

WHITE PAPER

The Evolution and Value of Purpose-Built Backup Appliances

Sponsored by: Symantec

Robert Amatruda

March 2013

IDC OPINION

Today, customers are still grappling with subpar backup performance as backups outstrip the allotted backup window time. Customer strategies for data protection and recovery continue to be dictated by aggressive SLAs, rapid recovery, and ease of integration in existing environments. As a result, firms are embracing more disk-based data protection technologies, including purpose-built backup appliances (PBBAs) to protect and recover data and applications. These appliances include features such as data deduplication, compression, encryption, and replication. Meanwhile, unabated data growth continues to pressure IT staff and protection and recovery processes, leading customers to consider alternative backup methods and targets.

The evolution of backup processes has provided storage administrators some relief, moving from tape for backup and recovery to virtual tape libraries (VTLs) to using general-purpose disk with software and now PBBA systems utilizing storage optimization such as data deduplication. IDC has discovered that PBBA systems provide measurable benefits to customers looking to improve their existing data protection and recovery processes. Furthermore, PBBAs, as the moniker suggests, are purpose-built, highly tuned, turnkey solutions that have preconfigured hardware, interconnects, interfaces, and software. Initially, PBBA systems were target based and still needed third-party backup application software to catalog, index, and schedule the backups from the host systems to the target appliance. Increasingly, customers demand integrated PBBA systems to provide a comprehensive solution that has backup software tightly integrated inside the system. Now, customers do not have to expend staff time, money, or resources to build a backup appliance from piece parts and test it prior to deployment. PBBAs are designed for ease of use and ease of deployment and mitigate the risks of using a solution that is cobbled together.

IN THIS WHITE PAPER

This IDC White Paper explores the increased use and adoption patterns of PBBAs, both integrated and targeted, and the utility these appliances provide to customers in their data protection processes. In addition, this White Paper illuminates the customer value that Symantec's Backup Exec and NetBackup appliances bring to the data protection and recovery process.

SITUATION OVERVIEW

The worldwide PBBA market comprises a wide array of products that are designed for backup and rapid recovery. PBBA solutions are designed for capacity optimization utilizing data deduplication that enables customers to manage their ever-growing volumes of backup data. We believe customers are gravitating toward deployment of PBBA solutions as a way to control their data protection capital and operational expenditures. Also, PBBA solutions provide customers a way to extend their existing investments in data protection hardware and software without having to change any existing processes. In many cases, a firm's investments in data protection software and hardware solutions are very costly and are not subject to rapid replacement cycles. Thus, PBBAs allow customers to have a great deal of coexistence with regard to software and hardware.

The worldwide PBBA market covers products designed for open system and mainframe environments. The open system PBBA systems are, in many cases, designed to supplant tape, while the mainframe PBBA systems are deployed with a large physical tape component. IDC defines a PBBA system as a standalone disk-based solution that utilizes software, disk arrays, server engine(s), or nodes that are used for backup and recovery. Furthermore, IDC defines and reports the following characteristics of the PBBA market.

PBBA Protocols and Interfaces

PBBA solutions can support multiple interfaces, including:

- ☒ **NAS (NFS/CIFS)**, which is used by a large number of external disk systems, devices, and backup applications. Data movement can be facilitated by using the Network Data Management Protocol (NDMP).
- ☒ **Virtual tape library (VTL)**, which uses Fibre Channel (FC)/iSCSI.
- ☒ **Others**, including Symantec's OpenStorage (OST) API as well as vendor-specific interfaces such as EMC's Data Domain Boost, HP's StoreOnce Catalyst, and Quantum's Accent. The vendor-specific interfaces facilitate storage optimization or deduplication across shared domains or media servers.

Furthermore, PBBAs fall into two distinct categories:

- ☒ **PBBA target systems.** These systems are used in conjunction with third-party backup software and designed to integrate with heterogeneous environments. PBBA targets require and do not include traditional data movement, scheduling, and/or media management components found in backup software.
- ☒ **PBBA integrated systems.** These systems are tightly integrated with backup software and have master or media servers built into the system to orchestrate the backup and movement of data to other systems or removable media such as tape.

