Samba and Btrfs

A Snapshot of Progress

Jim Mcdonough
jmcdonough@suse.com

David Disseldorp
ddiss@suse.com

Overview

- Btrfs
- Snapper
- Samba Enhancements
  - Transparent Compression
  - Server-Side Copy
  - Snapshots
    - Previous File Versions
    - Remote Share Snapshots

Btrfs and Snapper

Btrfs

- Next generation filesystem for Linux
- Resilient
  - Checksumming of data and meta-data
- Integrated redundancy
  - Multi-device support (mirrors and stripes)
- Snapshots
  - File range (clone) or subvolume granularity
- Transparent compression
- De-duplication

Snapper

- Separate utility for managing Btrfs snapshots
  - Timeline, application or user initiated
  - Stores associated metadata
- Compare and revert changes
- Authorization
- D-Bus interface

Samba Enhancements

Samba
Btrfs Integration

- SMB protocol can expose Btrfs features to clients
  - Compression
    - FSCTL_SET_COMPRESSION and FSCTL_GET_COMPRESSION
  - Server-Side Copies
    - FSCTL_SRV_COPYCHUNK and FSCTL_SRV_COPYCHUNK_WRITE
  - Snapshots
    - Previous File Versions
    - File Server Remote VSS Protocol (FSRVP)

- Windows Explorer can expose these features to users

Compression
Samba + Btrfs

- User flags compression at a file or directory level
  - Files inherit compression flags from parent directory

