

EXC01: To Backup or Not Backup: That Is The Question

Michael B. Smith

The Essential Exchange

michael@TheEssentialExchange.com



Who is Michael B.?

- Remember the B! 😊
 - And yes, that really is my name!
- Long-time Exchange guy
 - Since 1996 and Exchange 5.0
- Eight year Exchange MVP
- Consultant and migration expert
 - Exchange, Active Directory, System Center

Do you do social media?

- Then



Find us on
Facebook

- <http://facebook.com/ExchangeConnections>
- Discussions, announcements, and conversations about this and future Exchange Connection events

Agenda

- Introduction
- VSS through the ages
- Basic VSS
- Basic Disk
- Exchange Backups & Restores with VSS
- Disk Storage
- Summary
- Bonus: Dumpster 1.0 vs. 2.0 vs. Litigation Hold

Introduction

- VSS – Volume Shadow Copy Services
- Introduced in Windows XP/Server 2003
- Enhanced in Windows Vista/Server 2008
- Even better in Windows 7/2008 R2
- Not a uniquely Microsoft solution
- Available as a hardware feature on many SANs
- Snapshot vs. Clone

Current Status

- **As of mid-2011:**
 - OS installed base: about half Windows 2003, half something later
 - Exchange installed base: somewhat under 1/2 Exchange 2003, half something later
- **People are upgrading slowly**
- **Why so?**
- **Most common answer: no major reason to upgrade**

What's the problem on Exch2003?

- First version of the Exchange Writer to integrate with VSS (Writers will be discussed later)
- Can only back up a single storage group at a time
- All IOPS happen on the active server
- VSS preferred backup mechanism – but streaming backups not yet “deprecated”

What's the problem on Exch2007?

- Microsoft changed the “implied contract”
- With Exchange 2007, streaming backups are deprecated, remote backups impossible (API removed)
- Windows 2008 does not include ntbackup
- Windows Backup:
 - Shadow copy only
 - Doesn't speak Exchange until SP2
 - Full volume only (excepting system state)
- So... no native backup tool for Exchange

What's the problem on Exch2010?

- Windows Backup functionality is quite limited (same in Exch2010 RTM as in Exch2007 SP2)
- Entire server or volume by volume backup
- Entire application restore (either in place or to alternate location)
- If using RSG, everything is manual

Or.... Is there?

-is there a way to do native backups?
- If so, how?
- Windows 2008 includes new tools
- VSS SDK included similar tools for Win2003
- Think on that, while you learn about VSS...

Basic VSS

- Volume Shadow Copy Service – VSS
- Maintains the “state” of a volume
- Maintains multiple copies
- Can revert-at-once
- Crash-consistent
- By default, it is not Backup/Application-consistent
- Third parties can add significant functionality

Basic VSS Definitions

- **Provider**
 - Low-level interface to hardware, a shim that monitors and filters file system access to create and maintain shadow copies
- **Requester**
 - Backup or restore application
- **Writer**
 - Part of an application that coordinates its I/O with VSS (e.g., flushes buffers to disk)

Basic VSS – Hardware Support

- Windows comes with a basic provider; third parties (i.e., SAN vendors) can add their own
- VSS snapshot capture's the volume state – effectively a copy of the volume bitmap and MFT
- Copy-on-write / fragmentation increase
- Copy exists in parallel with live volume
- May be many copies
- Snapshot vs. clone

Basic VSS – Note

- The following slides represent a conceptual background for understanding how VSS and Exchange and Windows work together.
 - If you need to know exactly how bits-and-bytes work, head for MSDN and TechNet

Basic Disk – Volume Bitmap

Block #	0	1	2	3	4	5	6	7
	y	n	n	y	n	n	n	y
Block #	8	9	10	11	12	13	14	15
	n	y	n	n	y	n	n	y
Block #	16	17	18	19	20	21	22	23
	n	n	y	n	n	y	n	n
Block #	24	25	26	27	28	29	30	31
	y	n	n	y	n	n	y	n

The Volume Bitmap is a “simple” array containing a list of all blocks on a mass-storage device showing which are in-use and which are not

Basic Disk – Master File Table (MFT)

FILE:NTLDR.SYS								
10	11	122	18	24	36			
FILE:AUTOEXEC.BAT								
92	13	8	64	77				
FILE:IO.SYS								
1023	934	12	431	432	136	137	138	139
982	981	980	989					

Among other items, the MFT contains the file names on disk and the list of blocks that contain the file.