PBBAs provide a multiprotocol, scalable, and flexible approach to data protection and recovery. Each PBBA system type has advantages. However, integrated PBBAs provide customers a comprehensive solution that includes the application software, allowing them to scale out their environment to add master or media servers easily. Additionally, the integrated PBBA solutions have a centralized management console, allowing for ease of management and implementation. Integrated PBBA solutions are highly tuned and provide customers a truly turnkey, simplified way to solve their backup and recovery challenges.

PBBA systems provide more value and utility to customers than solutions cobbled together with general-purpose disk and software. Typically, customers expend more money and time with do-it-yourself disk systems used for backup and recovery. PBBA target systems allow customers to utilize existing application software that has already been deployed to provide indexing, cataloging, scheduling, and movement of data to and from the host and PBBA system. In integrated PBBA systems, the application software is tightly coupled with the system to allow customers to rapidly deploy the system and perform backup procedures with the innate ability to index, catalog, and schedule the movement of data to the system. In addition, integrated PBBA solutions allow customers to extend their data protection and recovery infrastructure quickly and easily by allowing them to add media servers. Integrated PBBA solutions typically have inclusive licensing, saving customers time and money.

The worldwide PBBA market experienced exceptional growth in 2011, with revenue totaling \$2.4 billion, a 43.4% increase over the prior year. The exceptional growth is due to increased use and customer adoption of PBBA systems for backup, recovery, and storage optimization with the use of data deduplication. The worldwide PBBA market is growing robustly, with revenue expected to have topped \$3 billion in 2012.

We believe the customer drivers for increased investment in PBBA solutions result from the need to improve backup window time, to provide faster restore and recovery times, and to enable seamless integration with existing backup applications. As a result, customers continue to aggressively move away from tape-based backup and recovery processes. This trend will continue for the foreseeable future as customers take advantage of PBBA's rich feature sets, particularly for virtual server protection, rapid recovery, and data deduplication. In addition, PBBA systems are increasingly used for remote office/branch office (ROBO) backup and recovery. We have discovered that ROBO backup and recovery can be a challenge for many organizations because of the physical handling of removable media at the remote site. PBBAs provide replication and capacity optimization such as deduplication to make backup of data more efficient and painless.

The long-term outlook remains very strong for PBBA market growth and adoption. We expect total worldwide PBBA revenue will grow robustly at a five-year compound annual growth rate (CAGR) of 19.4%, totaling nearly \$5.9 billion by the close of 2016. Customers that need to augment their existing data protection and recovery infrastructure or to enhance their backup infrastructure are driving the growth of the PBBA market as they look to take advantage of the technology's rapid recovery, data deduplication, and integrated management features. We expect the overall PBBA market to grow five times faster than the external disk market and four times faster than the total general-purpose disk and software markets used for data protection and recovery.

In the past year, there has been robust growth in the integrated PBBA market. Integrated PBBA systems are increasing at twice the rate of target PBBA systems. This growth is due to the simplified deployment model with application software being built into the integrated PBBA system. Customers realize operational improvements and cost savings by not having to manage a separate management or reporting console. Integrated PBBA revenue increased by 107% annually in 2011 while target PBBA revenue increased by 38% for the same period. In the first half of 2012, integrated PBBA revenue increased by 46% while target PBBA revenue increased by only 18%. Furthermore, the total integrated PBBA market represented 20% of the total open systems PBBA revenue in 2010 and increased to 27% of total revenue at the close of 2011. In the long term, we expect over a third of the total open system PBBAs will be integrated PBBA solutions as customers adopt a more cost-effective deployment model for data protection and recovery solutions. Furthermore, integrated PBBA systems allow customers to extend their data protection and recovery framework without additional licensing or maintenance costs.

