of Storage Replication**

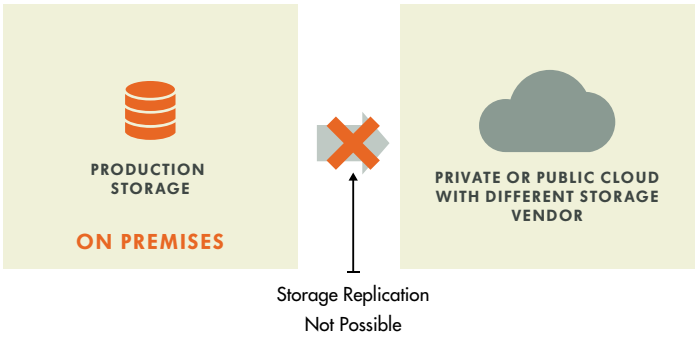

  - Often, DBAs don't trust storage level snapshots for database consistency, and as a result, they make their own copies on the side doing periodic dumps – further driving up infrastructure costs.

- **Host-based Replication or Continuous Data Protection (CDP):** CDP software is typically deployed inside hosts or hypervisors. Such products replicate changed blocks of data to remote sites and store data in its native format for faster recoveries.
  CHALLENGES:
  - CDP products maintain a journal so users can go back in time, typically a few days, to do point-in-time recoveries. CDP introduces copy-on-write operations either during replication or at the time of recoveries.

Figure 6: **Copy-On-Write Penalties of CDP**

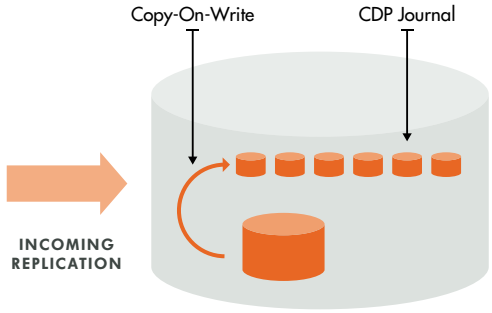

These copy-on-write operations introduce huge IOPS requirements on the target storage systems. For example, for every write operation on primary production database, the CDP software might trigger 3 to 5 times IOPS on the target storage. This causes enterprises to buy higher performance (expensive) storage that can deliver high IOPS for protection of very large databases.

  - Typically CDP products also need a lot of bandwidth. The reason: Such tools replicate every changed block constantly, meaning that changes to the same block, over and over, are replicated, which increases I/O and bandwidth costs. In contrast, point-in-time replication tools replicate the 'final' changed blocks in a given time period, and thus avoid replicating short-lived changes.

- Native database replication: Most database products come with native replication tools. They typically keep track of changed blocks and replicate to another 'live' database instance at remote DR site.
  CHALLENGES:
  - Such tools lack the ability to go back in time and recover from previous recovery points – they usually only maintain a current copy of the database. Thus corruptions in production at the application layer get replicated to DR site, rendering the DR copy useless.

Figure 7: **Data Is Useless Once Corruptions Get Replicated**

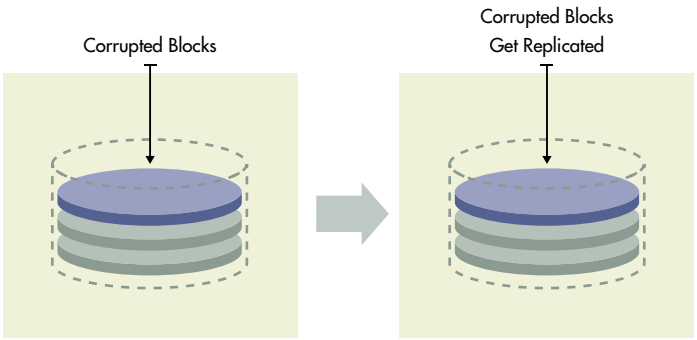

- Typically, database replication tools need a live OS and live database license running at the remote site, increasing licensing costs.