Basic VSS – Software Support

- VSS included in Windows XP/2003 and above
- VSS Writers work with the VSS Provider to ensure that their applications are backup consistent, not just crash consistent
- Writers include Exchange, Hyper-V, Registry, System, IIS, etc.
- Stages of a Snapshot:
 - PrepareBackup
 - Freeze
 - Thaw
 - PostBackup

Exchange Backup with VSS

- Involves stores, log files, and system files
- Exchange Writer ensures that disk buffers are flushed, and log files are flushed and the database header is consistent
- Backup is “dirty”, requires recovery to run (because backing store is not flushed)
- Must verify backup with `eseutil /k`
- Transaction Logs are flushed on success

Exchange Restores with VSS

- Involves stores, log files, and system files
- Per storage group
- All stores in storage group must be dismounted
- Must be careful with “old” log files
- Recovery runs when stores are remounted (makes RSGs a challenge)
- ...or, you can run recovery on a snapshot!

Simple VSS Backup Script

- Using Diskshadow

```
set verbose on
set context persistent
set metadata
c:\temp\20090320184921.cab
begin backup
add volume C: alias shadow_C
create
expose %shadow_C% g:
exec c:\temp\20090320184921.cmd
end backup
exit
```

Simple VSS EXEC file

- More complex. Multiple phases involved.
- For every storage group, copy all stores, log files, and system files from source to destination. If any copies fail, abort the entire backup
- For every store (mailbox and public folder) verify the store (`eseutil /k`). If any verify fails, abort the entire backup
- If everything is good, signal to delete log files

Where are we now?

- We understand, in detail, how backups work in modern Exchange
- This provides a base level of knowledge for understanding how to protect business data stored within Exchange
- Now, compare and contrast other features...

Disk Storage

- Some concepts
 - RAID-0
 - RAID-1
 - RAID-01 and RAID-10
 - RAID-5 and RAID-6
 - SAN
 - DAS
 - JBOD
 - DAG

Disk Storage: #2

- Total copies of data?
 - RAID-0
 - RAID-1 / RAID-01 / RAID-10
 - RAID-5 / RAID-6
 - SAN
 - DAS / JBOD
 - DAG

Disk Storage: #3

- Data retention
 - Backups
 - Dumpster 1.0
 - Dumpster 2.0
 - Litigation Hold
 - Single Item Recovery
 - Lagged DAG
 - DAG
 - Legal Requirements

To Back Up or Not to Back Up

- **What trumps?**
 - Legal requirements
 - Process requirements
- **Should time allow – Dumpster 1.0 vs. Dumpster 2.0 vs. Litigation Hold (after the break!)**

Presentation Availability

- You can download the slides from here:
- http://www.devconnections.com/updates/LasVegas_Fall_11/Exchange

Your Feedback is Important

Please fill out a session evaluation form
drop it off at the conference registration
desk.

Thank you!

Dumpster v1.0

- Let's take a poll...
- Can a user permanently remove something from their mailbox?
- No.....
- Yes.....

