International Journal for Multidisciplinary Research (IJFMR)
E-ISSN: 2582-2160 ● Website: www.ijfmr.com ● Email: editor@ijfmr.com

Protecting NetApp ONTAP Volumes with BlueXP Backup and Recovery to Azure

Venkata Raman Immidisetti

Infrastructure Architect, Raleigh, North Carolina
vimmidisetti@gmail.com

Abstract

Organizations increasingly rely on hybrid cloud infrastructures, making unified data protection essential for both on-premises and cloud environments. This paper presents a technical overview of protecting NetApp ONTAP volumes using BlueXP Backup and Recovery with Microsoft Azure as the exclusive backup target. We describe an architecture that integrates on-premises ONTAP systems and cloud-native ONTAP (Cloud Volumes ONTAP) under a single backup framework, leveraging Azure Blob Storage for reliable off-site backups. The solution utilizes NetApp snapshot technology and incremental replication to Azure, preserving storage efficiencies and enabling high-performance, secure backups without disrupting production workloads. We detail how the system architecture functions in hybrid and cloud-native deployments, explain the backup and restore workflows, and discuss design considerations such as network connectivity and data immutability. By focusing on Azure as the backup destination, this paper demonstrates a cohesive approach to safeguarding data across hybrid cloud environments while adhering to the 3-2-1 backup strategy. The result is a cost-effective, highly scalable, and resilient data protection solution for enterprise storage volumes, with rapid recovery capabilities and seamless management through the BlueXP unified control plane.

Keywords: NetApp ONTAP, BlueXP, Azure Blob Storage, Cloud Volumes ONTAP, Hybrid Cloud Backup, Data Protection, SnapMirror to Cloud, Incremental Backup, Disaster Recovery

I. INTRODUCTION

Data is a critical asset for enterprises, and protecting that data across on-premises and cloud environments is a paramount concern. NetApp ONTAP, a leading storage operating system, is widely deployed on dedicated storage appliances in data centers as well as in cloud-based instances (Cloud Volumes ONTAP) to manage cloud storage. As organizations adopt hybrid cloud strategies, they require a unified backup solution that can safeguard ONTAP-managed data wherever it resides. Traditional backup methods often involve disparate tools for on-premises and cloud data, leading to complexity, higher costs, and potential gaps in coverage. There is a need for an integrated approach that ensures consistent protection and centralized management of backups across hybrid environments.

BlueXP Backup and Recovery addresses this need by providing a cloud-integrated backup service tailored for ONTAP volumes. It leverages the efficiency of NetApp's snapshot and replication technologies to create point-in-time backup copies of volumes and stores them in low-cost cloud object storage. By exclusively using Microsoft Azure as the backup target, organizations can take advantage of Azure's scalable and durable storage infrastructure to maintain off-site backups of their critical data.

This approach aligns with the industry-standard 3-2-1 backup strategy (3 copies of data on 2 different storage types, with 1 off-site copy) by using on-site storage for primary data and Azure cloud storage for the off-site backup copy. Furthermore, Azure's global infrastructure and security features make it a robust repository for protected data, enabling geographic redundancy and compliance with data retention policies.

In this paper, we present an academic-style examination of the architecture and operation of BlueXP Backup and Recovery for NetApp ONTAP when Azure is the designated backup destination. We focus on two primary deployment scenarios: a hybrid deployment where on-premises ONTAP clusters are backed up to Azure, and a cloud-native deployment where Cloud Volumes ONTAP instances in Azure are backed up to Azure Blob Storage. We detail how the BlueXP system orchestrates backup and restore processes in these scenarios, how it maintains data consistency and security, and how it leverages Azure services to enhance reliability. Diagrams are included to illustrate the system architecture and data flow during backup operations. Through this analysis, we demonstrate that BlueXP Backup and Recovery offers a cohesive, efficient, and secure solution for protecting enterprise data across hybrid cloud landscapes, using Azure as a unified backup target.