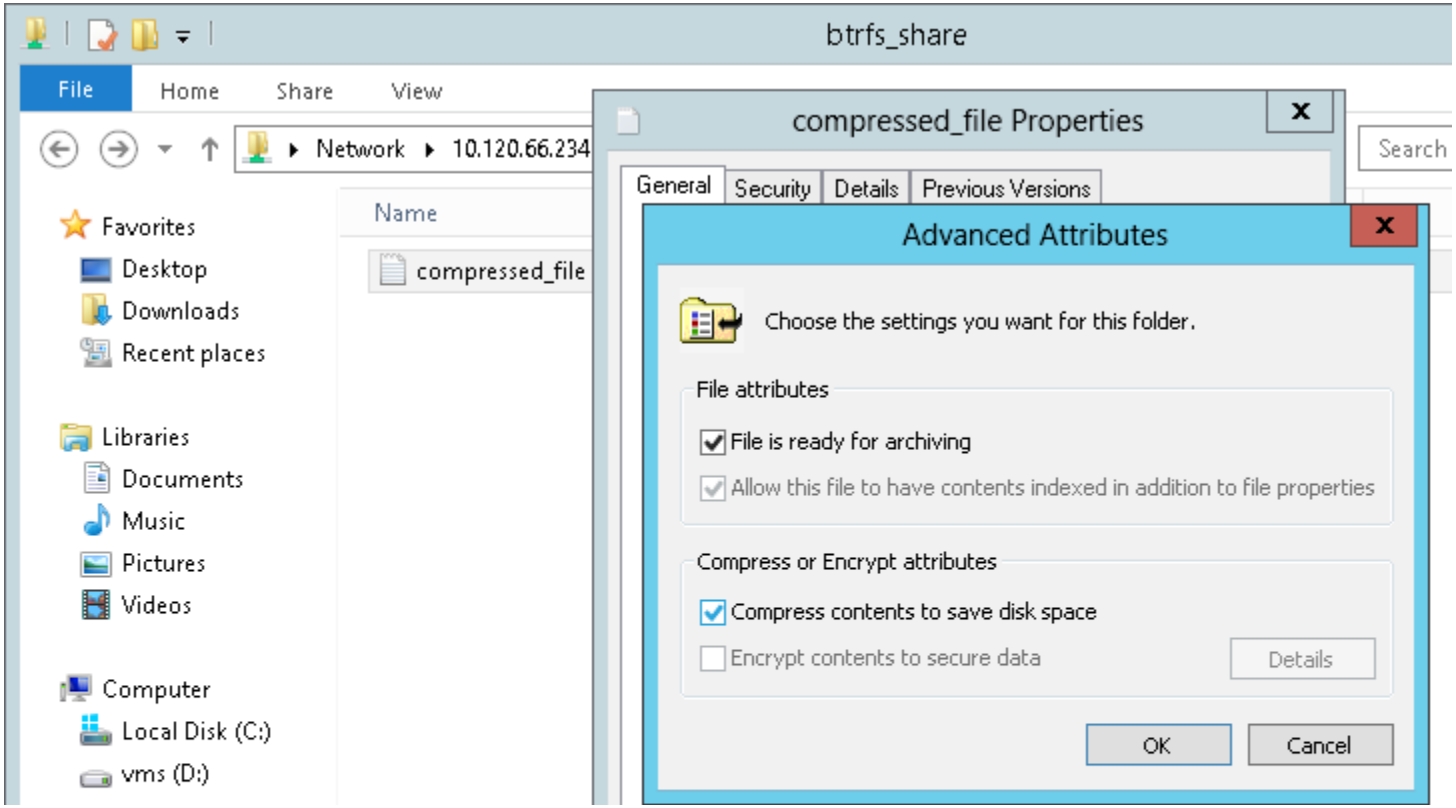

Compression
Samba + Btrfs

- FSCTL_SET_COMPRESSION sent over the wire
  - Samba's *btrfs* VFS module translates request into FS_IOC_SETFLAGS ioctl
  - Same for FSCTL_GET_COMPRESSION and FS_IOC_GETFLAGS respectively
- Btrfs compresses and uncompresses data during IO

