OPENBENCH LABS

Data Center Management

PERFORMANCE Analysis: Unified Backup: Virtual & Physical

Dell AppAssure 5
Backup, Replication, and Recovery Software
vs.
Symantec Backup Exec 2012

PERFORMANCE

Analysis: Unified Backup: Virtual & Physical

Dell AppAssure 5
Backup, Replication, and Recovery Software
vs.
Symantec Backup Exec 2012

Author: Jack Fegreus, Ph.D.
Managing Director
openBench Labs
http://www.openBench.com
August 10, 2012

Table of Contents

Executive Summary

Why read this document?

If you're currently using or considering either Symantec Backup Exec or AppAssure software to protect your Windows environment, you may already be aware that both products were reintroduced this year in significantly enhanced forms.

openBench Labs is pleased to be the first independent test organization to compare them so you can make an informed purchase decision. We pitted the performance of Symantec Backup Exec 2012 software (introduced in February) versus AppAssure 5 (released in June) in a well-appointed Windows lab environment. We found that while each offering has focused on ease-of-use through the addition of a new interface and has substantially enhanced its capabilities over its previous version, customers will judge these products primarily by how well they answer three time-tested questions:

- **How easy is it to use?**
- **How fast can it complete a backup, with all the functionality you need for recovery?**
- **What does it take to recover when you need it most?**

Findings & Benefits

In testing, we found AppAssure offered significant advantages in ease-of-use, backup speed, and recovery time. Here are several key takeaway findings and their impact on IT operations:

- **Finding: More backups mean more recovery point choices.**

AppAssure's incremental forever image-level backup technology can be set to execute every five minutes for 288 recovery points a day. More importantly, data transfer optimizations automatically set by AppAssure Smart Agents make this a deliverable level of performance for many transaction processing applications.

Benefit: IT operations managers can practically eliminate the potential for even minor data loss by creating very closely spaced recovery points for frequently updated mission-critical applications like Exchange, SQL and SharePoint.

- **Finding: Backup processing time for VMs hosted on vSphere consistently faster.**

With True Global Deduplication and Recovery Assure™ active, AppAssure was 5X faster executing a full backup of a VM running Exchange with 417GB of active data than Backup Exec with PureDisk deduplication, file and data-item Granular Recovery Technology (GRT), and no verification testing.

Benefit: Backup windows are effectively eliminated and every single backup is verified automatically for higher application recovery assurance.

• **Finding: Pipelining replication and recovery processes enhances fast recovery options with more recovery points.**

We know of no data protection software that matches AppAssure's ability to maintain a warm standby VM for either a virtual or physical production server while maintaining a five-minute RPO in addition to a five-minute RTO.

Benefit: Administrators have more recovery options with outstanding speed to meet today's increasingly stringent service level agreements.

• **Finding: One backup pass, not two, provides full protection for all VMs.**

With BE2012, common VM protection scenarios can require running two separate backups: once as a VM and once as a legacy server using file-based backup. AppAssure 5 requires only one backup pass to do the same job.

Benefit: Administrators spend less time on administration and less disk storage is required.

• **Finding: Near-zero time recoveries perceived by users.**

The most magical and impressive Smart Agent™ technology is without doubt AppAssure Live Recovery™, which provided AppAssure with a strong performance advantage in test after test.

Benefit: Users perceive a nearly instant recovery of a Windows application regardless of how data they store.

Building on Block Virtualization

"By implementing data deduplication as a standard feature, AppAssure 5 is able to transcend the issues of incremental, full, and synthetic backup and directly leverage the virtual block space created by applying True Global Deduplication to the data blocks transferred by Smart Agents installed on protected machines."

Under Test: Windows Server Data-Protection

Dell AppAssure 5 Value Proposition

1) **Maintain a Virtual Disk Block Space of all Protected Systems:** AppAssure 5 maintains a repository of virtual disk blocks rather than an archive of backup sets. ***Benchmark: With data deduplication, a full backup of a VM with 417GB of active data was 5X faster with AppAssure 5 than Backup Exec 2012.***

2) **Protect Physical and Virtual Machines with Changed-Block Tracking using Smart Agent™ Technology:** An AppAssure Smart Agent installs a filter driver to track all data changes by changes to disk blocks.

3) **Minimize Physical Storage Using True Global Deduplication:** AppAssure 5 stores a virtual space of globally deduplicated disk blocks of all protected systems in an object-oriented repository without the overhead level of an RDBMS. ***Benchmark: We averaged 20X storage savings on each of 5 protected virtual machines yielding a 100X reduction over traditional storage methods.***

4) **Simplify Application Validation with Recovery Assure™:** Smart Agents automatically discover key applications and configure default recoverability tests.

5) **Recover Files and Application Items via Universal Recovery™:** Retrieve files or individual data objects, such as an email messages and database tables, ad hoc by mounting and sharing any backup as a virtual disk.

6) **Minimize DR restore Time with a Standby VM:** Universal Recovery allows any backup to be used to automatically update system and data files on a warm-standby VM or to burn a hot-metal backup CD for any hardware platform.

7) **Provide Near-Zero Data Recovery Time with Live Recovery™:** Smart Agents reorder blocks transferred during recovery giving users immediate access to data. ***Benchmark: Using Live Recovery we restored an Exchange mailbox database for 500 users while passing a heavy-use Jetstress benchmark.***

8) **Minimize Data Loss with Near-Continuous Recovery Points:** AppAssure leverages changed block tracking to provide fast incremental backups, which can be scheduled in five minute intervals.

9) **Pipeline AppAssure Cores using True Scale™ Architecture:** IT can readily scale data protection processes by configuring multiple AppAssure Core servers in a pipeline to scale a data protection process linearly and enhance off-site availability of warm-standby virtual servers using a WAN for replication.

In this new Lab Report, openBench Labs examines Dell AppAssure Backup, Replication and Recovery software 5 with Symantec Backup Exec 2012. For protected physical and virtual machines (VMs), AppAssure implements changed-block tracking technology (CBT) to provide data protection for all servers running supported operating systems. More importantly, AppAssure 5 complements CBT with global deduplication and an object-oriented (OO) file system to create a unique repository architecture that presents IT administrators with a virtual disk space representing all virtual and physical systems protected by the repository's Core server.

For IT operations, the new OO architecture of AppAssure 5 provides a highly scalable infrastructure for a powerful "forever incremental" backup scheme that eliminates all of the traditional overhead for an incremental backup regime including both full and

synthetic backups. In contrast, Backup Exec 2012 uses a traditional file-based backup scheme to promote compatibility with the largest number of systems and only leverages CBT in a Virtual Infrastructure (VI).

AppAssure utilizes its own CBT technology to provide IT administrators with the same advanced features for a VM or a physical system. Key AppAssure features include True Scale™ Architecture to enable pipelining of Core server protection processes, True Global Deduplication to minimize the amount of repository data, Live Recovery™ for near-zero downtime; Recovery Assure™ for verifying object level recoverability of files and application-items in any backup; and Universal Recovery™ for cross-platform recovery of data volumes, including bare metal recovery, from physical to virtual (P2V), virtual to virtual (V2V), virtual to physical (V2P) or physical to physical (P2P) systems.