II. ARCHITECTURE

BlueXP Backup and Recovery is architected as a centrally managed service that spans the on-premises data center and the Azure cloud. The key components of the solution include the NetApp ONTAP storage systems (source volumes), the BlueXP management plane with its Backup and Recovery service, a BlueXP Connector for communication and orchestration, and Azure Blob Storage as the backup repository. The design ensures that data can flow securely from ONTAP volumes to Azure storage, under the coordination of BlueXP, with minimal administrative overhead.

BlueXP Management Plane and Connector: BlueXP is NetApp's cloud-resident management platform that provides a single control plane for multiple data services. The Backup and Recovery service within BlueXP is responsible for orchestrating backup schedules, retention, and restore operations for ONTAP volumes. To interface with customer environments, BlueXP uses a software proxy called the BlueXP Connector. The Connector can be deployed either on-premises (within the customer's data center) or in the Azure cloud (within the customer's Azure Virtual Network). Its role is to act as a communication broker between the BlueXP cloud service and local resources: it issues ONTAP API calls to the storage systems and facilitates data movement by coordinating with Azure services. The Connector is lightweight and can be installed as a virtual appliance or VM. For Azure deployments, a Connector can be launched from the Azure Marketplace for convenience. In all cases, secure connectivity (TLS 1.2 over HTTPS) is used between the Connector, the ONTAP systems, and the BlueXP cloud, ensuring that control commands and metadata exchanges are protected in transit.

NetApp ONTAP Source Systems: These are either on-premises ONTAP clusters (such as NetApp FAS or AFF systems, or ONTAP Select virtual appliances) or cloud-native ONTAP instances (Cloud Volumes ONTAP running in Azure). The source ONTAP system contains the primary data volumes that need to be protected. ONTAP provides inherent data protection features like *Snapshot copies*, which are immutable point-in-time images of a volume. BlueXP Backup and Recovery leverages these Snapshot copies as the foundation for backups, avoiding any performance impact on active data. In a typical

backup operation, the ONTAP system will create a new Snapshot of the volume (if not already taken recently per policy) and use that Snapshot as a consistent read-only source for the backup, ensuring that the data captured is crash-consistent and does not change during the transfer.

Azure Blob Storage (Backup Target): Microsoft Azure Blob Storage serves as the dedicated backup target in this solution. Blob storage is a highly scalable object storage service ideal for storing large amounts of unstructured data such as backup files. In the context of BlueXP Backup and Recovery, each ONTAP volume backup is stored as a set of objects (often referred to collectively as a backup file or image) in an Azure Blob Storage container. By using Azure Blob, the solution benefits from its durability (multiple replicas of data are maintained within the storage account), availability, and cost-effectiveness (with tiered storage options for cooler data). By default, the backup storage is configured with locally-redundant storage (LRS) within a region for cost optimization, but users have the option to configure zone-redundant storage (ZRS) or even geo-redundant options if higher resilience is required. Each backup is stored in Azure as an immutable object set – once written, the backup data is not altered, which provides protection against accidental or malicious modifications. Access to the backup containers is restricted through Azure's security (access keys or Azure AD integration), and data at rest in Blob Storage is encrypted using AES-256 encryption, meeting industry compliance standards.

Network Connectivity: A crucial aspect of the architecture is the network path between the ONTAP source systems, the BlueXP Connector, and Azure Blob Storage. There are two primary connectivity methods:

- **Public Endpoint Connectivity:** In this method, the ONTAP source sends data to Azure Blob over a secure HTTPS connection via the public Azure endpoint (internet). This requires the ONTAP cluster to have network access to Azure (either directly or through a corporate firewall allowing outbound HTTPS to Azure). The BlueXP Connector, whether on-premises or in Azure, communicates with Azure Blob service and the BlueXP cloud control plane over the internet as well. All communication is encrypted. This option is simpler to set up and is often used when no private network link exists between the data center and Azure. It is depicted in Figure 1 (left side) where the on-premises connector and ONTAP use public internet routes to reach Azure Blob Storage.

