International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 07 Issue: 04 | April - 2023 SJIF Rating: 8.176 ISSN: 2582-3930

Economical Approaches to NetApp Storage for Effective Database Management

Sethu Sesha Synam Neeli
sethussneeli@gmail.com
Sr. Database Administrator

Abstract:

Database storage is paramount in mission-critical applications, significantly influencing system performance metrics. In the contemporary technological landscape, the financial burden of procuring servers equipped with high-capacity storage solutions is considerable. To mitigate these expenses, organizations should evaluate alternative avenues, such as NetApp, a remarkably economical option.

NetApp has many features, including data cloning, deduplication, and advanced snapshot management, which optimize the storage architecture where databases are housed. This discourse will delve into best practices for integrating NetApp storage with an array of database management systems, emphasizing algorithms that enhance data retrieval efficiency and storage optimization.

The strategies proposed herein will furnish organizations with pragmatic methodologies to curtail storage expenditures while ensuring that critical performance measures and scalability are upheld. This approach guarantees that databases operate at their pinnacle of efficiency, all while remaining within predetermined budgetary parameters.

Keywords: NetApp storage, cost optimization, database management, deduplication, tiering, cloud integration, data lifecycle management.

1. Foundations and Structural Design

ONTAP Cloud is a software-centric manifestation of Data ONTAP, the data management operating system developed by NetApp for its physical storage appliances. This iteration of ONTAP has been meticulously tailored to function as an instance within the Amazon Elastic Compute Cloud (EC2), leveraging the cloud's robust infrastructure.

Key Characteristics of ONTAP Cloud:

1. Instantaneous, writable data replicas that require no supplementary storage capacity.
2. Rapid backup and recovery solutions for datasets of varying sizes (featuring NetApp snapshot copies, SnapMirror technology, and SnapVault).
3. Capability to simultaneously employ multiple protocols (CIFS, iSCSI, NFS) from a singular storage system.
4. Storage efficiencies that permit the utilization of reduced underlying storage capacity to meet data requirements (including thin provisioning, data deduplication, and data compression).
5. Space-optimized, user-friendly, bi-directional data transfer capabilities.
6. Data-at-rest encryption utilizing encryption keys securely stored on controlled key management systems.

7. Storage tiering to AWS S3 for cost optimization.

Layers NetApp ONTAP: High-Level Layered Design

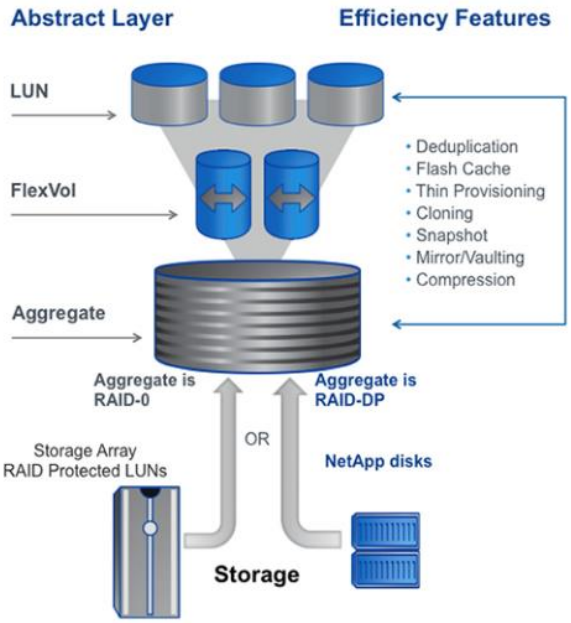

In AWS, an aggregate may encompass a maximum of 6 disks. The foundational disk configurations can consist of a singular EBS storage variant or a multilevel storage architecture comprising an EBS performance layer alongside an S3 capacity layer. The upper limit for disk size is 16 TB, facilitating an aggregate usable capacity of up to 85.05 TB.