openBench Labs Data Protection Testing Summary

Feature	Key Test Analysis	Dell AppAssure 5	Backup Exec 2012
Uniform Execution for Physical and Virtual Systems	•Uniform *execution of all commands from a single GUI lowers operation complexity*	•All AppAssure Core servers operate from a single GUI uniformly across all clients	•VM-based backups do not support synthetic or forever incremental backup schemes
Global Data Deduplication	*•D2D backup file size directly affects resource costs*	•Global data deduplication standard •Only stores deduplicated disk blocks •Highest data reduction ratio	•Data deduplication optional •Stores deduplicated data blocks and backup process files •Deduplication extends verification and restore
Backup Execution Performance	*•Minimum time to back up governs frequency and bounds an obtainable RPO*	•Fastest backups with deduplication •Parallel backups scale on single and multiple systems	•Shared common database required to centrally control parallel backups
Granular Recovery of Files and Application Objects	*•Recovery of application data objects localizes disruption of production processes*	•Discovers applications and adds validation tests automatically •Status of backed up data objects displayed in GUI •Mount and share a full volume associated with any restore point for application-level data recovery	•Scanning files during GRT catalog updates significantly extends backup processing for vSphere VMs •File-based synthetic backups exclude application-specific GRT •In VMs, SQL Server transaction logs truncated only using file-based backup
VMs and Near-zero RTO Disaster Recovery	*•VM clones of virtual and physical systems minimize RTO* *•Highly optimized proprietary features restore data with no perceived delay*	• VMs automatically configured as identical to production servers •Instantly restored an Exchange mailbox database while passing a Jetstress benchmark	• VMs must be rebuilt entirely, as continuous updates are not supported •No specialized features for near-zero DR are provided

Especially at small to medium business (SMB) sites, IT administrators often struggle to support two very different data protection schemes: a block-based solution for VMs and a file-based solution for physical systems. AppAssure's Smart Agent technology avoids this by providing a unified approach to data protection for all physical and virtual systems at a site. In this way, AppAssure holds down labor costs while enabling IT to support unmatched Recovery Time and Recovery Point Objectives (RTO and RPO).

Testing Overview

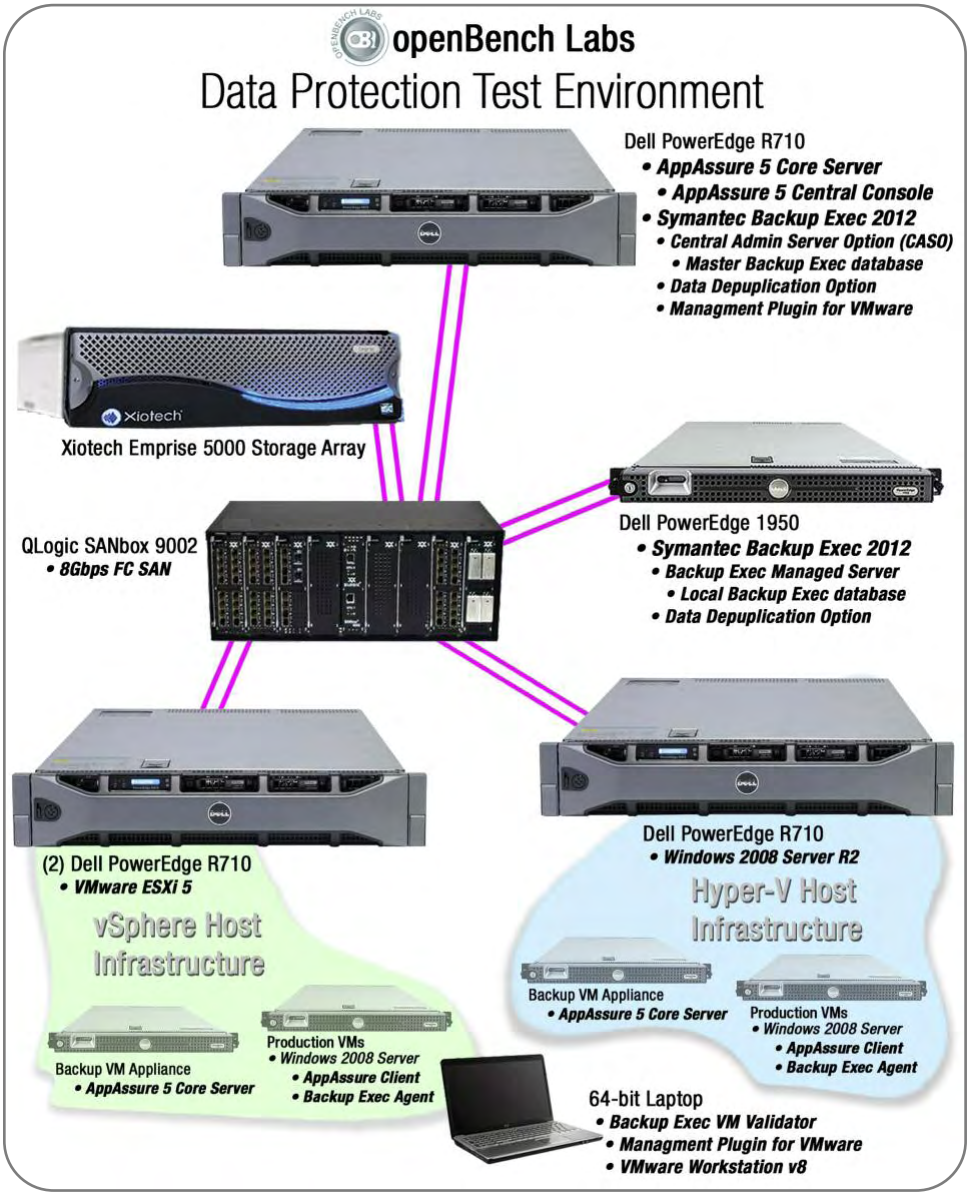

OpenBench Labs conducted tests on Dell AppAssure 5 Backup, Replication, and Recovery software and Backup Exec 2012 using a test bed populated with five physical servers, one laptop workstation, and two VM backup servers. We centered physical systems around a Dell PowerEdge R710 server, which was provisioned with dual 6-way core processors. This system served as the primary AppAssure 5 Core server running AppAssure Central Console and the primary Backup Exec 2012 server running the Central Admin Server Option (CASO).

For Backup Exec, which extensively leverages underlying Storage Area Network (SAN) resources, we used an 8Gbps Fibre Channel (FC) SAN to provision storage resources. To enhance Backup Exec scalability, we utilized quad-core Dell PowerEdge 1950 servers with 8GB of RAM and dual FC ports. On each server we installed a Backup Exec CASO managed

server, including options for data deduplication and VM client support. We also deployed a 64-bit laptop running VMware Workstation to boot VMs from Backup Exec backup files and validate recoverability.

For AppAssure, which is network-centric, we utilized 10GbE and 1GbE. In particular, we created virtual equivalents to the PowerEdge 1950 servers on each vSphere and Hyper-V host. We configured each VM-based AppAssure Core server with four CPUs, 8GB RAM, and an internal 10GbE virtual NIC.

"Both Backup Exec 2012 and AppAssure 5 implement global deduplication of the underlying disk block space of protected machine; however, the way in which each of these virtual block spaces is utilized is dramatically different and that results in equally dramatically different outcomes."

Within this environment, we set up data protection tests designed to assess each package's ability to add value through the simplification of IT management tasks and the reduction of operating expense (OPEX) costs. In particular we focused on three critical goals for IT operations:

- **Reduce backup time**
- **Speed recovery**
- **Reduce backup storage**

A critical success factor to achieving those goals is the ability to implement a highly optimized disk-to-disk (D2D) backup strategy that encompasses incremental backups of virtual and physical servers. A D2D strategy provides IT with a number of operating efficiencies versus disk-to-tape (D2T); however, an overall data reduction ratio of 20:1, which is equivalent to a 95% average storage reduction through data deduplication, is necessary for a D2D scheme to be cost effective.

From Meta Data to Real Data

To optimize a D2D scheme, data deduplication is an essential component. Both Backup Exec 2012 and AppAssure 5 implement global deduplication of the underlying disk block space of protected machine; however, the way in which each of these virtual block spaces is utilized is dramatically different and that results in equally dramatically different outcomes.

