IDC ANALYZE THE FUTURE

IDC TECHNOLOGY SPOTLIGHT

Sponsored by: Veritas

Critical data is increasingly spread across datacenter, cloud, and edge repositories. Edge workloads can contain highly valuable and sensitive information that requires special handling. Backup appliances both improve data protection and simplify backup operations at edge locations.

Using Purpose-Built Backup Appliances to Fully Protect Edge Workloads

September 2019

Written by: Phil Goodwin, Research Director, Infrastructure Systems, Platforms, and Technologies Group

Introduction: The Virtual Infrastructure Imperative

Data protection for edge locations, including remote offices and branch offices, has long been a challenge for organizations due to lack of skilled IT staff in the remote location, obsolete backup infrastructure, and rapidly changing business environments. By upgrading edge backup infrastructure to a purpose-built backup appliance (PBBA), IT organizations can simplify backup deployment, eliminate the need for onsite IT expertise, and rapidly adapt to changing business requirements for data availability.

The data volume at individual edge locations often does not justify full-time IT staff. Nevertheless, these locations have critical data, including sensitive information such as personally identifiable information (PII), that requires regulatory compliance. Moreover, when organizations have dozens or hundreds of edge locations, the collective volume of data can be substantial.

AT A GLANCE

KEY STATS

- 90% of organizations expect to use the cloud as part of their data protection strategy within 12 months.
- The most common architecture will be hybrid cloud, whereby data is located both on-premises and in the cloud.
- Organizations will tend to keep a primary copy of data on a purpose-built backup appliance (PBBA) in the compute location (core or edge) for optimized backup and rapid restore.
- A secondary copy of data will reside on a cloud-based PBBA (physical or software-defined) for assured data survival and disaster recovery staging.

Traditional tape backup at edge locations requires a knowledgeable administrator to ensure that tapes are regularly rotated, failed jobs are restarted, and data is restored from tape as necessary. Tape backup at edge locations is also typically labor intensive and error prone. Human errors can involve lost tapes or procedural mistakes, while technology errors involve broken tapes, bad tapes, and so on.

Data protection at edge locations may be further complicated by connectivity infrastructure (i.e., networks and the internet) that has limited bandwidth or inconsistent connectivity. Organizations may copy data from edge locations to a central datacenter or perhaps a cloud repository, and limited bandwidth may cause this transfer to take many hours that bleed into production worktime and negatively impact operations. Moreover, the longer the data transfer window, the more likely there is to be a job failure or error. Similarly, attempting to restore data across the network may take many hours (or even days) and impact the organization's ability to conduct business.

Virtual infrastructure (VI) now dominates the IT infrastructure landscape and, according to IDC research, accounts for more than 85% of x86 workloads. The deployment of VI has led to a consolidation of responsibilities within IT

organizations. For many, the VI administrator is tasked with managing not only the VI but also the backup environment. This creates two imperatives for backup solutions:

- They must operate within and protect VI environments.
- They must simplify backup operations to reduce the workload on VI administrators.

Purpose-built backup appliances have emerged as a cornerstone technology in simplifying data protection operations from the core to the edge while substantially improving the reliability of backup and recovery. PBBAs can eliminate the need for skilled staff at edge locations and reduce data transfer times when copying data to other locations, whether a central datacenter or the cloud.

Understanding PBBAs and Their Benefits

Purpose-built backup appliances have evolved into three primary categories: target, integrated, and virtual. Target appliances are specialized disk arrays that simply act as a substitute for tape systems. The IT organization must supply the backup/recovery software as well as any integration necessary. Integrated appliances are an integrated package of hardware and backup software, preconfigured and installed; integrated appliances are as close to "plug and play" as possible. Virtual appliances, or software-defined appliances, are software downloads onto common infrastructure or in the cloud.

Regardless of category, all PBBAs replace relatively less reliable tape infrastructure with more reliable hardware arrays. These appliances eliminate the tape-related errors mentioned previously, both human related and technology related. With fewer job-related errors, administrators spend less time on reactive problem solving and more time on productive projects. Restores are accomplished at disk speeds without the need to locate and load tapes.

