Technology Lab Test Report

UNIFIED BACKUP: VIRTUAL & PHYSICAL

AppAssure v4.7
Backup & Replication Software

vs.

Symantec
Backup Exec 2010 R2

PERFORMANCE

Analysis: Unified Backup: Virtual & Physical

AppAssure v4.7
Backup and Replication Software
vs.
Symantec Backup Exec 2010 R2

Author: Jack Fegreus, Ph.D.
Managing Director
openBench Labs
http://www.openBench.com
October 27, 2011

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.

Table of Contents

Executive Brief

"AppAssure's key advantage over Symantec Backup Exec is its ability to provide a unified approach to data protection that works with any physical or virtual server in exactly the same way to yield unmatched RTO and RPO levels."

Server Virtualization as the Norm

Under Test: Windows Server Data-Protection

AppAssure Value Proposition

1) **Live Recovery™ with Near-Zero Data Recovery Time:** AppAssure intelligently reorders blocks during volume recovery to permit immediate access to critical application data during the recovery of a full volume
2) **Assured Recovery™ for Automatic Application Validation:** Extend backup validation processes by automatically discovering key applications, including SQL Server and Exchange, on clients and then testing application recoverability in each backup
3) **Cross-platform Universal Recovery™ to any VM or physical server:** Any backup can be restored to a hot-metal backup CD for any hardware or used to build a VM on any Virtual Infrastructure host
4) **Universal Recovery™ for Disaster Recovery (DR):** Automatically update system and data files on a warm standby VM and user data on a hot standby physical server after every backup
5) **Granular Universal Recovery™ of Files and Application Items:** Retrieve files or individual data objects, such as an email messages and database objects by mounting and sharing any backup as a virtual disk
6) **Changed Block Tracking on All Backups:** AppAssure backup and replication client installs a filter driver that tracks all data changes by disk blocks on virtual and physical systems running Windows
7) **Minimal Data Loss with Near-Continuous Recovery Points:** AppAssure leverages block-level data change tracking to provide fast, small footprint, incremental backups with automatic roll up to synthetic full backups, which can be scheduled in five minute intervals
8) **Off-site DR via WAN and LAN Backup Replication:** Automatically replicate backup archives for enhanced off-site availability of warm, standby, virtual servers and hot, standby, physical servers

In this new Lab Report, openBench Labs examines AppAssure backup and replication software v4.7 with Symantec Backup Exec 2010 R2. AppAssure implements changed-block tracking technology to provide data protection for both virtual and physical servers in specific OS environments. In contrast, Backup Exec 2010 R2 uses traditional file-based backup to promote compatibility with the largest number of operating systems.

AppAssure leverages its changed-block tracking technology to provide IT operations with advanced features that include Live Recovery for near-zero downtime; Assured Recovery for testing and verifying object level recoverability for files and application-items in any backup; and Universal Recovery that permits cross-platform recovery of data volumes from physical to virtual, virtual to virtual, virtual to physical or physical to physical—P2V, V2V, V2P, P2P—including bare metal recovery to dissimilar hardware. In addition, these features also combine to provide a powerful, cost-effective, and unified Backup and Disaster Recovery alternative to Symantec Backup Exec.

The explosion in the number of applications and the amount of data residing in virtual and cloud environments was brought into the light with Symantec's 2010 State of the Data Center Report. The report, based on a worldwide survey of 1,780 IT sites, the report that identified 88% of all datacenters have a virtual infrastructure project planned or in progress—fueling future data growth concerns.

For a small to medium business (SMB) site, ensuring business continuity in a cross-platform environment with virtual and physical servers is especially difficult. Many SMB sites deal with the architectural differences between virtual and physical servers by adopting a separate data protection operations for virtual and physical systems. All too often, this strategy results in IT administrators struggling to support two very different data protection solutions: a block-based solution for virtual systems and a file-based solution for physical systems.

AppAssure's key advantage over Symantec Backup Exec is its ability to provide a unified approach to data protection that works with any physical or virtual server in the exact same way to yield unmatched RTO and RPO levels. Specifically, IT operations can implement a data protection regimen built on a foundation of near-continuous incremental backups that can be scheduled to occur in time intervals as short as five minutes.

Testing Overview

openBench Labs Data Protection Test Scenarios

1) Corporate Email Scenario:
- One virtual server running Windows Server 2008 R2 as a Domain Controller
- One virtual server running Exchange Server 2010
- One physical workstation running Windows 7 and Outlook 2007

2) Database-driven Business Process Scenario:
- One physicall server running Windows Server 2003 R2 as a Domain Controller
- One virtual server running Windows Server 2008 R2 and SQL Server 2008 R2
- One physical server running Windows Server 2008 R2 and SQL Server 2008 R2

3) End-User Desktop Workstation Scenario:
- One physical workstation running Windows 7
- One virtual workstation running Windows 7

OpenBench Labs conducted tests on AppAssure 4.7 backup & replication software and Backup Exec 2010 R2 using a test bed monitored with up.time 5.5 software and populated with four physical servers, one workstation, and a fibre channel storage area network (SAN). Within this environment, openBench Labs set up three different IT scenarios to host and evaluate our data protection tests.

Data protection tests began with an examination of each package to simplify IT management and lower operating expense (OPEX) costs by providing consistent uniform access to all combinations of data protection functions. Operating tests were conducted within the context of our three business scenarios and focused on a disk-to-disk (D2D) backup strategy employing continuous incremental backups of virtual and physical

servers. A D2D strategy provides IT with a number of operating efficiencies, versus a disk-to-tape (D2T) strategy; however, to be economically viable, storage utilization between production data and backup data must be on the order of 25 to 1.

openBench Labs Data Protection Testing Summary

Data Protection Test Suite	Key Test Metrics	Analysis	AppAssure v4.7	Backup Exec 2010 R2
Universal Backup Execution	•Unified execution of all commands	•Execution of all options from a single GUI lowers operation complexity	•All options run from a single management GUI	•Automatic incremental backup rollup requires True Image Restore •True Image does not support data item recovery
	•Backup time	•Backup time governs RPO	•Fastest backup times	•Data deduplication significantly extends both backup and verify time
	•Client CPU overhead	•Service overhead restricts backup window options	•Minimal CPU overhead	•Heavy CPU overhead
	•Backup file size	•Backup size directly affects resource costs	•Smallest incremental backup files	•Data deduplication needed for small backup files
Granular Application Recovery	•Extended validation of data objects	•Application validation essential for SLA support ot both RTO and RPO	•Automatically discovers protected applications and adds validation tests	•Application-specific backups excluded from synthetic backup
	•Restoration of item-level data from any backup	•Tested the ability to restore objects from a SQL Server database and email messages from an Exchange 2010 mailbox	•Status of databases and mailboxes shown in GUI •Any backup can be used to recover data objects	•Email retoration GUI lack integration with MAPI
Standby VMs and Near-zero Disaster Recovery	•Utilize VM clones of virtual and physical systems for DR	•Update standby VMs—OS and data—with production server disk snapshots from backups	•Standby VMs configured identically to production servers	•Optional module for Bare Metal backup and restore •Continuous updates of a standby VM not supported
	•Restore applications and data with no perceived delay	•Restore applications and data with no perceived delay	•Supported an Exchange mailbox restore while passing a Jetstress benchmark	•Near-zero DR unsupported

In our tests, backup and verification processing time, which sets the theoretical limit for an RPO, was consistently faster for AppAssure—***without compression or data deduplication, AppAssure completed full backups 7 times faster than Backup Exec.***

In addition, AppAssure imposed dramatically lower CPU overhead on clients***—in full backups of a dual-processor server AppAssure overhead peaked at 10% of a CPU, while Backup Exec consumed 90% of a CPU.*** Backup overhead is another important factor limiting the ability to implement an aggressive RPO. Storage space savings, which are essential for deferral of capital expense (CAPEX) costs, were statistically the same for both AppAssure and Backup Exec 2010 R2, when data deduplication was implented in our backup strategy. ***Without data deduplcation, incremental backups with Backup Exec ranged from 10 to 100 times the size of an AppAssure incremental backup.***

More importantly, the inability to apply key backup and restore options across all types of backups with Backup Exec makes it very difficult to compare the performance of features between AppAssure and Backup Exec. Specifically it is not possible to utilize incremental backups that will automatically be rolled up without manual intervention and which will also support restoration of individual files or data items.