With Backup Exec, PureDisk data deduplication is an optional component that is seen as a means to enhance a traditional scheme that archives each full and incremental backup file as a data protection entity. When a backup is processed, the data in the backup file is scanned for sequences of block patterns that have been previously discovered and saved as meta data. If a disk block pattern is matched, and a pointer is generated to the location of the existing data string and the pointer is substituted for the original string. This means that whenever a validation or a restore process is used in conjunction with a deduplicated backup file, that file must go through a process that rehydrates all of the original data.

"By essentially reversing the CBT-based backup process to accelerate recovery and support an aggressive RTO, AppAssure 5 entirely changes the dynamics of recovery, as well as backup."

To implement PureDisk global data deduplication, Backup Exec creates a specialized data deduplication volume that works somewhat like a virtual tape library. More importantly, each Backup Exec server can implement only one data deduplication volume. That restriction limits the ability to scale the capacity of deduplicated volume to the ability to extend a dynamic disk volume.

By implementing data deduplication as a standard feature, AppAssure 5 is able to transcend the issues of incremental, full, and synthetic backup and directly leverage the virtual block space created by applying True Global Deduplication to the data blocks transferred by Smart Agents installed on protected machines. Furthermore, by using an OO file system and standard OO techniques to underpin Core server repositories, AppAssure 5 assigns extensive characteristics, including recovery point information, to stored data blocks, and eliminates the need to save backup files.

What's more, the scalability of an AppAssure Core repository is just as impressive as its functionality. The encapsulated OO file system is built and extended by assigning standard NTFS directory files resident on up to 256 different volumes to a repository. That makes provisioning an AppAssure 5 repository an exercise in optimizing minimal, rather than maximal, increments in storage.

For IT operations, the payoff from the AppAssure repository starts with automatic support for a fast "incremental forever" backup scheme that is free of synthetic backups. Moreover, storage reduction ratios on the order of 20-to-1 are typically generated for each protected machine. Furthermore, since separate backup files are not saved for each protected machine, overall storage savings for a Core server is the sum of the savings calculated for each protected machine. On the AppAssure Core servers we established on VI hosts, overall storage savings typically ranged from 75- up to 100-to-1.

Our main concern, however, was recovery. In a business continuity context, backup is simply a means to an end.

Mitigating the amount of processed data likely to be lost when recovering from a computer outage requires closely spaced recovery points, which intrinsically links an aggressive RPO to the minimum time required to complete a backup. As a result, minimizing data loss during recovery brings us back full circle back to leveraging incremental backups. More importantly, any work to provide an aggressive RPO can be quickly negated by a prolonged recovery process.

Fast backup must be balanced with fast reliable recovery. By essentially reversing the CBT-based backup process to accelerate recovery and support an aggressive RTO, AppAssure 5 entirely changes the dynamics of recovery, as well as backup. Specifically, AppAssure 5 leverages the disk block repository and Smart Agents to roll data to

protected machines in a Live Recovery process that continuously reorders the blocks transmitted back to the protected machine based on meeting I/O requests of users.

Email Checking

To test and verify the interaction of Smart Agents with the OO repository on backup and restore operations, we set up multiple VMs running Windows Server 2008 R2 and various key applications, including Active Directory, SQL Server 2008 R2 and Exchange Server 2010 on vSphere and Hyper-V hosts.

Exchange VM Environment Baseline Performance

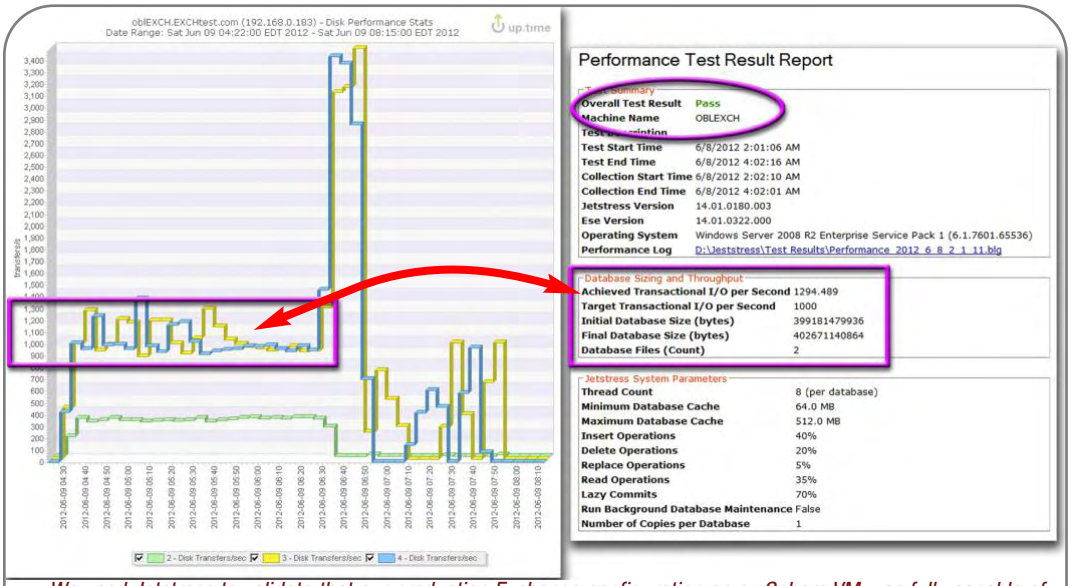

We used Jetstress to validate that our production Exchange configuration on a vSphere VM was fully capable of sustaining a heavy 1,000 TPS load. In our analysis of data protection performance, we focused on the ability to maintain this baseline level of heavy email processing in a backup scenario, and more importantly to restore processing back to this level in a DR scenario.

We made extensive use of a VM running Exchange 2010 and the Jetstress benchmark. To provide significant I/O stress, we configured an enterprise-class Exchange server with 1,000 email accounts using dedicated virtual disks: for email databases and Exchange transaction logs.

For benchmark analysis, email transactions are distributed into four groups as follows: 35% read, 20% delete, 5% replace, and 40% inserts. A normal email processing load is defined as 1 transaction per second (TPS) for 50% of the user mail boxes on a server, with an average response time of under 20ms. A heavy processing load extends the email load to 1 TPS to include all user mail boxes on a server. In our scenario, a heavy email activity load translated to 1,000 TPS spread over two Exchange mailbox databases.

Backup Performance

> "With True Global Deduplication and Recovery Assure active, AppAssure was 5X faster executing a full backup of a VM running Exchange with 417GB of active data than Backup Exec with PureDisk deduplication, GRT, and no verification."

Backup Software Architecture Differences

A central issue for comparing AppAssure with Backup Exec is the scope of the two products. AppAssure is designed for use with virtual or physical systems that are using its CBT technology. To set up AppAssure's generalized disk-block tracking mechanism, IT administrators install a Smart Agent for the AppAssure service on every client system to collect, package, and send block data back to an AppAssure Core server.

