



Assuring Successful Exchange Recoveries

By Mark Arnold

Abstract

Microsoft Exchange is often the single most important application for the businesses that use it. Should Exchange or Exchange data be unavailable, the entire organization comes to a halt until, at least temporary, adjustments are made. A consistent method of validating the recoverability of Exchange is paramount to any everyday or disaster recovery plan.

Implementing a consistent method can be difficult in IT shops that are strapped for resources; staff, time, or budget. Replay from AppAssure is designed to streamline Exchange backup, verification, maintenance, and recovery in terms of the time those processes take, the storage they consume, and the human resources they require.

“20% of unplanned Exchange downtime is caused by Exchange database failure.”

—Gartner

Verifying Backups

Most Exchange administrators spend all their time making sure the system stays up and running. They have little or no time to validate the integrity of the backups and, therefore, have no idea whether the backups will be useful if needed. Backups to tape have a verify option, which merely indicates that the data on the tape is what was sent to the tape. It does not verify that the data on the tape can either be taken off the tape or, that once back on the hard drive after a restore, is of any use.

In order to first verify the Exchange is properly backed up and can be restored and then actually restore Exchange, organizations typically invest in a combination of labor, hardware, and separate, expensive software solutions. The procedures required to recover Exchange using disjointed technology tend to be outside, or at best just within, the skills of the onsite administrator and the time available.

Recovery SLAs

Given all the current attention to RPO and RTO (recovery point and time objectives), implementing a system for Exchange recovery is more important than ever. The idea of an SLA (service level agreement) is becoming less relevant because businesses want their e-mail available - period. Ninety-eight percent availability is not good enough. Fast recovery of mission-critical systems, in this case Exchange, is the key.

“90% of Exchange administrators are unable to carry out preventative maintenance of the Exchange databases until after a disaster strikes.”

— Microsoft

The Exchange recovery challenge really lies with the RPO. Businesses have a range of requirements here. Some want a recovery to the previous night's condition, indicating that their business can withstand quite a large gap in their e-mail. Other organizations require an hour, fifteen minutes, or less. Still other businesses require absolute, up-to-the-moment-of-failure recovery. Short of complex clustering or (SAN) synchronous replication, what application can achieve that level of recoverability?

Probably the most important thing that administrators overlook when recovering information is that the Exchange databases must be validated when they have been restored. In the case of large databases this will take two, three, or even more hours to complete.

Then there is the Microsoft Exchange application. Little, yet important, changes are always being made to the operating system and Exchange. Transient files, which do not typically get included in the backup, are added and deleted. An application that can capture the full, current view of the entire Exchange server, transient files included, can reduce RTO and RPO to minutes, which is far less time than required by a complete tape-based recovery.

How can organizations preserve and improve Exchange service through efficient backup, verification, and maintenance processes without stretching or adding staff resources?

Replay for Exchange – The Concepts

Consider a Microsoft Exchange server running several thousand mailboxes. The Exchange administrator will spend at least an hour a day making sure that tape or disk drives will accommodate its backup, conducting the backup, and then verifying datastores. With more employees working flexible hours, IT departments are on the hook to offer uninterrupted service for more than the normal eight-hour day. An hour of interrupted Exchange service at any time during extended working hours is unacceptable. Users do not appreciate poor performance from their e-mail just because the processor is busy with the backup job.

And, since more e-mail is being stored within Exchange, the nightly maintenance processes take longer and require more processor time. So as Exchange datastores grow, so do the staff and process time requirements.

Replay answers those administrative challenges by eliminating the daily hit an Exchange server takes with the daily backup process. *Replay* spreads the backup load over the entire day and allows the administrator to recover e-mails that might have just been deleted, thereby eliminating delays in recovery.

“Our last Exchange recovery took over eight hours due to corruption, and e-mail was down for the entire company.”

