actifio

ORACLE Platinum Partner

SOLUTIONS

Critical Capabilities Checklist for Oracle Backup, Recovery & Cloning

As businesses have become data-driven, their need to backup, retain, and recover data on-demand has become more critical than ever. Similarly, the demand for cloning data for many internal users such as Dev, QA, UAT for testing new features, business intelligence users for analytics, and production support to accelerate root cause analysis has increased exponentially.

With data management capabilities for Oracle databases being so critical, there is an urgent need to ensure that data management solutions adhere to best practices and can be evaluated objectively by enterprises. After working with the top Fortune 2000 enterprises, we have created the following checklist to provide a comprehensive list of desired features.

Additionally, we provide a spreadsheet to assist your vendor evaluation process. The sheet allows you to assign a weighted score to each capability to help you better assess vendor offerings. The spreadsheet helps you compare and contrast vendors and decide the best features and solution for you.

Backup Features

1 **Block level incremental forever backup**

Traditional backup application performs full backups weekly, if not daily, and incremental backups on a daily basis. Unfortunately, full backups impact production database performance and backup window because they create a very high storage & network IO, CPU & memory utilization.

An ideal solution relies on block-level (not file-level) incremental forever backups thus backing up just the changed blocks and leverages RMAN to ensure database consistency is maintained with incremental copies. Even more important is the ability to synthesize a full image quickly and efficiently after each incremental backup.

2 **Parallel backup from multiple nodes in RAC**

In a RAC environment with multiple nodes, DBAs should have the flexibility to control the number of RAC nodes to participate in the backup process. Along with incremental forever backup, this feature can dramatically reduce the database performance impact and backup window.

Another benefit is transparent failover of RMAN backup from one failed node to another available RAC Node while continuing with incremental backups.

3 **Protocol flexibility**

Many environments have invested in Fibre Channel, and they prefer to leverage native Fibre Channel to perform backup, recovery & cloning, instead of impacting the IP network. Other enterprises prefer IP-based technologies like iSCSI and NFS.

It's essential for a backup solution to provide the flexibility to backup, recover, or clone using any of these underlying protocols.

4 **Storage reduction**

Many environments write RMAN dumps to staging disk and have backup software sweep the staging disk. With recurring full backups, this increases the storage consumption and the block storage TCO gets worse in public cloud.

A better approach would be for the backup software to do RMAN incremental forever backups with storage compression and/or deduplication to reduce storage consumption.

5 **DBA self-service, role-based access controls**

When eliminating RMAN dump to disk to reduce storage, it is very important for DBAs to have control over when and how often to backup, recover and clone databases.

Role-based access control is critical in ensuring security while delivering self-service capabilities via a very simple UI, API and CLI for DBAs.

6 **Flexible backup retention**

Based on an organization's data retention policies, an admin / DBA should be able to specify data retention for days, weeks, months or years at primary and DR sites.

Basic replication, CDP, or snapshot tools do not support this because they generally do not offer cost-effective long-term retention options.

7 **Store backups in cloud object storage**

For Oracle databases running on-premises or in the cloud, leveraging cost-effective cloud object storage such as AWS S3/S3-IAS, Azure Blob or Google delivers cost-effective, long-term retention for weeks, months or even years.

Another benefit is that DBAs don't have to worry about storage capacity management as cloud object storage can grow and shrink on demand.

8 **Flexibility to offload backups from Data Guard standby database**

Many environments have Active or non-Active Standby DataGuard architectures in place. Backup software should have the flexibility to backup the database from the primary or standby nodes.

9 **Wide platform support for Oracle Exadata, SuperCluster, ODA, OVMs, AWS, Azure, Google**

Enterprises evolve over time. They might start with OVMs and then buy other appliances from Oracle, or they might spin up Oracle databases in the cloud. It is vital that a backup software solution provides common capabilities across all platforms to simplify manageability and the ability to meet business SLAs.