- **Private Endpoint Connectivity:** For enhanced security, the architecture can utilize Azure Private Link to connect to Blob Storage. In this configuration, a Private Endpoint for the Blob Storage account is created within an Azure Virtual Network. If the organization has a VPN or Azure ExpressRoute connection from the on-premises environment to the Azure VNet, the ONTAP cluster's traffic to Blob Storage can be routed through this private connection, never traversing the public internet. The BlueXP Connector can also use the private endpoint to access Azure Blob. Figure 1 (right side) illustrates this scenario, showing a Private Endpoint within Azure and data transfer over VPN/ExpressRoute. This method provides a fully private, internal path for backup data, which may be required by organizations with strict security compliance needs. Regardless of method, the data transfer uses HTTPS, and only the minimal required ports (e.g., port 443 and Azure-specific endpoints) are opened, reducing the attack surface.

Hybrid and Cloud-Native Deployments: The architecture supports both hybrid and cloud-native use cases seamlessly:

- In a **hybrid deployment**, the source is an on-premises ONTAP cluster. The BlueXP Connector may reside on-premises (for local API access) or in Azure. The connector discovers the on-premises cluster (through its management IP interface) and orchestrates backup transfers from the cluster's networking interface to Azure. Optionally, organizations can also deploy a Cloud Volumes ONTAP instance in Azure as a secondary system for disaster recovery; the on-premises volumes might be replicated to this secondary ONTAP (using NetApp SnapMirror technology) in addition to being backed up to Blob storage. In this layered approach, the cloud ONTAP serves for quick failover or recovery, while Azure Blob serves for longer-term retention and recovery in worst-case scenarios. However, BlueXP Backup to Azure can be used independently even without a secondary ONTAP – the primary copy remains on-premises and the backup copy in Azure.

- In a **cloud-native deployment**, the source is a Cloud Volumes ONTAP (CVO) instance running in Azure. Here, the environment is entirely within Azure: the BlueXP Connector runs in the same or a nearby Azure region (often in the same VNet as the CVO instance), and the backup target is Azure Blob in that region. This setup benefits from Azure's internal high-bandwidth network, enabling efficient data transfer. CVO instances typically have direct network access to Azure services, and a Private Endpoint can be used by the Connector and ONTAP to reach Blob storage without leaving Azure's backbone. In this scenario, the backup architecture avoids any data movement over the public internet altogether. The proximity of CVO and Blob storage also means low latency and potentially faster backup and restore throughput. Cloud-native ONTAP deployments might use Azure managed disks for active data and Azure Blob for backups, achieving a fully cloud-based data protection cycle. If multiple CVO instances exist in different regions or accounts, BlueXP can manage all of them, backing up each to the designated Blob storage and even allowing backups to be restored across instances if needed (for example, restoring data from one CVO to another in a disaster recovery event).

The architecture is designed to be scalable and secure. It can protect many ONTAP volumes in parallel, as each backup operation is handled at the volume level and Azure Blob storage can ingest multiple streams concurrently. BlueXP manages this concurrency and ensures that system resources (network and storage throughput) are used efficiently. Security is enforced at multiple layers: ONTAP volumes remain secure in transit (TLS encryption) and at rest in Azure (encryption and optional object locking for immutability), and all control plane operations require authentication and authorization via BlueXP. This integrated design enables enterprise-grade data protection for ONTAP environments using Azure as a seamless extension of the backup infrastructure.