This means a highly optimized D2D backup strategy comparable to AppAssure simple is not possible. IT operations must tradeoff between backup file size, backup speed, and complexity of backup proceedures. We chose to minimize operations complexity and backup file size for D2D backup, which required running full backups with data deduplication using Backup Exec 2010 R2. ***As a result, our operational scenario left us comparing similar sized incremental AppAssure backups with compression and deduplication that took under two minutes to process, which was 40X faster than full backups with Backup Exec with data depuplication.***

With the growth of application-specific SLAs for business continuity, our Recovery testing suite included an assessment of each package's ability to support a deep validation of application recoverability within a backup file. Operating tests concentrated on two scenarios for data-item level recovery: the recovery of individual SQL Server databases and the recovery of Exchange mailboxes and messages. In addition we examined Backup Exec and AppAssurep for specialized recovery packages supporting recovery features for perceived near-zero recovery time.

In our third test suite, we examined the ability of AppAssure and Backup Exec 2010 R2 to extend data protection to disaster recovery functions. Beyond generating a bare metal installation CD, which typically leaves IT with several hours of work to configure a new system, we examined each package's ability to leverage a warm standby VM.

By aggressive exploiting its changed block tracking mechanism across all data protection operations, AppAssure was able to forge a measurable advantage in all three of our test suites. More importantly, AppAssure leveraged block-based tracking in the introduction of three key technologies: Live Recovery™ to meet near-zero RTO and five-minute RPO; Assured Recovery™ to automatically verify that both a backup is recoverable and that application data within the backup is recoverable; and Universal Recovery™ to support granular recovery of files and application items, as well as bare metal recovery that can be P2P, V2V, P2V, or V2P. As a result,AppAssure dramatically reduces both the capital and operating costs associated with data protection and DR.

File to Block Redirection

> "The biggest savings in capital and operating expenses for IT comes from the ability to readily extend AppAssure's Universal Recovery options to Disaster Recovery."

Disk Image Foundation

Traditional backup solutions, such as Symantec's Backup Exec, transfer server files to an alternate storage medium to protect data. Unfortunately, file-level backups do not extend easily to other objects. When IT needs to recover data items contained in files or recover virtual systems that host files, traditional backup software requires specialized modules, which increase capital expenses and add to operational complexity.

AppAssure I/O Redirection

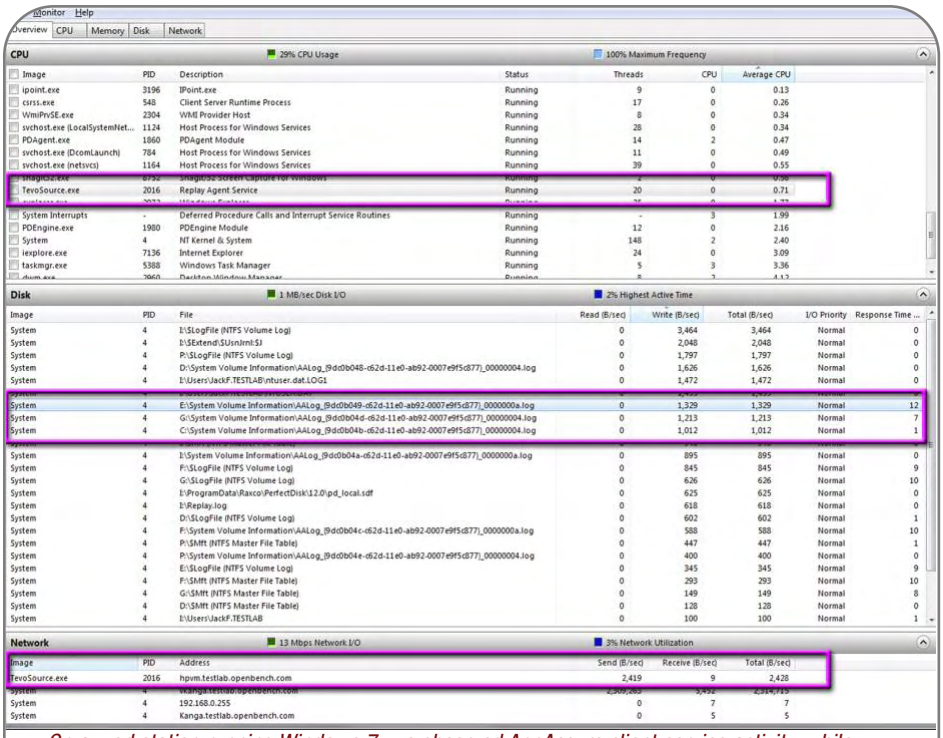

On a workstation running Windows 7, we observed AppAssure client service activity, while generating a high level of I/O. As file data changed, the client service logged meta data about disk block changes to an AppAssure (AA) log file in the disk's System Volume Information directory.

In contrast, AppAssure is built on a simple mechanism to track data changes to disk blocks rather than files. Using changes to disk blocks instead of files, AppAssure supports a unified, disk-imaging approach for all data protection and DR functionality.

To establish a generalized disk-block tracking mechanism on Windows-based systems, IT Administrators install an AppAssure service agent on every client system to collect, package, and send block data to a central AppAssure Core server. The AppAssure agent installs 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.

