

# Beyond CIFS

David Kruse

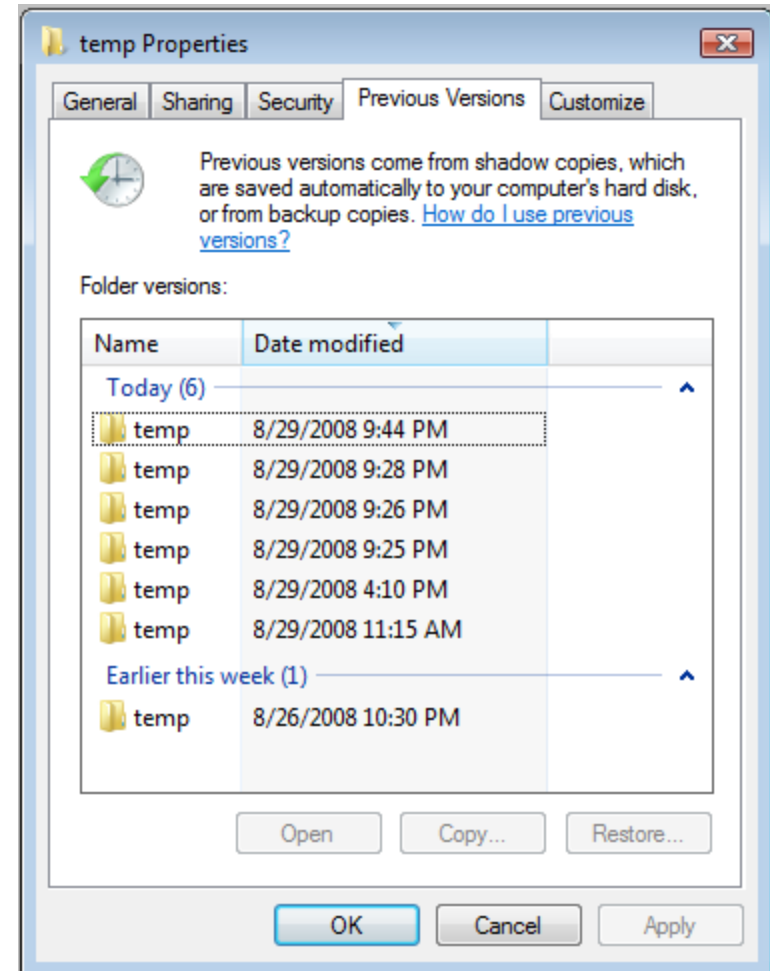
Microsoft

# Agenda

- ❑ Previous Version Access
- ❑ Access Based Enumeration
- ❑ Server Side Data Movement
- ❑ Symbolic Link Support in SMB2
- ❑ Durable Open Support in SMB2


# Previous Version Access

- ❑ Offers the user access to previous versions of files or folders.
- ❑ Reduces IT costs by enabling end-user to restore content.
- ❑ Native support in WinXP SP2+, earlier requires Shadow Copy Client (<http://technet.microsoft.com/en-us/windowsserver/bb405951.aspx>)