AppAssure Smart Agent Overhead

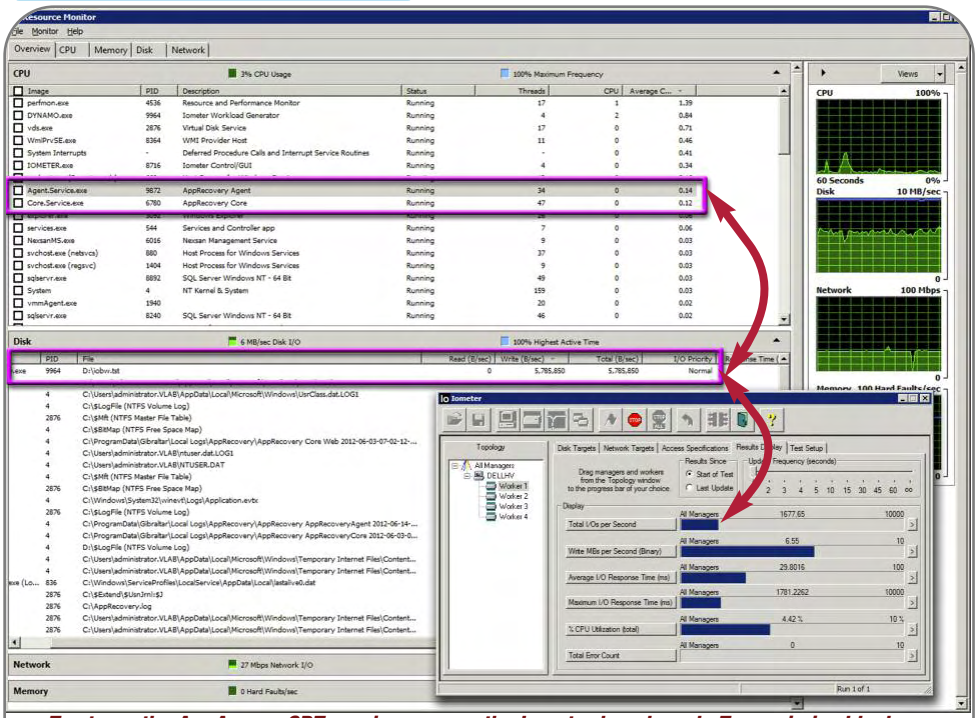

To stress the AppAssure CBT service, we ran the Iometer benchmark. To maximize block changes, we used 4KB data operations to write 1,700 blocks per second randomly over a disk. During this test, the AppAssure service processes consumed just 0.25% of CPU resources.

The AppAssure Smart Agent places a filter driver between the Windows file system and the Windows kernel to capture block-level changes on the underlying logical disk volume associated with data changes in files made by applications. Using the filter driver, AppAssure Smart Agents on client systems log meta data about disk-block changes for each logical disk. During a backup, the agent utilizes the meta-data logs to send an AppAssure Core server all of the block-level data changes that have occurred since the last backup.

The block data that is sent to an AppAssure Core server goes through a global data deduplication process—dubbed True Global Deduplication—and is maintained in a

repository built on an object-oriented file system. The AppAssure repository immediately incorporates the new block data tagged as volume recovery points and minimizes backup windows by creating a forever incremental backup scheme that enables an aggressive RPO and never requires a synthetic full backup.

AppAssure Unified Client Protection GUI

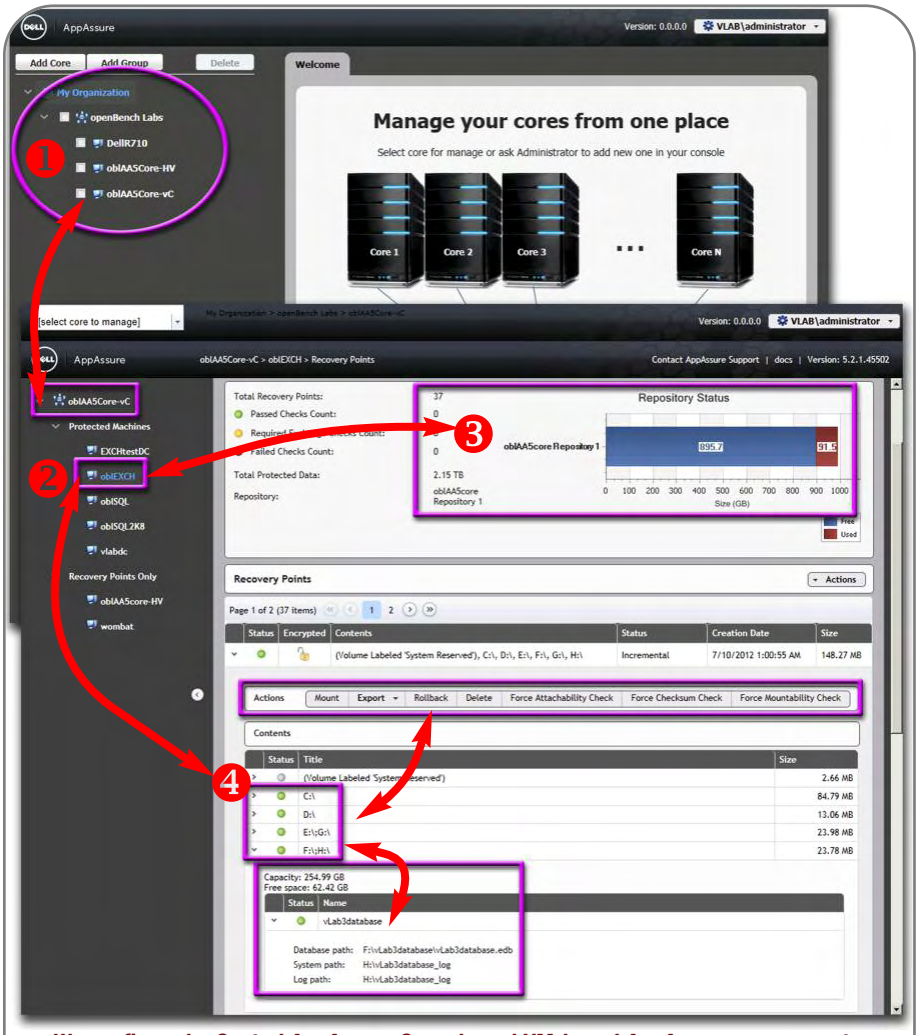

We configured a Central AppAssure Console and VM-based AppAssure servers on two ESXi and one Hyper-V host❶. On an ESXi-based AppAssure server, we protected a VM running Exchange❷. With respect to the Exchange server VM, 61GB of the 91.5GB in the core's repository represented 37 recovery points protecting 2.15TB of data❸. More importantly, at each recovery point AppAssure Smart Agents grouped related volumes and displayed their state as a single resource❹. In particular we drilled down and examined the state of specific mailbox data stores and their related log volumes and invoked commands to test and recover these volume groups as single resources.

To further control operations overhead, Smart Agents immediately check a new machine protected by an AppAssure Core server to discover protected applications. When an application is found by a Smart Agent, AppAssure implements custom verification tests, which may performed on every backup or on the final backup of the day for more highly involved tests.

In addition, AppAssure customizes the GUI to display the new clients configuration options. Not only are the disks used by the application grouped, but the Smart Agents drill down into the application and will create subgroups of disk volumes that will need to be treated as a unified resource. On our Exchange server, disk volumes used to protect a mailbox database and the logs for that database were grouped as separate resources. Meanwhile, the AppAssure GUI tailors the presentation of the total of the unique blocks of data in the repository as a set of recovery points for each protected machine. While presenting the total number of blocks in the repository as a representation of the recovery points for a protected machine is accurate, it is also very

conservative with respect to actual data reduction savings for that machine.

When we replicated the 37 recovery points for our Exchange server VM in our repository for all VMs, only 61GB of the 91.5GB in that repository were replicated. As a result, the real data reduction ratio for our Exchange server VM was 36 to 1. At the same time, 30.5GB of block data in our repository was providing 20-to-1 data reductions for four different VM servers, each of which had about 35GB of active data. For this group of four servers, the repository yielded a 92-to-1 reduction over traditional backup file storage.

What's more, backup processing time for VMs hosted on vSphere was consistently faster using AppAssure. With True Global Deduplication and Recovery Assure active, AppAssure was 5X faster executing a full backup of a VM running Exchange with 417GB of active data than Backup Exec with PureDisk deduplication, GRT, and no verification.