Impact of Virtualization on Data Protection and Recovery

The proliferation of virtualization has had an impact on how data and systems are protected, backed up, and recovered. Virtualization has evolved beyond test and development and is now primarily used for production applications. Thus, fast, reliable, and efficient backup and recovery is paramount. Increasingly, business-critical workloads such as ERP and CRM are being virtualized and are widely deployed in customer environments. The growth and adoption of virtual servers have exacerbated the problems associated with backup and recovery. Consequently, comprehensive and streamlined protection and recovery processes are more vitally important than ever. The adoption of virtualization on server infrastructure has raised several challenges, which include:

- ☒ Impact on host resources, limiting processing power to complete backups
- ☒ Limited host/network resources for backup
- ☒ Image-level recovery versus file-level recovery
- ☒ Hypervisor integration capabilities
- ☒ Explosion in storage capacity for backup data
- ☒ Network and/or I/O bottlenecks adversely impacting backups

A protection methodology using traditional backup, with a backup agent on each guest system, allows firms to use existing processes, incur limited incremental costs, and have the advantage of single file or object recovery. Typically, this is implemented as a backup-to-disk scheme. However, overhead associated with backup may warrant limiting the number of virtual machines per host. Another limiting factor of using traditional backup in a virtualized environment is supporting multiple guest systems with a single backup window. These approaches may not provide for image-level recovery of an entire virtual machine. Firms may need to utilize a service console or consolidate backup to gain any benefits in virtualized environments.

Firms that virtualize their environments need to consolidate backup by using either a backup agent at the service console or a consolidated backup method offered by one of the virtualization vendors. Either method is an alternative to the traditional approach of using a backup agent on each guest system. A virtualized protection strategy provides the tools needed for administrators to perform virtual machine-level recovery. Using a backup proxy server and VMware Consolidated Backup, firms can offload the physical host, reduce network traffic, and minimize backup windows.

Virtualization fundamentally changed server and storage architecture. Virtual servers are less protected than physical servers due to several factors: Easy provisioning of VMs has led to rapid growth and sprawl and inconsistent enforcement of data policies. Traditional data protection solutions have not worked well in virtual environments. More importantly, point products have not addressed customer virtualization needs; rather, they have only added to management complexity. When considering a data protection and recovery solution, customers need to be mindful of the following issues:

- ☒ Impact on host/network resources for backup
- ☒ Volume of data as well as the impact of that data on the host is an important part of storage architecture
- ☒ Image-level recovery versus file-level recovery
- ☒ Hypervisor integration capabilities

Additionally, customers are grappling with a multitude of different backup packages used for disaster recovery, ROBO protection, point-oriented VM software, and snapshot technology unique to a specific disk array vendor. Target PBBA systems don't resolve this without client deduplication or reducing redundant data at the source. Customers need tighter integration between the target PBBA system and backup application that allows for client deduplication, replication awareness, individual file recovery, load balancing, and disaster recovery. Some pertinent examples are Symantec's OpenStorage or EMC's Data Domain Boost. Customers using integrated PBBA solutions with client-side deduplication and application awareness will have more cohesive tools to solve the issues associated with explosive data growth and a rapidly growing VM server environment.

Data Deduplication Enabling More Efficient Data Protection

Customers continue to grapple with unabated data growth regardless of company size or application. We foresee this trend continuing with data growth fueled by new applications, the proliferation of virtualization, the creation of electronic document stores, document sharing, and the retention or preservation of digital records. Furthermore, customers have constrained IT budgets and need to reduce capital and operating costs. As a result, we see data deduplication as a technology that will accelerate storage efficiency and reduce costs, especially for customer environments that have physically constrained datacenters with limited infrastructure in terms of power, cooling, and floor space.

Additionally, deduplication addresses challenges associated with data management, backup, and network inefficiency. Deduplication reduces the data footprint, allowing firms to keep their staffing and budgets in line with explosive data growth. Furthermore, data deduplication technology can optimize available physical and virtual infrastructure by sending less data over local or remote network links. It can also improve service-level response times and help meet shrinking backup windows. Deduplication also makes use of random access media (disk), improving recovery times, data security, and reliability. Firms are deploying data deduplication in a number of places in the infrastructure stack to address these practical, real-world challenges. The benefits of deduplication include the following:

- ☒ **Driving down cost.** Deduplication offers resource efficiency and cost savings that include a reduction in datacenter power, cooling, and floor tile demands as well as storage capacity, network bandwidth, and IT staff.
- ☒ **Improving backup and recovery service levels.** Deduplication can significantly improve backup performance to meet limited backup windows. Deduplication technology also leverages random access disk storage for improved recovery performance compared with sequential access methods such as physical tape.
- ☒ **Changing the economics of disk versus tape.** Deduplication makes disk-based backup feasible for a wider set of applications. Tape has had a role in enterprise datacenters because of its economics and archival properties. However, price per gigabyte declines for disk when used with deduplication, resulting in disk costs that are equal to or less than tape costs.