Establishing the picture

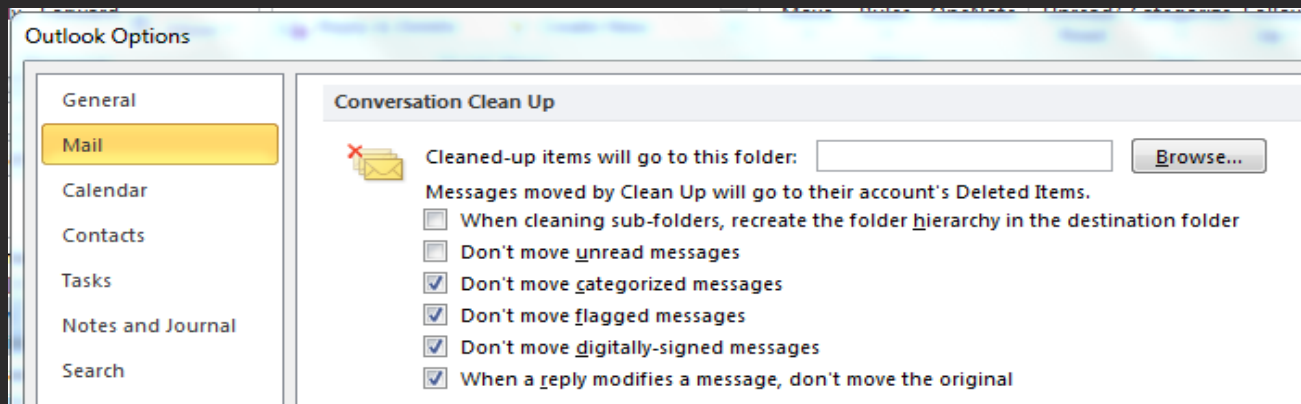
- **Some quick background:**
 - In Exchange, everything is either an item or a folder.
 - Items include mail messages, appointments, contacts, tasks, journal entries – you get the idea
 - Folders include mailboxes, user controlled folders (Inbox, Deleted Items, etc.) – plus a lot of folders a user can't see

Now, fill in the background

- Folders can contain items and additional folders (subfolders)
- Both folders and items have attributes that are used by Exchange and Outlook to determine how to treat those objects
 - Often referred to as a MAPI property
 - Key attributes for this discussion:
 - Hidden
 - Deleted


Dumpster v1.0

- When a user clicks “Delete” in Outlook, the item is moved to the mailbox’s Deleted Items folder
 - Outlook soft delete



Dumpster v1.0

- When a user clears the Deleted Items folder (manually or by a setting in Outlook)

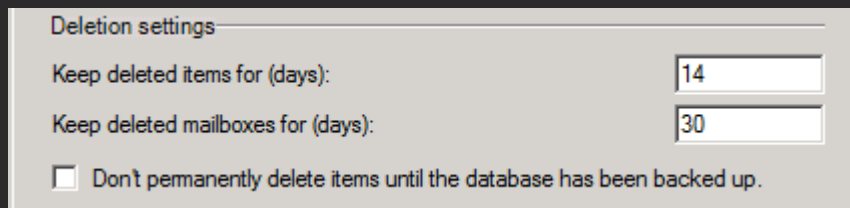


or does a shift-delete the item is “moved to the dumpster”

- Item has the “Deleted” attribute set on it
- Also ptagDeletedOnFlag (timestamp)

Dumpster v1.0

- Items with “Deleted” set are excluded from almost all Outlook folder & item views
- The `ptagDeletedOnFlag` is important
 - Key attribute associated with timeframe supported by Deleted Item Recovery



Deletion settings

Keep deleted items for (days):

Keep deleted mailboxes for (days):

Don't permanently delete items until the database has been backed up.

- DIR applies to deleted user-level folders
- Deleted objects are processed by store maintenance nightly

Dumpster v1.0

- Emptying Deleted Items or doing a shift-delete has two names:
 - Outlook Hard Delete
 - Store Soft Delete
- Before DIR has expired, items can be recovered by user (or administrator) using Recover Deleted Items in Outlook or OWA
- Will be restored to folder from which the deletion occurred

Dumpster v1.0

- After DIR is expired, recovering the object requires restoring from a backup
 - User must know timeframe when item was deleted
 - Admin must perform restore (whether to RSG or using third-party tool – almost certainly still an admin function)
- Technical name for DIR expiration:
 - Store hard delete – an object's DB pages are placed in available pool