A hallmark of PBBAs is the deduplication of data, often at rates of 20:1 or more. This deduplication has several benefits: lowering the cost per GB of data stored, reducing power/cooling per GB stored, and compressing the amount of data to transfer over the network. This latter factor can reduce the time to transfer data from an edge location to a datacenter or cloud by the same ratio (20:1) as the data deduplication rate. This is especially true for edge deployments where source-side deduplication (meaning at the edge location) reduces the amount of data to be transferred before the data ever enters the network.

A hallmark of PBBAs is the deduplication of data, often at rates of 20:1 or more.

Many PBBAs are equipped to transfer data to the cloud using common protocols (e.g., S3, OpenStack, SMB), a practice known as cloud tiering. Many organizations will choose to simultaneously replicate data from edge locations to the central datacenter as well as the cloud. Replicating to the cloud ensures data survivability in the event of a disaster and stages the data for disaster recovery. In addition, organizations may stage the cloud replication as an "air gap" to avoid propagating malware or ransomware from location to location and into the backup set. A third use case is to transfer data to the cloud for long-term archiving.

Another key benefit of PBBAs is a significant reduction in required human labor. In addition to managing fewer failed backups due to missed schedules, there is no need for any backup tape management with PBBAs. The devices can be managed remotely from a central location, meaning that no trained staff are needed at the edge location because there are no local administrative tasks. The central IT organization can create master policies that can be pushed out to the edge devices, simplifying the management of dozens or even hundreds of devices and ensuring that corporate data protection and governance policies are applied uniformly.

Key Trends

IDC research has found that 90% of organizations expect to use the cloud as part of their data protection strategy within 12 months. The most common architecture will be hybrid cloud, whereby data is located both on-premises and in the cloud. Organizations will tend to keep a primary copy of the data on a PBBA in the compute location (core or edge) for optimized backup and rapid restore with a secondary copy on a cloud-based PBBA (physical or software defined) for assured data survival and disaster recovery staging.

Our research also forecasts that 75% of data will require special handling by 2025. Examples include confidential information, personally identifiable information, financial information, and classified information. Organizations must have strategies and mechanisms for finding this information as well as tracking it and governing it appropriately, whether on primary storage or in backup data sets.

We also find the distinction between backup and recovery and disaster recovery to be diminishing. Organizations are driving to ever lower recovery point objectives (RPOs) and recovery time objectives (RTOs) and replicating data in near real time to the cloud. In combination with workload migration, recovery orchestration tools, and on-demand compute resources, application recovery within minutes is becoming realistic. Thus, organizations will be able to create a scenario where the data protection and disaster recovery infrastructure and processes are essentially the same.

Considering Veritas

IDC data shows that Veritas is among the market leaders by revenue for both data protection software and purpose-built backup appliances. The company's NetBackup Appliance family fits the integrated appliance category; it is a fully integrated bundle of appliance hardware and NetBackup software. Among this family of systems is the Flex 5150. Flex appliances are designed to be highly adaptable in order to reduce costs, simplify management, and rapidly respond to changing business requirements.

The Flex 5150 was designed specifically for edge locations. A common use case is large retail organizations that may have dozens or hundreds of stores with point-of-sale (POS) systems as well as inventory and financial applications, among others. It is simply not practical or cost effective to have IT staff at each location, yet the systems at each location must be backed up regularly and disaster prepared. Other industries such as restaurant and hospitality, financial services, and government have similar scenarios.

Physically, the Flex 5150 has a 1U form factor, meaning it has a minimal space requirement and can operate in an office environment. Its 14.55TB maximum usable capacity makes it an excellent fit for smaller edge locations, such as retail stores as well as remote offices and branch offices. Because of deduplication, the effective capacity of the device may be substantially more. The Flex 5150 is designed for easy upgrades as an organization or a location grows.