Furthermore, *Replay* allows organizations to gain availability and recoverability enhancements without additional expenditure in the procurement or enhancement of existing storage area networks (SAN). Their investment in tape backup solutions is not rendered obsolete. *Replay* works with existing storage infrastructures. When a SAN is in place, *Replay* stores continuous, consistent application snapshots on the SAN. Those snapshots are protected by disaster recovery systems that are already protecting the SAN. This feature is an improvement over SAN-based snapshots that are not application-aware and need further processing before they can be used for restoration. That further processing can take from ten to hundreds of hours depending on the size of the Exchange environment. When tape backup systems are in use, that infrastructure is directed to the *Replay* Server.

Replay also introduces the concept of continuous application protection (CAP) to the entire Exchange server. *Replay* examines each volume-block on the storage volumes and sends those to a *Replay* Server where they are stored in a manner that allows the Exchange server to be recovered to a clean, consistent state. The fact that *Replay* examines changed blocks from the volume rather than changed files or even changed file-blocks is key to its administrative benefits.

Figure 1 below shows the *Replay* Console. The console protects all of the volumes on the Exchange server or Exchange cluster node.

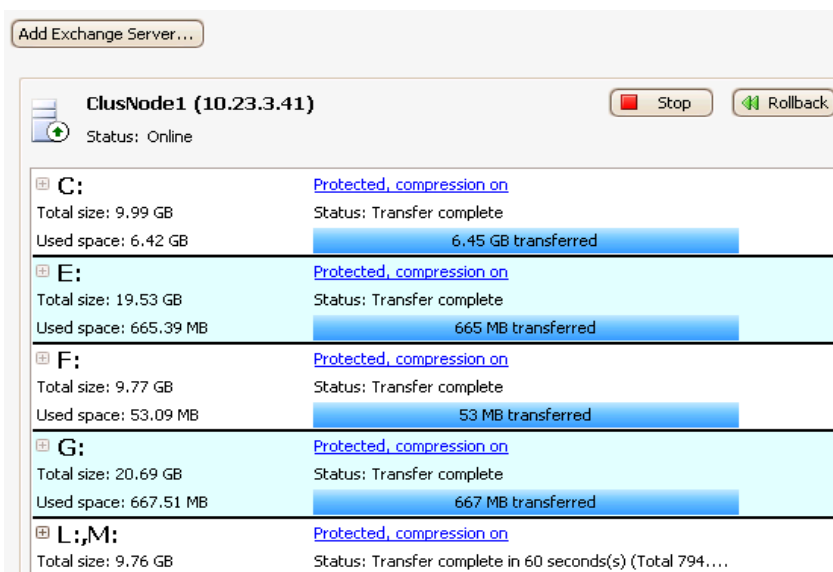


Figure 1 – *Replay* Console

The *Replay* Agent installed on the Exchange server dispatches only those small sections of data that have changed to the *Replay* Server. This function is in contrast with dispatching a 5MB/1MB log file that includes the space consumed by an individual e-mail arriving, or even an entire information store complete with white space. The *Replay* Agent also allows you to take a virtual image of the server operating system, sending the changed volume-blocks to the *Replay* Server in the same

“We don’t have a DR plan for our mission-critical Exchange environment – it’s too expensive and complex.”

way that Exchange data is transferred.

This feature provides several advantages. If your entire server fails you can bring up another similar server using PXE boot technology and use *Replay* to drop the entire contents of the server onto that new platform. This function is similar to how disk imaging applications work, in that blocks of data rather than individual files are pushed out to the new server. Of course, disk images can only deliver your server back to its last imaged state. And disks are rarely imaged. So relying on a disk image for backup means you might be returning to a long-ago state, maybe even to initial install and configuration.

Some imaging or image-based backup applications can only bring a server back to life as if it had “blue screened,” meaning that everything, from disks to applications, has to undergo a consistency check. *Replay* is able to use its state-aware logic to restore a server to the consistent state.

The frequency at which *Replay* captures information and secures the data can be adjusted from as often as every few minutes to an hour or more, enabling rollbacks to restore a server to an operational state in a short period of time. Obviously, a properly configured Exchange server will roll the transaction logs forward when an information store is reverted to a prior state. Once the Exchange data has been secured on the *Replay* Server, the roll forward can take place.

