

Original Article

Optimizing NetApp Storage Cost-Effective Strategies for Database Management

Balakrishna Boddu

Sr. Database Administrator, USA.

Received Date: 02 April 2023

Revised Date: 04 May 2023

Accepted Date: 03 June 2023

Abstract: Database storage plays a very critical role in the Real world of Critical applications, In Today's Environment Infrastructure is very costly when we procure a server with high-level storage, to reduce cost, we must focus on multiple options line NetApp which is a very cost-effective solution. NetApp provides multiple features like cloning, compression, and snapshot management of the disk where databases reside. We will explore best practices for integrating NetApp storage with various Database management systems. The proposed strategies offer organizations a practical approach to lowering storage costs without sacrificing performance or scalability, ensuring that databases operate at peak efficiency within budget constraints.

Keywords: Netapp Storage, Cost Optimization, Database Management, Deduplication, Tiering, Cloud Integration, Data Lifecycle Management.

I. INTRODUCTION AND ARCHITECTURE

ONTAP Cloud is a software-only version of Data ONTAP, which is the data management operating system from NetApp that is used on physical NetApp storage appliances. With ONTAP Cloud, the operating system has been customized to run as an Amazon Elastic Compute Cloud instance.

A. The features of ONTAP Cloud include:

- Instant, writable data clones that consume no additional storage capacity.
- Instant backup and recovery for data of all sizes. (NetApp snapshot copies, Snap Mirror technology, Snap Vault)
- Ability to use multiple protocols (CIFS, iSCSI, NFS) from the same storage system, at the same time.
- Storage efficiencies that enable you to use less underlying storage capacity for your data needs. (thin provisioning, data deduplication, data compression)
- Space-efficient, intuitive, bi-directional data transfer.
- Data-at-rest encryption using encryption keys that is stored on key managers under our control.
- Storage tiering to AWS S3 to reduce costs.




Figure 1: High-Level Architecture of NetApp ONTAP Layers



In AWS, an aggregate can contain up to 6 disks. The underlying disk type can be a single EBS storage type or a tiered storage configuration that includes an EBS performance tier and an S3 capacity tier. The maximum disk size is 16 TB, which enables an aggregate of up to 85.05 TB of usable capacity.

The total amount of storage purchased from AWS is the raw capacity. The usable capacity is less because approximately 12 to 14 percent is overhead that is reserved for ONTAP Cloud use. For example, if Cloud Manager creates a 500 GB aggregate, the usable capacity is 442.94 GB.

II. RESEARCH METHOD

NetApp has a huge number of options that can implement multiple database flavors like SQL Server, Oracle, MySQL, VMware, SAP, and Postgres with Different Cloud Solutions. Get predictable performance and the flexibility to run a variety of workloads. NetApp solutions combine simplicity of management, scalability, and data protection. The biggest benefit would be the storage efficiency of using NetApp ONTAP data management software. You are also able to use something called a Flex Clone during test and development, so you can clone your database which Will take up any capacity until you change the clone. The scope, scale, and complexity of today's data-driven world create new demands for agility in the data center.

A. Speed Backup and Increase Data Protection:

With NetApp integrated data protection solutions, you get quick, space-efficient backups, and a comprehensive set of Windows PowerShell cmdlets for easy scripting of backup/restore workflows. With federated backup, you can add multiple database instances and databases to the same federated group and take a NetApp Snapshot™ copy to back up all databases in that group at the same time. You can restore to a marker, so you can tie recovery to a business event, which is particularly useful for applications that span multiple databases across multiple database instances. You also can enable NetApp thin replication technology to protect your critical databases database data, including database Availability Groups, against disasters.

B. Provide Multitenancy and Qos:

Customers can use the functionality that ONTAP offers to provide data compartmentalization for all tenants, with each granted access only to its designated storage container. Rich support based on Windows PowerShell for both ONTAP and NetApp integrated backup empowers administrators to develop appropriate workflows that support the needs of your multitenant environment. In addition, ONTAP provides storage QoS workload management using policies that specify a throughput limit on NetApp ONTAP.

C. Achieve Intelligent Management

NetApp integrated backup helps you streamline administrative tasks such as backup, restore, cloning, and disaster recovery so that administrators can focus on more strategic tasks and business initiatives. With Snap Manager federated backups you can simplify data protection of Database applications where data is distributed across multiple databases. In addition, the combination of ONTAP and Windows PowerShell scripting enables DBAs to automate the management of LUNs and volumes.

D. Improve Manual, Slow, and Space-Inefficient Test and Development Processes

Accelerate the release and improve the quality of new applications based on databases with NetApp Flex Clone technology. With Flex Clone, you can create space-efficient copies of Databases data for development, testing, reporting, and disaster recovery testing in seconds. You can also perform clone lifecycle management and on-demand, periodic, and customized staging and user-acceptance testing services. You can do this quickly through a policy-based clone automation framework and without causing errors to other DBAs and end users

E. Speed Deployment of Business-Critical Applications with Converged Infrastructure

FlexPod-validated solutions combine storage, networking, and server components into a single, flexible architecture for business workloads such as Databases. FlexPod solutions speed the deployment of infrastructure and business-critical applications while reducing costs, complexity, and risk. With a validated design and centralized management, you can quickly and easily deploy and host multiple instances of databases, all with cooperative support from Microsoft, Cisco, and NetApp.