Backup Exec Agent Trust

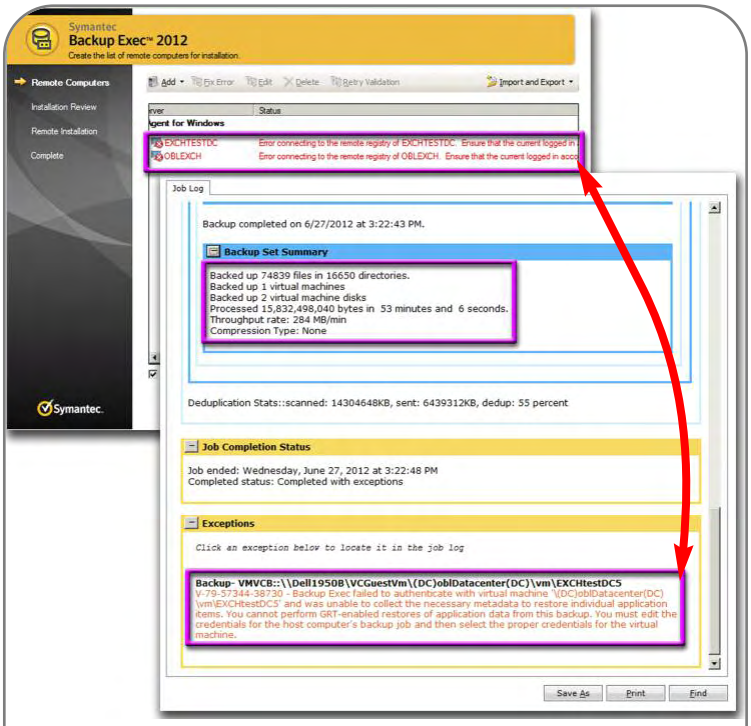

Using Backup Exec, a remote agent must be installed on a VM in order to establish a trust relationship with every Backup Exec server that will run a backup of the VM. More importantly, to install the agent, our Backup Exec server had to be in the same Windows domain as the client. We could not install the agent across trusted domains. More importantly, for full GRT functionality, the client also needed to be backed up by a server in the same domain. When we backed up a Domain Controller from a Backup Exec server in a different domain, GRT cataloging ran for over 45 minutes, however only GRT file-level recovery was available from the backup. Active Directory items were excluded from the GRT recovery features.

What really makes AppAssure 5 radically different as a data protection technology, however, is the aggressive exploitation of CBT across multiple data protection operations, not just backup. AppAssure leverages CBT in three key restore technologies: Live Recovery to meet near-zero RTO; Recovery Assure to verify that a backup, including the application data within a backup is recoverable; and Universal Recovery to support granular recovery of files, application items, as well as P2P, V2V, P2V, and V2P recovery options. As a result, AppAssure forges measurable advantages in our recovery scenarios.

Backup Exec, which was created for a legacy environment with physical systems, could not make assumptions for advanced features based on data structures exclusive to a Virtual Infrastructure (VI). More importantly, advanced features, such as data deduplication, Granular Recovery Technology (GRT) and the ability to scale and manage multiple servers via CASO are dependent on Backup Exec trust relationships that need to work with physical and virtual systems. While IT administrators can easily perform an agentless backup of a VM with Backup Exec 2012, an agentless backup with

Backup Exec 2012 precludes the ability to invoke advanced recovery functions.

With AppAssure, Smart Agents automatically discover protected applications, such as Exchange and SQL Server, and configure specialized recovery validation tests before the first backup. The AppAssure Core server is always aware of special applications resident on a protected machine. On the other hand, Backup Exec servers must scan and discover special applications on every backup for GRT processing.

The heavy-weight trust relationship required between Backup Exec servers and clients proved problematic in our test bed, which encompassed two mutually trusted Windows domains. We were unable to install an agent on a client that was not in the same Windows domain as the Backup Exec server. As a result, we needed to install a full Backup Exec server in each Windows domain we intended to protect with GRT.

SQL Server Log Issues

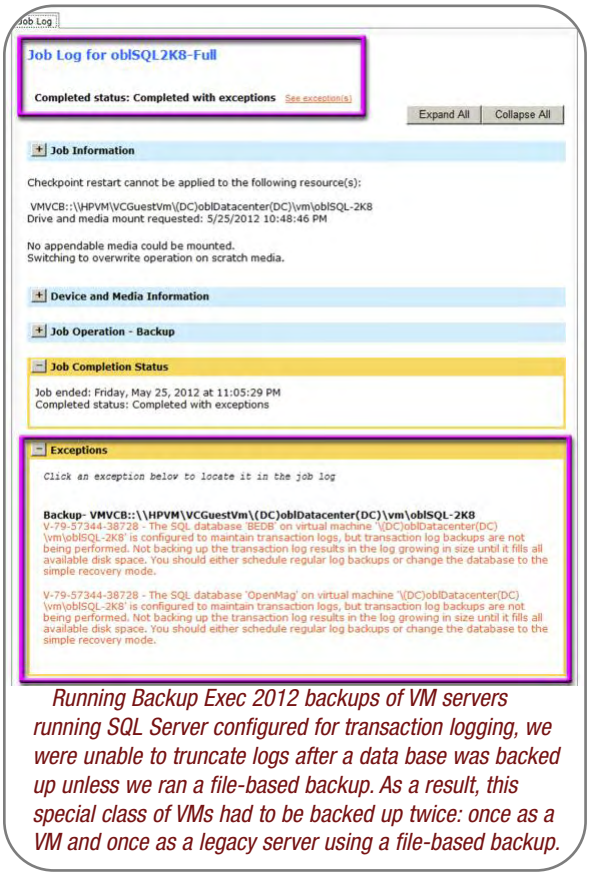

Running Backup Exec 2012 backups of VM servers running SQL Server configured for transaction logging, we were unable to truncate logs after a data base was backed up unless we ran a file-based backup. As a result, this special class of VMs had to be backed up twice: once as a VM and once as a legacy server using a file-based backup.

Once a Windows agent was installed on a client system, however, we were able to leverage existing Windows domain trusts to configure a Backup Exec trust relationship between the Windows client and any Backup Exec server. With both the Windows domain trust and Backup Exec 2012 trusts in place, we were then able to run a GRT-enabled backup, albeit not all GRT functionality could not be invoked. Only general file-level data was available from our GRT backup. We could not access data items associated with Active Directory.

We also ran into agent limitations with Backup Exec 2012 running backups of VMs with SQL Server installed. While AppAssure has a has a standard default to truncate database logs when validating a backup of a VM or physical system running SQL Server, Backup Exec has no complimentary feature for a VM backup.

Backup Exec categorizes backups of a VM running SQL Server with transaction logging as successful with an exception. To truncate the database transaction logs, we needed to run an additional file-based backup of every VM running SQL Server configured to use transaction logging in order to have an option to truncate SQL Server logs.

What's more, enabling GRT on backups of vSphere-hosted VMs consistently added about 32 minutes to either a full or incremental backup as the file directories were scanned and cataloged after the backup completed. On Hyper-V, this GRT cataloging process added about 2 minutes. This relatively fixed amount of time for vSphere-hosted VMs was

particularly troubling for incremental backups, which would otherwise complete in about 5 minutes without GRT. As a result, the minimum time for an incremental Backup of a vSphere VM using Backup Exec with GRT was about 35 minutes.

GRT functionality, however, is only one of many features that require support from Backup Exec agents. As a result, Backup Exec is constrained in dealing with the two contextual personalities of a VM, which can be viewed either as an application running on a host's hypervisor or as a logical system running an OS and a set of applications.