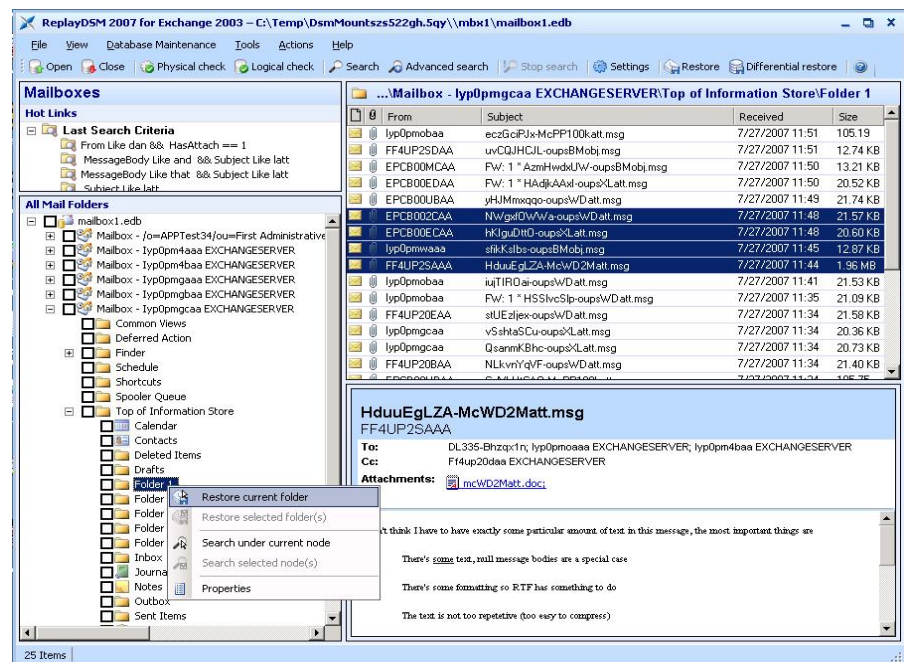


Figure 2 - *Replay*DSM Console

Using the data blocks it has captured, *Replay* internally formulates a complete information store, carries out a verification of that database and marks it as consistent. This backup is then available for any purpose from searching to restoration through the DSM console (Figure 2 above).

Finally, *Replay* minimizes the storage space that is required to process recovery through the creation of virtual recovery points. Traditionally, an Exchange

server has to maintain sufficient space to carry out restores to the recovery storage group (RSG) as well as space within the information store for deleted items retention (items users have deleted from their deleted items folder but are retained on the server to allow recovery). This additional space can be considerable and, since Exchange 2007 permits five, unlimited size, information stores, *Replay's* approach meets Exchange 2007 requirements without adding storage space to the already stretched IT budget.

Replay Helps Maintain Service Levels

Replay takes a near constant stream of information and intelligently stores it in a repository so that an Exchange information store can be rolled back to a previous state quickly. For example, let's say database corruption has occurred and a complete restoration is required. The information store is offline and users are demanding a speedy return to service. Under conventional backup methods the entire information store would have to be restored from disk or tape and then checked for consistency. This process will take close to two hours from start to finish. Figure 3 compares the RTO, RPO, and corruption detection capabilities of different DR solutions.

	Continuous Application Protection	SAN-Based Replication	File-Based Replication	Traditional Backup
RTO	Under 5 minutes	Under 1 hour	Under 1 hour	Days
RPO	15 minutes	0	Minutes	Up to 24 hours
Corruption detection	Yes	No	No	Yes
Point in time	Yes	No	No	No
Granularity	E-mail message through entire server	Entire server	Exchange data only	Exchange data only
Test ability	Automatic	Difficult	Difficult	Difficult
Scope	Entire server	Entire server	Exchange data only	Exchange data only
Cost	Thousands	Millions	Thousands	Thousands

Figure 3 - DR Solutions Compared

With *Replay*, the situation is different. The management console will help an administrator determine which recovery point is the most appropriate to restore. From there it will determine which data blocks to place back onto the volume that stores the database in order to recover the Exchange information store to the required previous state.