# Previous Version Access

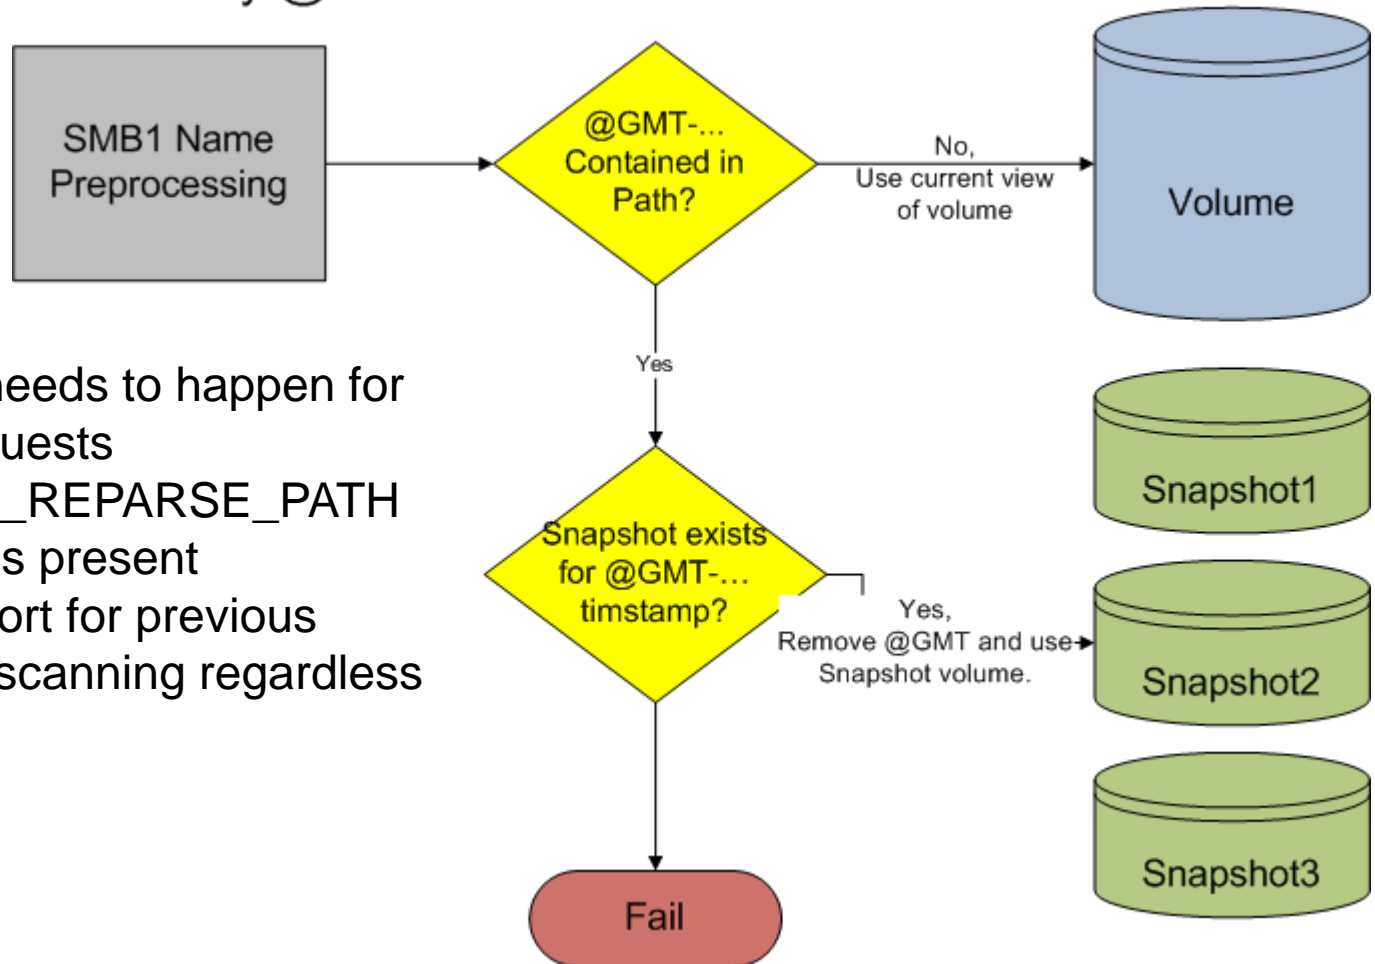
0	1	2	3	4	5	6	7	8	9	1	0	1	2	3	4	5	6	7	8	9	2	0	1	2	3	4	5	6	7	8	9	3	0	1
NumberOfSnapShots																																		
NumberOfSnapShotsReturned																																		
SnapShotArraySize																																		
SnapShots (variable)																																		
...																																		



- Implementation is aligned with volume snapshot features, but true versioning can be supported
- Client must expect that snapshots can be removed at any time