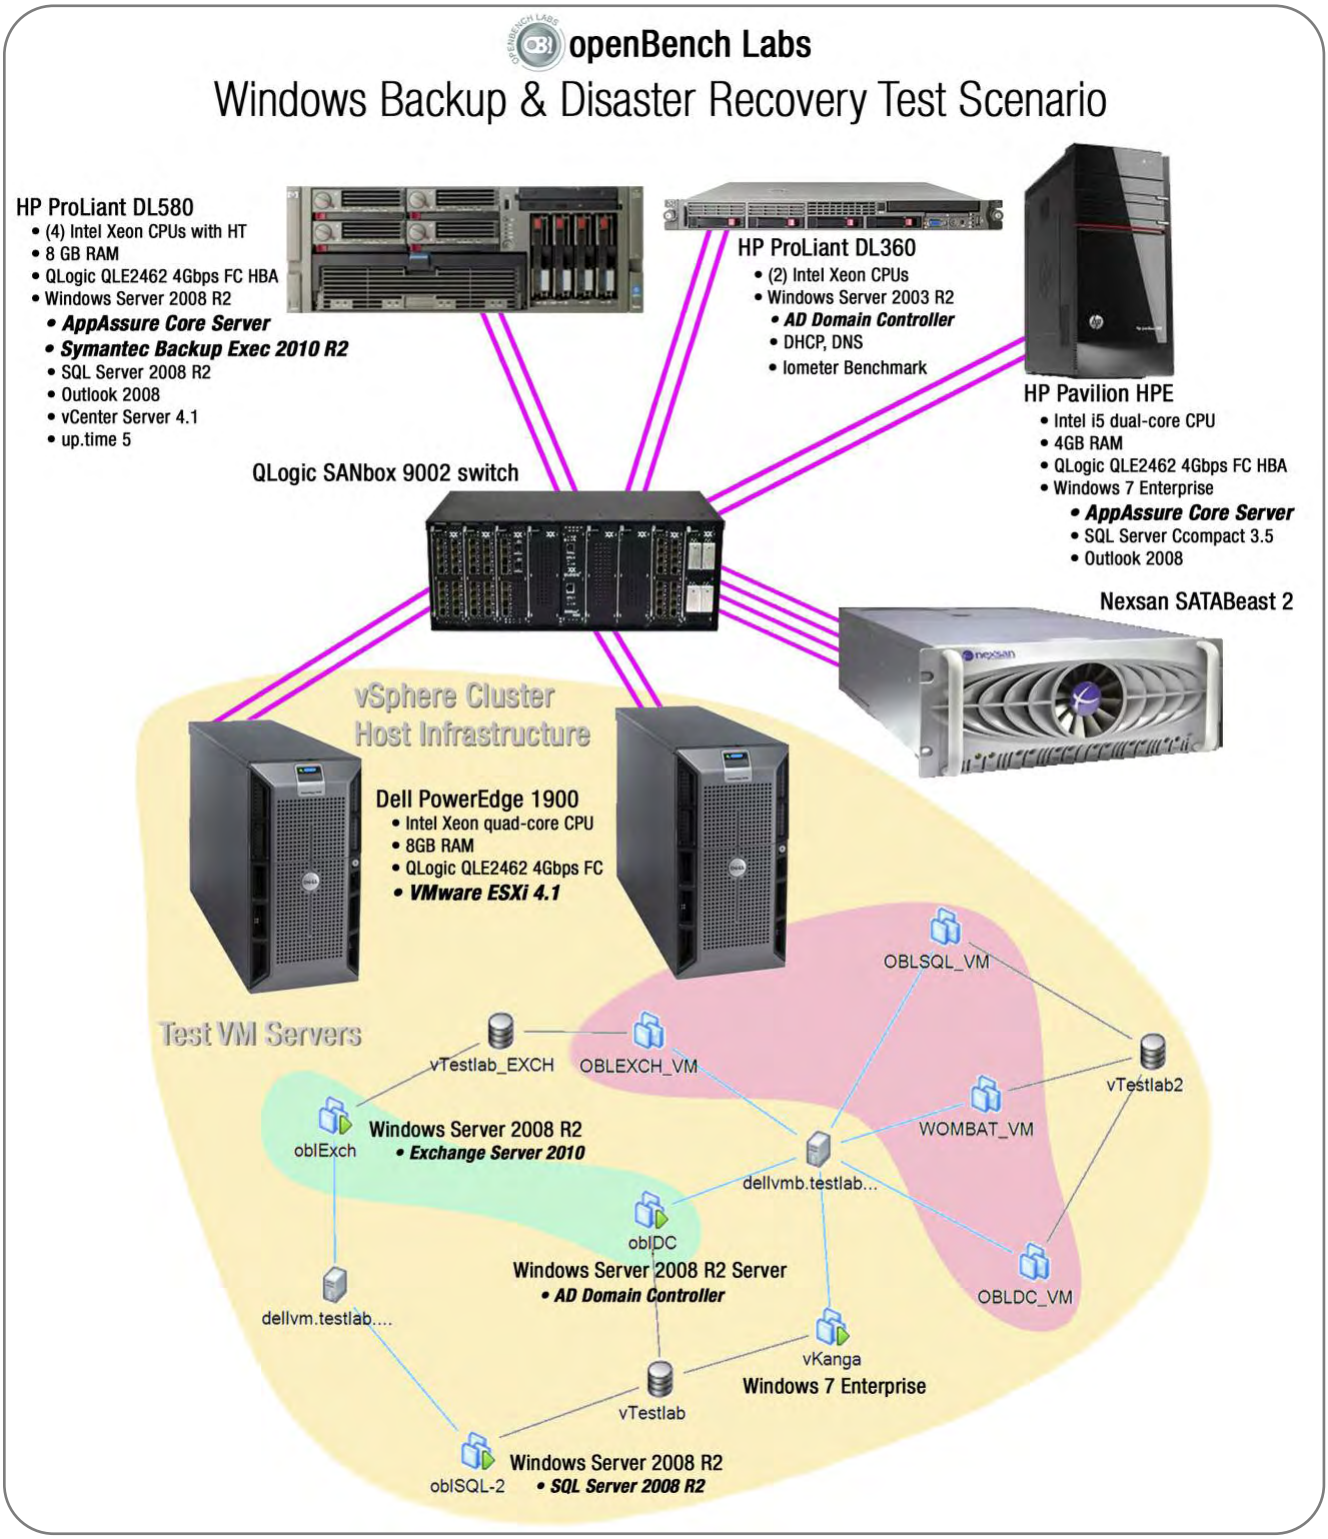

Using the filter driver, AppAssure service agents on client systems log meta data about disk-block changes in the System Volume directory for each logical disk. During a

backup, the agent utilizes meta-data logs to send the block-level data changes that have occurred since the last time an AppAssure Core server requested a backup data transfer.

Test Bed Basics

The test bed was set up with five physical systems, which included four servers and one workstation. The physical systems were centralized around two key servers. An HP ProLiant DL360 ran Windows Server 2003 R2 and played the role of an Active Directory Domain Controller. The second server was a quad-processor HP ProLiant DL580 with 8GB of RAM to serve as our primary AppAssure Core server and as our VI infrastructure controller running vSphere 4.1 to manage a VMware HA cluster with two ESXi hosts.

To test and verify granular item-level recovery, we installed the Express version of SQL Server 2008 R2 and a MAPI-compliant email client on the AppAssure Core server. In addition, we used a workstation running Windows 7 as a secondary AppAssure Core to test high availability (HA) for DR based on replication of backup stores.

On two ESXi hosts, we set up three VMs running Windows Server 2008 R2 as clients of our AppAssure Core server. On one VM, we installed a full version of SQL Server 2008 R2 along with Microsoft's AdventureWorks database. On the other VMs, we set up our email server scenario: The first VM acted as an Active Directory domain controller and the second VM ran Exchange Server 2010.

To provision storage for all openBench Labs tests, we utilized a 4 Gbps Fibre Channel (FC) Storage Area Network (SAN). All virtual disk volumes were created on a Nexsan SATABeast 2 array server. The SATABeast 2 featured two independent FC controllers with dual 4Gbps ports on each controller. In addition, Nexsan supports Asymmetric Logical Unit Access (ALUA), which enables hosts to direct I/O exclusively to active ports connected to the controller servicing a targeted logical volume. As a result, we measured total read and write throughput rates of 1GB per second on each physical system.

Test 1: Universal Execution for All Backup Options

> "AppAssure, unlike Backup Exec, permits IT administrators to implement all data reduction features, including data compression and deduplication, with either an incremental or a synthetic backup."

Backup Software Architecture Differences

Of particular concern for openBench Labs is the ability of software packages to simplify IT operations. For data protection software, a critical factor is the ability to minimize operating complexity by allowing IT administrators to implement all backup and recovery options using a minimal set of GUI commands.

When comparing the ease of use of AppAssure with that of Backup Exec 2010 R2, a central issue is the intended scope of the two products. AppAssure is designed explicitly for use in environments that enable implementation of its block-level tracking, while Symantec Backup Exec 2010 R2 is designed to support the largest possible collection of servers in a wide range of environments.