The total storage acquired from AWS is referred to as the raw capacity. The usable capacity is diminished by approximately 12 to 14 percent due to overhead reserved for ONTAP Cloud operations. For instance, if Cloud Manager provisions a 500 GB aggregate, the net usable capacity would amount to 442.94 GB.

2. Research Approach:

NetApp offers a vast array of solutions capable of seamlessly integrating diverse database architectures, including SQL Server, Oracle, MySQL, VMware, SAP, and Postgres, across various cloud environments. This facilitates consistent and predictable performance while providing the flexibility to accommodate a myriad of workloads.
NetApp solutions epitomize a harmonious blend of management simplicity, scalability, and robust data protection protocols. A significant advantage lies in the storage efficiency provided by NetApp ONTAP data management software, which optimizes system performance metrics through advanced algorithms. Additionally, the incorporation of FlexClone technology enhances the development and testing phases by allowing for efficient database cloning without immediate storage overhead, as clones utilize minimal capacity until modifications are made.
In light of the expanding scope, scale, and intricacy inherent in today's data-centric landscape, these capabilities foster enhanced agility and responsiveness within the data center.

1. Expedited Backup Operations and Elevated Data Protection:

With NetApp's integrated data protection solutions, users benefit from rapid and space-efficient backup processes, supported by a comprehensive suite of Windows PowerShell cmdlets that facilitate the automation of backup and restore workflows through advanced scripting capabilities. Leveraging federated backup architecture, administrators can incorporate multiple database instances and their respective databases into a unified federated group, enabling simultaneous creation of NetApp Snapshot™ copies for holistic backup operations across all databases within that group.

The ability to restore to a specific marker allows for precise recovery aligned with critical business events, a particularly advantageous feature for applications that rely on multiple databases across diverse instances. Furthermore, by implementing NetApp's thin replication technology, organizations can safeguard vital database data—including components within Database Availability Groups—against potential disasters, thereby ensuring robust data integrity and resilience in system performance metrics.

2. Delivering Multitenancy and Superior QoS:

Customers can leverage the functionalities provided by ONTAP to facilitate data compartmentalization across all tenants, ensuring that each tenant has exclusive access to its designated storage repository. The robust integration of Windows PowerShell support for both ONTAP and NetApp integrated backup solutions empowers administrators to architect sophisticated workflows tailored to the requirements of a multitenant environment.

Furthermore, ONTAP enhances system performance metrics through advanced storage Quality of Service (QoS) algorithms that manage workload distribution. By utilizing policies that enforce throughput limits, administrators can ensure optimal resource allocation and performance consistency across diverse database workloads on NetApp ONTAP, thereby maintaining an efficient and equitable operational landscape.

3. Harnessing Intelligent Management Practices

NetApp's integrated backup solutions facilitate the optimization of administrative operations, including backup, restoration, cloning, and disaster recovery functionalities, thereby allowing administrators to concentrate on higher-level strategic initiatives and business objectives. Through the utilization of Snap Manager's federated backup capabilities, data protection for database applications is significantly simplified, particularly in scenarios where data is dispersed across multiple database environments. Moreover, the synergistic integration of ONTAP with Windows PowerShell scripting empowers Database Administrators (DBAs) to automate the orchestration of Logical Unit Numbers (LUNs) and volume management, enhancing operational efficiency and effectiveness in data governance.

4. Enhancing Inefficient Manual Test and Development Procedures

Accelerate the deployment and enhance the quality of new applications reliant on databases through the utilization of NetApp FlexClone technology. By leveraging FlexClone, organizations can generate space-efficient replicas of database datasets for development, testing, reporting, and disaster recovery validation within mere seconds.

This technology not only facilitates robust clone lifecycle management but also supports on-demand, scheduled, and tailored user acceptance testing services. Utilizing a policy-driven automation framework for cloning expedites these processes while ensuring minimal disruption to other Database Administrators (DBAs) and end users, thereby maintaining system performance metrics and operational integrity.