Challenges–In Summary

With so many point tools, and so much extra storage and other infrastructure servicing backup copies, it's no wonder why so much of enterprises' IT budgets are sunk into protection of these critical databases.

Fortunately, there is a better approach, 'Copy Data Virtualization', that can deliver instant recoveries for VLDBs with lower costs.

THE SOLUTION

Let's think through what an ideal solution would look like by dividing the problem into three buckets.

1. What's the best way to CAPTURE data?
2. What's the best way to MANAGE data?
3. What's the best way to RECOVER data?

What's the Best Way to Capture Data?

Once a full backup from production database is done, an ideal solution would be to never do a full backup again. After an initial ingest of production databases, the solution should do 'incremental forever' backups and each backup should be database-consistent. From these incremental backups, the solution should offer the ability to synthesize a point-in-time virtual full backup by merging in changed blocks in an automatic and transparent fashion. This concept of synthesizing a 'virtual full' out of incremental database changes is a foundation of Copy Data Virtualization technology.

Figure 8: **Incremental Forever Backups**

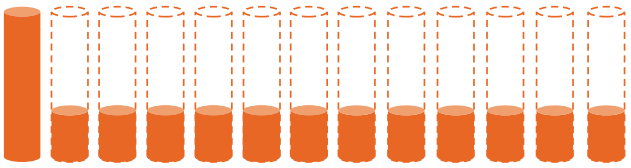

For example: If a production database is 100TB in size, the full backup would be 100TB. If the average change rate (at a block level) is 2% per day, the incremental backup size would be 2TB. Even though a 2TB incremental backup was performed, the solution should be able to recover the 100 TB database without having to manually apply all the incremental backups.

The advantages of such an incremental forever backup solution are:

1. It significantly reduces (up to 10x) the storage I/O, network traffic, and CPU and memory utilization on production database environments. This enables DBAs to deliver consistently high performance for production workloads, which is one of DBAs top two concerns, year-after-year (the other one is protection).
2. It reduces backup windows and increases the probability of success for backup jobs.
3. Such a smaller backup window opens up the possibility of many backups per day, thus reducing the RPO, even for a 100+ TB database.

What's the Best Way to Manage Data?

Outside of regulated industries with long-term data retention compliance needs, typical enterprise IT operational requirements are to retain data for a few weeks or months locally and at a DR/secondary site. The natural inclination is to deduplicate the data so that there is minimum storage consumption. However, instant recovery (in minutes) from deduplication storage systems is impossible. And even if there was a solution to emulate an instant recovery volume from deduplicated storage, the IO performance will be horrible after recovery because of the scatter-gather pattern of IO on underlying deduplicated storage.

Thus an alternate mechanism to manage data, and the best solution, for such VLDB is to:

1. Store data in its native format (instead of deduplicated and proprietary backup formats) so that the recoveries are instant,
2. Automatically synthesize and create "point-in-time virtual full" backups with sparse retention to satisfy retention needs such as keeping daily copies for a few days, weekly copies for a month, and monthly copies for a few months. As the figure below indicates, this ensures a minimum of storage consumption with the incremental forever strategy to satisfy retention needs–cost effectively.

Figure 9: **Sparse Retention with Incremental Forever Backups**

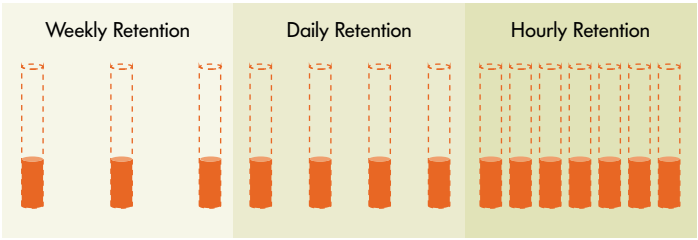

3. Compress data at rest to further reduce storage consumption,
4. Replicate data to remote clouds or sites with only changed blocks and compression to reduce bandwidth consumption. See figure below.