AppAssure Unified Client Protection GUI

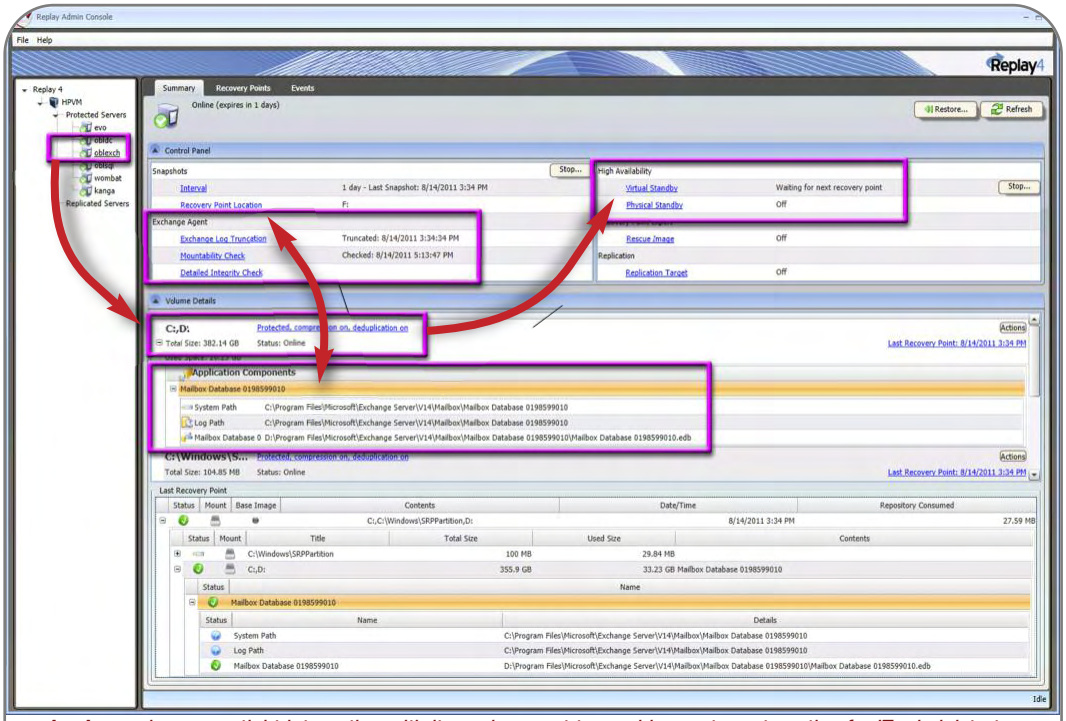

AppAssure leverages tight integration with its environment to provide greater automation for IT administrators. When we added a server running Exchange Server 2010 R2 to our AppAssure Core server for protection, AppAssure automatically recognized that Exchange was running and configured item-level protection for every mailbox database. In addition, we easily extended DR protection by setting up a warm virtual standby VM, which was automatically updated following any backup of the protected server.

To meet the general demands of an enterprise-class IT site, Backup Exec typically requires the addition of optional modules to the base package. Key optional features include data deduplication, synthetic backups with true image restore, and a VM agent for a VMware VI. These add-on components also need to

implement advanced functionality across multiple architectures, which results in adding complexity and operational overhead to day-to-day activities for IT administrators.

Deciphering Deduplication

For disk-to-disk (D2D) backups, Backup Exec leverages tape library integration to implement an ersatz tape library structure with fixed-size media for use in D2D backups. For an IT operations manager optimizing advanced D2D features, such as data deduplication, this faux library and cartridge scheme obscures the details of key features, such as Backup Exec's powerful PureDisk enterprise-class data deduplication scheme.

Backup Exec Media Scheme for D2D Backups

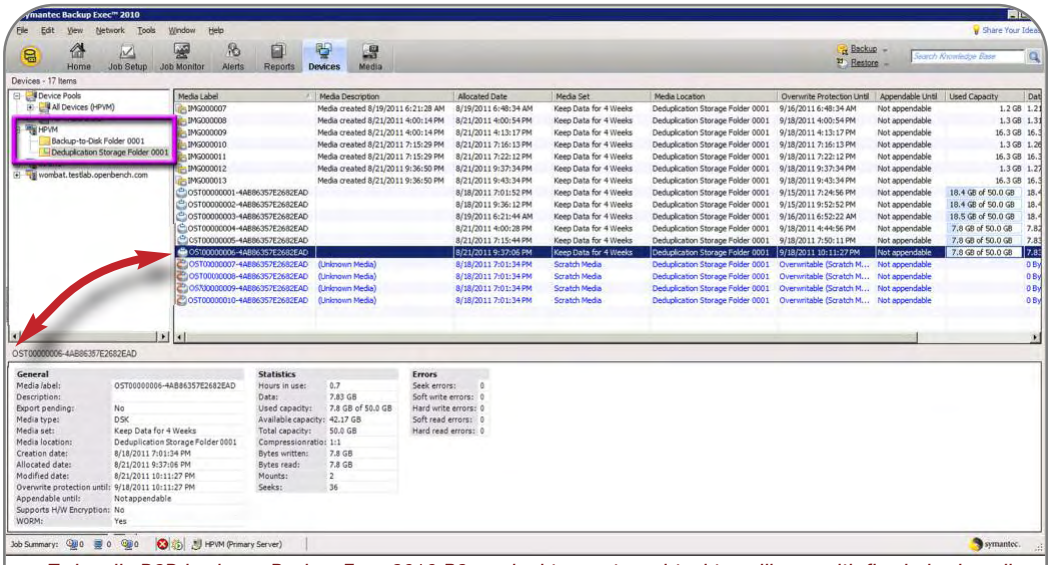

To handle D2D backups, Backup Exec 2010 R2, we had to create a virtual tape library with fixed-sized media files acting as tape cartridges. Worse yet, to handle deduplication, we had to create another distinct library with media for single instance data storage and media for backups.

PureDisk deduplicates data from multiple logical storage volumes on a system by discovering data segments that repeat in one or more backups. During a backup, PureDisk systematically breaks the data being streamed into segments, which are compared to hash-coded fingerprints that identifies a stored copy of segments already saved. If no match occurs, Backup Exec stores the new segment as a single data instance and creates a fingerprint to identify and point to the new data instance. If a match occurs, PureDisk replaces the known segment in the backup file with a copy of the matching fingerprint pointer.

To run the PureDisk deduplication scheme, IT administrators must create a special D2D media library that contains two distinct groups of media files: one group to store backup data and one group to store global single-instance segment data. More importantly, software compression of D2D backups significantly disrupts PureDisk efficiency and Backup Exec prohibits the use of data deduplication in synthetic backups. Backup Exec 2010 R2 also excludes the restoration of data items, such as SQL databases or Exchange mailboxes from a synthetic backup. As a result, to match the ease of use, automation, and storage savings of AppAssure with Backup Exec 2010 R2, we needed to implement a regime of full backups with data deduplication.

Deduplication Overhead

In our second set of tests, we focused on reducing the amount storage needed for D2D backups. While there is a strong consensus for data deduplication and compression as the options of choice, architectural differences between applications, such as Backup Exec 2010 R2 and AppAssure, result in dramatic differences in how these options are applied.

The goal of a traditional enterprise data deduplication scheme, such as PureDisk in Backup Exec 2010 R2, is to physically store backed-up data only once. To accomplish this, a backup file is divided into segments and replaces known segments with a pointer to a saved single-instance of the data. While Backup Exec uses this method to make significant data storage reductions, it also requires a great deal of CPU processing, which we restricted to our media server. In a file-based backup environment, any change to a file triggers the inclusion of the entire file in an incremental backup: For example a single edit in a 30-page report will need to have nearly all of the data replaced with pointers.

Storage Reduction: Full Backup & Verification

Exchange Server 2010: Active data 33.4GB

Software Options	Backup Software	Backup File	Backup Window
No Data Reduction	AppAssure	33.4GB	5 minutes
	Backup Exec	33.4GB	35 minutes
Software Compression	AppAssure	9.2GB	18 minutes
	Backup Exec	19.1GB	40 minutes
Deduplication & Compression (full)	AppAssure	8.4GB	19 minutes
Deduplication & Compression (incremental)	AppAssure	3.7MB	2 minutes
Deduplication (initial)	Backup Exec	17.4GB	127 minutes
VM Deduplication (initial)	Backup Exec	11.0GB	31 minutes

Now consider our incremental file scenario in a changed-block backup environment, such as AppAssure. In a changed-block backup, incremental backups only the blocks changed in a file are processed. There are no unchanged data blocks that need to be deduplicated. The only possible duplicated data sources are multiple instances of the same change within the changed blocks. That's why AppAssure can restrict data deduplication to just the current backup and eliminate the CPU overhead and extended wall clock time associated with data deduplication of a file-based backup.