# Previous Version Access

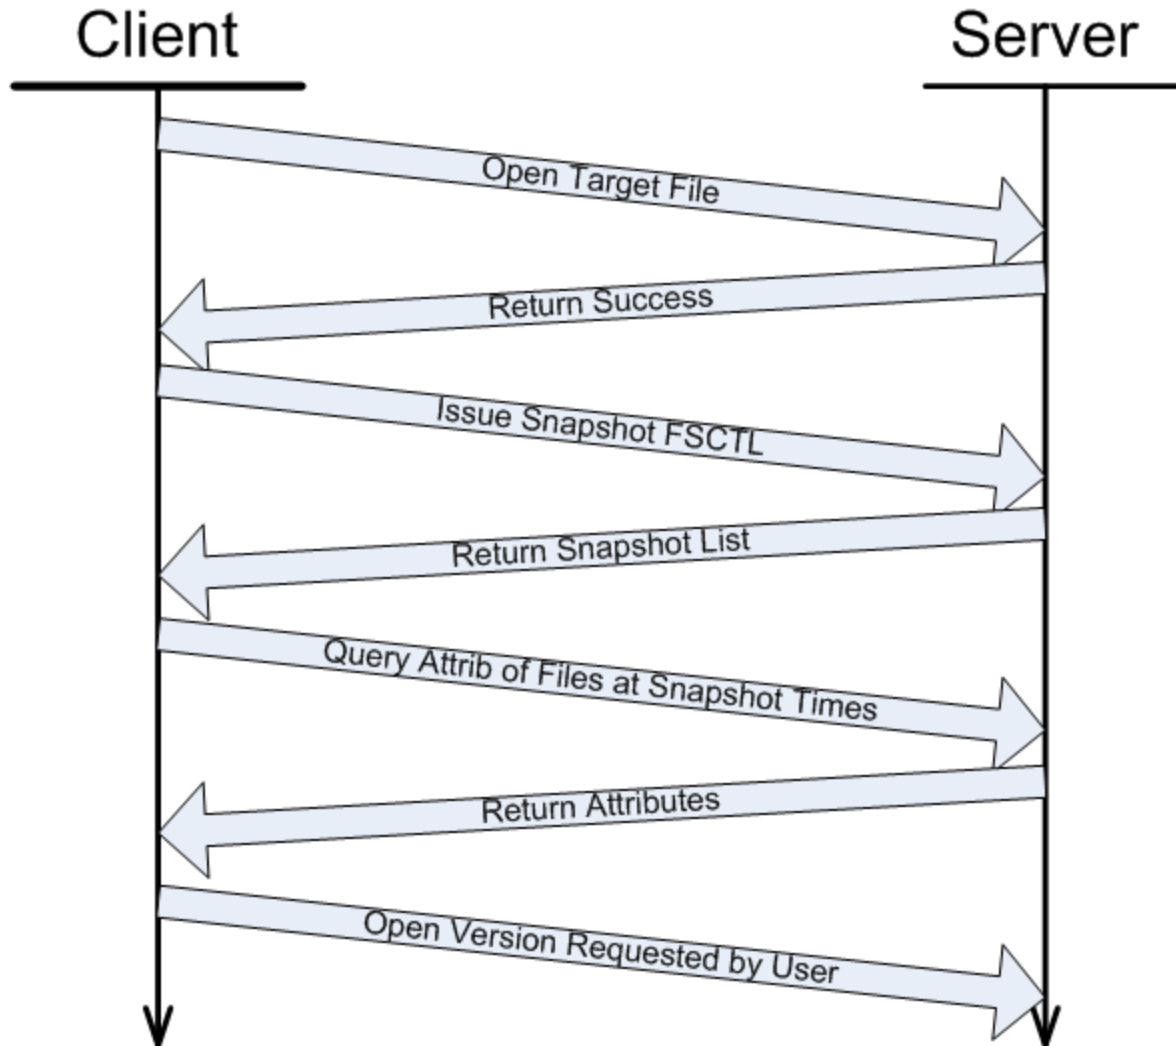
\Directory\@GMT-2008.09.25-08.30.00\File.txt



- Preprocessing needs to happen for all path-based requests
- SMB2\_FLAGS2\_REPARSE\_PATH hints that @GMT is present
- Downlevel support for previous versions requires scanning regardless of flag

- SMB2 Support is simpler:
  - Create Context contains @GMT token converted to LARGE\_INTEGER time value
  - Single create path and handle based command set mean that only create logic needs to be Previous Version aware.

# Previous Version Access



# Previous Version Access

- Downlevel (Win9x/NT4) Support
  
  
  
  
  
  
  
  
  
  
- <downlevel enumeration example here>



# Access Based Enumeration


- ❑ Provides a filtered view of the namespace based on user access
- ❑ View filters out files and directories which the user does not have READ\_DATA/LIST\_DIR access to *unless* the caller is privileged. (backup privilege)
- ❑ Enabled/Disabled at share granularity
- ❑ <http://www.microsoft.com/windowsserver2003/techinfo/overview/abe.mspx>

# Access Based Enumeration

- ❑ SMB/SMB2 server should filter directory enumeration results as desired
  - ❑ Always allow “.” and “..”
  - ❑ Handle wildcard and non-wildcard queries
  - ❑ Consider error mapping on create requests
- ❑ Server **MUST** set bit in tree connect response
  - ❑ Technologies like Offline Files require knowledge that they are seeing a filtered namespace