Dumpster v1.0

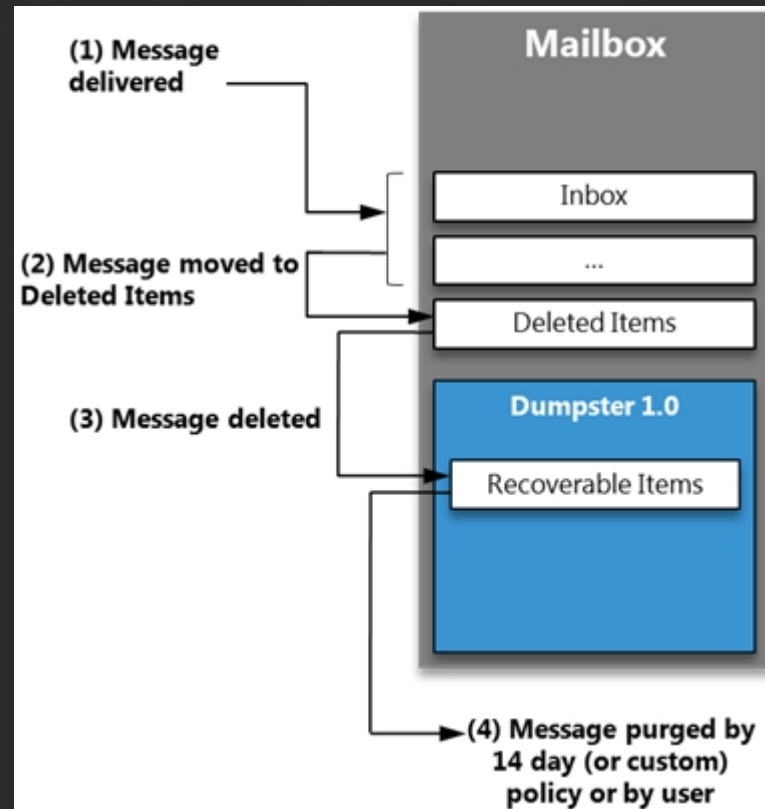
- **What's the problem with this scenario?**
 - Within Recover Deleted Items, end-user can purge items currently in the Dumpster
 - !!!!
 - Yep.
 - Provides an inappropriate level of plausible deniability
 - User can “prove” they didn't send or receive item
 - Admin (anyone with FC MB access) could potentially cause havoc

Dumpster v1.0

- Any other problems with v1.0?
- Sure:
 - You lose dumpster data when you move a mailbox
 - Dumpster data can't be searched or indexed by native products
 - Users can modify objects ("Edit" in Outlook) and again have an inappropriate level of plausible deniability

Dumpster v1.0

- Discuss general process



Dumpster v2.0

- Designed to address the deficiencies associated with Dumpster v1.0:
 - User interface has not changed
 - Translation between v1.0 and v2.0 semantics and behavior is provided by the Client Access Server
 - A user is not aware that they have lost plausible deniability
 - Higher overhead – must be enabled on a per mailbox basis

Dumpster v2.0

- Provides a legal compliance experience that should meet the needs of most companies (a/k/a “legal hold”)
 - Dumpster data moves with the mailbox
 - Is now indexed and discoverable / searchable (by admin)
 - Dumpster has a quota, separate from mailbox
 - Dumpster is per mailbox (not per folder)
 - User’s cannot purge their data
 - Item modifications are tracked

Dumpster v2.0

- In order to provide those features:
 - Dumpster v2.0 completely re-architected
 - No longer simply a database index/view
 - Now a folder named “Recoverable Items”
 - Located in NON_IPM_Subtree of mailbox
 - Hidden
 - Three subfolders exist:
 - Deletions
 - Versions
 - Purges