In addition, Backup Exec's file-based backup method puts a number of constraints on data deduplication starting with the need to continually make a large number of data segment comparisons. Our initial full backup of an Exchange server took just over 2 hours and yielded a 2 to 1 reduction in data storage. ***After several more backups, deduplication ratios grew as high as 400 to 1; however, the time to complete a full backup with Backup Exec remained an hour and 45 minutes.***

PureDisk Time Trap

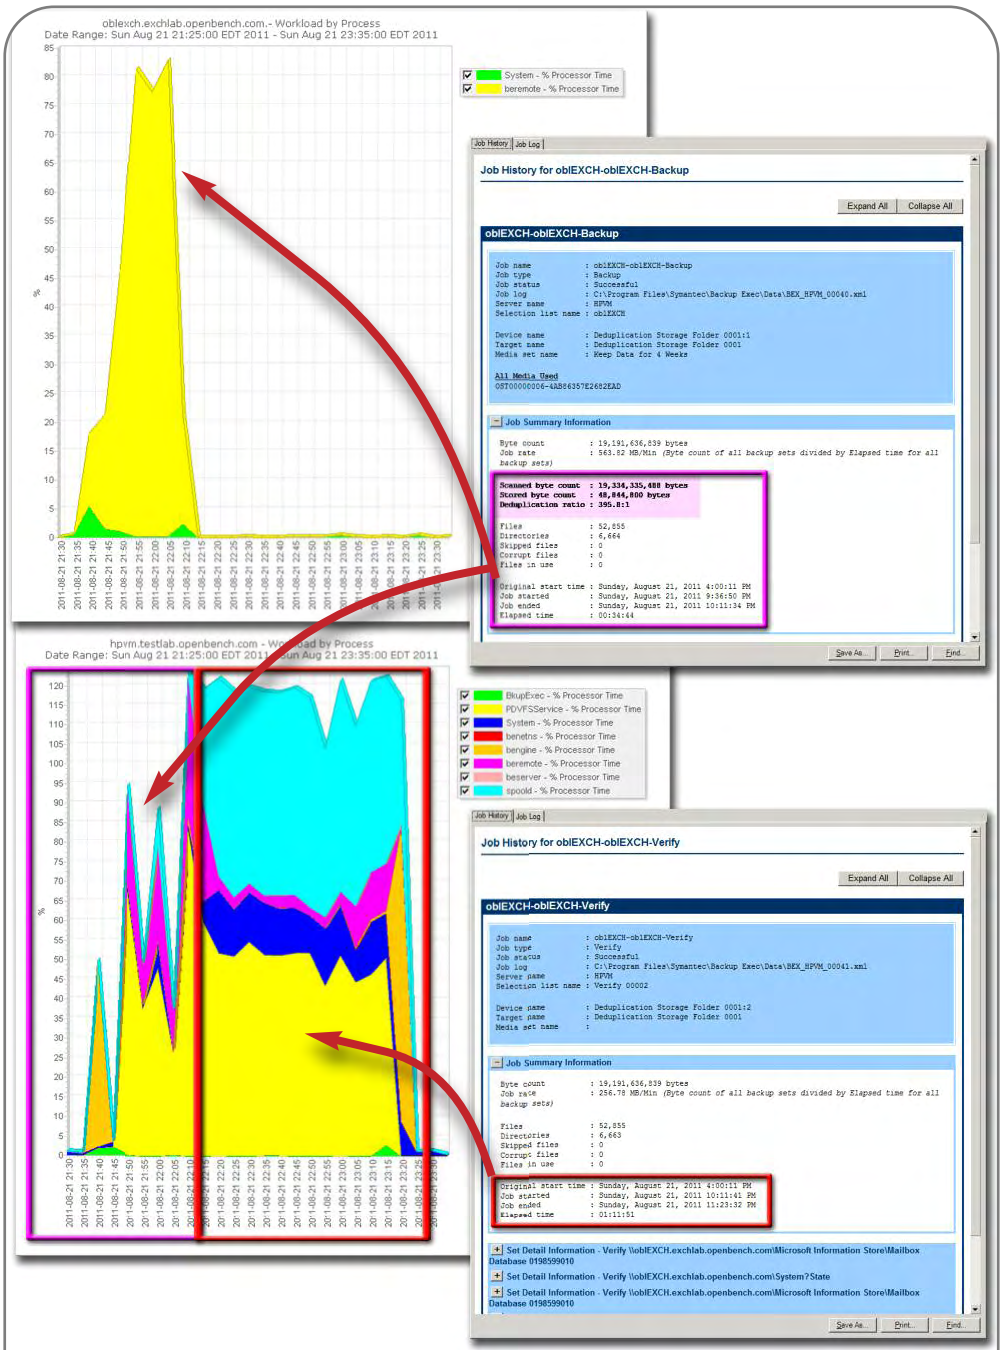

After multiple full backups of our Exchange Server using Backup Exec 2010 R2 with data deduplication, it was not uncommon for stored backup data files to be comparable in size to an incremental backup with AppAssure. In our tests, deduplication ratios on the order of 400-to-1 were not unusual, as stored backup files often fell to less than 50MB. Nonetheless, ***reducing 33.4GB of active data to a 48.8MB backup file with Backup Exec 2010 R2 required 1 hour and 47 minutes to deduplicate the backup and then rehydrate the backup file to verify the contents.***

Incremental backups shared this same problem, since full files, not just changes, are backed up in an incremental backup. In contrast, when we used Backup Exec with an identical configuration running on a VM, Backup Exec leveraged the changed-block vStorage API to reduce processing time for a full backup by 75%.

On the other hand, verification time increased as data deduplication increased in all test scenarios. To verify a backup, the stored data must be read and rehydrated. More importantly, IT operations constraints within Backup Exec 2010 R2 extended beyond data deduplication issues.

Backup Exec 2010 R2 dubs restore operations using synthetic backups that

automatically roll up an incremental backup into a full backup for restoration as "True Image." incremental IT administrators must use a special policy to run True Image, which does not support the inclusion of meta data containing system or application state information. As a result, recovery of data items, such as individual email messages or tables in a database are not possible using a True Image synthetic backup with Backup Exec 2010 R2.