Single-Server Backup with File- & Object-level Restore
File & Application-Object Restore

Protected Machine Configuration	Backup Software Options	Backup Window (hh:mm:ss)
Physical Server Windows Server 2003 Active Directory PDC 30.5GB active data 250MB average change	AppAssure 5 Compress, Dedupe & Test	00:03:29 Full 00:00:19 Incremental
	Backup Exec 2012 Dedupe & GRT	00:37:30 Full 00:07:33 Incremental
vSphere 5 VM Windows Server 2008 R2 Active Directory PDC 19.8GB active data 250MB average change	AppAssure 5 Compress, Dedupe & Test	00:02:17 Full 00:00:28 Incremental
	Backup Exec 2012 Dedupe & GRT	00:57:58 Full 00:45:38 Incremental
vSphere 5 VM Windows Server 2008 R2 SQL Server 2008 R2 27.5GB active data 250MB average change	AppAssure 5 Compress, Dedupe & Test	00:02:04 Full 00:00:28 Incremental
	Backup Exec 2012 Dedupe & GRT	00:50:02 Full 00:34:48 Incremental
vSphere 5 VM Windows Server 2008 R2 SQL Server 2008 R2 37.1GB active data 250MB average change	AppAssure 5 Dedupe & Test	00:02:35 Full 00:00:36 Incremental
	Backup Exec 2012 Dedupe & GRT	00:51:59 Full 00:36:58 Incremental
vSphere 5 VM Windows Server 2008 R2 Exchange Server 2010 417GB Active Data 250MB average change	AppAssure 5 Compress Dedupe & Test	00:31:44 Full 00:00:56 Incremental
	Backup Exec 2012 Dedupe & GRT	02:48:40 Full 00:36:35 Incremental
Hyper-V VM Windows Server 2008 R2 SQL Server 2008 R2 29.7GB active data 250MB average change	AppAssure 5 Compress, Dedupe & Test	00:02:15 Full 00:00:37 Incremental
	Backup Exec 2012 Dedupe & GRT	00:06:07 Full 00:04:52 Incremental

In particular, IT can only implement synthetic full backups from generic file-based incremental backups of VMs or physical servers. For a VM backed up as a VM, the inability to create a synthetic full backup precludes the ability to implement a forever-incremental backup scheme. This makes it very important to verify incremental backups and terminate chains of incremental backups with a full backup as soon as possible in order to hedge against the loss of integrity.

Furthermore, Backup Exec 2012 Validator, which utilizes VM Workstation to test the ability to boot and recover a VM directly from a backup file, can only be used to validate a full backup. The only way to verify an incremental backup is to rehydrate the data and run a full checksum.

Recovery and Replication Performance

> "Not only were we able to use AppAssure 5 to maintain a VM running Exchange server with an aggressive RPO of 20 minutes and an RTO of 5 minutes, we did so while the original server processed 1,000 email transactions per second."

The Road to Zero RTO

For IT, backup has long been a necessary daily activity centered on backing up data, applications, and OS files in a minimal amount of time. From the perspective of a Line of Business (LoB) executive, however, the value of IT data-protection operations rests entirely in recovery processes. As a result, data-protection Service Level Agreements (SLAs) between IT and LoB units focus on RPO and RTO targets and leave backup as simply a means to provide a recovery point.

More importantly, LoB executives are beginning to link computer downtime to more than lost sales revenue: They equate computer outages to potential losses in customer confidence and market share. That's why growing numbers of senior executives expect IT to meet RTOs for critical business applications measured in minutes. With traditional system restore techniques taking hours to complete, the challenge for IT decision makers is to assuage the business-continuity fears of corporate executives within budgets that exclude capital expenditures on costly components, such as fault-tolerant servers.

Next, we measured the time needed to extract and rehydrate VM data files from a backup file, write the data to a vSphere host, and register the VM. The results, consistently paralleled the original backup time. A traditional recovery test of a VM server running Exchange 2010 using Backup Exec 2012 took 2 hours and 44 minutes, to prepare a new VM that was ready to boot into the production environment.

Standing by Godot

To resolve the problems associated with meeting strict RPO and RTO requirements, both AppAssure and Backup Exec 2012 provide IT with the ability to maintain a warm-standby VM of an existing system as part of the backup process. In addition, AppAssure further improves recovery functionality and performance by using Smart Agent technologies to recover an entire disk with a zero-perceived restore time.

Both Backup Exec and AppAssure provide IT Administrators with the ability to create a warm-standby VM. When Backup Exec 2012 is used, however, both the process of creating a warm-standby VM, which is completely disjunct from the creation of a backup file, and the resultant VM are significantly limited. First, the process deletes any existing standby VM before creating a new one. As a result, a full file-based backup is required to create a warm-standby VM. In essence, Backup Exec's "Conversion to a VM" process is analogous to a full file-level restore.

Conversion to a Standby VM

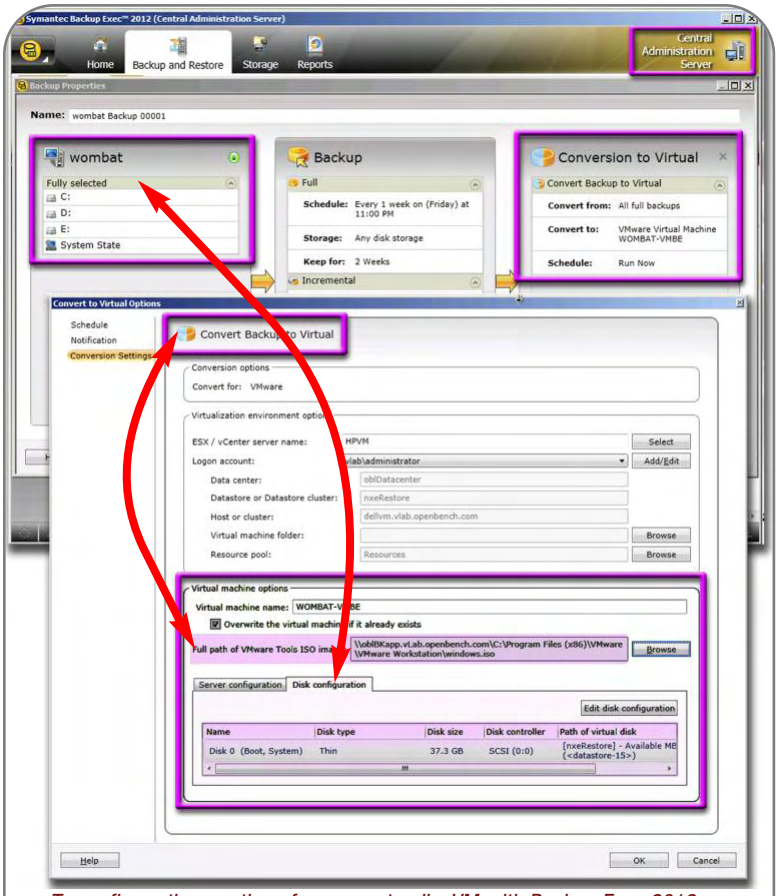

To configure the creation of a warm standby VM with Backup Exec 2012, we had to point to the VMware tools ISO for Windows on the media server. In addition, the disk conversion process only recognized the system disk (disk 0) in the conversion process.

Several configuration requirements further underscore the impression that the creation of a warm-standby VM is in reality a P2V conversion process that leverages a P2P backup process. In particular, an IT administrator must configure a path on the Backup Exec 2012 server pointing to the location of a VMware tools ISO file. That's a problem because the file in question is a hypervisor component rather than a vCenter component.

The ISO file can also be found in a VMware Workstation installation, which is required for the Backup Exec Validator option. Nonetheless, we could not simply copy that file to the Backup Exec server and point to its location. In order to be able to point to the ISO file within the Backup Exec 2012 GUI, we needed to install VMware Workstation on the Backup Exec server. We did not have to activate the software:, but we did have to install it.

Handling the VMware tools ISO was not our only problem. More consequential was the exclusion of all disks other than the system/boot disk from the conversion process.

When a full file-level backup of the original server included multiple disks, the Backup Exec VM conversion configuration only recognized the system/boot disk. We needed to configure and provision all ancillary disks manually after the conversion completed. Worse yet, disks were not the only hardware resources missing. While all of our physical systems included an Intel Gigabit (e1000) network adapter, none of the standby VMs created with Backup Exec 2012 had a network adapter configured. Along with data disks, network connections had to be configured manually in a follow-up process.