F. Respond Quickly to Data Growth:

With seamless scaling, you can quickly and easily scale NetApp solutions to meet your changing business needs. With thin provisioning provided by NetApp FlexVol technology, you can cost-effectively expand and reallocate storage while increasing efficiency through the deduplication of active data—all without interfering with your business operations. ONTAP allows you to

non-disruptively move your data within your storage cluster to meet the needs of SQL Server data growth and increased and changing database application workloads.

III. STORAGE DESIGN AND IMPLEMENTATION:

Architecture:

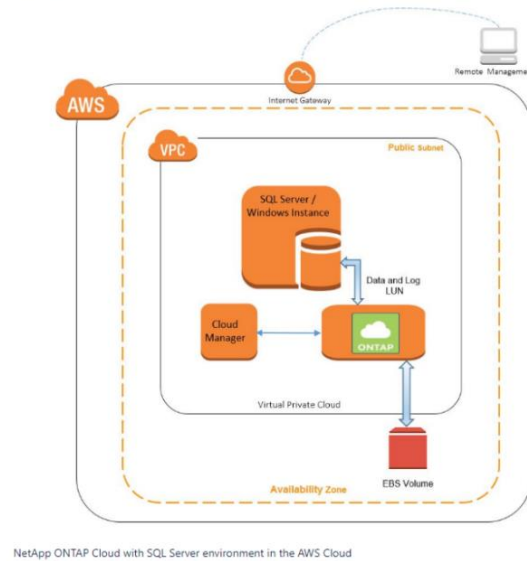


Figure 2: NetApp ONTAP Cloud with SQL Server Environment in the AWS Cloud

A. Basic Design of a NetApp in AWS

SnapCenter provides application-consistent data protection for ONTAP systems. SnapCenter allows taking Snapshot copies of applications, databases, file systems, and virtual machines. Use SnapCenter to back up to a secondary storage system within the same data center or to a storage system located in a different data center.

Example configuration can be used for SQL Server instances that require basic performance and contain multiple small databases. The database storage design has the following characteristics:

- Contains one aggregate for SQL Server instances. Microsoft SQL Server 2017 on NetApp ONTAP
- Uses a dedicated volume and LUN for the SQL Server system databases, including the tempdb database.
- Uses a dedicated LUN for each database.
- Uses a single volume for both data and log.
- Uses dedicated SMB shares for both data and log (if using SMSQL for backup).




Figure 3: Basic Design of a NetApp in AWS

IV. SHARED INSTANCE VERSUS DEDICATED INSTANCE

If an application has many schemas/stored procedures, this could potentially affect other apps that share the same SQL Server instance. Instance resources could potentially become divided/locked, which in turn causes performance issues for any other apps with databases hosted on the shared SQL Server instance. Troubleshooting performance issues can be complicated because you must figure out which instance is the root cause. This question is usually weighed against the costs of operating system and SQL Server licenses. If application performance is paramount, then a dedicated instance is highly recommended. Microsoft licenses SQL Server at the server level per core and not per instance. For this reason, database administrators are tempted to install as many SQL Server instances as the server can handle, to save on licensing costs, which can lead to major performance issues later. NetApp recommends choosing dedicated SQL Server instances whenever possible to get high performance.

V. BENEFIT OF CLOUD VOLUMES ONTAP ON NETAPP

A. Reduce costs:

Reduce your TCO by consolidating your SQL Server databases onto NetApp storage systems. Our support for FC, FCoE, iSCSI SAN, and SMB (CIFS) means you can use your existing infrastructure to cut storage outlay. Our unified architecture helps you increase storage and server utilization while simplifying data management. Particularly, NetApp's built-in support for SMB 3.0 greatly simplifies the provisioning of database clones to use in test/dev/test.

B. Provide SnapCenter advantages

NetApp SnapCenter software delivers a unified, scalable platform and plug-in suite for application-consistent data protection and clone management. IT infrastructure admins can empower application and database administrators to self-manage their own copies and data protection. They retain centralized control through powerful policies and reporting capabilities—across the data.

With SnapCenter, data protection, and clone management tasks become:

- Simple: A centralized GUI supports monitoring, notification, logging, reporting, and scheduling for all application and database plug-ins, including custom plug-ins.
- Scalable: You can add SnapCenter servers with transparency for high availability and load balancing, supporting thousands of applications.
- Empowering: Role-based access control (RBAC) provides self-service for application admins; SnapCenter infrastructure administrators retain oversight.