Low Impact AppAssure Backup with Warm VM Standby

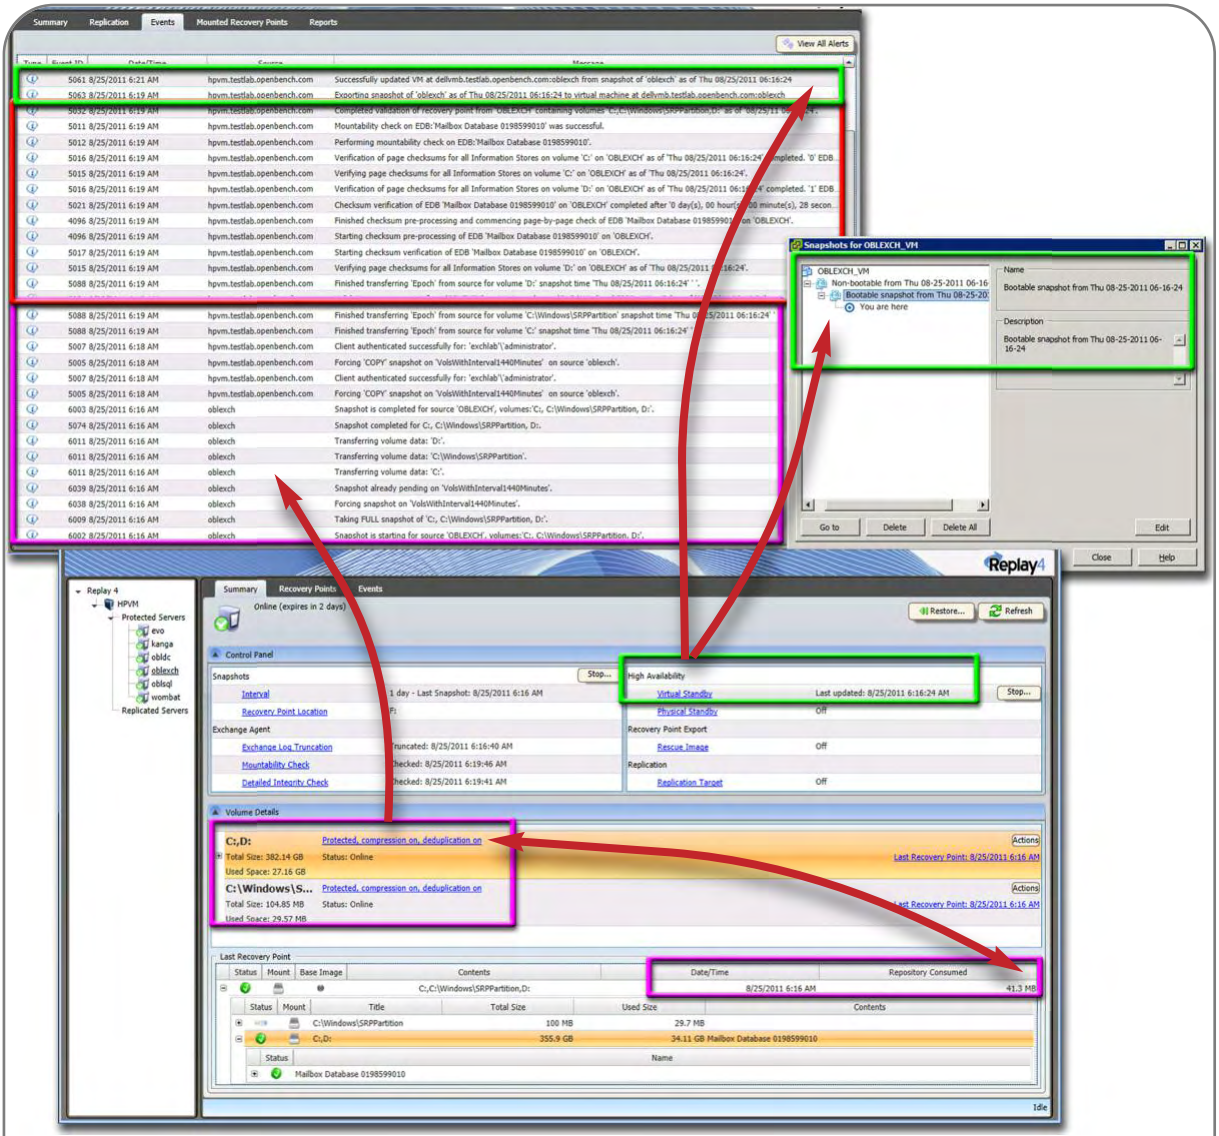

Using AppAssure with incremental updates in 10-minute intervals, we easily met aggressive RTO and RPO levels while placing minimal overhead on IT infrastructure. Every 10 minutes, we ran an incremental backup process, which included data compression and deduplication. The backup took about 4 minutes to complete—3 minutes to back up and 1 minute to verify the data. Backup files for this process were typically well under 50MB. What's more, we were also able to automatically update a warm standby VM within that 10 minute window by applying snapshot updates. ***In about 7 minutes, we processed an incremental backup of our Exchange 2010 server and updated a standby VM server for DR.***

Without True Image, our incremental backup file would indicate the presence of all files at the recovery point. Nonetheless, when we recovered an incremental file without first running a manual roll up, we discovered that we had restored pointer place-keeping files for any files skipped as unchanged during the incremental backup. To achieve similar functionality to AppAssure—a small backup foot print, data item recovery, and no manual file rollups—we needed to implement full backups with deduplication. As a result, our backups using Backup Exec 2010 R2 included all of the overhead that an in line data deduplication scheme places on CPU and memory resources.

Option Simplification

AppAssure, in contrast, enables IT administrators to implement all data reduction features, including data compression, and data deduplication, with either a full or incremental backup. To enable unrestricted functionality, AppAssure implements a simplified data deduplication scheme that only compares the data contained in an active backup. This means an AppAssure deduplicated backup file is totally self-contained and does include pointers to external data files.

"For an SMB site with constrained IT resources, AppAssure's ability to automatically discover critical applications and configure Universal Recovery of data items cannot be over-estimated."

Self-containment allows IT administrators to turn space-saving features, such as data compression and deduplication, on or off at will and replicate backup files between AppAssure Core servers for enhanced high availability and robust DR planning. Senior IT executives can also leverage AppAssure's operational flexibility for key strategic initiatives. CIOs can rely on AppAssure to automate a synthetic backup at any recovery point to support aggressive Service Level Agreements (SLAs) with respect to RTO and RPO goals.

What's more, fine-grained data item recovery is available for applications such as Exchange, SharePoint, and SQL Server. In addition to automating synthetic backups, AppAssure automatically modifies the presentation of backed up volumes when it discovers a protected application. To ensure fine-grained retrieval of data items with Universal Recovery, AppAssure groups all of the disks utilized by the application and presents them as a logical resource for recovery. For an SMB site with constrained IT resources, AppAssure's ability to automatically discover critical applications and configure Universal Recovery of data items cannot be over-estimated.

Test 2: Universal Recovery of Granular Items

"With Live Recovery, the restoration of a volume starts with the transfer of the Master File Table (MFT), which provides users with immediate access to the drive and the perception that the drive has been fully restored."

The Road to Zero RTO

IT traditionally viewed data protection as necessary activity centered on backing up data, applications, and OS files in a minimal amount of time. From a Line of Business (LoB) perspective, however, the value of data-protection operations rests entirely on recovery. That's why Service Level Agreements (SLAs) between IT and LoB units focus on RPO and RTO targets. Today, backup is simply a means to provide a recovery point.

For LoB executives, an RTO represents the maximum amount of time before an interruption to a core business process negatively impacts an enterprise. For IT administrators, resolving such an outage can range from restoring a database table to restoring a full system from OS to applications. The Symantec 2010 DR survey reported the total amount of time spent recovering a system averaged approximately five hours.

File and Volume Recovery with Deduplication
Windows Server 2003 R2

Software Options	Backup Software	Recovery Target	Time mm:ss
Deduplication & Compression	AppAssure	5GB *file*	3:05
Deduplication	Backup Exec	5GB *file*	7:30
Deduplication & Compression	AppAssure	33GB *volume*	31:00
Deduplication	Backup Exec	33GB *volume*	35:23
Deduplication & Compression Live Recovery	AppAssure	33GB *volume*	Perceived Instant

Most restore operations, however, are far from critical. A typical restoration involves retrieval of an end-user file lost or corrupted in day-to-day use. As a result, we began recovery testing with the restoration of a single 5GB user file.

Using Backup Exec 2010 R2, recovery of a file is solely an IT task, which begins when an IT administrator selects a recovery point that includes the desired file. This seemingly simple task is more complex than it appears. Only a recovery point associated with a full or true image synthetic backup is guaranteed to contain all of the files present on the system: With other recovery points, the restored file might be a simple place-holder file. For our single file restored from a full backup with deduplication, recovery process took 7 minutes and 30 seconds to extract, rehydrate, and transfer the data.

Zero Restore Time Legerdemain

File- and volume-level recovery is a very different process with AppAssure. For the restoration of any non-system volume, AppAssure provides IT with Live Recovery, a volume restoration feature that leverages block data to provide end users with zero perceived restore time.