# Access Based Enumeration

## ABE Enabled vs. Not Enabled

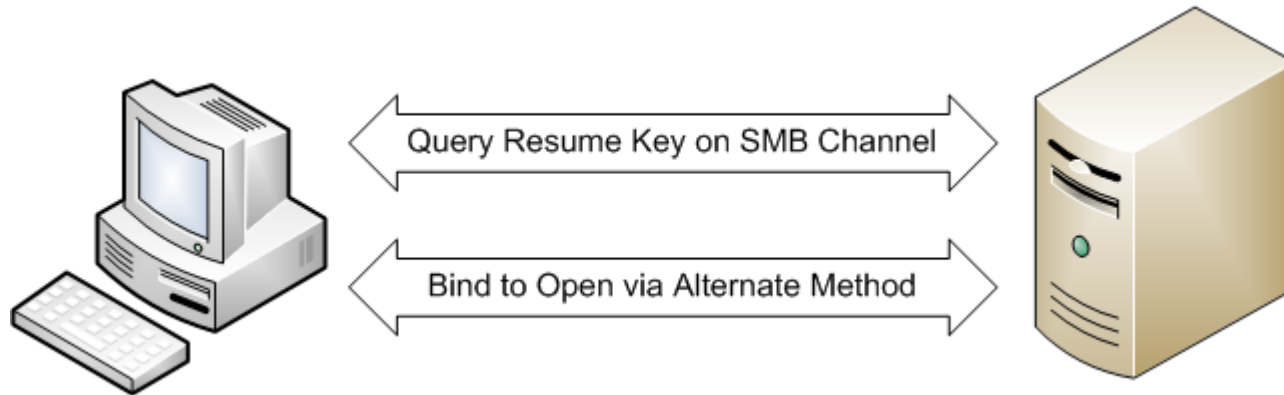


# Server Side Data Movement

- ❑ Resume Key interaction provides a portable identifier for an existing open to allow referencing in new scenarios
- ❑ Server Side Copy enables clients to request data movement between source and destination files without pulling the data to the client

- ❑ Resume Key Format
  - ❑ 24 byte opaque value, defined by you
  - ❑ Locally unique to the server
    - ❑ May be assembled from GUID, FileId, Open Time, Process Id, signature, etc.
  
- ❑ Why not simply use Fid?

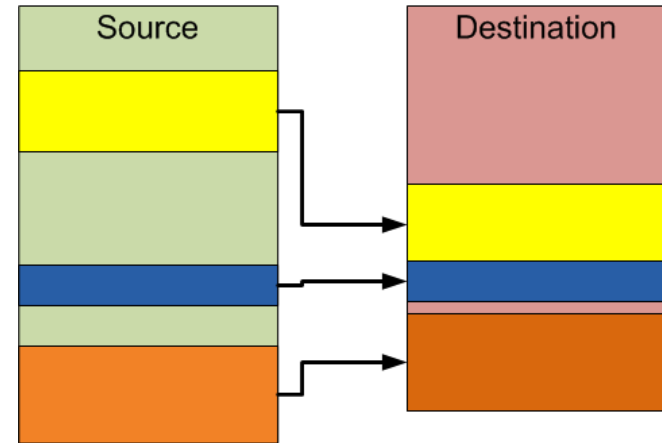
# Server Side Data Movement



<Perf Data Here>

# Server Side Data Movement

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



On Success or Failure, returns information on written data.

If Request exceeds configured limits, returns configured limits and `INVALID_PARAMETER`.

- Current Client Usage:
  - Vista+ for Remote<->Remote CopyFile
  - Vista+ for Offline Files Bitmap Upload
  
- <insert CopyFile perf chart>



# Symbolic Link Support for SMB2

- ❑ Allows SMB2 clients to process absolute and relative symbolic links.
  - ❑ (You may have heard of them before)
- ❑ Symlink evaluation always done by the client
- ❑ Vista client default policy locks down remote symlink evaluation
  - ❑ Local Computer Policy \ Computer Configuraton \ Administrative Templates \ System \ NTFS Filesystem \ “Selectively allow the evaluation of symbolic links”
  - ❑ Enable “Remote->Remote” and potentially “Remote->Local”

# Symbolic Link Support for SMB2

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
SymLinkLength																															
SymLinkErrorTag																															
ReparseTag																															
ReparseDataLength																Reserved															
SubstituteNameOffset																SubstituteNameLength															
PrintNameOffset																PrintNameLength															
Flags																															
PathBuffer (variable)																															
...																															

# Symbolic Link Support for SMB2

□ <Symlink format examples>

# Durable Open Support for SMB2

- ❑ Enables a client to re-establish an open held on the server after a disconnect has occurred.
  - ❑ Client requests this behavior during Create
  - ❑ Server can grant it as appropriate
  - ❑ Reconnect is not necessarily guaranteed by the server.
- 
- ❑ Supported on Vista+ clients.

# Durable Open Support for SMB2

- ❑ SMB2\_CREATE\_DURABLE\_HANDLE
  - ❑ Sent on create to request durability
  - ❑ Acknowledged in successful create response if durability is available
- ❑ SMB2\_CREATE\_DURABLE\_HANDLE\_RECONNECT
  - ❑ Sent on new connection to rebind to durable handle.
  - ❑ SessionId, TreeId must be valid

# Durable Open Support for SMB2

- ❑ Client is responsible for replaying operations in flight for which response was not received.
- ❑ Vista Client: Reconnect is retry based, not timeout based (Reconnect attempted 5 times)
- ❑ Server 2k8 holds durability after disconnect while oplock is present

# What else is there?

- SMB I
  - Reauthentication on Kerberos Ticket Expiry
  - Signing Key Protection
  - Extended Responses (with Maximal Access)
  - 32-bit Read/Write lengths
  - Native NT information levels via Pass-thru

# Questions?