C. Storage:

With Cloud Volumes ONTAP, you get advanced data management that enhances service levels, saves time for IT and DevOps, and reduces storage management and associated costs. The following are the benefits of using Cloud Volumes ONTAP.

- Cost savings with storage efficiencies. Cloud Volumes ONTAP can save up to 90% on storage capacity with space-efficient technologies: data deduplication, compression, thin-cloning, and Snapshot copies that don't affect storage footprint.
- High availability. Achieve high availability with a two-node solution that supports multiple Availability Zones and enables business continuity for your critical production workloads and databases with no data loss (RPO=0) and short recovery times (RTO < 60 secs).
- Data protection and disaster recovery. Recover from data corruption or loss with efficient data Snapshot copies and disaster recovery copies which are easily configured, cost-effective, and support seamless failover, failback, restore, and recovery processes that meet minute-level SLAs.
- Hybrid and multi-cloud environments. Save time and money by using the same storage and advanced NetApp ONTAP data management software across hybrid and multi-cloud environments, including DR, HA, Dev/Test and DevOps, sandbox, reporting, data tiering, workload hosting, and training.
- Data mobility. Migrate, replicate, and synchronize your data securely, using efficient data Snapshot copies to transfer only incremental changes and recover from any point in time using NetApp's Snap Mirror.
- Cloning technology for developers. Increase DevOps agility by cloning writable volumes from Snapshot copies so data can be shared simultaneously across organizations and regions with zero capacity and performance penalties using NetApp Flex Clone.
- Interoperability. Use multi-protocol support (iSCSI and SMB) for your data and file shares and meet the demands of SQL Server workloads.
- Flexible licensing. There are multiple Cloud Volumes ONTAP solutions from hourly priced options to

longer-term subscriptions and Bring Your License (BYOL) options. • Enhanced security. In addition to security and privacy features offered by the hyper scaler, it can provide NetApp-managed encryption which gives you the ability to manage encryption keys on your premises.

VI. UP TO 50% LOWER STORAGE COSTS FOR DATABASES

FSx for ONTAP brings several key improvements to reduce database storage capacity and lower your costs:

- FSx for ONTAP storage is accessed over high-throughput network bandwidth. As such, you're able to use smaller and less-pricey Amazon EC2 instance types for your database's computing needs.
- Spend less on database licenses without sacrificing resiliency or protection. For example, with FSx for ONTAP, MS SQL Server databases can get the resiliency and protection that an Always-on Availability Group (AOAG) provides, all while using the Standard MS SQL license. Standard is much less expensive than the Enterprise licenses required to use AOAG.
- Thin cloning keeps your costs low because lightweight copies consume nearly zero storage space, instead of the duplicate storage costs that full copies incur.
- With thin cloning, new database environments are created instantly, speeding up your DevTest cycles, team productivity, and TTM.
- Storage efficiency features, including thin provisioning, deduplication, compression, and compaction, significantly reduce the amount of physical storage consumed, lowering your costs. This data footprint reduction applies to the production data, the DR, and backup data, as well as to the data transfer costs.

VII. CONCLUSION

Organizations typically face a series of significant challenges in their effort to increase the return on their Server investments and optimize their infrastructure to support business and IT requirements. They must Accelerate new database implementations or migrations and lower the risk of these operations, make sure that the underlying storage infrastructure is fully optimized to support SLAs, including performance, scalability, and availability, consolidate existing databases and infrastructure to lower costs, reduce complexity and simplify IT infrastructure, Increase the productivity of IT personnel. To handle these challenges, the architects, sysadmins, or DBAs are looking to deploy their databases and storage infrastructure based on proven best practices and technology.

VIII. REFERENCES

- [1] "Performance Evaluation of Data Deduplication Techniques for SAN Storage" by Xiangyu Lei and Willy Zwaenepoel (2013). This article explores data deduplication, a key storage efficiency technique mentioned in the paper.
- [2] "A Survey of Data Deduplication Techniques for Network Attached Storage" by Wei Xu, Lingling Wang, and Wei Li (2011). This is another article on data deduplication techniques.
- [3] "Storage Tiering for Cloud Data Management: A Survey and Future Directions" by Yiming Hu, Cheng Li, David Wang, Rajiv Ranjan, and Natalia Anerousis (2016). This article explores storage tiering, another storage efficiency technique mentioned in the paper.
- [4] "Data Center Storage: Designing, Deploying, and Operating High-Performance Storage Infrastructure" by Jeffrey Hill (2014). This book provides a comprehensive overview of data center storage concepts and technologies relevant to NetApp storage.
- [5] "NetApp Data ONTAP Deep Dive: Configuration and Administration Guide" by Timothy McGeeney (2017). This book offers an in-depth exploration of NetApp's Data ONTAP storage operating system.
- [6] "SQL Server High Availability and Disaster Recovery" by William Stanton (2019). This book explores high availability (HA) and disaster recovery (DR) concepts mentioned in the paper, focusing on SQL Server.