Live Recovery Optimized Data Block Restoration

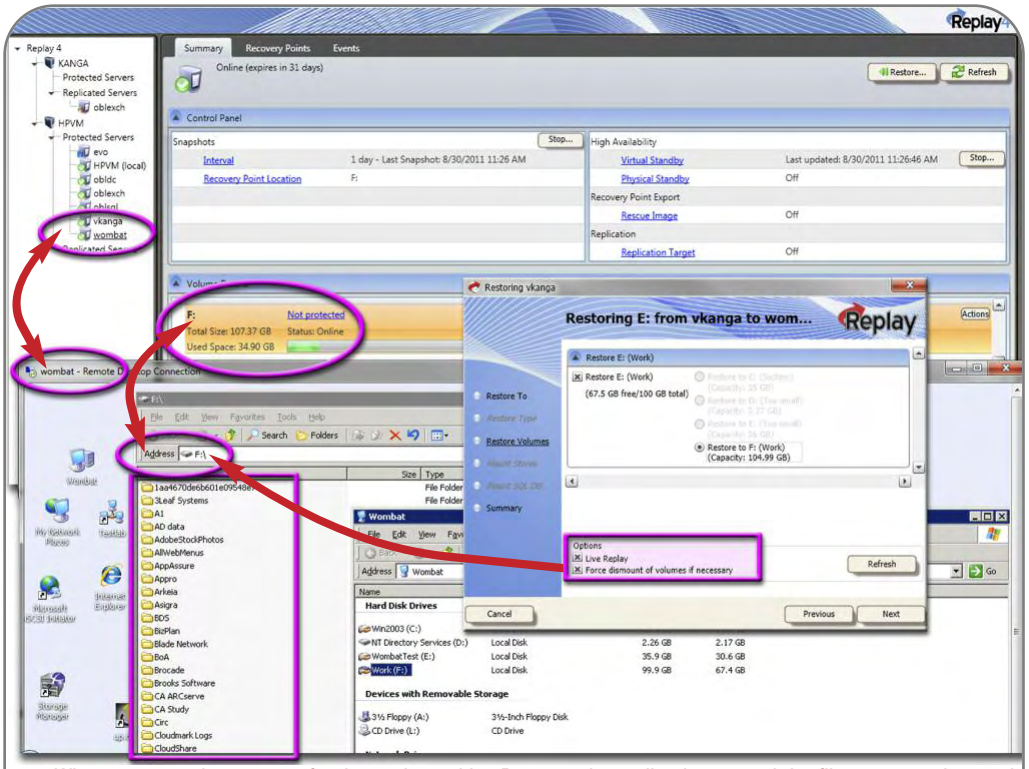

When we started a restore of a data volume, Live Recovery immediately restored the file structure data and enabled access to the volume. When we accessed any file during the restoration process, the AppAssure Core server immediately reprioritized the block data flow and sent the blocks that we needed. As a result we could instantaneously access any file or even an Exchange mailbox database.

With Live Recovery, the restoration of a volume starts with the transfer of the Master File Table (MFT), which provides users with immediate access to the drive and the perception that the drive has been fully restored. Moreover, if a user tries to access a file or data item still queued on the AppAssure Core server, the requested data blocks are transferred immediately. As a result, a large volume file server or Exchange mailbox database with hundreds of gigabytes of data can be restored instantaneously from an end user's perspective. ***With Live Recovery, we restored an Exchange database supporting 100 user accounts while sending each account one transaction per second in a Jetstress benchmark.***

Simplified Data Item Recovery

With the capability to mount and share any Assure backup file as a logical disk volume, IT can turn file-and item-level data recovery requests into self-serve end-user processes. A recovery point be mounted and shared in seconds, which opens the door for end users to manipulate application-level data items with existing application tools.

As a test, we mounted and shared a backup file containing the SQL Server 2008 R2 databases. We were then able to mount the backup shared from our Core server and

using SQL Server Management Studio we were able to attach any database contained in the backup and retrieve fine grain items such as views and queries.

Exchange Message Recovery

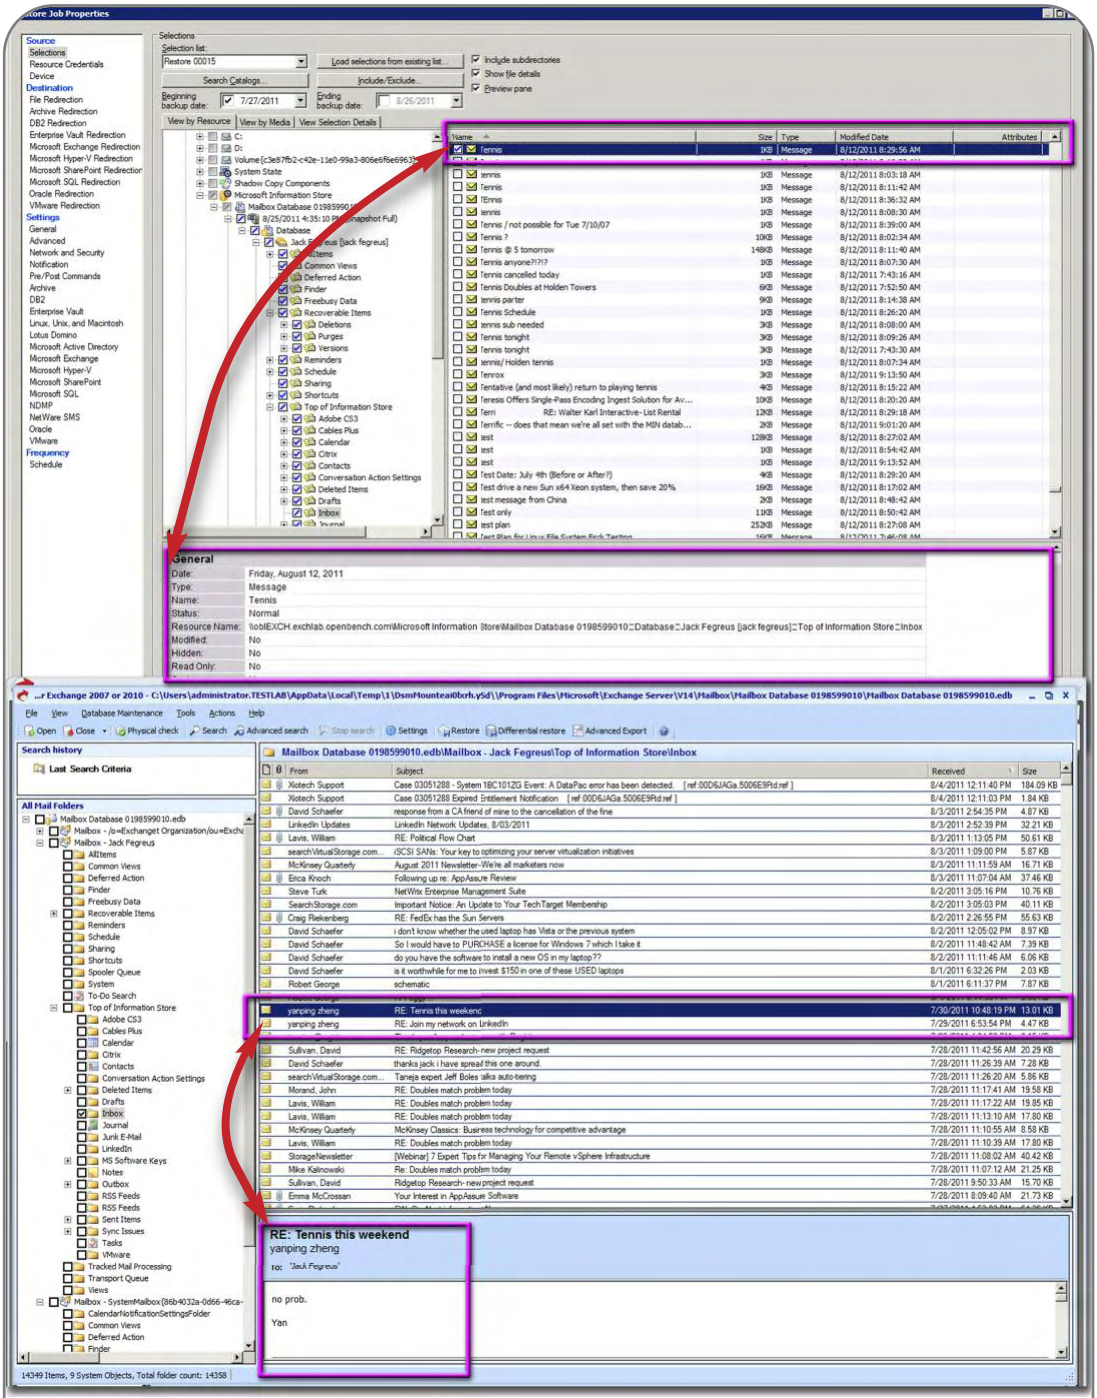

While we were able to restore individual mail messages from a full backup of an Exchange database with Backup Exec and any AppAssure backup. AppAssure made message identification much easier by leveraging MAPI to provide the sender name, message arival time, and full text of all messages in an email database.