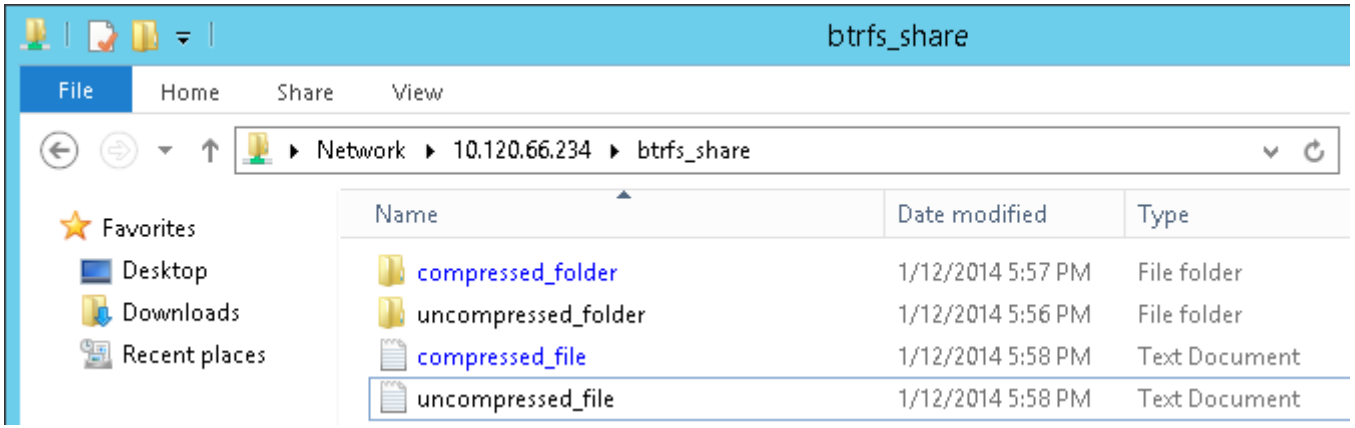

Server-Side Copy

Samba + Btrfs

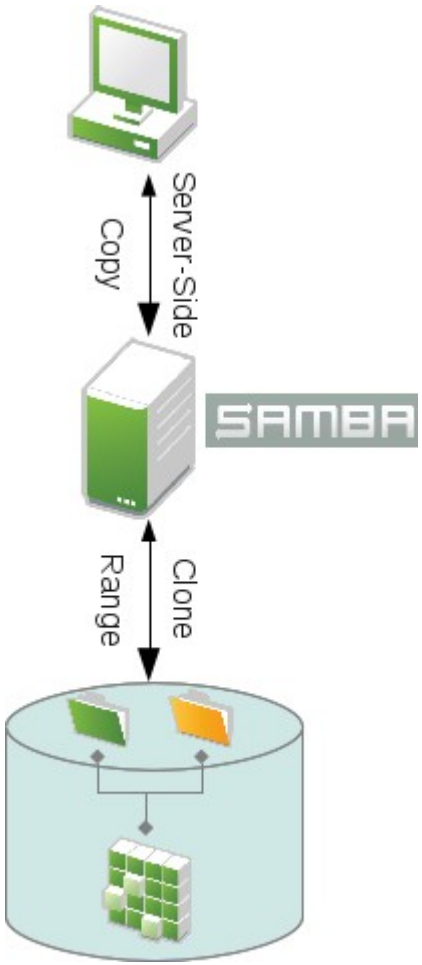

- FSCTL_SRV_COPYCHUNK sent over the wire
  - Avoids network round-trip of file data done by a traditional copy
- Samba's *btrfs* VFS module translates the request into a BTRFS_IOC_CLONE_RANGE ioctl
  - Avoids disk round trip of file data
  - No duplication of file data on disk

Server-Side Copy
Samba + Btrfs

- Dependent on client support
  - Windows Explorer in Windows 8 and Windows Server 2012
  - Robocopy in Windows 7 and Windows Server 2008
  - Linux Kernel CIFS client 3.13.0+
    - Via CIFS_IOCTL_COPYCHUNK

- Constrained by ioctl requirements
  - Btrfs block-size alignment
    - Performs local copy on fallback
    - Network round-trip still avoided
  - Copy-chunk requests only issued within a single share

Previous File Versions

Samba + Snapper + Btrfs

- Allow users to view and restore files from snapshots
  - Windows Explorer associates existing paths with matching snapshot paths as Previous Versions
    - Only displays versions with unique modification times
- Samba / Snapper configuration
  - Snapper to take periodic snapshots of share path
  - *snapper* Samba VFS module enabled for share
    - Communicates with Snapper daemon via D-Bus
  - Users permitted snapshot list and folder traversal permission

Previous File Versions

Samba + Snapper + Btrfs

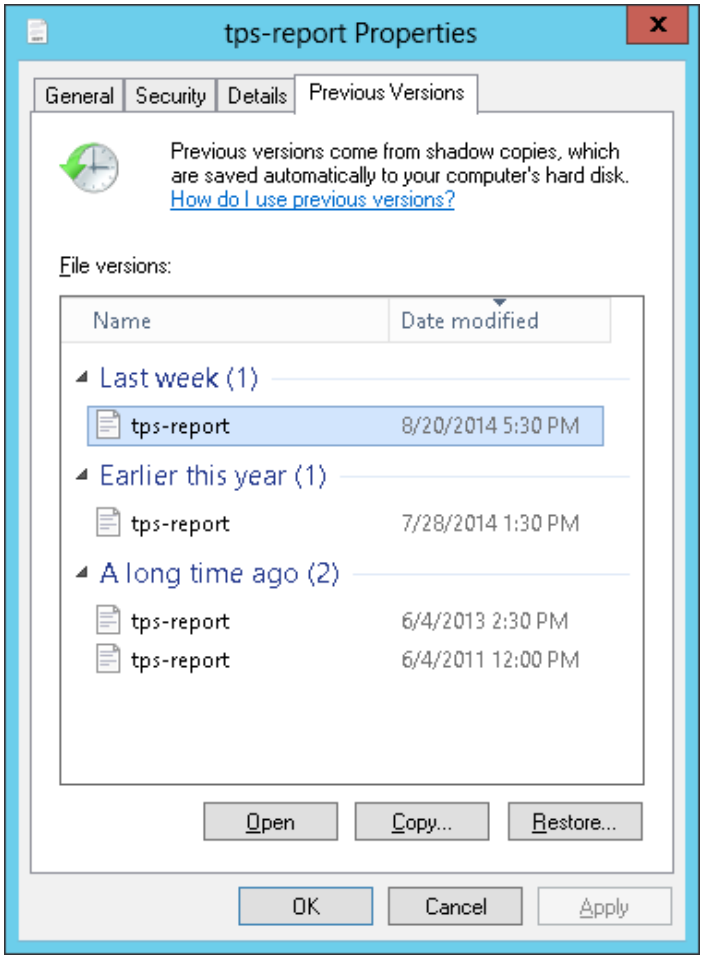

- Expose snapshots via SMB
  - Client enumeration via FSCTL_SRV_ENUM_SNAPS
- Intercept paths that correspond to snapshots
  - CIFS *@GMT* path token
  - SMB2 timewarp (*TWrp*) create tag