The need for IT administrators to manually finalize the configuration of a standby VM can be very significant with respect to RPO and RTO. First, the requirement to use a full file-based backup compromises the strategic purpose of a warm-standby VM, which is to maintain a copy of a production system in a state that is as close as possible to that of the original system. Second, the need for Backup Exec to recreate the standby VM on every "update" precludes the ability to make persistent changes to optimize RTO.

Pipelining a Standby VM With an Aggressive RPO

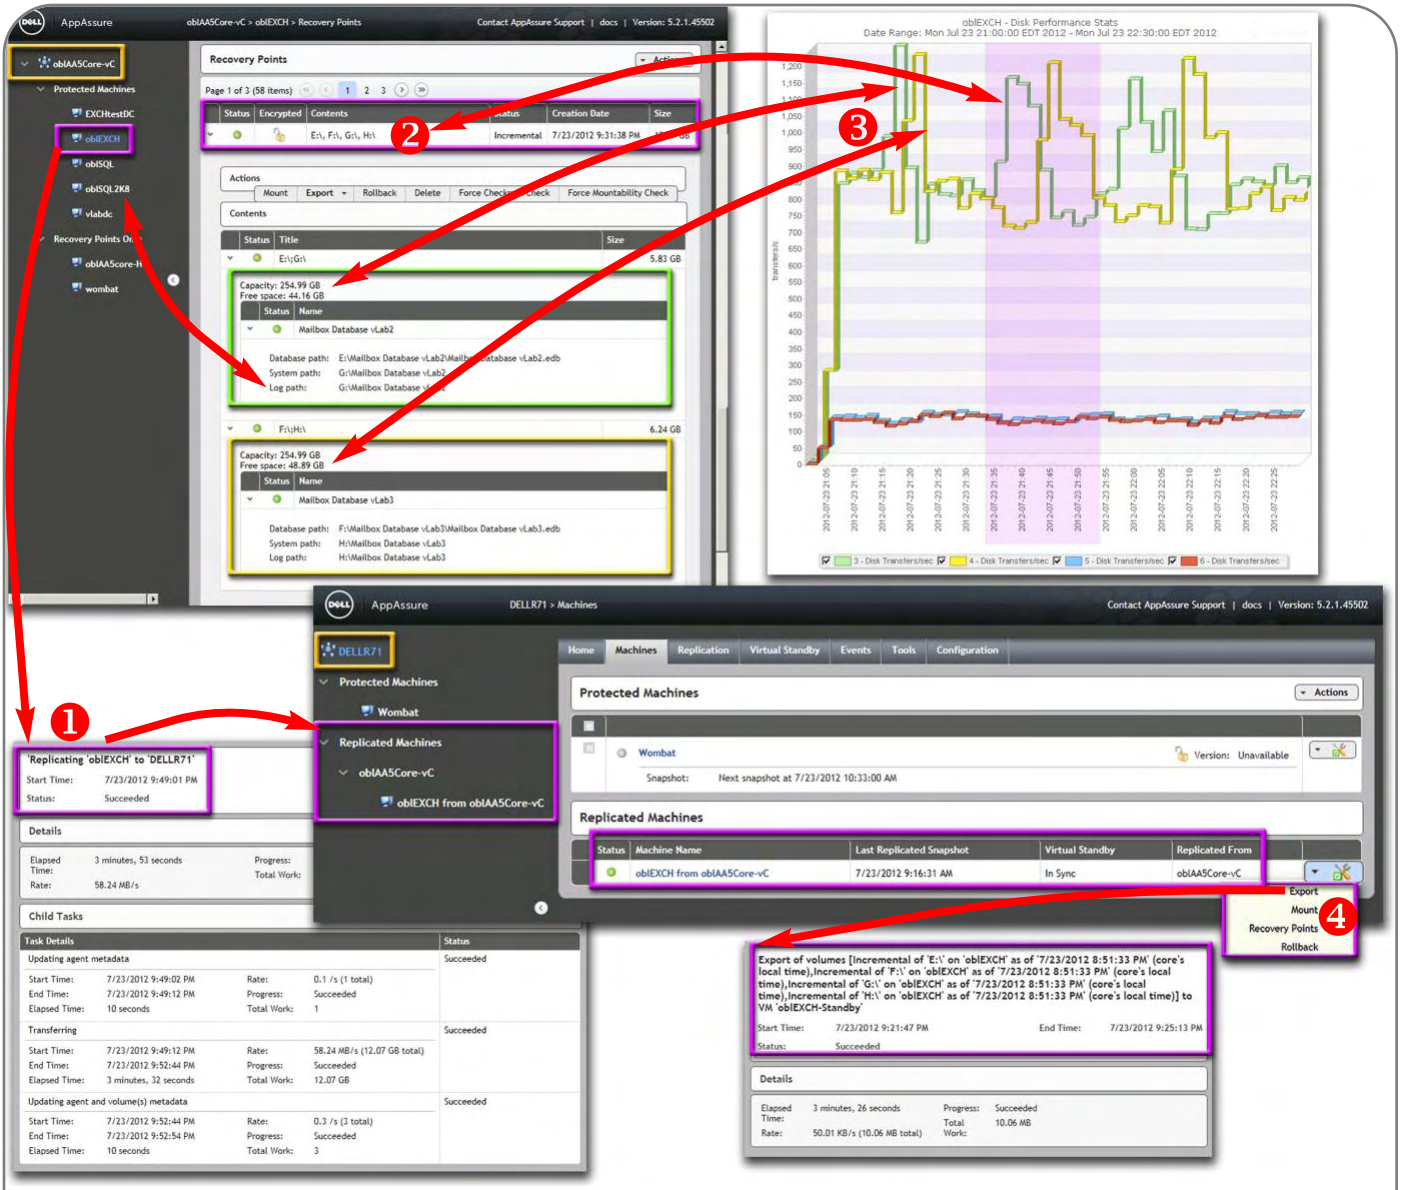

To provide a 20-minute RPO and a 2-minute RTO, we set up a pipline❶ between the AppAssure Core on the Exchange VM's host and the Core server on our central Dell PowerEdge R710. While our Exchange VM processed 1,000 TPS generated by Jetstress, we ran incremental backups on the VM-based Core server❷ every 20 minutes. During backups of disks containing mailbox databases and logs, Smart Agents increased throughput on the disks being backed up to ensure our Exchange sever continued to deliver the required 1,000 TPS❸. When the backup completed, the AppAssure Core VM replicated the Exchange VM data in its repository to the Central AppAssure Core and we were able to run any recovery task from either Core server. In particular, we maintained a standby VM❹ for an enhanced RTO on the Central AppAssure Core asynchronously from the RPO-oriented backup tasks running on the AppAssure VM Core server.

In sharp contrast, AppAssure leverages the ability to update a VM not in a running state by generating a disk snapshot for each logical volume after every incremental

backup of the original machine. More importantly, AppAssure 5 extends this capability by using Smart Agents to optimize the transfer of data between AppAssure protected machines clients and AppAssure Core servers, as well as the ability to replicate backup repositories among AppAssure Core servers. AppAssure implements all of this technology with an extraordinary level of performance.

In particular, we leveraged Smart Agents to safely run an incremental backup every 20 minutes, while the Exchange server VM processed 1,000 email transactions per second. Each incremental backup typically transferred 20GB of changed block data to a VM-based AppAssure Core server on the same vSphere host.

Using True Scale architecture, we decoupled updates of a standby VM on another host by automatically replicating the Exchange server portion of the backup server's repository to a second AppAssure Core after each backup. The second Core was able to use the replicated repository to create and transfer disk snapshots to a standby VM asynchronously from the backup of the original VM.

Not only were we able to use AppAssure 5 to maintain a standby VM running Exchange server with an aggressive RPO of 20 minutes and an RTO of 5 minutes, we did so while the original server processed 1,000 email transactions per second. The key to this legerdemain in business continuity was the ability to leverage Smart Agents during all aspects of the data protection process.