5. Accelerate Deployment of Mission-Critical Applications Through Converged Infrastructure

FlexPod-validated architectures integrate storage, networking, and server components into a cohesive, adaptable infrastructure tailored for enterprise workloads, including databases. These solutions significantly accelerate the deployment of infrastructure and mission-critical applications while concurrently mitigating costs, complexity, and associated risks. With a rigorously validated design and centralized management framework, organizations can

swiftly and seamlessly deploy and host multiple database instances, all benefiting from synergistic support provided by industry leaders Microsoft, Cisco, and NetApp.

6. Agile Solutions for Managing Data Growth:

With seamless scalability, organizations can rapidly and effortlessly augment NetApp solutions to align with evolving business requirements. The thin provisioning capabilities offered by NetApp FlexVol technology facilitate cost-effective expansion and reallocation of storage resources, while enhancing operational efficiency through the deduplication of live data—achieved without disrupting ongoing business processes.

Moreover, ONTAP empowers users to non-disruptively migrate data within their storage clusters, accommodating the demands of SQL Server data expansion and adapting to the fluctuating workloads of database applications. This strategic flexibility ensures optimal system performance metrics and effective resource management in response to dynamic data growth.

3. Storage Infrastructure Design and Execution:

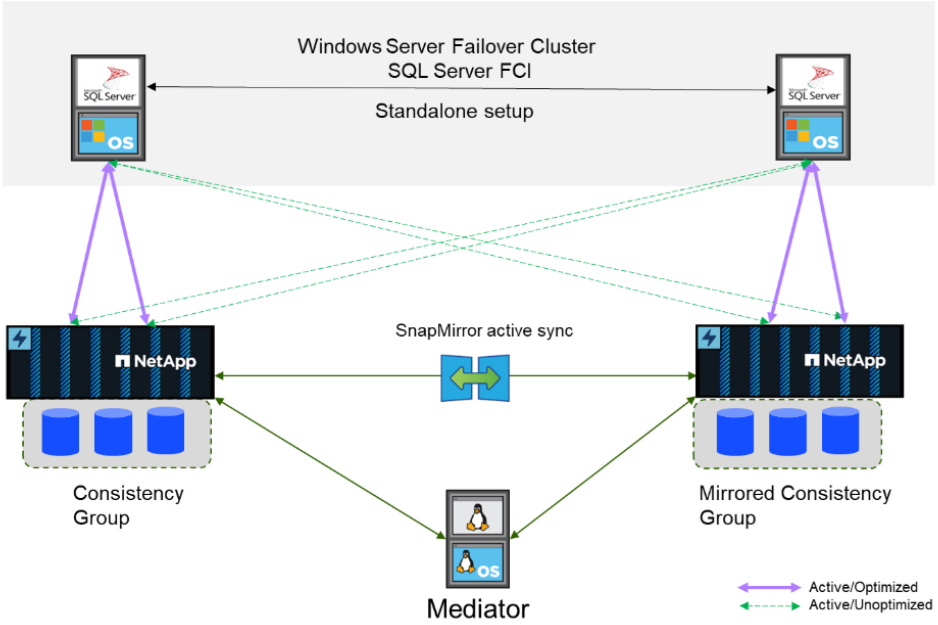

Basic Design of a NetApp in AWS

SnapCenter offers robust application-consistent data protection tailored for ONTAP environments, leveraging sophisticated algorithms for optimal database integrity. This powerful tool facilitates the creation of Snapshot copies for a diverse array of entities, including applications, databases, file systems, and virtual machines, ensuring minimal disruption during backup processes. With SnapCenter, users can efficiently orchestrate backups to a secondary storage system, whether housed within the same data center or strategically located across different geographic regions, all while maintaining high system performance metrics and maximizing data availability.