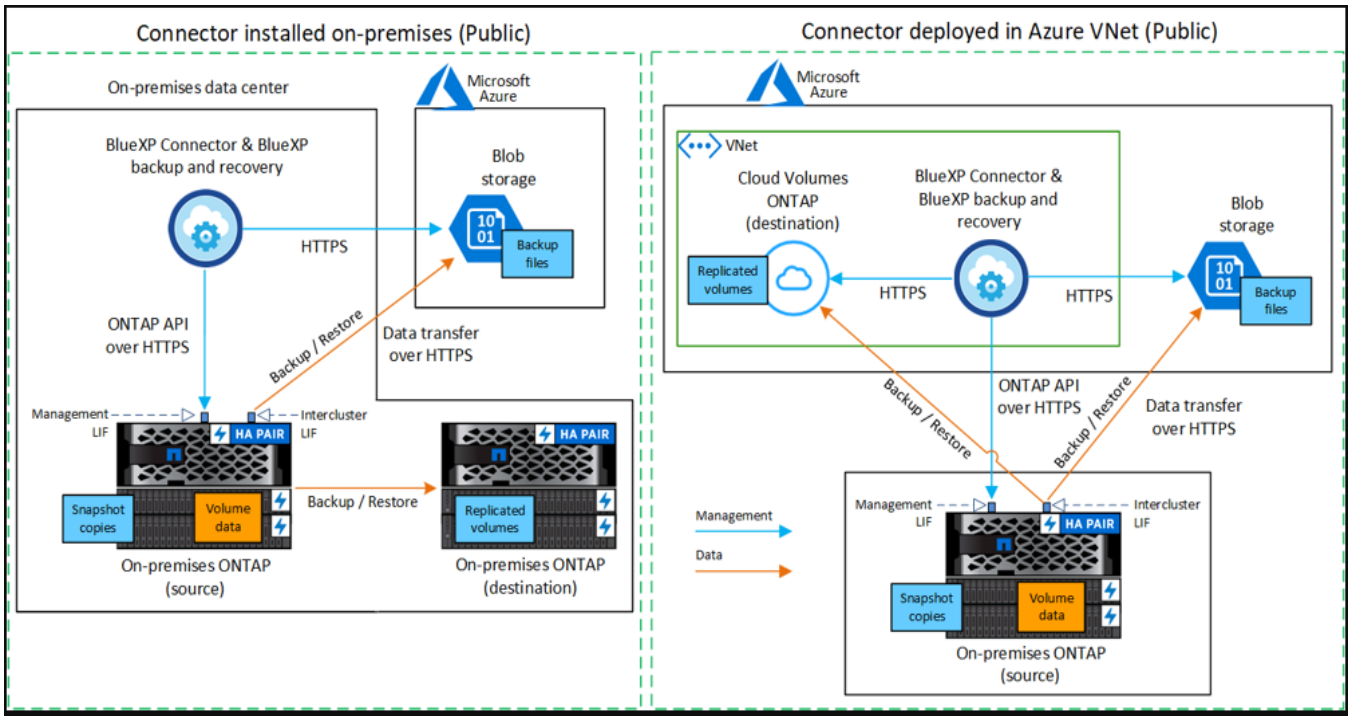

FIGURE1: BlueXP architecture with connector on-premises and cloud

III. HOW BLUEXP BACKUP AND RECOVEY WORKS

BlueXP Backup and Recovery operates by creating and maintaining incremental, point-in-time copies of ONTAP volumes in Azure Blob Storage. The process is largely automated once configured, relying on ONTAP's robust snapshot mechanism and NetApp's SnapMirror-to-Cloud technology to efficiently transfer only changed data blocks to the cloud. Below, we outline the backup workflow and the restore workflow, explaining how data moves through the system and how consistency and integrity are preserved.

Backup Workflow: The backup process can be scheduled or triggered on-demand through the BlueXP interface. Administrators define backup policies in BlueXP, which include the backup frequency (e.g., hourly, daily, weekly) and retention (how many backup versions to keep or how long to retain them). Once initiated, a typical volume backup proceeds through the following steps:

1. **Snapshot Creation:** BlueXP instructs the ONTAP source to create a new Snapshot copy of the target volume (if one has not been taken recently according to policy). ONTAP snapshots are quick, metadata-based operations that mark a consistent state of the file system without duplicating the data, thus incurring minimal performance impact. Users can also integrate application-consistent snapshot creation (using NetApp's SnapCenter or other tools) if a quiesced state is required for critical databases, but in most cases crash-consistent snapshots are sufficient for file services.

2. **Data Transfer Initialization:** The BlueXP Connector, upon confirmation of the snapshot, orchestrates the transfer of data from the ONTAP system to Azure Blob Storage. Under the hood, this leverages NetApp's SnapMirror technology extended for cloud targets (often referred to as

SnapMirror Cloud). SnapMirror identifies the new or changed blocks in the volume since the last backup (using the snapshot delta between the new snapshot and the previous backup snapshot) – this delta technique means that after the first full backup, subsequent backups send only incremental changes. This approach eliminates redundant data transfer, significantly reducing bandwidth and time required for backups. It also preserves ONTAP storage efficiencies: blocks that have been deduplicated or compressed on the source remain in that state during transfer and in the stored backup, avoiding re-inflation of data.