Smart Zero RTO

The most magical and impressive Smart Agent technology is without doubt AppAssure Live Recovery. Live Recovery is a unique volume restoration feature. During a recovery, a Smart Agent monitors block data requests from users and instructs the backup server to reorder the data being transferred to meet user requests. As a result, users perceive all data as having been restored immediately, as they have immediate access to any data on any volume.

With Live Recovery, the restoration of a logical volume starts with the transfer of the volume's Master File Table (MFT). With the MFT in place, users immediately see all of the files that were on the drive at the time of the restore point. As a result, end users have the impression that the drive has been magically restored to full operational status in seconds. Nonetheless, in a database environment, there is the distinct risk of corrupting one or more critical internal tables or causing considerable delays and problems for mission-critical, transaction-processing applications.

When an application attempts to access data queued on the AppAssure backup server, the requested data addresses are transferred to the backup server, which reorders the block stream queued for the client so that the requested blocks are at the top of the queue. As a result, a large volume file server or an Exchange mailbox database can be restored instantaneously from the perspective of an end user.

AppAssure Live Recovery Rollback

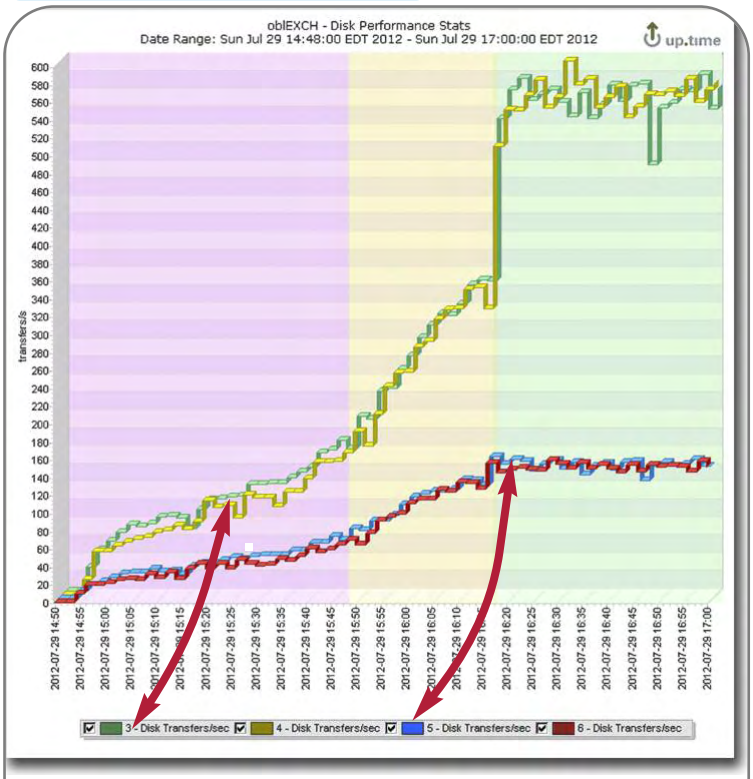

To stress an AppAssure Live Recovery we rolled back two volumes used for an Exchange mailbox database and logs for 500 users, while we simultaneously ran Jetstress. During the process, Live Recovery allowed Exchange server to successfully respond to all Jetstress requests. The I/O response rate increased slowly until 65% of the data had been restored, when both the response rate and rollback throughput rapidly increased.

To test Live Recovery, we began a rollback of one of our mailbox resource groups, which consisted of the disk volume supporting a mailbox database for 500 users and the disk volume supporting the logs for that database. When we began our rollback, our biggest problem turned out to be the Smart Agents that optimized the transfer of the minimal amount of data needed to return to the chosen recovery point. Unless we either went far enough back in time or deliberately destroyed the entire volume, the rollback finished before we could set up a Jetstress benchmark to stress the recovery.

After we began a full restore process for both volumes—202GB of data—we started a Jetstress benchmark that required the presence of the missing database and its log files. From the start of the test, the AppAssure client coordinated the order in which mailbox data was transferred from the AppAssure backup server. As the transaction load stressed the Exchange mailbox database and log files that we were recovering, AppAssure ensured that the process never failed an I/O request. During the test period, the IOPS rate continuously escalated as the amount of data available locally increased.

Test Summary

"No data protection software tested by openBench Labs has been able to match AppAssure's ability to provide a warm VM for both virtual and physical production servers while maintaining a five-minute RPO in addition to five-minute RTO."

Rewriting the Rules of Business Continuity

Concerns expressed by Line of Business (LoB) executives over business continuity are helping to drive the next wave of IT projects. In a competitive 24x7 environment, computer downtime represents more than lost revenue to sales and marketing executives. These executives equate lengthy computer outages with potential losses in customer confidence and market share. As a result, senior LoB executives expect IT to meet an RTO that is measured in hours rather than days and an RPO that is close to zero.

AppAssure Feature Benefits

1) **Repository Stores a Virtual Disk Block Space:** AppAssure 5 uses an object oriented (OO) file system to maintain a virtual space of globally deduplicated disk blocks in place of an archive of backup sets.
2) **Agent-based Changed Block Tracking:** IT administrators install agents on client systems that utilize a filter driver to track block-level changes to greatly reduce the processing of incremental backups.
3) **Minimize Physical Storage Using True Global Deduplication:** AppAssure 5 leverages an object-oriented repository in a global deduplication scheme that reduces storage consumption without incurring the overhead level of an RDBMS
4) **Pipeline AppAssure Cores using True Scale™ Architecture:** IT can readily scale data protection processes by configuring multiple AppAssure Core servers in a pipeline to scale a data protection process linearly.
5) **Live Recovery Restores Volumes and Key Application Files with Zero Perceived Delay:** Live Recovery immediately restores directory details and reorders the transfer of blocks based on user access of files during recovery.
6) **Recovery Assure Automatically Tests and Verifies Application Recoverability:** AppAssure discovers key applications, such as SQL Server and Exchange, on clients and applies special item-level verification tests for these applications.
7) **Universal Recovery to Any Hardware for DR:** Recover systems using a bare metal recovery CD, a hot standby server or a warm standby VM.
8) **Universal Recovery at Any Granularity:** Recover systems, disk volumes, files, or application-level data items, such as email messages and database tables using standard application management tools.

AppAssure 5 Backup, Replication and Recovery software is designed to leverage changed disk blocks in place of changed files on any Windows system—physical or virtual—to ensure RPO and RTO objectives. More importantly, AppAssure 5 introduces an OO file system that underpins a globally deduped repository, which represents a virtual disk space for all systems protected by the AppAssure 5 Core Server. For IT operations, the new OO architecture of AppAssure 5 provides a highly scalable infrastructure for a powerful "forever incremental" backup scheme that requires a minimal amount of computer processing time.

Using AppAssure 5, IT is empowered to implement a DR plan with aggressive RTO

and RPO levels to support any SLA required by Line of Business executives. Moreover, IT can double down on the advantages garnered in incremental backup processing to utilize very frequent automated incremental backups to provide minimally spaced recovery points for mission critical systems.

Specifically, IT operations managers can leverage AppAssure incremental backup advantages by running frequent incremental backups for closely spaced recovery points of mission critical systems and use standby VM servers to double down on an aggressive DR plan with an RTO measured in minutes. No data protection software tested by openBench Labs has been able to match AppAssure's ability to provide a warm VM for both virtual and physical production servers while maintaining a five-minute RPO in addition to a five-minute RTO.

Jack Fegreus is Managing Director of openBench Labs and consults through Ridgetop Research. He also contributes to InfoStor, Virtual Strategy Magazine, and Open Magazine, and serves as CTO of Strategic Communications. Previously he was Editor in Chief of Open Magazine, Data Storage, BackOffice CTO, Client/Server Today, and Digital Review. Jack also served as a consultant to Demax Software and was IT Director at Riley Stoker Corp. Jack holds a Ph.D. in Mathematics and worked on the application of computers to symbolic logic.