SYMANTEC: MARKET LEADER IN DATA PROTECTION AND RECOVERY

As the market share leader in data protection and recovery software, Symantec continues to bring innovation to the data protection and recovery market with key enhancements to its flagship software products — Backup Exec and NetBackup. More than two years ago, Symantec added source (client-side) and target deduplication and distinguished itself from point product vendors by offering solutions with both physical and virtual server support. Symantec continues to enhance and refine its core data protection and recovery portfolio to provide its NetBackup and Backup Exec customers more flexibility, scale, and ease of use. Symantec is delivering on its plan to offer its customers a set of cost-effective and comprehensive tools to more efficiently protect and manage their data from cradle to grave. Symantec's enhancements to data protection provide customers a comprehensive approach using deduplication, snapshots, and physical and virtual protection to control capital and operating costs while decreasing the complexity of recovery. Exiting 2011, the most recent full year for which IDC data is available, Symantec remained the market leader in the data protection and recovery software market with a nearly 30% share.

Two years ago, Symantec entered the PBBA market with a series of integrated PBBA solutions built around its Backup Exec and NetBackup software products. Today, Symantec has a broad portfolio of integrated and target PBBAs to meet the needs of

midmarket customers and enterprises looking to update their infrastructure and simplify their data protection and recovery environment. Symantec's integrated PBBA solutions represent a new and complementary set of offerings in the data protection and recovery market. Overall, the PBBA market is growing rapidly, and Symantec holds a leading position in the data protection and recovery software market as well as the know-how and expertise to provide customers with a comprehensive set of tools to solve their data protection and recovery challenges. More importantly, Symantec's integrated PBBA systems are a new delivery vehicle for Backup Exec and NetBackup software for customers that need to cost-effectively extend their infrastructure, and the systems provide a strong value proposition for new customers looking to easily deploy new instances of NetBackup and Backup Exec.

Symantec integrated PBBA solutions include:

☒ **Symantec Backup Exec appliance.** Backup Exec 3600 is a simple, effective, and low-risk solution.

- ❑ Appliance form factor enables Symantec to penetrate larger SMBs and midmarket customers that Backup Exec 2012 would not address with a single instance of software alone. The Backup Exec 3600 customer value propositions are reduce risks, save money, and simplify the backup and recovery processes.
- ❑ Specs: 1U form factor supporting 2 x 80GB 2.5in. SSD disks for OS (RAID 1), 4 x 2TB hot-swappable 3.5in. HDD (RAID 5) with 5.5TB usable capacity, and SAS card for tape out.
- ❑ Common hardware platform with two software licensing configurations to meet specific use cases. Both include unlimited licensing for physical and virtual system, application, and database protection up to the capacity of local storage. The higher-end configuration includes replication, 10-drive tape library support, and central administration to permit management of multiple Backup Exec appliances and/or instances of Backup Exec software from a single console.
- ❑ Hardened and secure, includes Symantec Critical System Protection

☒ **Symantec integrated NetBackup appliances.** The NetBackup 5200 series appliances are designed for enterprise scalability and performance.

- ❑ Specs: The NetBackup 5200 series provides 4TB of usable capacity in a node-based configuration that is expandable to 72TB. The 5200 series features built-in source and media server deduplication. Each node is a 2U device, while the optional storage shelf is 3U and comes in 24TB or 36TB configurations. Both feature redundant power supplies, redundant fan modules, and hot-pluggable disks. The NetBackup 5200 series supports both 1Gb and 10Gb Ethernet ports as well as 8Gb Fibre Channel ports.

- ❑ NetBackup 5200 series provides a tape out option via a gigabit dual channel Fibre Channel host bus adapter. This option is ideal for organizations that wish to use a hybrid approach for data protection consisting of disk and tape backups.
- ❑ The Symantec NetBackup 5200 series provides multiple use cases including remote office backups, virtual machine backups, and datacenter backups. The built-in source-side deduplication allows bandwidth-efficient centralized backups of remote offices while reducing dependence on tape and unskilled IT personnel at remote offices.
- ❑ Scalable, cost-effective, turnkey approach to data protection and recovery.
- ❑ Allows customers to control operating costs.