3. **Transfer over Network:** The identified changed blocks are transmitted from the ONTAP cluster to Azure Blob over the network. The route depends on the connectivity setup (public internet or private link as discussed in the Architecture section). The data is sent securely (HTTPS). ONTAP uses an intercluster LIF (logical interface) for SnapMirror data transfers, which in this case connects to the Azure endpoint. The BlueXP Connector monitors the transfer progress but typically the data flows directly from ONTAP to Azure, not through the connector, to optimize throughput. Azure Blob Storage receives the data and writes it into the designated container as objects. The backup data may be stored as a bundle of objects representing the volume's logical content or as a single large object; in either case, it is associated with metadata identifying the source volume, snapshot time, and other information.

4. **Backup Confirmation and Cataloging:** After successful transfer, BlueXP verifies the backup integrity (ensuring all expected blocks were transferred and possibly validating checksums). It then catalogues the backup: the backup instance is recorded in the BlueXP Backup and Recovery catalog, which tracks available backup versions for each volume, their timestamps, and retention status. This catalog is crucial for browse and restore operations. Additionally, if the Search & Restore feature is enabled, at this point BlueXP may trigger an indexing process. This process involves scanning the file system metadata within the backup to build a searchable index of file names and attributes. NetApp employs an indexing service (using Azure Data Lake Storage and Azure Synapse Analytics in the backend) to allow users to later search for individual files across all backups. This step is optional and runs in the background without affecting the completion of the backup.

5. **Retention Management:** BlueXP enforces the retention policy by marking older backups for deletion when they fall outside the retention window or when the maximum number of backups (up to 4,000 per volume) is reached. When a backup version expires, BlueXP will prune it by deleting the associated objects from Azure Blob Storage. Importantly, the retention logic and deletion of old backups do not affect active backups or the live data; only the objects corresponding to expired recovery points are removed, maintaining cost efficiency on the storage usage.

This incremental forever backup mechanism means that after the initial full copy, each subsequent backup operation is relatively quick, often completing in minutes for modest changes, and efficient in storage consumption. There is effectively no traditional "backup window" where the application needs to be taken offline; ONTAP volumes remain online and serving data, since the backup reads from a

snapshot and transfers in the background. By using Azure's elastic storage and networking, the solution can scale to protect petabytes of data and thousands of volumes without a linear increase in backup time.

Restore Workflow: Restoring data from Azure backups is a critical capability of the system, and BlueXP provides flexible restore options to address different recovery scenarios:

- **Full Volume Restore:** This is used in scenarios such as data corruption or loss of a primary volume, where the entire volume needs to be recovered to a previous state. The administrator selects a specific backup version from the BlueXP catalog. BlueXP then orchestrates the restore by instructing the ONTAP system to retrieve the backup data from Azure. If restoring to the original source cluster, ONTAP can either overwrite the existing volume (typically by reverting to the snapshot of the backup) or, for safety, restore into a new volume. The ONTAP system connects to Azure Blob (using the same mechanism as backup but in reverse) and downloads the needed blocks. Thanks to the incremental nature of backups, ONTAP knows which blocks have changed between the backup point and the current volume; it can request only the blocks needed to reconstruct the volume to the backup state. Azure Blob serves the data, and ONTAP writes it to the volume, rehydrating the volume to the exact contents of the snapshot. This process is efficient because it, too, avoids transferring unchanged data — if restoring to a new volume, it effectively pulls the full backup image; if overwriting an existing volume, it may only fetch differences, similar to a SnapMirror update. Once completed, the volume is available in the state it was when the backup was taken.