The proposed configuration is optimized for SQL Server instances necessitating fundamental performance metrics while accommodating multiple petite databases. The database storage architecture encompasses the following intricate characteristics:

- It supports a singular aggregate dedicated to SQL Server instances, specifically utilizing Microsoft SQL Server 2017 within the NetApp ONTAP framework.
- A dedicated volume and LUN are allocated for SQL Server system databases, inclusive of the pivotal tempdb database, ensuring superior performance and rapid data retrieval.
- Each individual database is provisioned with its own dedicated LUN, facilitating enhanced management and isolation of resources.
- A unified volume is employed to house both data and transaction logs, streamlining data organization and access.
- Dedicated SMB shares are established for both data and log files when employing SMSQL for backup processes, thereby optimizing network throughput and reliability through advanced communication algorithms.

This configuration not only aligns with best practices in database management but is also engineered to maximize efficiency and resilience in high-demand environments.

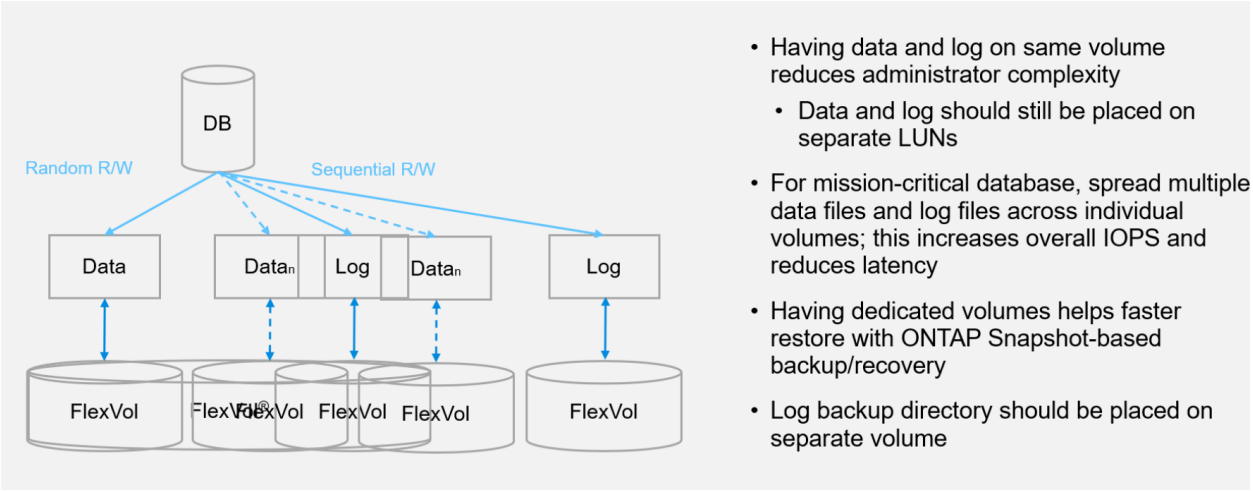

4. Shared Instance Versus Dedicated Instance:

When an application encompasses numerous schemas and stored procedures, it can significantly impact the performance of other applications operating within the same SQL Server instance. Resource contention may arise, resulting in divided or locked instance resources that lead to performance degradation for any applications utilizing databases hosted on the shared SQL Server environment. Diagnosing these performance concerns can become intricate, as it necessitates identifying the specific instance contributing to the issues.

This consideration is often juxtaposed with the financial implications of operating system and SQL Server licensing costs. In scenarios where application performance is of utmost importance, it is advisable to opt for a dedicated SQL Server instance. Microsoft licenses SQL Server on a per-core basis at the server level rather than per instance, which often entices database administrators to maximize the number of SQL Server instances installed on a single server to optimize licensing expenditures. However, this approach may culminate in significant performance challenges down the line.

NetApp advocates for the selection of dedicated SQL Server instances whenever feasible, as this strategy enhances system performance metrics, mitigates resource contention, and ensures optimal algorithmic efficiency across database operations.