Dumpster v2.0

- By changing to a folder-based storage location, instead of a view-based storage location, three objectives are immediately easily achieved:
 - Simple to move dumpster data when you move the mailbox
 - Easy to index and search / discover data contained in dumpster (all subfolders)
 - Recover Deleted Items is now per mailbox, not per folder

Dumpster v2.0

- **Deletions subfolder**

- Equivalent to “Recover Deleted Items” view in Dumpster v1.0
- When a user does an Outlook hard-delete or empties their Deleted Items folder, the affected objects are moved to “Recoverable Items → Deletions”

Dumpster v2.0

- **Purges subfolder**

- If a user attempts to purge an item when using Recover Deleted Items, instead of the object actually being purged, the item is moved from “Recoverable Items → Deletions” to “Recoverable Items → Purges”.
- Plausible deniability is gone
- Other uses will be discussed later

Dumpster v2.0

- **Versions subfolder**

- Each time a user edits an item, the original item is stored in the Versions subfolder (copy-on-write)
 - For messages and posts, copy-on-write will capture changes in the subject, body, attachments, senders/recipients, and sent/received dates
 - For other items, copy-on-write will occur for any change to the item except for moves between folders and read/unread status changes
 - Drafts are exempt from copy-on-write

Dumpster v2.0

- Dumpster Quotas

- Default to the store quotas
- Can be set individually

- RecoverableItemsQuota

- Specifies the size limit for the Recovery Items folder. When you reach the quota limit, you can't put any more items in the Recovery Items folder. The “hard limit”.

- RecoverableItemsWarningQuota

- Specifies the quota for when a warning event is entered in Event Viewer. The “soft limit”.

Dumpster v2.0

- **Dumpster quotas continued:**
 - Separate limits for primary and archive mailboxes
 - Separate dumpsters for primary and archive mailboxes
 - Quotas for dumpster are intended to prevent denial-of-service on MB servers
- **DumpsterAlwaysOn has no current effect**

Dumpster v2.0

- To enable:

```
Machine: ex2010sp1.test.local
[PS] C:\>Set-Mailbox test.user -SingleItemRecoveryEnabled $true
WARNING: The single item recovery setting may take up to 60 minutes to take effect.
[PS] C:\>
```

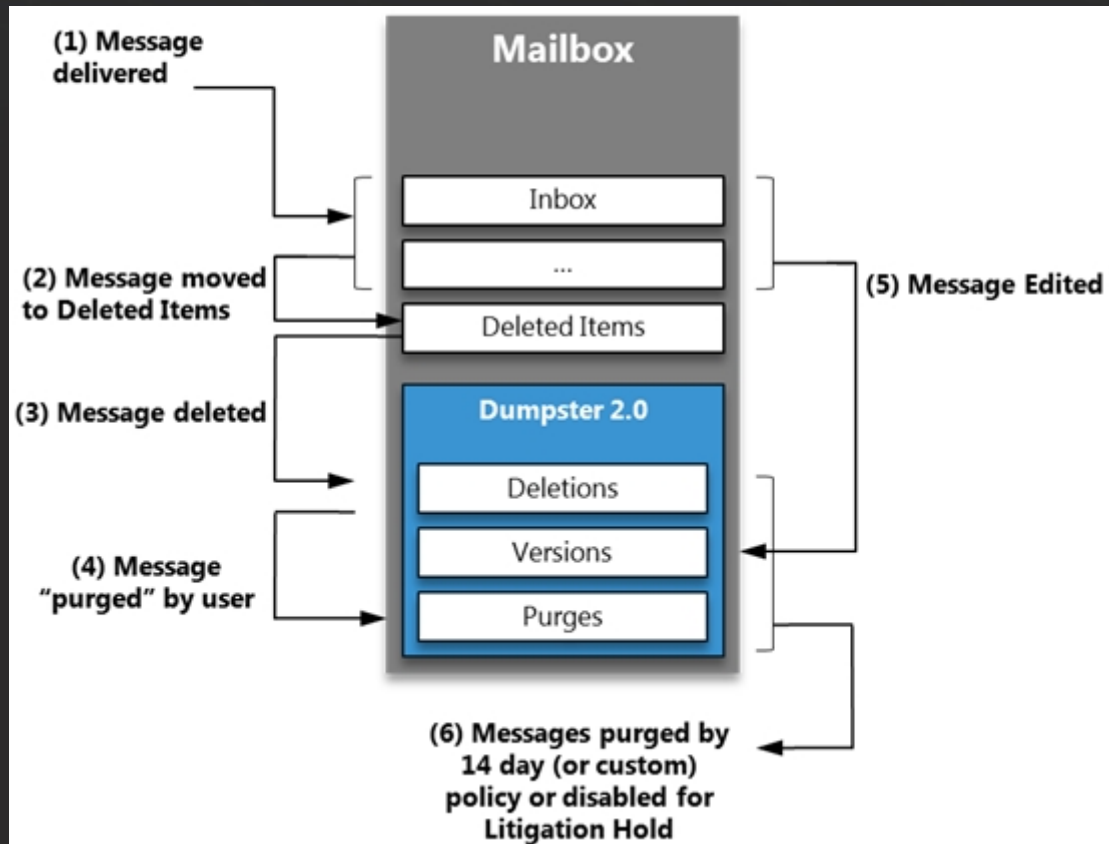
- The “60 minutes” is pretty-much bogus. However long AD replication takes in your environment is the real answer.
- Deleted Item Recovery still applies when SingleItemRecovery is set