- **Restore to Alternate Location:** BlueXP allows restoring a backup to a different ONTAP system than the source. For example, an on-premises volume's backup in Azure could be restored to a Cloud Volumes ONTAP instance in Azure. This is useful for disaster recovery (if the on-premises site is down) or for testing (restoring a copy of production data into a test environment in the cloud). In this case, BlueXP coordinates access between the backup storage and the alternate ONTAP target. The target ONTAP will authenticate to Azure Blob (via the credentials or access setup in BlueXP) and pull the backup data to create a new volume. This cross-system restore showcases the benefit of having backups in a universally accessible cloud object store: any authorized ONTAP instance that can connect to Azure can potentially pull the backup and recover the data. The data remains encrypted in transit and at rest, ensuring security even during cross-location restores.

- **File-Level Restore:** Often users do not need an entire volume restored, but just a specific file or folder that was accidentally deleted or modified. BlueXP Backup and Recovery provides a mechanism to retrieve individual files from backups. Using the indexed catalog (if enabled) or by mounting a backup snapshot, BlueXP allows browsing the backup content. Under the covers, the system can instantiate a temporary restore operation of the snapshot (either on the source or on an auxiliary space) to access the file system structure. The user can then select the needed file, and BlueXP will extract that file from the backup objects and deliver it to the ONTAP system, where it can be copied to the original volume or delivered to the user. This granular restore avoids having to make a full volume copy when only a small item is needed, saving time and resources. The process still ensures data integrity: the file retrieved is exactly as it was at the time

of the backup. File restores also benefit from the fact that metadata was captured in the backup; attributes like permissions and timestamps are preserved.

Throughout the backup and restore processes, several features ensure data integrity and security. Every backup is essentially a read-only snapshot image, which means it is inherently consistent (there are no partial backups or open-file issues since ONTAP snapshots freeze the image at a point in time). The backups stored in Azure are logically immutable – they appear as objects that are not altered until deletion (and Azure can be configured with write-once-read-many policies if regulatory immutability is required). Data integrity checks (checksums) are maintained from ONTAP through to the object store, ensuring that a restore will not proceed with corrupt data. Additionally, access control is maintained: BlueXP requires credentials to connect to ONTAP and to Azure storage, meaning only authorized operations can pull backups or perform restores. This guards against unauthorized data exfiltration.

Optimization and Performance: BlueXP Backup and Recovery is optimized to reduce both the backup footprint and the performance impact on systems:

- Because it uses block-level incremental updates, even large volumes with minimal daily changes will incur very small backups. For instance, if only 5% of a volume's data changes between daily backups, only that 5% (at the block level) will be sent to Azure.
- ONTAP's ability to preserve storage efficiencies means if the volume data was compressed or deduplicated, the backup will not expand those blocks. This can dramatically save space in the Azure Blob container and reduce egress bandwidth. Effectively, the backup in the cloud is as space-efficient as the source storage.
- Backups can be performed frequently (even hourly) without performance degradation, because ONTAP snapshots are instantaneous and the incremental transfer can run in the background. This allows enterprises to achieve low RPO (Recovery Point Objective) values – they can restore data from a point just an hour or less before an incident, minimizing data loss in the event of corruption or deletion.
- Multiple volumes can be backed up in parallel. BlueXP manages task scheduling to ensure that concurrent transfers do not overwhelm the network. Azure's elasticity means that scaling the number of backups primarily impacts the available bandwidth and throughput from the ONTAP side. In practice, NetApp ONTAP is capable of high-throughput transfers, and Azure Blob can intake data very quickly, so the limiting factor is often the network link (for on-premises, the WAN or internet bandwidth, and for cloud, the VM's bandwidth limits). Administrators can tune the number of simultaneous backups or throttle bandwidth if needed to avoid saturating links during business hours.

Security Considerations: Given that backups represent a last line of defense (especially against threats like ransomware), the solution incorporates strong security measures:

- All communications between BlueXP, the Connector, ONTAP, and Azure are encrypted (HTTPS with TLS). Additionally, ONTAP supports encryption of data in flight (using TLS) natively for SnapMirror transfers.