5. Advantages of Cloud Volumes ONTAP on NetApp Solutions:

1. Cost Reduction:
Lower your Total Cost of Ownership (TCO) by consolidating your SQL Server databases onto NetApp storage systems. Our comprehensive support for Fibre Channel (FC), Fibre Channel over Ethernet (FCoE), iSCSI SAN, and SMB (CIFS) protocols allows for the seamless integration of your existing infrastructure, thereby minimizing storage expenditures. The unified architecture of NetApp enhances both storage and server utilization while streamlining data management processes. Notably, NetApp's native support for SMB 3.0 significantly simplifies the provisioning of database clones, facilitating rapid deployment for development, testing, and staging environments. This optimization not only enhances operational efficiency but also leverages advanced algorithms to ensure high system performance metrics across your database ecosystem.

2. Provide SnapCenter advantages
NetApp SnapCenter software provides a cohesive, scalable platform along with a comprehensive suite of plug-ins designed for application-consistent data protection and efficient clone management within complex IT ecosystems. IT infrastructure administrators can empower application and database administrators to self-manage their own data copies and protection strategies while maintaining centralized governance through robust policies and advanced reporting capabilities across diverse data environments.
With SnapCenter, data protection and clone management tasks are transformed into:
Simplicity: A centralized graphical user interface (GUI) supports real-time monitoring, notifications, logging, reporting, and scheduling for all application and database plug-ins, including those that are custom-built. This user-friendly interface enables administrators to efficiently manage system performance metrics and streamline operations.
Scalability: The architecture allows for seamless integration of additional SnapCenter servers, enhancing high availability and load balancing capabilities, thereby supporting thousands of applications without compromising performance or reliability.
Empowerment: Role-based access control (RBAC) facilitates self-service options for application administrators while allowing SnapCenter infrastructure administrators to maintain oversight. This balance ensures that the delegation of responsibilities does not detract from the governance and security of the data protection environment. Overall, SnapCenter not only simplifies the management of databases and application data but also leverages sophisticated algorithms to optimize performance, ensuring that organizations can effectively scale their data management solutions while preserving system integrity.

.3. Storage Solutions with Cloud Volumes ONTAP:
Utilizing Cloud Volumes ONTAP offers sophisticated data management capabilities that significantly enhance service levels, save time for IT and DevOps teams, and reduce the complexities and costs associated with storage management. The following delineates the advantages of employing Cloud Volumes ONTAP in your IT infrastructure:
Cost Efficiency through Storage Optimization: Cloud Volumes ONTAP can achieve up to 90% reduction in storage capacity requirements by leveraging advanced space-efficient technologies, including data deduplication, compression algorithms, thin-cloning, and Snapshot copies that do not impact the overall storage footprint.
High Availability and Resilience: Attain unparalleled high availability through a robust two-node architecture that spans multiple Availability Zones, ensuring business continuity for critical production workloads and databases. This design guarantees zero data loss (RPO=0) and rapid recovery times (RTO < 60 seconds), thus optimizing performance metrics for mission-critical applications.
Comprehensive Data Protection and Disaster Recovery: Efficiently recover from data corruption or loss utilizing agile data Snapshot copies and disaster recovery copies that are easily configured and cost-effective. These features

support seamless failover, failback, restoration, and recovery processes that align with stringent minute-level Service Level Agreements (SLAs).

Support for Hybrid and Multi-Cloud Environments: Streamline operations and reduce costs by deploying the same advanced NetApp ONTAP data management software across hybrid and multi-cloud environments. This versatility extends to disaster recovery (DR), high availability (HA), development/testing (Dev/Test), DevOps, sandboxing, reporting, data tiering, workload hosting, and training initiatives.