Symantec target PBBA solutions include:

- ☒ **Symantec target NetBackup appliances.** The NetBackup 5000 series appliances are complete deduplication solutions that provide the benefits of NetBackup's deduplication technology in an easy-to-use appliance form factor. An open and flexible architecture allows midsize and enterprise environments efficient long-term retention of data on disk, providing faster and more reliable backup and recovery of data.
- ❑ Specs: NetBackup 5000 series appliances can scale up in increments of 32TB to a maximum usable deduplication capacity of 192TB. Each NetBackup 5000 series appliance node is a 4U device and features redundant power modules, redundant fan modules, and hot-pluggable disk modules. The NetBackup 5000 series supports both 1Gb and 10Gb Ethernet ports as well as 8Gb Fibre Channel ports.
- ❑ Scalable, cost-effective, turnkey approach to data protection and recovery.
- ❑ Allows customers to control operating costs.

FUTURE OUTLOOK

The long-term outlook for PBBAs appears to be very robust for the next five years. We expect total PBBA revenue will grow at a five-year 19.4% CAGR, totaling nearly \$5.9 billion in 2016. The open systems portion of the PBBA market is expected to grow at a 21.2% CAGR. We believe that integrated PBBA systems will be the engine of growth as customers look to deploy turnkey, easy-to-use, and easy-to-manage solutions. In addition, we expect PBBA solutions, particularly integrated PBBAs, to provide channel partners with a vehicle to add more value-added services such as hosted services and cloud options. This model will enable channel partners not only to upsell but also to position themselves as trusted advisors to customers.

We expect vigorous customer adoption of PBBA solutions for the foreseeable future. IDC believes Symantec will continue its market leadership as a result of delivering solutions today that:

- ☒ Demonstrate the utility and value of using PBBA solutions in concert with existing data protection infrastructure — software and hardware.
- ☒ Provide nondisruptive integration of PBBA solutions to customers' existing data protection and disaster recovery frameworks so that customers need not change any procedures and policies.
- ☒ Support seamless data movement from PBBA systems to physical tape to support archive and disaster recovery.
- ☒ Support industry-standard interfaces, APIs, and application software to ease deployment.
- ☒ Add features that optimize storage efficiency and security and provide a more cohesive disaster recovery approach such as data deduplication, encryption, or replication.
- ☒ Support virtualized server environments enabling customers to use the same tools to safeguard both physical and virtual infrastructure.
- ☒ Allow for cloud-enabled storage models.

CHALLENGES/OPPORTUNITIES

In the near term, Symantec's greatest challenge will be established systems or hardware vendors aggressively holding or maintaining their share positions. Symantec may encounter very aggressive pricing from competitors or significant bundling of hardware, software, and compute infrastructure. In addition, Symantec must put greater emphasis on education and awareness of the benefits that integrated PBBA solutions bring to a customer environment. This would include demonstrating the cost savings that companies can realize by using an integrated solution rather than building, testing, and deploying heterogeneous backup and recovery solutions by themselves. Further, Symantec is subject to macro-economic pressures that could adversely impact the company's ability to maintain a solid growth trajectory in the PBBA market.

CONCLUSION

Symantec has extended its leadership position in the data protection and recovery software and hardware markets with continued innovation in full-featured, integrated PBBAs. The company boasts an extensive set of backup and recovery solutions that span a broad array of customers, company sizes, and use cases. The market success that Symantec has achieved with its PBBA solutions validates its strategy of pursuing a new and innovative approach to solving customer data protection challenges in a simplified, cost-effective manner. Furthermore, Symantec's emphasis on integrated PBBA solutions provides greater value and utility to help customers manage their data protection challenges today and in the future. We believe the

company is extending its core competencies and participating in an adjacent market that is proving to be higher growth than the traditional on-premise, packaged data protection and recovery market. Symantec is well positioned to be an industry leader in the integrated PBBA market.

Copyright Notice

External Publication of IDC Information and Data — Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2013 IDC. Reproduction without written permission is completely forbidden.