As an integrated appliance with NetBackup pre-installed, the Flex 5150 is designed to work seamlessly with a central copy of NetBackup (i.e., in the core datacenter). Flex 5150 can be managed entirely from the central NetBackup console. When numerous Flex 5150 appliances are deployed, policies can be pushed from the central location to all Flex 5150s simultaneously, ensuring that backup and retention policies are consistently applied. Data at the edge locations can optionally be copied to the central office or tiered to the cloud as secondary or tertiary copies; Veritas CloudCatalyst provides the cloud tiering functionality. However, if communications with the central office are not available, the Flex 5150 will continue to operate autonomously to process backup jobs as scheduled. When communications are restored, data will be automatically copied or tiered according to the established policies.

All of this may make the Flex 5150 sound very simple, but that is the intention of the design. Veritas has engineered the installation process to be completed in a matter of minutes with a single screen for application deployments. Complete NetBackup data protection (both master and server) is built in. The result is a system designed to provide seamless data protection from core to edge.

Challenges

All NetBackup appliances are highly dependent on NetBackup software, making them ideal for existing NetBackup customers who want to simplify or expand their backup infrastructure. However, for non-NetBackup customers, it requires buy-in for the entire package. This is both an obstacle and an opportunity for Veritas. These customers must choose to replace an existing backup vendor or add Veritas as another vendor. Nevertheless, most organizations have multiple backup vendors and choose to deploy them on an application or use case basis. Organizations with edge workloads, especially with numerous locations, may find an opportunity to simplify and improve backup operations with NetBackup appliances.

Conclusion

Edge workloads are generating more and more data, and much of that data is mission critical. Certainly, edge locations, such as retail and hospitality, cannot operate without data availability. It is important that edge locations have the same enterprise-class data protection capabilities that are in the core datacenter.

The old-school methods of protecting data at remote office and branch office locations using manual systems and processes are both costly and error prone. Purpose-built backup appliances are designed to be self-contained with simple deployment and centralized management. With the addition of cloud-tiering capabilities, organizations can deploy enterprise-class solutions across far-flung locations and deliver the highest level of data availability. The NetBackup appliances and Flex 5150 edge appliance represent an opportunity for widely distributed organizations to create just such an environment.

MESSAGE FROM THE SPONSOR

Veritas Flex 5150 Appliance

The Veritas Flex 5150 appliance is designed to bring NetBackup enterprise-class data protection the edge of the enterprise network. It is a complete NetBackup data protection solution in a self-contained, compact, easy-to-use appliance. The Flex 5150 appliance brings the following benefits to branch and remote offices at the edge of the network:

- Maximize performance right out of the box with optimized integration of NetBackup Software in a compact 1U form-factor including software and storage.
- Easily configure a data protection deployment in minutes with a menu of NetBackup services available in the Flex 5150 appliance.
- Remotely identify problems before they happen with AutoSupport that monitors Flex 5150 status and alerts a Veritas support team if a problem is detected.
- Save storage costs and efficiently tier to the cloud with NetBackup CloudCatalyst.
- Minimize downtime with built in redundancy.

For more information visit the ***Flex Appliance Webpage***.

About the analyst:

Phil Goodwin, *Research Director, Infrastructure Systems, Platforms, and Technologies Group*

Phil Goodwin is a Research Director within IDC's Enterprise Infrastructure Practice, covering research on data management. Mr. Goodwin provides detailed insight and analysis on evolving industry trends, vendor performance, and the impact of new technology adoption. He is responsible for producing and delivering timely, in-depth market research with a specific focus on cloud-based and on-premises data protection, business continuity and disaster recovery, and data availability.

IDC Custom Solutions

IDC Corporate USA

5 Speen Street
Framingham, MA 01701, USA
T 508.872.8200
F 508.935.4015
Twitter @IDC
idc-insights-community.com
www.idc.com

This publication was produced by IDC Custom Solutions. The opinion, analysis, and research results presented herein are drawn from more detailed research and analysis independently conducted and published by IDC, unless specific vendor sponsorship is noted. IDC Custom Solutions makes IDC content available in a wide range of formats for distribution by various companies. A license to distribute IDC content does not imply endorsement of or opinion about the licensee.