Data Mobility and Synchronization: Securely migrate, replicate, and synchronize your data while utilizing efficient Snapshot copies to transfer only incremental changes. With NetApp's SnapMirror technology, you can recover data from any point in time, ensuring minimal disruption and enhanced operational agility.

Cloning Technology for Enhanced Development: Elevate DevOps agility by enabling the cloning of writable volumes from Snapshot copies, allowing for concurrent data sharing across organizations and regions without incurring capacity or performance penalties, thanks to NetApp FlexClone technology.

Interoperability Across Protocols: Benefit from multi-protocol support (iSCSI and SMB) for data and file shares, effectively catering to the demands of SQL Server workloads and optimizing access patterns for improved system performance.

Flexible Licensing Options: Choose among various Cloud Volumes ONTAP solutions tailored to your needs, ranging from hourly pricing options to longer-term subscriptions and Bring Your Own License (BYOL) plans, ensuring cost-effective scaling in alignment with organizational requirements.

Enhanced Security Features: In addition to the robust security and privacy measures provided by the hyperscaler, Cloud Volumes ONTAP offers NetApp-managed encryption, empowering you to control encryption key management directly on-premises.

By employing Cloud Volumes ONTAP, organizations can realize a transformative approach to data management that maximizes performance while reducing overhead and risk across their IT landscapes.

6. Up to 50% lower storage costs for databases

Up to 50% Reduction in Database Storage Costs with FSx for ONTAP

FSx for ONTAP introduces significant enhancements designed to optimize database storage capacity and significantly reduce associated costs. The following outlines the pivotal improvements:

High-Throughput Accessibility: FSx for ONTAP storage is accessible via high-throughput network bandwidth, enabling organizations to leverage smaller and more cost-effective Amazon EC2 instance types for their database computing requirements. This optimization allows for improved resource allocation and cost efficiency while maintaining performance metrics.

Cost-Effective Licensing: With FSx for ONTAP, you can minimize database licensing expenditures without compromising on resiliency or data protection. For instance, MS SQL Server databases benefit from the robust availability and protection features of an Always-On Availability Group (AOAG) while utilizing the Standard MS SQL license—significantly more affordable than the Enterprise licenses typically required for AOAG.

Thin Cloning Technology: The implementation of thin cloning drastically reduces costs by allowing for lightweight copies that consume negligible storage space. This approach eliminates the substantial storage costs associated with traditional full copies, thereby optimizing budget allocation for database management.

Rapid Environment Provisioning: Thin cloning facilitates instantaneous creation of new database environments, accelerating Development and Testing (DevTest) cycles and enhancing team productivity. This expedited provisioning contributes to a reduced time-to-market (TTM) for new applications and features.

Advanced Storage Efficiency Features: FSx for ONTAP incorporates a suite of storage efficiency technologies, including thin provisioning, deduplication algorithms, compression mechanisms, and compaction techniques. These

innovations significantly minimize the physical storage footprint, thus lowering costs across production data, disaster recovery (DR), backup data, and data transfer expenses.
By harnessing the capabilities of FSx for ONTAP, organizations can achieve a transformative reduction in database storage costs while enhancing performance and operational efficiency through advanced data management strategies.
.

7. Conclusion

Organizations frequently grapple with a multitude of substantial challenges as they strive to enhance the returns on their server investments and refine their infrastructure to effectively align with business and IT requirements. They must expedite the implementation of new databases or facilitate migrations while mitigating associated risks, ensuring that the foundational storage infrastructure is meticulously optimized to adhere to Service Level Agreements (SLAs) encompassing critical parameters such as performance, scalability, and availability. Furthermore, they aim to consolidate existing databases and infrastructure to reduce expenditures, diminish complexity, and streamline the IT landscape while simultaneously enhancing the productivity of IT personnel. In response to these intricate challenges, architects, system administrators, and database administrators (DBAs) are seeking to deploy their databases and storage infrastructure in accordance with established best practices and innovative technologies to ensure optimal performance and efficacy.

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.