Figure 10: **Onsite Sparse Retention, Replication of Changed Blocks, Offsite Sparse Retention**

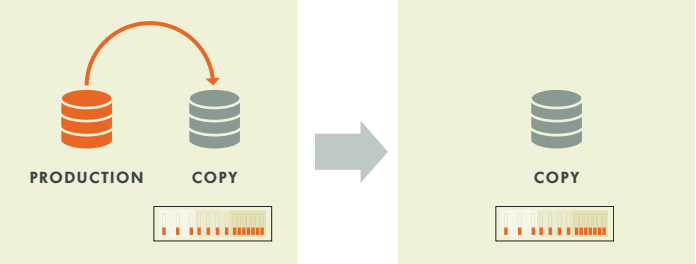

5. For retention needs that span to multiple months or years, compress and vault data to on-premises or cloud object storage. Object storage is the most scalable and cheapest form of storage at scale and is ideally suited for long-term data retention needs. Further, if data can be stored in object storage in its native application format (not a proprietary backup format), that data can be accessed instantly for use cases such as on-demand data analytics, or data warehousing in cloud. An example of such architecture is shown in the figure below.

Figure 11: **On-Demand Data Warehousing in the Cloud**

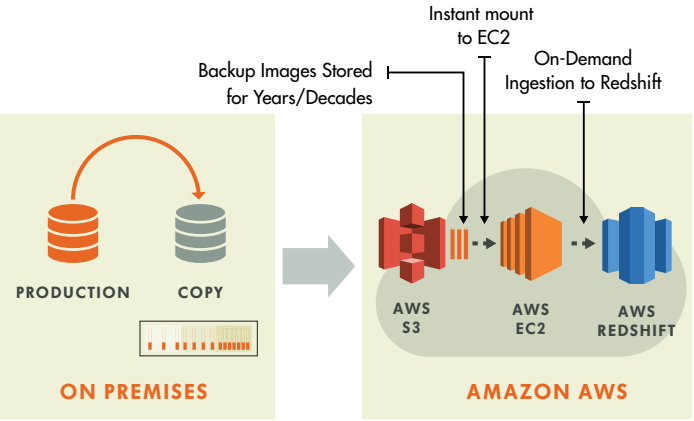

This concept of managing sparse retention of point-in-time virtual full images, and storing data in object storage with instant access capability is another important aspect of "Copy Data Virtualization" technology.

What's the Best Way to Recover Data?

To recover VLDB instantly with a low RTO, an approach that moves data from a backup repository to a server where recoveries are being done, is not a viable. For example, assuming a tape library or a dedup storage can provide a sustained 3 Gbps of throughput; it would take almost 72 hours to recover a 100TB database.

Thus, to provide instant recovery, the only other alternative (see Figure 12 below) is to dynamically provision or mount a virtual storage LUN from a backup system to the server where the database needs to be recovered for disaster recovery or a DR test. This storage LUN, provisioned over fiber channel or IP, would consist of a point-in-time version of the database, inclusive of files and logs. Thus the database instance can point to these files and bring up the database online.

Figure 12: **Instant Recovery with Virtual LUN**

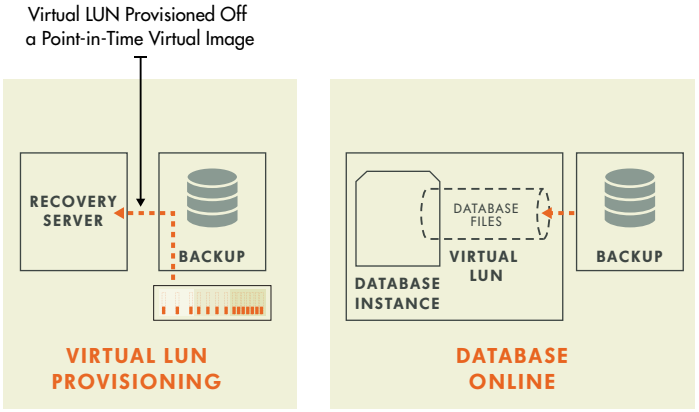

Such a solution does not have to move any data and hence delivers rapid, instant recoveries straight off the backup images. This concept of provisioning virtual storage over fiber channel or IP is the most important aspect of copy data virtualization because it gives architects the freedom to use any storage underneath the data virtualization appliance.