The data to be recovered might still be in the order of gigabytes, especially on a large information store. Since *Replay* works with volume-based blocks, the transfer over the network is far more efficient than other products that work with file blocks or complete files. Recovery of a physical information store that is perhaps 100GB in size can, therefore, be achieved by the transfer of a few hundred megabytes over the network. Since *Replay* has already ensured that the data it is transferring back to the production server is consistent, the store is

guaranteed to mount immediately. Thus, the timeframe from failure to recovery is reduced from hours to minutes.

The same procedure is adopted where the information store is affected by a virus infection. Restoring the store to the point just before the virus entered the Exchange server, updating the anti-virus signatures before mounting the store, and playing the log files will allow the anti-virus application to remove the infection from the store as the transactions are played into it. This feature ensures the users experience no data loss.

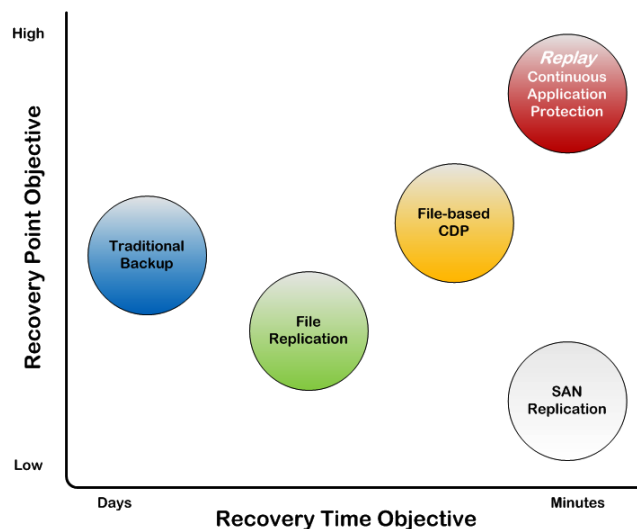
Beyond Application Protection

Since *Replay* works with volume-blocks it is able to protect the entire Exchange application, Exchange data, and the operating system, something other solutions cannot do.

For example, in a conventional SAN-protected environment the Exchange server always boots up in a dirty condition, necessitating a *chkdsk* process and introducing another delay in the recovery of the system partition.

When an Exchange server suffers a catastrophic failure, the normal process of recovery takes a considerable length of time following the delivery or replacement hardware. The operating system has to be installed, the service packed and patched, and then the Microsoft Exchange Server has to be reinstalled using the disaster recovery switch before the actual information stores can be restored.

Replay works together with a PXE bootable network card to stream the entire data set back to the new server. Again, this function occurs at the volume-block level, which enables the server to be returned to service much faster than it would be by using a conventional recovery procedure. And because *Replay* makes sure that the databases are in a consistent state, the information store will mount at the end of the recovery process.



Conclusion

Replay is the fastest, most flexible and cost-effective solution that can keep a Microsoft Exchange server protected, keeps the server available, and returns the server to production as quickly as possible - whether an Exchange outage was planned or not.

Replay is the only solution on the market today that provides a recovery time objective in the same timeframe as SAN-based replication techniques and to granular recovery point that simply cannot be matched by any other applications or technologies.

Replay is also the only solution that captures the application (Exchange), the operating system, and associated data (Exchange data store) in a consistent manner, whether the server is booted from local disks, from a LUN on the SAN, or within a VMware environment. And *Replay* will always inform you of detected corruption in the data it is protecting, giving you ample opportunity to take corrective action.

About the Author

Mark Arnold is an independent messaging consultant working in the Philadelphia, Pennsylvania area, advising customers on Microsoft Exchange designs and implementations for medium and large businesses. He is an MCE+M and has been a Microsoft MVP in the Exchange server discipline since 2001.

AppAssure Software, Inc.
1925 Isaac Newton Square East, Ste 440
Reston, VA 20190
Fax: 866-602-7934
Phone: 703-547-8686

For more information or to download your free trial,
visit us online: www.AppAssure.com

© 2008 AppAssure Software. All rights reserved. The information contained in this datasheet is subject to change without notice. AppAssure Software assumes no responsibility for any errors that may appear. AppAssure Software and the AppAssure Software logo are all trademarks of AppAssure Software. All other trademarks and registered trademarks are the property of their respective owners.