Database Recovery Features

10 **Instant Recovery in minutes**

Traditional backup and RMAN dump approaches take a lot of time to recover a database because of the need to copy the entire backup image from a backup server or RMAN dump location into the Oracle database. This process is typically compute, I/O and network intensive which further adds to the burden.

Some solutions offer the ability to instant mount from deduplicated storage. The mount gives "access" to the RMAN dump files instantly. However, the recovery will be slow because of the "physical copy" from the mount point to the Oracle server.

With the above approaches, the larger the database, the longer the recovery time.

Thus, it's very critical to have a capability that can allow DBAs/admins to recover a 1TB or 50+TB database instantly, in just minutes, which is typically achieved by mounting the backups instantly over Fibre Channel / iSCSI / NFS

11 **High performance post-recovery using any storage**

Some solutions store Oracle backups in native format. This is a better approach than RMAN dumps because users can instantly mount a multi-TB backup image and bring the database online. However such solutions are at a significant performance disadvantage post recovery. These systems deliver low I/O performance because all the reads and writes are delivered from lower performance storage that struggles to deliver performance with very high random I/O workloads.

An ideal solution would be to store the backup in native Oracle format (instead of RMAN dump format) and allow users to pick and choose any storage for their backups based on their post-recovery performance needs. For example, use AWS EBS magnetic disk (low cost) for low-performance needs, and AWS EBS SSD (higher cost) for high-performance needs.

12 Instant Recovery in minutes from older points in time

Traditional recovery mechanisms rely on a full recovery followed by the application of incremental restores and/or transaction logs. The older the point-in-time, the longer the recovery time.

An ideal solution would have the ability to enable DBAs to recover data, in minutes, from a previous point-in-time with full automation. For example, recovering a 20 TB database from a day ago should take the smae amount of time as a copy from two weeks ago—minutes.

13 Archive log backup, purge, retention, and recovery

A backup solution should not only have the ability to do block-level incremental forever backups but also provide the capability to backup archive logs in between these incremental backups.

Thus the ideal solution would allow DBAs to recover, in minutes, to any point-in-time and roll forward archive logs with full automation.

14 Instant Recovery in minutes to ASM environments

Many DBAs prefer ASM in their production environments because of very rich storage management functionality. Thus it is critical that a backup solution provide the ability to recover databases back to ASM environments in minutes.

Beware of solutions that claim backup from ASM, but can't do instant mount and recovery to ASM environments.

15 ASM rebalance integration

Instant recovery, in minutes, using ASM format on backup storage is great. Post recovery, all transactions happen in the recovered database running on backup storage. Once the production storage is available, how does somebody get their database back to production storage?

Integration with the Oracle ASM rebalance feature ensures that data syncs seamlessly from backup storage to production storage and upon completion switches the database to run transparently off the production storage, without any downtime.

16 Recover a full machine with Oracle database in AWS/Azure/Google cloud

Many enterprises are leveraging public cloud for on-demand DR. Many Oracle databases run on on-premises physical servers or in VMs.

An ideal solution would have the ability to recover the entire on-premises physical machine in the cloud and provide instant mount and recovery of the Oracle database to deliver low RTO.

17 Instant mount and recovery even from cloud object storage

Cloud object storage offers a compelling method to store database backups for many months, years or even decades. However, it's equally important for organizations to be able to spin up a DB in the cloud straight from the cloud object storage to enable instant data access.

Compare this approach with solutions that would take many hours or days to restore data from object storage to block storage and then make the database accessible.

18 Leverage Flash Array snapshots on a production storage array

Modern Flash Arrays provide advanced snapshot capabilities. Ideally, a backup product would be able to manage native array snapshots with application consistency, capture archive logs, and capture changed blocks between the snapshots and replicate them to a remote site for DR purposes.

The backup product should also leverage storage snapshots to mount and recover Oracle databases instantly in minutes. Mounting over Fiber Channel also has the advantage of delivering high performance over a dedicated link, post recovery.