Both AppAssure and Backup Exec are able to restore messages from an Exchange mailbox. AppAssure, however, leverages MAPI to provide far more data about each message. With the AppAssure GUI, we had all of the data found in an email client, such as Outlook. With Backup Exec, however, selecting an email message was quite difficult without the aid of rich MAPI data.

Test 3: Extend Universal Recovery to DR

"When a protected AppAssure client crashes, a warm standby VM server can be booted in seconds and brought online processing data in minutes to satistfy the most stringent business continuity SLA with respect to an aggressive RTO."

DR: The New Backup and Restore

DR VM Standby & Replication Topology

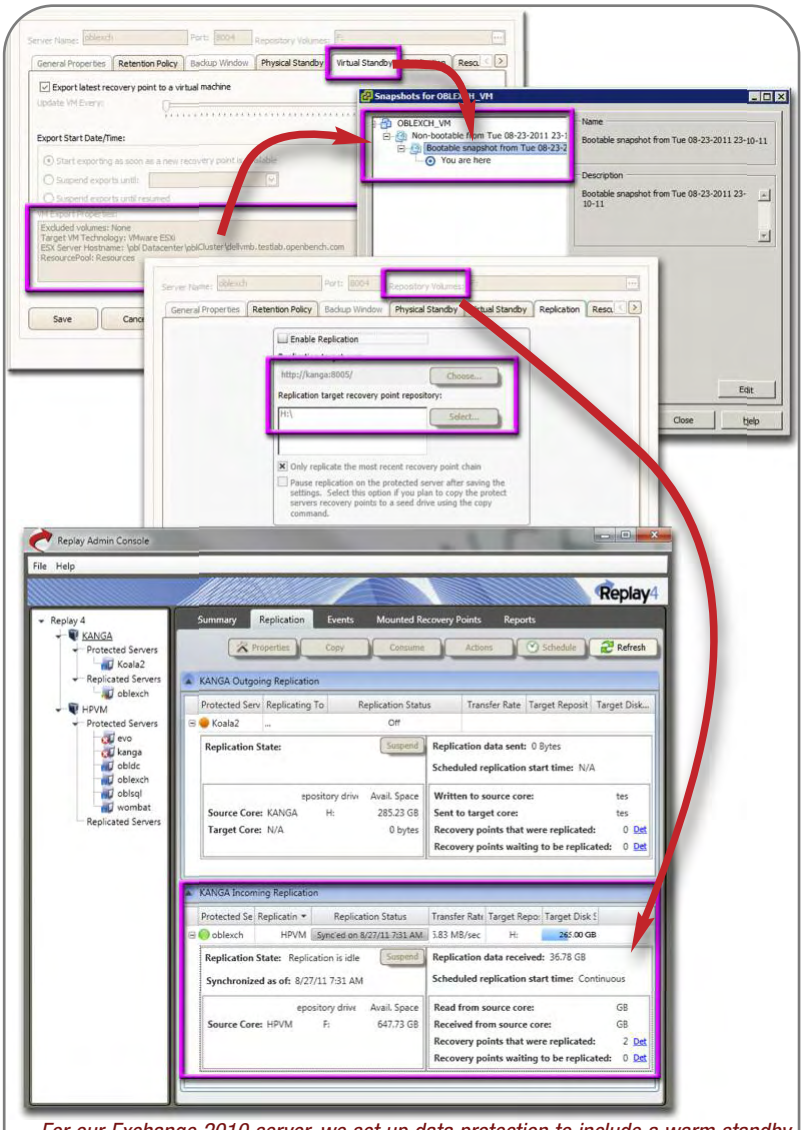

For our Exchange 2010 server, we set up data protection to include a warm standby VM. on an ESXi host. Following each incremental backup, AppAssure sent an update snapshot for each disk on a standby VM to its ESXi host.

AppAssure can automatically turn any incremental backup into a bare-metal recovery CD or update user files on a hot standby physical server. Similar DR functionality can found within Backup Exec; however, AppAssure offers a unique and more powerful option derived through the use of block-based disk imaging.

ESXi and ESX hosts can update VMs that are not in a running state. That allows an AppAssure Core server to use incremental backups of any physical or virtual system to generate disk snapshots for a warm standby VM, which is a complete copy of the original system down to the state of the OS.

The standby VM boots directly from a standard host datastore, immediately takes the identity of the original system, and exhibits full disk and network I/O performance. When a protected AppAssure client crashes, a warm standby VM server can be booted in seconds and brought online processing data in minutes to satisfy the most stringent business continuity SLA with respect to an aggressive RTO.

A key to keeping a warm

standby VM updated via disk snapshots is the minimal amount of CPU overhead imposed on the AppAssure Core server when processing incremental updates. The entire process of sending an incremental backup from the client to processing, and saving the update on the AppAssure Core server can be repeated over intervals measured in minutes. As a result, the AppAssure Core server can quickly send a corresponding disk snapshot to the host of the warm standby VM. As a result, the RPO for a business continuity SLA can be just as aggressive as the RTO.

What's more, AppAssure backups are self contained, which allows backups on one Core server to be replicated to another Core server for an HA scenario. Replicating the small, compressed and, deduplicated incremental backups puts minimal stress on a LAN infrastructure and is very well suited to WAN infrastructure. As a result, IT garners greater efficiency by replicating incremental backups to an off-site secondary AppAssure Core server via a WAN and using the off-site server to generate disk snapshots for an off-site, warm standby VM.

"IT garners greater efficiency by replicating incremental backups to an off-site secondary AppAssure Core server via a WAN and using the off-site server to generate disk snapshots for an off-site, warm standby VM."

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) **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
2) **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
3) **Assured Recovery 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
4) **Universal Recovery to Any Hardware for DR:** Recover sytems using a bare metal recovery CD, a hot standby server or a warm stanby VM
5) **Universal Recovery at Any Granulariy:** Recover sytems, disk volumes, files, or application-level data items, such as email messages and database tables
6) **Automate incremental Backups For Near-Lossless RPO:** Incremental backups with both deduplication and compression can be automated to run in intervals as short as every 5 minutes for minimally spaced restore points
7) **Compress and Deduplicate Data on any Backup Process:** AppAssure can apply both compression and data deduplication to incremental backups

AppAssure backup and replication 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. With the installation of a client service that includes a filter driver that updates a changed-block log in the system volume directory of a disk, AppAssure leverages disk image technology to generate fast, low-impact, incremental backups. The incremental backups are maintained as restore points on an AppAssure Core server and automatically rolled up to a full synthetic backup when accessed for a recovery process.

Most important, the block-based architecture of AppAssure enables the universal application of a functions and features with all backups, while Backup Exec 2010 R2 has a list of features and functions that are incompatible. For an SMB site, feature and function incompatibility can excessively burden IT administrators with complex data

protection procedures.

Moreover, AppAssure requires a minimal amount of computer processing time to support and optimize incremental backups that are up to 40X faster and 100X smaller by using changed disk blocks in place of a file-based backup. With minimum overhead IT is empowered to implement a DR plan with aggressive RTO and RPO levels to support any SLA required by Line of Business executives.

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.