Remote Snapshots
Samba + Snapper + Btrfs

- Allow users to remotely manipulate share snapshots
  - Using the File Server Remote VSS Protocol (FSRVP)
    - Introduced by Molly Brown at SDC 2011
  - Create, delete and expose share snapshots

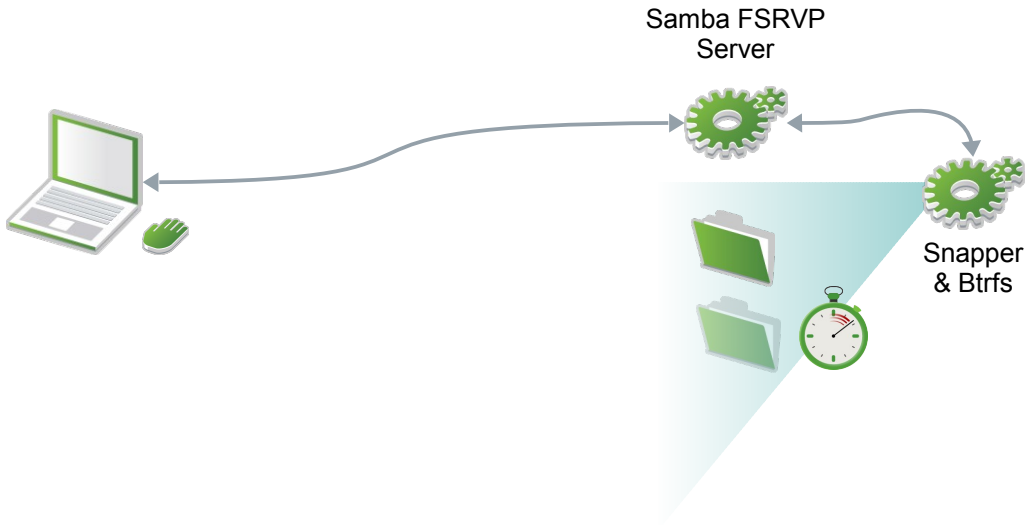

Remote Snapshots
Samba + Snapper + Btrfs

- Samba / Snapper configuration
  - Samba's *snapper* VFS module and FSRVP server enabled
  - Registry shares enabled
  - Users granted appropriate permissions
    - Snapper (ALLOW_USERS) and Samba (Administrator or backup operator)

- Clients
  - Remote backup applications
  - Samba rpcclient
  - Windows Diskshadow.exe
  - Microsoft System Center Data Protection Manager

Integration Testing

- Part of Samba's continuous integration test suite
  - Analyse underlying filesystem
    - Enable *btrfs* VFS module
    - Run server-side copy and compression with and without module
- Fake snapshots for previous file versions and FSRVP
  - Simply copy share contents into created snapshot
  - Enumerate
  - Test all FSRVP operations against Samba server

More information:
www.samba.org

Thank you.


SUSE®
We adapt. You succeed.

Title of No-Logo Slide Here (32pt)

- First-level bullet (24pt)
  - Second-level bullet (20 pt)
    - Third-level bullet (16pt)
      - Fourth-level bullet (14pt)

Unpublished Work of SUSE. All Rights Reserved.

This work is an unpublished work and contains confidential, proprietary and trade secret information of SUSE. Access to this work is restricted to SUSE employees who have a need to know to perform tasks within the scope of their assignments. No part of this work may be practiced, performed, copied, distributed, revised, modified, translated, abridged, condensed, expanded, collected, or adapted without the prior written consent of SUSE.
Any use or exploitation of this work without authorization could subject the perpetrator to criminal and civil liability.

General Disclaimer

This document is not to be construed as a promise by any participating company to develop, deliver, or market a product. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. SUSE makes no representations or warranties with respect to the contents of this document, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. The development, release, and timing of features or functionality described for SUSE products remains at the sole discretion of SUSE. Further, SUSE reserves the right to revise this document and to make changes to its content, at any time, without obligation to notify any person or entity of such revisions or changes. All SUSE marks referenced in this presentation are trademarks or registered trademarks of Novell, Inc. in the United States and other countries. All third-party trademarks are the property of their respective owners.