- Data at rest in Azure is encrypted using Microsoft-managed keys by default. If required, customers can use Azure's options for customer-managed keys to have control over encryption keys for the backup storage account.
- Role-based access control in BlueXP ensures that only authorized users can initiate restores or modify backup settings. Fine-grained permissions can limit, for example, a junior admin to only perform restores but not delete backups.
- Immutability settings can be leveraged on the Azure side (via Blob object locking or Azure Vault policies) to prevent deletion of backups for a prescribed retention period. This can protect backup data from malicious tampering or premature deletion in case of a security breach.
- The BlueXP Backup service itself is monitored and maintained by NetApp (if using the SaaS control plane), which means updates and security patches are regularly applied to the service and connector, reducing the management burden on the user's side and ensuring the latest security standards are met.

In summary, the "How It Works" can be characterized by automation and efficiency. The heavy lifting is done by ONTAP's proven data replication engine, while BlueXP orchestrates and keeps track of everything. Azure serves as a robust and scalable storage backend. This synergy allows administrators to "set and forget" backup policies, knowing that their ONTAP volumes — whether on-prem or in the cloud — are continuously protected and that any needed recovery, from an individual file to an entire data volume, can be executed swiftly through a single pane of glass.

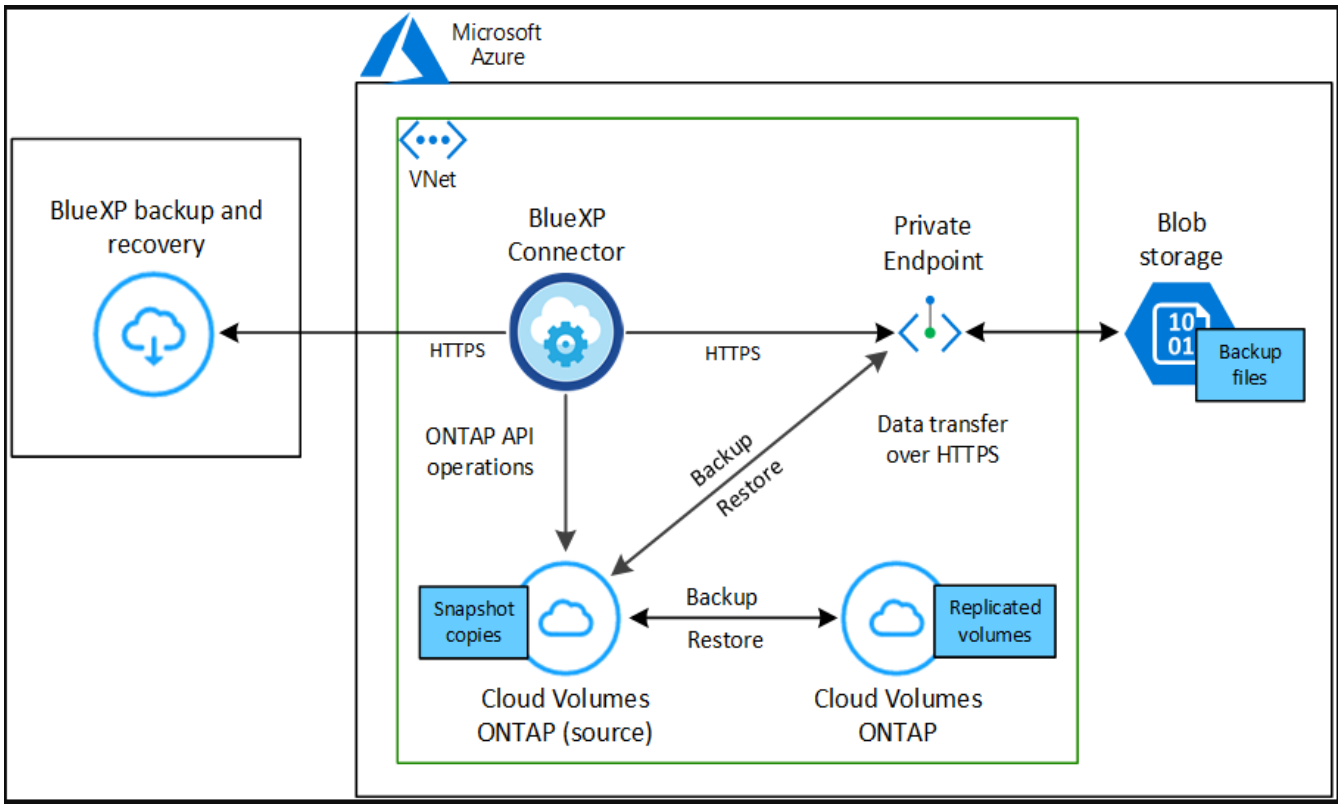

FIGURE 2: BlueXP backup and restore flow in azure

IV. CONCLUSION

This paper has presented a unified approach to protecting NetApp ONTAP volumes—both on-premises and cloud-native—using BlueXP Backup and Recovery with Azure as the exclusive backup target. By leveraging ONTAP snapshots and SnapMirror-based incremental replication, the solution efficiently transfers data to Azure Blob Storage while preserving storage efficiencies and minimizing performance impact. The architecture supports both public and private network paths and enables seamless management through a single control plane.BlueXP facilitates scalable, secure, and policy-driven backup workflows with granular restore capabilities, including full-volume and file-level recovery. The ability to restore data across environments enhances disaster recovery readiness, while Azure's durability and flexibility ensure long-term retention and cost optimization.In summary, this cloud-integrated backup solution effectively addresses hybrid cloud protection requirements, aligning with enterprise data resiliency goals and modern 3-2-1 backup strategies. It exemplifies how cloud-native services can extend traditional storage systems to deliver comprehensive, unified data protection across hybrid IT environments.