Once the recoveries are done, all read/write I/O from the database will happen off the storage behind the backup system. Thus the virtual LUNs provisioned are re-writable. All writes to the database will be stored on the storage behind the backup system. This is ideal for on-demand DR testing, test/dev, analytics, data warehousing - anywhere, on premises or off-premises in a private cloud from an MSP, or public cloud like Oracle cloud, AWS, or Azure. Following such an instant recovery during a real recovery process, while users and applications are using the-just recovered database from the backup systems, IT teams can redeploy production storage at the target site. Then, tools such as Oracle ASM rebalance can be invoked to migrate transactions off the backup system to production storage.

Comparison

Even though this article focuses on a 100+ TB Oracle database as an example, this applies to any multi-TB database or file system. Actifio's copy data virtualization technology enables the architectural concepts above. It's a cloud, storage, and database agnostic solution, providing enterprises freedom to pick and choose their storage and private-public-hybrid cloud vendor of choice.

A summary comparison of an "As-Is" approach vs. Actifio Copy Data Virtualization is shown in Table 1 below:

Conclusion

Copy Data Virtualization is a refreshingly simple and better approach to database protection that simplifies the data management architecture for small, medium, large, and super large 100+ TB datasets. Actifio Copy Data Virtualization is a cloud-, storage-, and database-agnostic solution that simplifies data management using a capture-manage-recover architecture, and delivers instant recovery for mission critical applications anywhere in private, public or hybrid cloud infrastructures.

Table 1: **Current Backup Approach vs. Actifio for VLDB Protection**

"AS-IS" APPROACH	ACTIFIO DATA VIRTUALIZATION	BENEFITS
Recurring full backups, at least once a week, plus partial dailies at night	Incremental forever approach	10x less Storage IO, Network IO, CPU, Memory utilization, and backup window.
24 hr RPO (Recovery Point objective)	1 hr to 24 hr RPO flexibility	Reduced TCO by eliminating other snapshot and replication products.
High RTO (Recovery Time Objective) with huge amounts of data movement	Low RTO in minutes with instant recovery	Peace of mind: reduce downtime for mission critical applications reduced from days or hours to minutes.
Many point tools: one for backup and retention, one for replication, another for snapshots and recoveries	Single platform for backup, retention, replication, and instant recovery	Better SLAs with lowest possible TCO
Barrier to private-public-hybrid cloud adoption with vendor lock-in; Continued dependency on storage and dedup appliances	Cloud-agnostic, database-agnostic, and storage-agnostic solution	Enables you to be cloud-ready right away
Exceptionally long restores when recovering point-in-time data that is many years old	Instant access to data in object storage that is many years or even decades old	Data accessible 'instantly' from anywhere eliminates the fear of not being able to restore for legal and compliance requests

About Actifio

Actifio virtualizes the data that's the lifeblood of businesses in more than 30 countries around the world. Its Virtual Data Pipeline™ technology enables businesses to manage, access, and protect their data faster, more efficiently, and more simply by decoupling data from physical storage, much the same way a hypervisor decouples compute from physical servers. For enterprise-class backup modernization, self-serve instant data access, or service provider business transformation, Actifio is the first and only enterprise class copy data virtualization platform.

For more information, on how Actifio can help you build higher quality applications faster, improve business resiliency and availability, and enable your enterprise hybrid cloud, please visit Actifio at **www.actifio.com** or contact Actifio at **info@actifio.com** or **855.886.8997**.

actifio®

Radically Simple