19 Leverage Flash Array snapshots on a backup storage array

#18 talks about creating and managing snapshots on a production array but creates risk since a failure of array hardware could take down production and the snapshots. A better solution would be one that could capture data in an incremental forever manner from production snapshots while storing data on a separate array. Recovery would leverage snapshots from the external array for instant access. This approach would mitigate the risk of production hardware failure while delivering the ability to mount any protected snapshot, recover instantly and deliver native flash performance post recovery.

Beware of the backup products that can manage storage snapshots on production array but not on a backup flash array.

20 Inter-Region Cloud Recovery in minutes

Enterprises running Oracle databases in one cloud region typically want to replicate to another region for disaster recovery purposes. It's important for the solution to backup in one region, enable rapid recovery in that region and replicate to another region where it can also deliver rapid recoveries.

21 Multi-Cloud Recovery in minutes

Some enterprises want to protect their Oracle databases in one cloud and replicate to another cloud provider for DR purposes. This is is similar to inter-region recovery except that the recovery option spans public clouds. (e.g., run in Amazon and have DR in Azure). Multi-cloud recovery enables users to avoid cloud vendor lock-in and hedge their risks and costs between multiple cloud vendors.

Database Cloning Features

22 **Reuse backups to provision database thin-clones in minutes**

Dev, QA and UAT users need copies of production databases for testing new features; BI users need copies for analytics; production support teams need copies for root-cause analysis. Creating physical copies for so many users consumes significant time and storage.

An ideal solution would enable users to reuse backups to provision database thin-clones to all users in minutes. Note that these clones are thin, i.e., they consume no extra storage. Thin database clones only consume storage when changes are written to them, thus keeping the storage consumption low.

23 **Provision database thin-clones in minutes in ASM format**

If production is ASM, DBAs and testers want ASM in their UAT and test environments as well. An ideal solution would provision DB thin-clones with ASM format in test environments as well.

24 **Self-service database cloning in minutes**

Dev, QA, UAT and BI users want self-service so they can provision DB clones, on demand, from any point-in-time. Moreover, different users would like to refresh the clones at different intervals. This flexibility with self-service is critical to eliminate wait cycles and inter-dependencies amongst users.

25 **Role based access control**

Self-service and the ability to instantly provision database clones are very beneficial. However, the solution must have role-based access control to ensure that clones can only be mounted on authorized test machines and accessed by authorized users.

26 **Automation for sensitive data masking**

Some environments may have sensitive data in the databases that needs to be masked before clones are provisioned to testers. The solution must have the ability to leverage any masking scripts or masking tools that DBAs already have to automate the masking process.

27 **Efficient incremental refreshes**

Once the database clones are provisioned to testers, the test database diverges from production over time. After each test cycle (in hours, days, or weeks) testers, typically, want to refresh their clones with changes from production so they can test on the latest data in the next test cycle.

The solution must offer the capability to refresh just the incremental blocks from production to test environments to accelerate the refresh process.

28 **API integration to popular DevOps, CI/CD tools**

It's not uncommon for enterprises to use Ansible, Chef, Puppet, ServiceNow, Jenkins or other similar tools to automate infrastructure and code provisioning in test environments. The solution should have sample plugins so architects can integrate via APIs to their existing toolchain.

Analyst Ratings

29 **High Analyst Ranking for Databases**

Gartner does a thorough job of compiling critical capabilities for backup and disaster recovery. Their latest report analyzes many workloads including databases. It's a good checkpoint to consider whom they rank high for Databases, Self-service, Cloud integration. You can access their document here.

30 **Customer Reviews for Databases**

Gartner also has a peer insights website where their enterprise customers have rated their experiences with various vendors. Check out which vendors get good reviews for databases.

Also, don't forget to download this spreadsheet so you can assign priority rank and capability rank to each feature while evaluating a data management solution for Oracle databases.

We hope this was helpful. If you have any questions, please do not hesitate to contact us.

About Actifio

Actifio virtualizes the data that's the lifeblood of businesses in more than 35 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.

actifio®

Radically Simple

18188B