REFERENCES

[1] Kumar, K. Praveen. "The Discussion on Banking System in Rural Area through Cloud Computing." *Globus An International Journal of Management & IT* 6, no. 1 (2014): 51-53.

[2] Patterson, R. Hugo, and Stephen Manley. "{SnapMirror}:{File-System-Based} Asynchronous Mirroring for Disaster Recovery." In *Conference on File and Storage Technologies (FAST 02)*. 2002.

[3] Pandey, Anoop Kumar, Amit Kumar, Nilesh Malviya, and Balaji Rajendran. "A survey of storage remote replication software." In *2014 3rd International Conference on Eco-friendly Computing and Communication Systems*, pp. 45-50. IEEE, 2014.

[4] Lin, David, and Tom Ledoux. "Optimizing Data Storage and Management for Petrel Seismic Interpretation and Reservoir Modeling." (2009).

[5] Wang, Yanlong, Zhanhuai Li, and Wei Lin. "Rwar: A resilient window-consistent asynchronous replication protocol." In *The Second International Conference on Availability, Reliability and Security (ARES'07)*, pp. 499-505. IEEE, 2007.

[6] Azagury, Alain, M. F. Factor, Julian Satran, and William Micka. "Point-in-time copy: Yesterday, today and tomorrow." In *NASA CONFERENCE PUBLICATION*, pp. 259-270. NASA; 1998, 2002.

[7] Zhao, Zhenhai, Tingting Qin, Fangliang Xu, Rui Cao, Xiaoguang Liu, and Gang Wang. "CAWRM: A remote mirroring system based on AoDI volume." In *2011 IEEE/IFIP 41st International Conference on Dependable Systems and Networks Workshops (DSN-W)*, pp. 99-104. IEEE, 2011.

[8] Xiang, Xiaojia, Hongliang Yu, and Jiwu Shu. "Storage virtualization based asynchronous remote mirror." In *2009 Eighth International Conference on Grid and Cooperative Computing*, pp. 313-318. IEEE, 2009.

[9] Curtis-Maury, Matthew, Vinay Devadas, Vania Fang, and Aditya Kulkarni. "To waffinity and beyond: A scalable architecture for incremental parallelization of file system code." In *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*, pp. 419-434. 2016.

[10]https://bluexp.netapp.com/blog/cvo-blg-snapmirror-in-the-cloud-new-uses-for-netapp-data-replication

[11]https://docs.netapp.com/us-en/bluexp-backup-recovery/concept-ontap-backup-to-cloud.html#features

[12] https://docs.netapp.com/us-en/bluexp-backup-recovery/concept-backup-to-cloud.html#how-bluexp-backup-and-recovery-works

[13] https://docs.netapp.com/us-en/bluexp-backup-recovery/task-backup-to-azure.html#quick-start