Dumpster v2.0

- **SingleItemRecovery** is designed to address short-term retention of data
 - This is, data contained within the user's mailbox currently and that which is within the deleted item recovery timeframe
- **Deleted Item Recovery** defaults to the database value, but can be set per MB
 - Set-Mailbox -RetainDeletedItemsFor

Dumpster v2.0

- Discuss process:



Dumpster v2.0

- Demo the process
 - Get-Mailbox test.user | Select SingleItem*
 - Open Outlook for test.user
 - Four messages, each uniquely identifiable, then:
 - #1 – normal delete
 - #2 – Outlook hard delete
 - #3 – Edit message
 - #4 – Purge message from Recover Deleted Items

Dumpster v2.0

- **Open MFCMAPI**
 - <http://mfcmapi.codeplex.com>
- **Compare and contrast MFCMAPI**
 - In cached mode (CAS translation)
 - In online mode (no CAS translation)
 - Examine “Recoverable Items” and each subfolder and their contents

Dumpster v2.0

- For long-term retention of data, you use **LitigationHold**:
 - Set-Mailbox test.user
 - LitigationHoldEnabled \$true
 - Quotas no longer apply
 - Items in Deletions will be automatically moved to Purges when DIR expires
 - At that point, user can't recover, but admin can
 - All else is the same

Dumpster v2.0

- **Recovering purged data**

- Need membership in “Discovery Management” role
- Search and identify target item with ECP (Discovery Search) or EMS (Search-Mailbox)
- Export-Mailbox (RTM) or New-MailboxExportRequest (SP1)
- Need membership in “Mailbox Import Export” role (which needs to be assigned to a role group)

Quick Sidebar

- How to assign the Mailbox Import Export role to a management role group
- Can't be done from GUI
 - New-ManagementRoleAssignment -Name "Mailbox Import Export_Discovery Management" -SecurityGroup "Discovery Management" -Role "Mailbox Import Export"

Dumpster v2.0

- **Additional Considerations:**
 - Enabling SingleItemRecovery has an impact on size of Exchange database, and thus on backup and recovery timeframes
 - Consider Crowley's Law: if you keep a month in DIR you'll probably never need to do a mailbox recovery
 - **Ed Crowley – long-time Exchange MVP**
 - **Paraphrased**

Summary

- **Dumpster v2.0**
 - Adds significant functionality over v1.0
 - Provides assurance on information integrity
 - Allows for auditing and collection of information regarding user attempts to invalidate data integrity
 - Provides a full litigation hold experience
- **Questions?**